NBER WORKING PAPER SERIES

THE VALUE OF STATISTICAL LIFE FOR SENIORS

Jonathan D. Ketcham
Nicolai V. Kuminoff
Nirman Saha

Working Paper 33165
http://www.nber.org/papers/w33165

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2024

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, Lisa Robinson, Dan Silverman, Kerry Smith, Jonathan Skinner, David Slusky, Greg Veramendi, Jeffrey Wooldridge, and audiences at Camp Resources, WCERE, AERE, ASHEcon, WEAI, RFF Workshop on Improving Federal Environmental Regulation, Arizona State University, Carnegie Mellon University, Georgetown University, Hong Kong Polytechnic University, Michigan State University, Stanford University, University of Birmingham, University of British Columbia, University of Florida, University of Georgia, University of Hawai'i-Manoa, University of Illinois Urbana-Champaign, University of Kansas, University of Nova School of Business and Economics, and University of Southern California. This research was supported by the National Science Foundation (award #2049902) and the U.S. Environmental Protection Agency (award #84018401). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Value of Statistical Life for Seniors
Jonathan D. Ketcham, Nicolai V. Kuminoff, and Nirman Saha
NBER Working Paper No. 33165
November 2024
JEL No. Q51

ABSTRACT

We develop a new revealed preference framework to estimate the value of statistical life (VSL). Our framework starts from a hedonic model of health care in which heterogenous individuals choose how much to spend on medical services that reduce mortality risk. Their choices generate an equilibrium survival function that can be differentiated to recover their marginal willingness to pay for mortality risk reduction. Our IV estimator uses survey data on Americans over age 66, linked to their federal administrative records. The mean VSL is approximately $1 million at age 67 and increasing in health, income, education, and life expectancy.

Jonathan D. Ketcham
Earl G. and Gladys C. Davis Distinguished
Research Professor in Business
Department of Economics
W.P. Carey School of Business
Arizona State University
300 E. Lemon Street
Tempe, AZ 85287-4106
ketcham@asu.edu

Nicolai V. Kuminoff
Department of Economics
Arizona State University
P.O. Box 879801
Tempe, AZ 85287
and NBER
kuminoff@asu.edu

Nirman Saha
School of Economics
University of Surrey
Stag Hill, University Campus
Guildford GU2 7XH
Surrey
United Kingdom
n.saha@surrey.ac.uk

1 Introduction

Mortality rates are affected by government activities such as regulating air pollution, mandating job safety, and funding health care programs. Evaluating the equity and efficiency of these activities requires weighing their benefits, including mortality reductions, against their costs. A conventional approach to monetizing mortality reductions is to multiply a change in the number of premature deaths avoided by a value per statistical life (VSL) between \$10 and \$15 million (Evans and Taylor, 2020; Banzhaf, 2022; Cropper et al., 2024). The resulting benefit measures often dominate cost-benefit analyses. For example, Lee and Taylor (2019) report that survival gains represent up to 70% of all monetary benefits calculated for all federal regulations in the United States.

VSL measures reflect how much individuals are willing to pay to reduce mortality risk. They are typically derived from compensating differentials paid to workers to perform jobs with higher risks of accidental death. The workers whose choices generate this evidence are almost entirely under age 65, whereas the benefits of policies targeting premature mortality are often concentrated among people over 65. This discrepancy is particularly stark for policies targeting air pollution and climate change, where mortality effects are concentrated among the elderly (Aldy et al., 2022; Carleton et al., 2022). Discrepancies between the age group used to calculate VSL and the age group to whom VSL is applied may yield substantial mismeasurement of private benefits. Theory predicts that VSL will evolve over the life cycle with health, wealth, and remaining life expectancy (Arthur, 1981; Ferranna et al., 2023). However there is virtually no revealed preference evidence on this evolution after age 65. Our study fills this knowledge gap.

We develop evidence on how willingness to pay to reduce mortality risk evolves with health and life expectancy after age 65. Our evidence comes from data on the rates at which people choose to consume medical care relative to other goods and services. We view these choices through the lens of a hedonic model of health care that builds on ideas from Grossman (1972) and Rosen (1988). We consider individuals who differ in their preferences, age, health, income, and exposure to environmental amenities. They choose how much to spend out-of-pocket on a bundle of medical services that affects mortality risk. Their choices equate the marginal private costs and benefits of reducing mortality risk, generating an equilibrium survival function. Differentiating that function to recover an individual's marginal effect of spending on their probability of survival reveals how much that individual is willing to pay for a marginal increase in their survival probability. Individual willingness to pay can be aggregated to calculate VSL measures by age, health, and income.

We derive VSL measures for a random sample of Americans over age 66 who participated in the longitudinal Medicare Current Beneficiary Survey from 2005 to 2011. We link each person's survey responses to federal administrative data on their birth date, death date, annual medical spending, medical history, and residential location history. These linked data provide the most comprehensive and accurate information that exists on medical spending by Americans over 65. The data track out-of-pocket expenditures, as well as expenditures paid by public health insurance programs and private health insurance plans. The survey provides additional information on each person's income, education, self-assessed health, health behaviors, and limitations on activities of daily living.

We use the linked data to estimate the marginal effect of spending on the probability of survival. A key threat to identification is that spending may be correlated with latent health. This will bias VSL measures if, for example, people who are sicker in unobserved ways spend more on health care and die sooner. We mitigate this threat using a multimarket hedonic identification strategy (Heckman et al., 2010; Banzhaf, 2021). Specifically, we derive an instrumental variable for medical spending from variation in the supply of health care across geographic markets. Intuitively, identical people who face different menus of treatment options will choose to spend different amounts. We use data on people who moved between markets to derive a measure of variation in spending that reflects individual optimization against spatially-varying menus. We provide indirect evidence on this instrument's validity by showing that it is statistically independent of observed health.

Results from our main IV specification imply a mean VSL just under \$1 million at age 67 (year 2024 dollars). The mean VSL declines as a near–monotonic function of age and is estimated with sufficient precision to distinguish our estimates from the conventional range of VSL measures based on compensating wage differentials (\$10 to \$15 million). The 95% upper confidence band on our estimate for the VSL-age function is approximately \$5 million in the late 60's, \$2 million in the mid 70's, and \$0.5 million in the early 90's.

Another way to summarize our results is to annuitize VSL estimates to calculate age-specific values per statistical life year (VSLY). We find the mean VSLY is approximately \$70,000 in the late 60's, \$40,000 in the mid 70's, and \$20,000 in the early 90's. Confidence bands include values up to \$400,000 in the late 60's and \$100,000 in the early 90's. This range is consistent with values obtained by annuitizing conventional wage-based estimates for VSL over workers' lifetimes.

Conditional on age, we find that VSL measures increase in health. This is true whether we focus on self-reported health, the ability to perform daily living activities, or medically diagnosed illnesses. For example, the mean VSL at age 67 is just under \$2 million for people reporting their health as "excellent" compared to \$0.6 million for those in "good" health. We also find that VSL measures increase in education and income, conditional on age. Our evidence on how VSL measures vary with age, health, education, and income reflects heterogeneity in individuals' willingness to pay to reduce mortality risk. This information can be combined with group-specific weights to evaluate social costs and benefits of policies.[1] However we take no stance on the normative question of how to define social welfare weights.

We conduct additional analyses to test whether our primary VSL estimates are confounded by agency problems or information frictions. Specifically, we estimate

[1]For example, the U.S. Office of Management and Budget suggests that federal agencies may choose to assign higher weights to lower income groups (U.S. OMB, 2023).

specifications that allow VSL measures to vary based on whether individuals had assistance making decisions (e.g. from a spouse or adult child) and whether they understood institutional features of health care markets. Our results are consistent with the hypothesis that agents tend to steer patients toward higher spending, which inflates VSL measures. On the other hand, we find that information frictions tend to attenuate VSL measures. However both effects are small, modifying our primary estimates by less than 10% on average. Further, our main results persist across a wide range of alternative model specifications.

Our findings have three broad implications for evaluating activities that affect mortality risk. First, using VSL estimates for working-age adults to evaluate policies that reduce premature mortality among the elderly can overstate private benefits by an order of magnitude. This is especially important for policies targeting air pollution and climate change, but also relevant for a broader set of public programs and private actions (Cropper et al., 2011; Evans and Taylor, 2020; Banzhaf, 2022; Cropper et al., 2024). Second, our findings validate an alternate approach to valuing mortality risk among the elderly that combines annuitized VSL measures with remaining life expectancy (e.g. Deryugina et al., 2019; Hollingsworth and Rudik, 2021; Carleton et al., 2022). Finally, improving seniors' health and human capital can increase their willingness to pay to reduce mortality risk. This suggests that VSL measures may be endogenous to some of the policies they are used to evaluate, such as policies regulating air pollution (Aldy et al., 2022).

Our study builds on literature that uses individuals' market choices to estimate their willingness-to-pay to reduce mortality risk. Prior studies focused mainly on wage premia for risk of death on the job (e.g. Smith et al., 2004; Viscusi and Aldy, 2007; Aldy and Viscusi, 2008; Kniesner et al., 2012; Deleire et al., 2013; Lee and Taylor, 2019) and, to a lesser extent, on price premia for automobile safety features (e.g. Li, 2012; Rohlfs et al., 2015; O'Brien, 2018). The subset of these studies that examine how VSL varies with age mostly report inverse u-shape functions from age 18 to 65 (e.g. Aldy and Viscusi, 2008; O'Brien, 2018). VSL estimates from this literature have been widely used for policy evaluation, but their credibility has been questioned due to concerns about imperfect agency, information frictions, and selection into markets for labor and automobiles based on age, health and human capital (Ashenfelter, 2006; Kniesner and Viscusi, 2019; Greenberg et al., 2021; Banzhaf, 2022; Lavetti, 2023; Cropper et al., 2024).

Our study advances this literature in three ways. It is the first to analyze a random sample of people over 65 and to examine how their VSL measures vary with age, health and human capital. Second, we strengthen the literature's revealed preference mapping from choices to preferences by focusing on medical decisions where mortality risk is salient. Finally, we investigate how VSL estimates are affected by agency in decision-making and information frictions. These advances leverage information we obtain by linking survey data to federal administrative records.

2 Data and Measures

The Medicare program provides near universal health insurance for Americans over age 65. The U.S. Centers for Medicare and Medicaid Services (CMS) maintains records on each individual's birth date, residential history, medical history, and death date. It supplements these administrative data with surveys. We link longitudinal data on individuals who participated in the Medicare Current Beneficiary Survey (MCBS) from 2005 to 2011 to their administrative records through 2012.

The MCBS is a rotating panel survey that is administered to approximately 16,000 randomly chosen Medicare beneficiaries each year. It provides a nationally representative sample of people over age 65.[2] Each respondent is interviewed for up to four 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. Survey modules provide annual data on socioeconomic status, knowledge of Medicare

[2]While all Americans over 65 are entitled to Medicare benefits, eligibility for those under 65 is determined by illness or poverty. Therefore, we cannot obtain a nationally representative sample of people under 65.

programs, utilization of assistance in making medical decisions, and self-assessed health. Importantly, the linked data provide a comprehensive measure of annual medical spending starting in the second year of survey participation.

2.1 Sample construction

The linked data contain 51,191 person-years of annual spending from 2005-2011 by people who survived to the end of the calendar year. We make two sample cuts. First, we drop 730 observations in which respondents declined to answer questions about their socioeconomic status or health, or reported total spending of zero or over \$100,000.[3] Second, we drop 5,764 observations where the respondent was employed at the time of their MCBS interview. Dropping workers sharpens our focus on the tradeoff between medical spending and mortality risk by allowing us to abstract from the potential effect of medical spending on future labor income.[4]

Our main sample is comprised of 22,206 individuals and 44,697 person-years. Table 1 reports summary statistics. Individuals are observed for one, two or three years. The panel is unbalanced because some people die while enrolled in the MCBS and others' enrollment cycles extend beyond our study period. The minimum age is 67 – the youngest age at which the MCBS reports a full year of expenditures.[5] The one-year mortality rate is approximately 5%.

2.2 Medical expenditures

Americans over 65 in the Medicare program have various sources of insurance, including "traditional Medicare" run by the government, Medicare Advantage plans

[3]Dropping extreme tails of the expenditure distribution reduces the scope for outliers to affect our estimates. Labor market studies of VSL make similar cuts based on extreme tails of the wage distribution, e.g. Kniesner et al. (2012).

[4]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. Section C.1.1 shows that our results are robust to including workers in estimation.

[5]An individual who joins the MCBS between the ages of 65 and 66 is 67 in their second full year of survey participation when medical spending is first recorded.

Table 1: Summary Statistics

Measure	Summary statistic	Source
1-year mortality (%)	5	admin
Gross annual medical spending ($2024)	16,544	MCBS-admin
Out-of-pocket annual medical spending ($2024)	2,617	MCBS-admin
Health		
mean age	78	admin
female (%)	59	admin
ever smoked (%)	58	MCBS
underweight BMI (%)	4	MCBS
mean HCC health risk 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 basic activities of daily living (%)	30	MCBS
has a Medicare Advantage insurance plan (%)	25	MCBS-admin
has a Medigap insurance plan (%)	63	MCBS-admin
receives Medicaid benefits (%)	12	MCBS-admin
Socioeconomic characteristics		
White, not-Hispanic (%)	85	admin
African American (%)	8	admin
Hispanic (%)	5	admin
education: high school degree (%)	31	MCBS
education: some college (%)	22	MCBS
education: college degree (%)	21	MCBS
married (%)	52	MCBS
has living children (%)	93	MCBS
number of people	22,206	
number of person years	44,697	

Note: Spending measures are adjusted to year 2024 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.

operated by private insurers, and wraparound policies offered through employers or bought in "Medigap" markets. MCBS spending data include all of these public and private forms of insurance as well as expenditures paid entirely out-of-pocket. They are the most comprehensive data on medical spending for Americans over 65.

The MCBS reports total and out-of-pocket spending during the second, third and fourth years of survey participation. Survey staff work with respondents to record medical events in calendars and keep records and receipts. CMS then reconciles these records with Medicare administrative files. Finally, CMS reports how each respondent's annual expenditures were divided across payers.[6]

This accounting generates person-by-year measures of medical spending. Table 1 shows the average person spent \$16,544 annually, of which \$2,617 was paid out-of-pocket.[7] These out-of-pocket expenditures were equivalent to 7% of per capita income in Census data. We use m_{it} to denote an individual's total spending in year t and γ_{it} to denote their coinsurance rate; i.e. the fraction of expenditures paid out-of-pocket.

2.3 Health

The average individual is 78 years old and 59% are female. Since life expectancy varies with age and sex, these variables may proxy for health. Table 1 summarizes additional measures of health that our models control for. First, we use survey responses indicating whether people have a history of smoking (58%) or are underweight based on their body mass index (4%). Second, we use administrative data to identify if and when each person was first diagnosed with 61 chronic illnesses.[8] We use CMS's hierarchical conditions categories health risk score (i.e. HCC score) to synthesize chronic illnesses diagnoses into an overall index of morbidity.[9]

[6]Appendix A.1 provides additional background on the MCBS spending measures.

[7]These statistics are for 12 months of spending. To measure per capita expenditures consistently we exclude calendar years in which people die.

[8]The average person has 7 chronic illnesses. The data are drawn from CMS Chronic Conditions Warehouse files. Conditions include 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]CMS uses this index to make capitation payments to Medicare Advantage plans. The index predicts variation in medical spending based on age, gender, low-income status, and past medical diagnoses. We follow Finkelstein et al. (2016) in adjusting raw HCC scores for spatial and temporal trends. Table 1 reports the mean of the adjusted measure. Additional details are provided in

We also use subjective measures of health. First, 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?*". Table 1 shows the distribution of responses is slightly left-skewed with 79% of people reporting their health is good, very good, or excellent. MCBS respondents are also asked about their ability to perform various activities of daily living (ADL). Approximately 28% say they have difficulty performing at least one "instrumental" ADL that affects their ability to live independently, such as managing money, doing housework, using the phone, or preparing meals. Approximately 30% report difficulty performing at least one "basic" ADL such as bathing, dressing, eating, or walking. Finally, we control for whether people receive Medicaid benefits, which can be triggered by disability or low income, and whether they were enrolled in Medicare Advantage plans or Medigap plans. These enrollment indicators serve to control for adverse or advantageous selection into different types of insurance plans.

As MCBS respondents age over the course of the survey they tend to be diagnosed with more chronic illnesses and become more likely to experience ADL restrictions.[10] We define the vector H_{it} to include age, sex, age-by-sex, and all other subjective and objective measures of health in Table 1.

2.4 Socioeconomic Characteristics

We also control for a vector of socioeconomic characteristics, D_{it}, that may help to explain variation in life expectancy. These variables, which are reported in the MCBS, include indicators for race, educational attainment, marital status, and presence of living children. Marital status and living children help to control for the role of informal care while educational attainment may modify the return to medical spending through adherence to treatment (Goldman and Smith, 2002). Table 1 shows that about half of the individuals in our sample are married and 93% having

Appendix A.4.

[10] Their medical expenditures also increase. These trends are shown in Appendix A.2.

living children. The distribution by sex, race and education approximately matches US Census data.[11]

2.5 Environmental Exposures

Residential environments can also modify life expectancy (Deryugina and Molitor, 2020; Finkelstein et al., 2021; Bishop et al., 2024). Therefore, we track exposure to a vector of measures for environmental quality and access to health care across 306 Hospital Referral Regions (HRRs). We denote this vector by X_{it}.

HRRs provide a convenient way to describe geographic variation in residential environments. They were designed to approximate regional markets for health care and, therefore, provide a natural delineation of the U.S. into local markets within which seniors may choose among physicians and treatments.[12] At the same time, HRRs are similar in size to metropolitan areas which are often used to delineate variation in exposure to environmental amenities.

The health care variables we use include annual Dartmouth Atlas per capita measures for the numbers of acute care hospital beds, primary care physicians, and medical specialists. In addition, we use two indices of HRR-specific health care quality: CMS's "hospital compare" index of hospital quality and the discharge rate for ambulatory care-sensitive conditions that can often be prevented through better outpatient management. For environmental amenities we use annual averages of minimum daily winter temperature, maximum daily summer temperature, particulate air pollution smaller than 2.5 microns in diameter, homicide mortality, and automobile mortality. We also use the urbanization rate, the high school graduation rate, and median household income to control for unobserved features of HRRs that could be correlated with latent health through residential sorting. Summary statistics and data sources are provided in Appendix A.3.

[11] 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.

[12] The Dartmouth Atlas defines HRRs to represent regional markets for medical care. Each HRR contains at least one hospital that performs major cardiovascular procedures and neurosurgery.

3 The Demand for Mortality Risk Reduction

We develop a simple model of the demand for mortality risk reduction to fix ideas and motivate our empirical approach. Retirees enter the model with endowments of health and wealth. Each period they choose how much wealth to allocate to medical spending. Increasing medical spending increases the expected quantity of life, while reducing consumption of other goods and services. The way that individuals respond to this tradeoff reveals their marginal willingness to pay for mortality risk reduction under assumptions that are similar to, but arguably weaker than, the standard identifying assumptions in the wage-hedonic literature on job-related mortality risk.

3.1 Survival

Consider a retired individual i who lives in location j in period t. The individual's health is defined by a single index, h_{it}. This index captures how the health stock depends on idiosyncratic factors such as age, gender, past health behaviors, education and other socioeconomic characteristics, and past environmental exposures. A second index, x_{it}, captures features of the individual's current residential environment that may affect mortality risk (e.g. air pollution, extreme temperatures). The individual's health evolves as a deterministic function of their previous health, previous environmental exposure, x_{it-1}, and previous medical spending, m_{it-1}. At the beginning of each period, the individual also experiences random shocks to their health and residential environment, measured by ϵ_{it} and ε_{jt} respectively. Examples of environmental shocks include extreme weather and effects of government policies. Thus, $h_{it} = f(h_{it-1}, x_{it-1}, m_{it-1}) + \epsilon_{it}$ and $x_{jt} = x_{jt-1} + \varepsilon_{jt}$.

The probability of surviving to period $t+1$ is denoted by s. It is assumed to be a continuous and differentiable function of medical spending, the health index, and environmental exposures:

$$s_{ijt} = s(m_{ijt}, h_{it}, x_{jt}). \tag{1}$$

The marginal effect of medical spending on the probability of surviving to the next period, $\frac{\partial s_{ijt}}{\partial m_{ijt}}$, may vary with h_{it} and x_{it}. Conditional on those factors, $\frac{\partial s_{ijt}}{\partial m_{ijt}}$ may also vary across locations due to spatial variation in medical treatment options, for example, due to Roy sorting by physicians, peer effects among physicians, and institutional features of local health care markets. This supply-side variation in the survival function's shape is represented by the j subscript on m_{ijt}. It is important for our identification strategy in Section 5.2.

3.2 Budget constraint

We abstract from credit markets and require retirees to maintain non-negative assets. Equation (2) shows the intertemporal budget constraint.

$$w_{it+1} = (1 + r)w_{it} + y_i - c_{it} - \gamma_{it}m_{ijt} \;\geq\; 0 \;\forall t. \tag{2}$$

The individual's assets in period $t + 1$ are equal to assets retained from the prior period, w_{it}, which grow at interest rate r, plus fixed income y_i from pensions, social security and other sources, less expenditures on non-medical consumption, c_{it}, and health care.[13] Health care is subsidized by the government so that out-of-pocket costs are $\gamma_{it}m_{ijt}$. The coinsurance rate in period t, γ_{it}, varies across people depending on their previously determined income, health, and insurance coverage.

3.3 Preferences and optimization

In principle, an individual could modify their survival probability in Equation (1) by adjusting medical spending or by moving to a new residential location. However, empirically, virtually all retirees purchase health care and very few move.[14] Therefore, we treat locations as predetermined. However, our model does not preclude

[13]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.

[14]Retirees migrate at low rates relative to the general population and most of their moves are local. In our estimation sample, for example, we only observe 1% of individuals moving between hospital referral regions. See Mathes (2024) for a dynamic model of retirees' migration decisions.

residential sorting. Retirees may be stratified across the $j = 1, ..., J$ locations by h_{it}, γ_{it}, y_i, and w_{it} at time t due to residential sorting prior to retirement.[15]

Conditional on their prior location choice, the individual's flow utility is assumed to be a function of their non-medical consumption and health: $u(c_{it}, h_{it})$. The individual decides how much of their assets to allocate to medical and non-medical consumption each period, subject to the budget constraint. This optimization problem can be expressed as the following Bellman equation:

$$V_t(w_{it}, h_{it}, x_{jt}) = \underset{\{c_{it}, m_{ijt}\}}{max} u(c_{it}, h_{it}) + \alpha_i \ s(m_{ijt}, h_{it}, x_{jt}) \ E[V_{t+1}(\cdot)]. \tag{3}$$

Each period, the individual chooses c_{it} and m_{ijt} to maximize expected utility over the remaining lifetime, with discount factor α_i. The expectation operator is taken with respect to the following period's health stock and environmental conditions. The maximization problem is subject to the budget constraint in (2) and the survival function in (1).

Solving the optimization problem in period t, combining the first-order conditions, and rearranging terms implies that the utility-maximizing individual will choose the level of medical spending each period to equate the marginal cost and benefit of reducing mortality risk:

$$\frac{\alpha_i E[V_{t+1}(\cdot)]}{u_c(c_{it}, h_{it})} + \alpha_i \frac{s_{ijt}}{u_c(c_{it}, h_{it})} E\left[\frac{\partial V_{t+1}(\cdot) \frac{\partial h_{it+1}}{\partial m_{ijt}} \frac{\partial m_{ijt}}{\partial s_{ijt}}}{\partial h_{it+1}}\right] = \frac{\gamma_{it}}{\partial s_{ijt} / \partial m_{ijt}}. \tag{4}$$

The expression to the left of the equality is the private benefit of marginally increasing the survival probability, expressed in dollars.[16] The first term is the present discounted value of remaining life expectancy conditional on health. The second term is the present discounted value of medical spending's effect on future health.[17]

[15] Our econometric model addresses residential sorting by controlling for individual fixed effects in Section 5.2.

[16] Both terms are divided by the marginal utility of income to convert utils to dollars.

[17] $\frac{\partial h_{it+1}}{\partial m_{ijt}} \frac{\partial m_{ijt}}{\partial s_{ijt}}$ tracks how the increase in medical expenditures that is used to marginally increase the survival probability affects future health which, in turn, influences both the quality of life and

Together, the two terms decompose the private benefit of mortality risk reduction into a mortality effect and a morbidity effect. These effects are typically conflated in empirical measures for the value of statistical life (Gentry and Viscusi, 2016), which we discuss in Section 4.2.

The expression to the right of the equality in Equation (4) is the private cost of marginally increasing the survival probability. It is proportional to the cost of saving a statistical life. For example, if a \$10,000 increase in gross medical spending increases the survival probability by 0.001, then the total cost of avoiding one statistical death among type i retirees at age t is \$10 million. If $\gamma_{it} = 0.25$ then the private cost to those retirees is \$2.5 million.

4 The Value of Statistical Life

Labor market estimates for the cost of reducing on-the-job mortality risk are commonly used to approximate workers' willingness-to-pay for statistical life extension and construct VSL measures (Evans and Taylor, 2020; Lavetti, 2023). The revealed preference logic is that workers can reduce mortality risk by sorting into safer jobs that pay lower wages. The opportunity for seniors to reduce mortality risk by increasing their out-of-pocket medical spending is analogous.

Figure 1 shows how our stylized model adapts the VSL literature's equilibrium hedonic structure to the health care setting. The supply side is comprised of healthcare providers who differ in their abilities to reduce a patient's mortality risk, given the patient's age, health, and environmental exposures. The demand side is comprised of patients who differ in their wealth and preferences and consume medical care until its marginal private cost equals the marginal benefit. The collective choices made by patients and healthcare providers trace out an equilibrium survival function. The points at which patients' indifference curves are tangent to the survival function illustrate the optimality condition in equation (4).[18] Thus, differentiating

the survival probability in future periods.

[18]The hedonic model's revealed preference logic does not require taking a stance on the nature

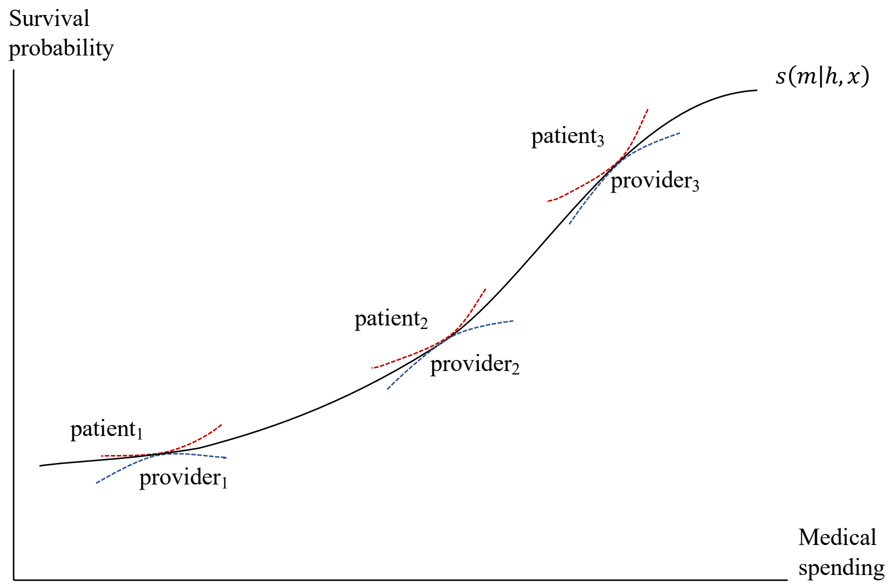

Figure 1: VSL Interpretation of the Return to Medical Spending

Note: The figure depicts patients' indirect indifference curves and healthcare providers' production functions in survival-spending space. The solid line depicts the survival probability as a function of medical spending conditional on age, health, and environmental exposures.

the survival function at a patient's point of consumption recovers their willingness to pay for a marginal reduction in mortality risk.[19] This revealed preference logic allows the private cost of saving a statistical life to be interpreted as a VSL measure.

Figure 1 depicts a single geographic market. The survival function may vary across markets for at least three reasons. First, it depends on the distribution of healthcare provider skills and beliefs, which may vary over space due to Roy sorting and peer effects. Second, it depends on the distribution of patient preferences and income, which may vary over space due to Tiebout sorting.[20] Third, it depends on

of competition in the health care sector, or the shape of its production function (Bajari and Benkard, 2005). Equation (4) simply requires seniors to make utility-maximizing decisions against the equilibrium survival function they face in the market where they live.

[19]An alternative way to estimate the willingness to pay for mortality risk reduction would be to assume parametric forms for the utility, health, and survival functions, and calibrate model parameters, similar to studies that have calibrated life-cycle models to consider the benefits of investments in health and longevity (e.g. Murphy and Topel, 2006; Hall and Jones, 2007; Hugonnier et al., 2013; Aldy and Smyth, 2014; Bauser et al., 2018; St-Amour, 2024). Another alternative is to design surveys to ask how much respondents are willing to pay for hypothetical reductions in mortality risk (Cropper et al., 2011; Evans and Taylor, 2020).

[20]The scope for Tiebout sorting during our study period is minimal: 99% of individuals stay

environmental factors that vary over space, such as pollution and climate.

4.1 Choice frictions

VSL literature generally abstracts from choice frictions such as discrete menus, incomplete information, and imperfect agency. Nevertheless, these frictions could undermine a VSL interpretation of equation (4). We address this concern in Section 7.1 by using survey modules to identify potential frictions and explore how they affect our results. To motivate that analysis we briefly highlight where our model abstracts from frictions, discuss their relevance for health care relative to the literature's canonical labor market setting, and preview our empirical findings.

First, continuity of the survival function in equation (1) and Figure 1 assumes seniors can adjust mortality risk through marginal changes to medical spending. We believe this provides a relatively good approximation to health care markets where patients are free to adjust consumption in small increments by choosing among numerous physicians and treatments. By contrast, a worker's menu of safety-wage options is constrained by discrete job opportunities.

Second, the optimization problem in equation (3) assumes that patients make choices that maximize their own utility. This is arguably a strong assumption for seniors' health care. Unlike the labor market setting, a senior's medical decisions may be made by an agent, such as a spouse or adult child, if the senior is medically incapacitated or cognitively impaired. However, we show that VSL estimates are fairly insensitive to whether medical decisions are made by the patient or an agent.

Third, the tangency between indifference curves and the survival function in equation (4) and Figure 1 implies that decision-makers know how medical spending affects short-term mortality risk. Full information is a strong assumption for any decision affecting mortality risk. Nevertheless, we expect the mapping from consumption to mortality risk to be more salient to older adults purchasing health care

in the same market. Nevertheless, the distribution of seniors' preferences and income could vary substantially across space due to Tiebout sorting at younger ages.

than to younger adults applying for jobs. For one thing, reducing mortality risk is a primary reason for purchasing health care. Further, physicians are tasked with informing patients and their agents about mortality risks of medical conditions, and the costs and benefits of treatment options. We find that VSL measures only vary slightly with patients' stated knowledge of health care markets.

4.2 Interpreting the VSL with endogenous future health

Our model also highlights the fact that VSL measures derived from market choices include effects of risk-reducing interventions on expected future health. For example, a medical treatment that reduces risk of a fatal stroke may simultaneously reduce risk of a stroke that is non-fatal but debilitating. The second term in Equation (4) formalizes how lowering the risk of future morbidity creates a co-benefit that can increase a senior's willingness to pay for statistical life extension. Such co-benefits are ubiquitous in VSL literature. The labor market analog is a safety precaution that reduces risk of both on-the-job fatalities and non-fatal injuries (Gentry and Viscusi, 2016). The automobile analog is a technology, such as airbags, that reduces risk of both fatal and non-fatal injuries in a crash. Indeed, VSL estimates derived from quasi-random changes to job safety and automobile safety embed their effects on expected future health (Kniesner et al., 2012; Rohlfs et al., 2015; Lee and Taylor, 2019). While our study is no different in this regard, the novelty of our setting makes it worth reconsidering the implications for interpreting VSL measures.

Under fairly weak assumptions, a VSL measure with endogenous future health will exceed a hypothetical health-neutral VSL. Specifically, if flow utility is strictly increasing in health, and expected future health is weakly increasing in medical spending, then a VSL measure derived from equation (4) will exceed a VSL measure that ignores the expected effect of medical spending on future health. This mechanism works against our empirical finding that the VSL for seniors is below conventional estimates for younger individuals.

An alternative hypothesis is that accounting for endogenous future health will reduce the VSL because decision-makers expect medical spending to impair future health, for example, through undesirable side effects of treatment.[21] We provide evidence against this hypothesis in Section 6.4 by showing that VSL measures decline with the onset of conditions that can be avoided by medical treatment.

5 Econometric Model

Equation (4) suggests a simple strategy to estimate VSL. First estimate the survival function. Then differentiate it with respect to individual medical spending. Finally, multiply the reciprocal by the coinsurance rate and aggregate over individuals to calculate VSL.

However, latent health poses a threat to identification. For example, seniors who are sicker in unobserved ways may spend more on health care and die sooner, biasing the estimator. We address this threat by using a two-stage control function approach that uses instrumental variables to isolate the effect of medical spending on survival (Heckman, 1979; Heckman and Robb, 1986).

5.1 Survival function

Let s^*_{ijt} be a latent variable that determines survival, defined so that person i in HRR j lives through year $t+1$ iff $s^*_{ijt} > 0$. We model survival as a linear function of medical spending, vectors of controls for health, socioeconomic status, and the residential environment, and an error, μ_{it}:

$$s^*_{ijt} = \beta + \beta_m m_{ijt} + \beta_H H_{it} + \beta_D D_{it} + \beta_X X_{jt} + \beta_C C_{it} - \mu_{it}. \tag{5}$$

In Section 2, we defined the control vectors H_{it}, D_{it}, and X_{jt}. H_{it} includes age, sex, age-by-sex, the HCC morbidity index, and indicators for whether the individual

[21] For example, prescription statins can reduce the risk of a heart attack or stroke, but increase the risk of developing diabetes among some patients.

ever smoked, is underweight, has limitations on activities of daily living, self-reported health, and insurance coverage. D_{it} includes indicators for race, educational attainment, marital status, and presence of living children. X_{jt} includes minimum daily winter temperature, maximum daily summer temperature, particulate air pollution smaller than 2.5 microns in diameter, homicide mortality, automobile mortality, urbanization rate, high school graduation rate, median household income, indices of local hospital quality, and per capita measures for the numbers of acute care hospital beds, primary care physicians and medical specialists. Finally, C_{it} is a vector of indicators for the states where individuals live. These indicators help to absorb variation in residential environments and state policies. We measure m_{ijt} and all other covariates in year t. Thus, we aim to identify how medical spending in year t affects survival through $t+1$.

5.2 Identification and estimation

Equation (5) is unlikely to yield a consistent estimator for β_m because μ_{it} and m_{ijt} are unlikely to be independent. Despite the extensive controls, μ_{it} may include aspects of latent health that are correlated with both m_{ijt} and s^*_{ijt}.[22] We mitigate this threat using a multimarket hedonic identification strategy (Heckman et al., 2010; Banzhaf, 2021). Specifically, we use ancillary data on seniors who moved between HRR markets to isolate variation in m_{ijt} that is driven by how individuals adjust their consumption of health care to between-market variation in treatment options. This strategy is analogous to way Kniesner et al. (2012) use job-to-job transitions to identify VSL measures for workers.

5.2.1 Constructing an instrument for medical spending

The institutional background for our instrument starts from the stylized fact that the supply of health care varies across the U.S. For example, Cutler et al. (2019) doc-

[22] For example, we observe if and when each person is diagnosed with various cancers, but not the stage at which they are diagnosed.

uments spatial variation in physician practice style and Chandra and Staiger (2007) highlights productivity spillovers and Roy sorting among physicians. Finkelstein et al. (2016) develops a method to decompose between-HRR variation in per/capita spending into demand side factors (e.g. health, income) and supply side factors (e.g. treatment options). We use this decomposition to construct an instrument for m_{ijt} that purges the effect of health on medical spending. Intuitively, we isolate variation in m_{ijt} that is driven by patients choosing to adjust consumption to local menus of treatment options, conditional on health.

Equation (6) shows the fixed-effects regression we use to derive the instrument.

$$m_{ijt} = \phi_j + \sigma_i + \chi_t + \psi W_{it} + o_{ijt}. \tag{6}$$

The dependent variable is annual medical spending for a person in HRR j. The variable of interest is the HRR-specific fixed effect, ϕ_j. The covariates are designed to absorb variation in m_{ijt} due to patient health so that $\hat{\phi}_j$ isolates variation in m_{ijt} driven by patients adjusting to HRR-specific menus of treatment options. Specifically, σ_i is a patient fixed effect that absorbs time-constant health, χ_t is a year fixed effect that absorbs trends in average spending, W_{it} is a vector of time-varying patient characteristics that absorbs changes in health, and o_{ijt} is an orthogonal error. We define the instrument as $z_{it} = \hat{\phi}_j - \hat{\phi}_k$ for a patient in region j in year t, where k is an arbitrary reference location.

Thus, our instrument is derived from changes in spending that coincide with moves. A natural concern is that moves may cause health shocks or be caused by health shocks, causing the instrument to be invalid. We mitigate this concern in three ways. First, we define W_{it} to include dummies for 5-year age bins and dummies for the current year relative to the move year. These dummies absorb trends in spending around moves that may reflect health shocks. Second, we exclude movers who were newly diagnosed with any new medical condition during their move year. Finally, the survival function controls for HRR environmental conditions in the covariate vector X_{jt}. Thus, we derive the instrument from differential changes in spending among individuals in the same age bin who moved between different HRR pairs during periods of stable health, while controlling for observed differences in residential environments.

A large sample of movers is required to estimate equation (6) precisely.[23] We meet this requirement by incorporating ancillary CMS administrative data. We use 3.2 million person-years of data on 484,000 seniors who were enrolled in traditional Medicare and changed their address from one HRR to another exactly once between 1999 and 2013. We extracted these data from a 10% random sample of all Medicare beneficiaries used in Bishop et al. (2023). Section 6.1 provides indirect evidence on validity of the constructed instrument by showing that it is uncorrelated with observed measures of health.

5.2.2 Control function estimation

We use the instrument defined above to create a control function for latent health from the residuals to the following regression.

$$m_{ijt} = \pi + \pi_H H_{it} + \pi_D D_{it} + \pi_X X_{jt} + \pi_C C_{it} + \pi_Z Z_{it} + \xi_{it}. \tag{7}$$

The control function is obtained by regressing m_{ijt} on the covariate vectors and Z_{it}, a vector of instruments. Z_{it} contains four elements: the instrument defined above and interactions of that instrument with the HCC morbidity index and indicators for whether the individual has restrictions on basic or instrumental activities of daily living. This set of instruments is designed to capture variation in how healthier and sicker individuals adjust their medical spending in response to geographic variation in treatment options.

We make two key identifying assumptions. The first is a relevance condition that equation (7) is correctly specified with $\pi_Z \neq 0$. The second is a validity condition

[23] This is infeasible for our main estimation sample. The number of people who move between HRRs in that sample (244) is smaller than the number of HRRs (306).

that the survival function error, μ_{it}, is independent of the instruments, Z_{it}, and the controls, H_{it}, D_{it}, X_{jt}, and C_{it}, given ξ_{it}. Under these assumptions, the residuals to equation (7) capture the effect of latent health on spending. We add the residuals as a control function to the survival equation: $CF_{it} = \hat{\xi_{it}}$.

Finally, we assume that the modified survival equation error is an *iid* draw from a Type I extreme value distribution. This implies the probability of survival, $s_{ijt} = Pr(s^*_{ijt} > 0)$, takes the complementary log-log form:

$$s_{ijt} = \exp[-\exp(\beta + \beta_m m_{ijt} + \beta_H H_{it} + \beta_D D_{it} + \beta_X X_{jt} + \beta_C C_{it} + \beta_{CF} CF_{it})]. \quad (8)$$

This form is an intuitive choice for modeling death among seniors because the function's potential asymmetry allows the probability to approach 1 (survival) slowly relative to the rate at which it approaches 0 (death).[24] It also allows the shape of the survival function to vary across time and space.

The control function estimator defined by equations (7) and (8) nests the canonical two-stage least squares estimator in the special case where the survival probability is linear in parameters. In non-linear models like ours, the estimators are not equivalent, but the intuition is similar. Our maximum likelihood estimator provides a consistent estimator for the β's as long as the controls are exogenous, the instruments are partially correlated with medical spending, and the instruments are exogenous (Wooldridge, 2015; Palmer, 2024).

In summary, estimation proceeds as follows. First we use CMS admin data to estimate equation (6) and construct the instrument. Then we use the linked MCBS-admin data to estimate equation (7) and create the control function. Finally, we estimate the survival function (8). We calculate standard errors using a bootstrap that repeats estimation of (7) and (8) after resampling one thousand times with replacement and clustering at the HRR level (Abadie et al., 2023)). Section 6.1

[24] Rescaling the dependent variable to equal one in the event of death yields the Gompit model, so named because of its similarity to the Gompertz mortality function used in Finkelstein et al. (2021) and Bishop et al. (2024). Section C.1.5 shows that our main findings are robust to estimating a Gompertz function.

provides evidence on instrument validity, and Sections 7 and Appendix C show that our results are robust to alternative specifications of (6)-(8).

6 Results

6.1 First-stage results and instrument validity

The IV captures substantial variation in individual medical spending.[25] For example, a standard deviation increase in the IV corresponds to a $1,024 increase in spending (6% of the mean).[26] In our preferred econometric specification, a $1 increase in the IV is associated with a $0.63 increase in spending for the average individual, with larger effects for those with higher HCC morbidity scores and restrictions on basic activities of daily living.[27] This is consistent with the hypothesis that sicker individuals are more likely to take advantage of opportunities for additional treatments that prolong life (Bauser et al., 2018).

While we cannot directly test instrument validity, we provide indirect evidence by examining how it modifies the partial correlation between medical spending and observed health, following Altonji et al. (2005). Figure 2 illustrates this by presenting standardized coefficients from regressing measures of observed health on medical spending and the instrument. The open circles depict coefficients from regressing each health measure on spending, absent controls.[28] Intuitively, all four morbidity measures are associated with higher spending. Closed circles show that the asso-

[25] The IV is also correlated with survey-based measures of spatial variation in treatment options from Cutler et al. (2019). Specifically, a standard deviation increase in the IV is conditionally associated with a 0.16 standard deviation increase in an HRR's share of "cowboy" physicians who consistently recommend intensive care beyond clinical guidelines and a 0.17 standard deviation decrease in the share of "comforter" physicians who consistently recommend palliative care for the seriously ill.

[26] Appendix figure B.1 shows the distribution of HRR estimates.

[27] First-stage regression coefficients from the main specification of our model in Table 2 are reported in Appendix Table B.1. The R-squared is 0.23 and the F-statistic on the excluded instruments is 11, indicating adequate statistical power.

[28] We use OLS regressions for the HCC index and indicators for restrictions on basic and instrumental activities of daily living, and an ordered probit model for the categorical measure of self-reported health status, scaled so that higher values indicate worse health.

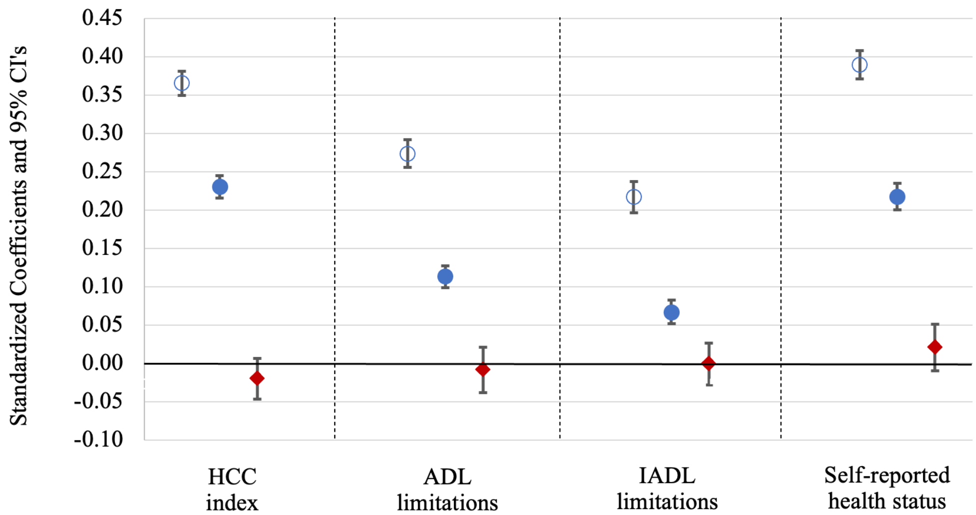

Figure 2: Selection on Observable Health

Note: The open circles depict coefficients from univariate regressions of the health measures shown on the horizontal axis on individual medical spending. Closed circles depict coefficients from regressions that add all other covariates from our main specification of the survival function in column 5 of Table 2. Closed diamonds depict coefficients from analogous regressions of each health measure on the instrument for individual medical spending. Tic marks define 95% confidence intervals. All variables are standardized. Health measures are defined in Section 2.

ciations persist in diminished form when we condition on the remaining covariates defined by H_{it}, D_{it}, X_{it}, and C_{it}. This motivates our concern that latent health is also likely to be correlated with residual variation in medical spending. The diamonds show how the IV mitigates this concern. They show that conditional associations between the IV and observed morbidity measures are close to zero. This provides indirect evidence in support of our maintained assumption that the IV is uncorrelated with latent health.

6.2 Survival functions

We find that increasing annual medical spending by $1,000 reduces the probability of death the following year by an average of 0.189 percentage points (pp) or 4% of the sample mean. The size of this effect increases with age, which causes VSL measures to decline with age. Our results suggest a mean VSL of $930,000 at age 67 and \$123,000 at age 87. To illustrate how these results depend on our identification strategy, we present the average marginal effect (AME) of medical spending on mortality and the resulting VSL measures from five specifications of the model described in Section 5.

The first column of Table 2 begins with an associative model of medical spending, health, and mortality. The next four columns address potential confounding by instrumenting for medical spending and incrementally adding the socioeconomic and environmental covariates defined by D_{it}, C_{it}, and X_{it} in equation (5). The final column presents our preferred specification.

Table 2: Average Marginal Effects of Medical Spending on Mortality

	(1)	(2)	(3)	(4)	(5)
Mortality effect of \$1,000 in spending	0.046	-0.182	-0.147	-0.202	-0.189
	(0.004)	(0.048)	(0.048)	(0.078)	(0.091)
mean VSL at age 67 (\$1,000)		916	1,156	864	930
mean VSL at age 77 (\$1,000)		320	409	303	324
mean VSL at age 87 (\$1,000)		124	155	115	123
health covariates	x	x	x	x	x
instruments for medical spending		x	x	x	x
socioeconomic covariates			x	x	x
state dummies				x	x
environmental covariates					x
F-statistic on instruments		13.9	15.3	10.6	11.0
number of observations	44,697	44,697	44,697	44,697	44,697
number of individuals	22,206	22,206	22,206	22,206	22,206

Note: The table reports average marginal effects of a \$1,000 increase in medical spending in year t on the probability of death in year $t+1$. Results are expressed as percentage point changes. All models control for age, sex, age-by-sex, HCC score, restrictions on basic and/or instrumental activities of daily living, self-reported health status, and insurance coverage. Columns (2)-(5) instrument for medical spending. Column (3) adds socioeconomic covariates indicating race, educational attainment, marital status, and presence of living children. Column (4) adds state dummies. Column (5) adds controls for the local residential environment including: CMS's hospital compare index, per capita measures of the numbers of acute care hospital beds, primary care physicians, medical specialists, and hospital admissions for ambulatory care-sensitive conditions, as well as measures of automobile mortality, homicide mortality, fine particulate matter, mean winter low temperature, mean summer high temperature, share urban, median income, and high school graduation rate. Standard errors are based on 1,000 bootstrap repetitions and clustered by hospital referral region.

Column (1) in Table 2 shows the association between medical spending and mortality conditional on observed health. Covariates include the measures of health represented by H_{it} in equation (5). The result indicates that a \$1,000 increase in spending is associated with a 0.046 pp *increase* in the one-year mortality rate. If this result were used to calculate the VSL, it would be negative.

Column (2) adds the medical spending control function. This causes the estimated AME of medical spending to change sign. This change is consistent with the result in Column (1) being driven by positive correlation between residual variation in medical spending and latent morbidity. Intuitively, people who are unobservably sicker tend to spend more on health care and die sooner. Instrumenting for medical spending addresses this confounding. The result in Column (2) indicates that a $1,000 increase in spending causes the one-year mortality rate to decline by an average of 0.182 pp.

Columns (3), (4), and (5) extend the IV specification in (2) by incrementally adding socioeconomic covariates, indicators for the states where individuals reside as a coarse control for environmental quality and, finally, covariates describing localized measures of environmental quality in the HRRs where individuals reside. The AME point estimates for these four specifications are all within a standard error of each other. Further, the corresponding age-specific VSL measures all start close to $1 million at age 67 and decline close to $0.1 million by age 87. The stability in these results as we incrementally add covariates provides some additional indirect support for our maintained assumption that the residual identifying variation in the instrument is uncorrelated with latent features of health and environmental exposures that affect mortality (Altonji et al., 2005).

The AME in Column (5) indicates that increasing annual medical spending by $1,000 reduces the one-year probability of death by an average of 0.189 percentage points.[29] This average reflects considerable heterogeneity. In particular, the AME increases with age and morbidity.[30] For example, among people who describe their health as "good" the AME increases in absolute magnitude from -0.07 at age 67 to

[29] Appendix Figure B.2 shows that this specification closely approximates age-by-sex mortality rates in the data, and Appendix Table B.2 reports AMEs for health, socioeconomic, and environmental covariates. The health covariates are strongly predictive of mortality.

[30] Appendix Section B.4 summarizes heterogeneity in AMEs with age and observed health.

-0.15 at age 77 to -0.33 at age 87. And among people aged 77, the AME is -0.06 for those in “excellent” health compared to -0.52 for those in “poor” health.

6.3 VSL measures by age

Table 2 reports the mean VSL at ages 67, 77, and 87 for each IV specification of the survival function. We calculate these measures in three steps. First, we divide each individual's coinsurance rate by the estimated marginal effect of spending on their survival probability. This recovers the marginal private benefit of spending, as shown in Equation (4). Then we take age-specific means of these marginal values. Finally, we scale the age-specific means to measure the aggregate willingness-to-pay to avoid one statistical death among seniors at a given age. This calculation produces VSL measures that decline with age because the marginal effect of spending on the survival probability increases with age whereas the coinsurance rate is relatively invariant to age.[31] Intuitively, the decision not to spend more on one's health when the marginal effect on survival is higher reveals that the VSL is lower.

Column (5) of Table 2 shows that our preferred specification of the survival function yields a mean VSL of \$930,000 at age 67. The mean VSL declines to \$324,000 at age 77 and \$123,000 at age 87. Figure 3a plots the mean VSL at every age from 67 to 97, along with shaded 95% confidence bands that account for statistical imprecision in our estimates for the return to spending.[32] The bands are asymmetric around the point estimates because the VSL is a nonlinear function of the return to spending.[33]

Figure 3a highlights three important features of the results. First, the willingness to pay for statistical life extension by seniors in their late 60's implies a mean VSL just under \$1 million. Second, the mean VSL declines with age in a near-

[31] Medical spending increases in age, as shown in Appendix Figure A.1. However, it does not increase enough to reduce the marginal return to further spending. Appendix Figure B.5 illustrates this by plotting the mean coinsurance rate and return to spending by age.

[32] Appendix Table B.3 reports the values underlying Figure 3.

[33] Appendix Figure B.3 plots the mean return to spending by age and 95% confidence bands.

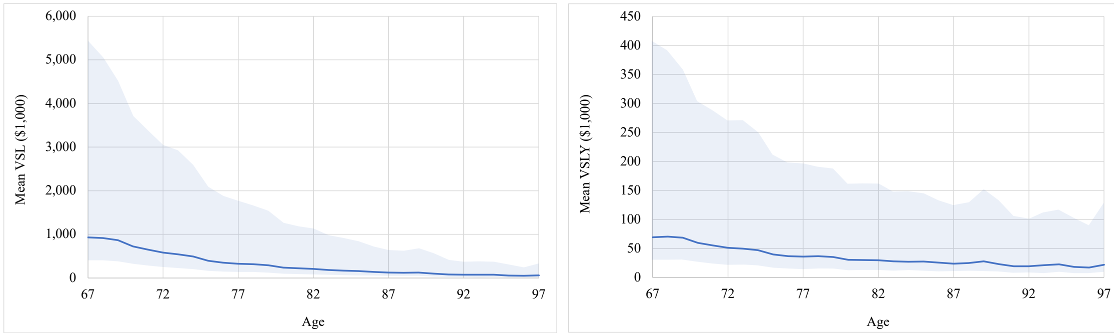

(a) value of a statistical life, by age (b) value of a statistical life year, by age

Figure 3: VSL Measures from Age 67 to Age 97

Note: The figure shows point estimates for the VSL and VSLY at each age. Values are reported in year 2024 dollars. The shaded areas are 95% confidence bands derived from 1,000 bootstrap repetitions with errors clustered by hospital referral region.

monotonic fashion. Third, the VSL-age profile is estimated with enough statistical precision to clearly differentiate it from the 10 to 15 million dollar estimates derived from compensating wage differentials (Kniesner et al., 2012; Lee and Taylor, 2019; Banzhaf, 2022; Lavetti, 2023; Evans and Taylor, 2020). While the upper confidence band includes values up to $5 million for seniors in their late 60's, it falls below $2 million after age 75.

Another way to summarize the results is to annuitize VSL estimates to calculate age-specific values per statistical life year (VSLY). Figure 3b converts our results to VSLY measures using remaining life expectancy from the US life tables and a 3.1% discount rate recommended for cost-benefit analyses by the US Office of Management and Budget (U.S. CDC, 2014; U.S. OMB, 2023). The resulting measures start just below $70,000 at age 67 and decline to $24,000 by age 87. This decline is consistent with evidence that the VSLY for workers aged 18 to 62 is an inverse u-shaped function of age (Aldy and Viscusi, 2008). At the same time, confidence bands include a mean VSLY between $100,000 and $200,000. This is consistent with values used in some studies to measure benefits of reducing exposure to air pollution and extreme temperatures among the elderly (e.g. Deryugina et al., 2019; Hollingsworth and Rudik, 2021; Carleton et al., 2022).

6.4 VSL measures by health, education, and income

Conditional on age, we find that VSL measures vary with many of the individual characteristics in Table 1.[34] Figure 4 provides four examples. It shows age-specific means for VSL, stratified by measures of health, education, and income.

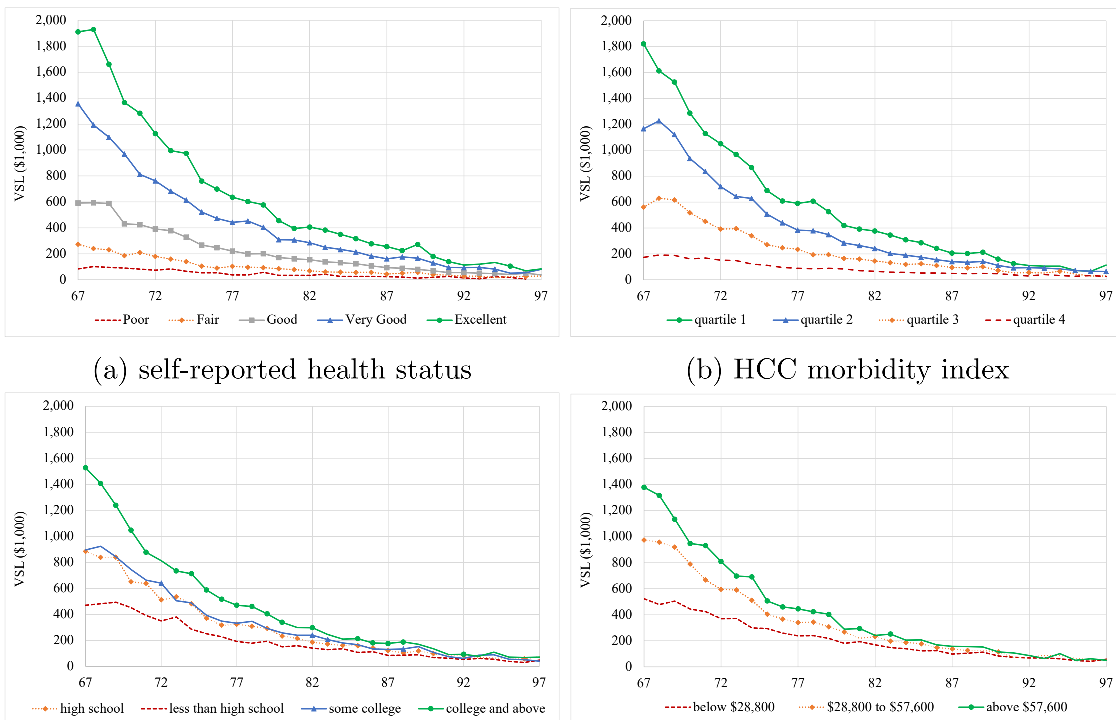

(a) self-reported health status

(b) HCC morbidity index

(c) educational attainment

(d) income

Figure 4: VSL Measures by Age, Health, Education, and Income

Note: Each panel shows the mean age-specific VSL in $1,000 (2024) dollars stratified by demographics. Markers along each line denote ages at which the VSL measure exceeds the VSL measure for the lower adjacent line in at least 95% of 1,000 bootstrap repetitions of the model in col (5) of Table 2.

Figure 4a shows that the VSL-age profile increases with subjective health. The VSL for 67-year-olds who state they are in excellent health for their age is just under $2 million and approximately 20 times larger than the VSL for those in poor

[34] Appendix Figure B.6 illustrates this variation by showing the distribution of VSL values at age 70 across different combinations of covariates. Approximately 75% of individuals have covariate vectors that correspond to type-specific VSL measures below $1 million, 16% have values between $1 and $2 million, and 9% have values over $2 million.

health.[35] The stratification between health categories declines in age as relatively healthier groups experience sharper declines. Figure 4b shows a similar pattern when VSL measures are instead stratified by objective health using quartiles of the HCC morbidity index. These stratification patterns reflect how our estimates for the AME of medical spending increase with morbidity.[36]

Figure 4c shows that the VSL-age profile also increases with educational attainment. This is unsurprising in light of Figures 4a and 4b because educational attainment is positively associated with health conditional on age (Goldman and Smith, 2002). Nonetheless, the magnitudes are striking. Between the ages of 67 and 87, the mean VSL among people with a college degree is, on average, approximately 50% larger than for those with a high school degree. In contrast, there is virtually no difference between people who finished high school and did not attend college and people who attended some college but did not obtain a degree.

Finally, Figure 4d shows that the VSL-age profile increases in income. This is not driven by an estimated effect of income on survival. Our models exclude income because we expect its effect on survival to be fully mediated by health and medical expenditures. Nevertheless, we can inform the associative relationship between VSL and income by stratifying the VSL-age profile by three income bins derived from the MCBS. Focusing on the minimum difference between the top and bottom bins defines upper bounds on the income elasticity of 1.64 at age 67, 0.87 at age 77, and 0.60 at age 87.[37] An upper bound below one is notable because compensating wage differentials imply an elasticity at or above one (Cropper et al., 2011; Banzhaf, 2022). Our findings are consistent with predictions from a calibrated life-cycle model in Aldy and Smyth (2014) that the VSL income elasticity will decline late in life due to declining differences in remaining life expectancy.

[35] People in poor health 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%).

[36] Appendix Figure B.4 shows that the age-specific return to spending increases in subjective and objective measures of health.

[37] For example, if we assume that the difference in income between people in the "below $28,800" and "above $57,600" bins is approximately $28,800, then doubling income at age 67 is associated with multiplying VSL by 2.64, yielding an upper bound on the elasticity of 1.64.

Appendix Section B.5 shows that VSL-age profiles vary similarity with other characteristics that proxy for health and remaining life expectancy. For example, VSL measures tend to be lower for ever-smokers compared to never-smokers, men compared to women, people diagnosed with more chronic medical conditions, and people with restrictions on basic and/or instrumental activities of daily living.

7 Validation Tests and Sensitivity Analysis

7.1 Main validation tests

Table 3 presents three validation tests of our main specification. First, we test whether our VSL estimates are confounded by heterogeneity in agency. For example, our estimates could be inflated if physicians steer patients toward costly treatments that do little to reduce mortality risk. Alternatively, our estimates could be attenuated if family caregivers steer patients away from such treatments. We test these hypotheses by adding interactions between medical spending and indicators derived from an MCBS question that asks whether an individual usually makes health insurance decisions on their own (67.6%), receives help from someone else (27.6%), or relies on others to make decisions for them (4.8%).[38] Column (2) reports the resulting VSL measures only for individuals who make their own decisions. The VSL-age profile is slightly flatter than our main specification (repeated in Column (1) for convenience). This is consistent with the hypothesis that agents steer patients toward more expensive treatments. However, the differences in age-specific means are less than 11%. Thus, our mean VSL estimates are fairly insensitive to whether medical decisions are made by the patient or their agent.

We further evaluate the revealed preference logic for our VSL calculations by interacting medical spending with an indicator for whether we have reason to suspect

[38] In cases of Alzheimer's disease or other impairments, the proxy who makes health insurance decisions also responds to the MCBS. Proxy decision-makers are almost always family members. The sample size declines because not all individuals respond to the question.

Table 3: Validation Tests

	(1)	(2)	(3)	(4)
Mortality effect of $1,000 in spending	-0.189 (0.091)	-0.204 (0.084)	-0.200 (0.093)	-0.191 (0.090)
mean VSL at age 67 ($1,000)	930	832	955	1,021
mean VSL at age 77 ($1,000)	324	298	342	316
mean VSL at age 87 ($1,000)	123	126	138	115
modification to main specification				
interact spending with agency indicator		x		
interact spending with information friction indicator			x	
add workers to estimation sample				x
F-statistic on instruments	11.0	10.3	11.0	10.4
number of observations	44,697	39,271	44,697	50,336
number of individuals	22,206	19,581	22,206	24,255

Note: The first column repeats our main results from Table 2. The next three columns report results from alternative specifications that are designed to test identifying assumptions that underlie our main specification. Standard errors are clustered by hospital referral region. See the note to Table 2 and main text for further details.

that an individual's health care decisions could be impaired by information frictions. We derive this indicator from an MCBS module that evaluates respondents' knowledge of health insurance programs. The indicator is based on whether one or more of the following statements about the individual is true: (1) does not make their own health insurance decisions, (2) has assistance managing money, (3) does not realize that out-of-pocket costs vary across Medicare 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.[39] These criteria lead us to classify 82% of all observations as potentially being affected by information frictions. This classification does not mean that revealed preference logic necessarily fails for these observations, only that we have reason to suspect that it might.

Column (3) of Table 3 reports VSL measures only for the 18% of individuals for whom there is no evidence of information frictions. The measures are 3% to

[39] The prescription drug 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?"

12% larger than in our main specification. This is consistent with the hypothesis that less informed individuals have higher marginal returns to spending, e.g. due to under investment in effective treatments. However, it can also be explained by the fact that the individuals making more informed choices are healthier (e.g. 36% of a standard deviation reduction in the HCC score). In any case, the relatively small effect on VSL measures suggests that heterogeneity in information frictions is unlikely to substantially attenuate our main estimates for the VSL.

Finally, we test whether the external validity of our results is likely to be compromised by excluding workers from our main estimation sample. Excluding workers could cause us to understate the mean VSL due to selection into labor participation by younger, healthier individuals. On the other hand, including workers could cause our estimator to overstate the mean VSL if workers spend more on health care in order to increase lifetime income by avoiding health shocks that would accelerate retirement. Column (4) shows that adding workers to the estimation sample has the largest effect on mean VSL for individuals in their late 60's. This is unsurprising because one third of all workers are aged 67 to 69. However, even for this age group, the mean VSL increases by less than 10%. Thus, our main results are also robust to including workers.

7.2 Additional sensitivity analysis

Our broad conclusions about the VSL-age profile persist when we modify our main specification to use: (1) alternative instruments for medical spending designed to allow for different forms of selection on unobserved health; (2) specifications that add interactions between medical spending and health, education, or age; (3) specifications that exclude interactions between the instrument and observed health; and (4) a Gompertz mortality function instead of the complementary log-log specification. We estimate survival functions and calculate VSL measures for every combination of these modifications to each specification in Table 3. We present results from this sensitivity analysis in Appendix C. In summary, our preferred specification is near the middle of the range of estimates. Moreover, across all specifications the mean VSL is below \$2.6 million in the late 60's and below \$1 million in the late 70s.

8 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 the resulting VSL measures by age, health, education, income, agency, and knowledge of market institutions. Our main results suggest a mean VSL close to \$1 million in the late 60's, and lower values for older ages. The order of magnitude difference from conventional \$10 to \$15 million VSL figures based on compensating wage differentials is striking. One explanation is that the youngest individuals we study are more than 25 years older than the average worker, and in worse health. This is consistent with our finding that VSL measures increase with health and remaining life expectancy. Annuitizing our VSL estimates implies values per statistical life year that are more consistent with values obtained by annuitizing conventional wage hedonic estimates over workers' lifetimes.

Our finding that VSL measures increase in health, income and remaining life expectancy suggests that VSL measures may be endogenous to the some of the policies they are used to evaluate, such as regulations on air pollution (Aldy et al., 2022). Multiplying a VSL measure by the number of premature deaths avoided by a policy will bias benefit measures toward zero for policies that simultaneously reduce morbidity. Such policies may trigger a virtuous cycle in which premature deaths are averted directly, but health is also improved, the VSL 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 further research.

References

Abadie, A., Athey, S., Imbens, G. W., and Wooldridge, J. (2023). When should you adjust standard errors for clustering? *Quarterly Journal of Economics*, 138:1–35.

Aldy, J. E., Aufhammer, M., Cropper, M., Fraas, A., and Morgenstern, R. (2022). Looking back at 50 years of the Clean Air Act. *Journal of Economic Literature*, 60(1):179–232.

Aldy, J. E. and Smyth, S. J. (2014). Heterogeneity in the value of life. *NBER Working Paper 20206*.

Aldy, J. E. and Viscusi, W. K. (2008). Adjusting the value of a statistical life for age and cohort effects. *The Review of Economics and Statistics*, 90(3):573–581.

Altonji, J. G., Elder, T. E., and Taber, C. R. (2005). Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools. *Journal of political economy*, 113(1):151–184.

Arthur, W. (1981). The economics of risks to life. *American Economic Review*, 71(1):54–64.

Ashenfelter, O. (2006). Measuring the value of a statistical life: problems and prospects. *Economic Journal*, 116(510):C10–C23.

Bajari, P. and Benkard, C. L. (2005). Demand estimation with heterogeneous consumers and unobserved product characteristics: a hedonic approach. *Journal of Political Economy*, 113(6):1239–1276.

Banzhaf, H. S. (2021). Difference-in-difference hedonics. *Journal of Political Economy*, 129(8):2385–2414.

Banzhaf, H. S. (2022). The value of statistical life: A meta-analysis of meta-analyses. *Journal of Benefit-Cost Analysis*, 13(2):182–197.

Bauser, D., Lakdawalla, D., and Reif, J. (2018). Mortality risk, insurance, and the value of life. *NBER Working Paper*, 25055.

Bishop, K. C., Ketcham, J. D., and Kuminoff, N. V. (2023). Hazed and confused: The effect of air pollution on dementia. *Review of Economic Studies*, 90(5):2188–2214.

Bishop, K. C., Kuminoff, N. V., Mathes, S. M., and Murphy, A. D. (2024). The marginal cost of mortality reduction: evidence from housing markets. *Forthcoming in Journal of Urban Economics*.

Carleton, T., Jina, A., Delgado, M., Greenstone, M., Houser, T., and et al. (2022). Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits. *Quarterly Journal of Economics*, 137(4):2037–2105.

Chandra, A. and Staiger, D. O. (2007). Productivity spillovers in health care: evidence from the treatment of heart attacks. *Journal of Political Economy*, 115(1):103–140.

Cropper, M., Hammitt, J. K., and Robinson, L. A. (2011). Valuing mortality risk reductions: progress and challenges. *Annual Review of Resource Economics*, 3(1):313–336.

Cropper, M., Joiner, E., and Krupnick, A. (2024). Revisiting EPA's value per statistical life. *Review of Environmental Economics and Policy*, 18(2):193–211.

Cutler, D., Skinner, J. S., Stern, A. D., and Wennberg, D. (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, T., Kahn, S., and Timmins, C. (2013). Roy model sorting and nonrandom selection in the valuation of a statistical life. *International Economic Review*, 54(1):279–306.

Deryugina, T., Heutel, G., Miller, N. H., Molitor, D., and Reif, J. (2019). The mortality and medical costs of air pollution: Evidence from changes in wind direction. *American Economic Review*, 109(12):4178–4219.

Deryugina, T. and Molitor, D. (2020). Does when you die depend on where you live? *American Economic Review*, 110(11):3602–3633.

Evans, M. F. and Taylor, L. O. (2020). Using revealed preference methods to estimate the value of reduced mortality risk: Best practice recommendations for the hedonic wage model. *Review of Environmental Economics and Policy*, 14(2):282–301.

Ferranna, M., Hammitt, J. K., and Adler, M. D. (2023). Age and the value of life. In *The Routledge Handbook on the Economics of Ageing*, pages 566–577. Routledge.

Finkelstein, A., Gentzkow, M., and Williams, H. (2016). Sources of geographic variation in health care: Evidence from patient migration. *The Quarterly Journal of Economics*, 131(4):1681–1726.

Finkelstein, A., Gentzkow, M., and Williams, H. (2021). Place-based drivers of mortality: Evidence from migration. *American Economic Review*, 111(8):2697–2735.

Gentry, E. P. and Viscusi, W. K. (2016). The fatality and morbidity components of the value of statistical life. *Journal of Health Economics*, 46:90–99.

Goldman, D. and Smith, J. P. (2002). Can patient self-management help explain the ses health gradient? *Proceedings of the National Academy of Sciences*, 99:10929–10934.

Greenberg, K., Greenstone, M., Ryan, S. P., and Yankovvich, M. (2021). The heterogeneous value of a statistical life: Evidence from US Army reenlistment decisions. *NBER Working Paper 29104*.

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

Hall, R. E. and Jones, C. I. (2007). The value of life and the rise in health spending. *The Quarterly Journal of Economics*, 122(1):39–72.

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica*, 47(1):153–161.

Heckman, J. J., Matzkin, R. L., and Nesheim, L. (2010). Nonparametric identification and estimation of nonadditive hedonic models. *Econometrica*, 78:1569–1591.

Heckman, J. J. and Robb, R. (1986). Alternative methods for solving the problem of selection bias in evaluating the impact of treatments on outcomes. In *Drawing inferences from self-selected samples*, pages 63–107. Springer-Verlag.

Hollingsworth, A. and Rudik, I. (2021). The effect of leaded gasoline on elderly mortality: evidence from regulatory exemptions. *American Economic Journal: Economic Policy*, 13(3):345–373.

Hugonnier, J., Pelgrin, F., and St-Amour, P. (2013). Health and (other) asset holdings. *Review of Economic Studies*, 80(2):663–710.

Kniesner, T. J. and Viscusi, W. K. (2019). The value of a statistical life. *Oxford Research Encyclopedia of Economics and Finance.*

Kniesner, T. J., Viscusi, W. K., Woock, C., and Ziliak, J. P. (2012). The value of a statistical life: Evidence from panel data. *Review of Economics and Statistics*, 94(1):74–87.

Lavetti, K. (2023). Compensating wage differentials in labor markets: Empirical challenges and applications. *Journal of Economic Perspectives*, 37(3):189–212.

Lee, J. M. and Taylor, L. O. (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.

Li, S. (2012). Traffic safety and vehicle choice: quantifying the effects of the 'arms race' on American roads. *Journal of Applied Econometrics*, 27:34–62.

Mathes, S. (2024). The dynamics of residential sorting and health: Implications of climate change in the U.S. *Forthcoming in International Economic Review.*

Murphy, K. M. and Topel, R. H. (2006). The value of health and longevity. *Journal of Political Economy*, 114(5):871–904.

O'Brien, J. (2018). Age, autos, and the value of a statistical life. *Journal of Risk and Uncertainty*, 57:51–79.

Palmer, C. J. (2024). An IV hazard model of loan default with an application to subprime mortgage cohorts. *Forthcoming in Journal of Finance.*

Rohlfs, C., Sullivan, R., and Kniesner, T. (2015). New estimates of the value of a statistical life using air bag regulations as a quasi-experiment. *American Economic Journal: Economic Policy*, 7(1):331–359.

Rosen, S. (1988). The value of changes in life expectancy. *Journal of Risk and Uncertainty*, 1:285–304.

Smith, V. K., Evans, M. F., Kim, H., and Taylor Jr, D. H. (2004). Do the near-elderly value mortality risks differently? *Review of Economics and Statistics*, 86(1):423–429.

St-Amour, P. (2024). Valuing life over the life cycle. *Journal of Health Economics*, 93:102842.

U.S. CDC (2014). U.S. Centers for Disease Control and Prevention, National vital statistics report, volume 63, number 7.

U.S. OMB (2023). U.S. Office of Management and Budget, Circular no. A-94: Guidelines and discount rates for benefit-cost analysis of federal programs.

Viscusi, W. K. and Aldy, J. E. (2007). Labor market estimates of the senior discount for the value of statistical life. *Journal of Environmental Economics and Management*, 53(3):377–392.

Wooldridge, J. M. (2015). Control function methods in applied econometrics. *Journal of Human Resources*, 50(2):420–445.

9 Supplemental Appendices

Appendix A Data

A.1 MCBS Spending Measures

The Medicare Current Beneficiary Survey (MCBS) Cost and Use files provide comprehensive measures of each respondent's total and OOP spending during their second, third and fourth years of survey participation. These are the best available measures of medical spending for the Medicare population. They account for all payments by Medicare as well as payments by Medicaid, Medigap, employer-sponsored plans, and other third-party payers. 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 plans, and spending in Medigap plans. Further, the Cost and Use files include expenditures paid entirely out-of-pocket (OOP) with no claim submitted to an insurer, such as purchases of generic drugs. Equally important is the fact that the reconciled spending measures provide a detailed accounting of how expenditures were divided across payers 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.

CMS's 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."

A.2 Evolution of Health and Medical Spending

Figure A.1 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 for 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 A.1 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

Figure A.1: Evolution of Health and Medical Spending Over MCBS Years 2 to 4

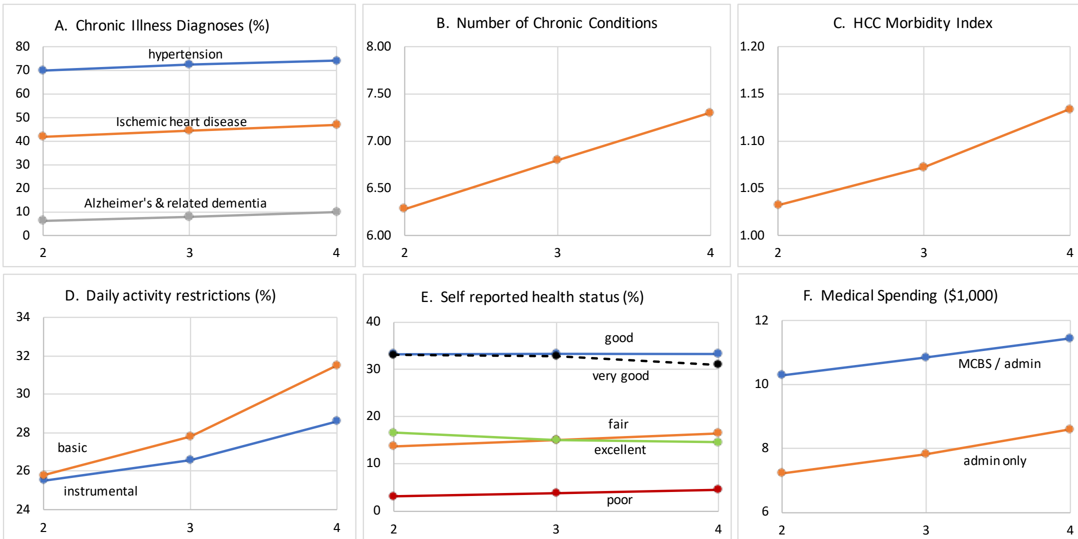

Note: The figure summarizes the evolution of health and medical spending during years two through four of the Medicare Current Beneficiary Survey, the period for which we observe comprehensive spending measures. The figure is constructed from data on the subset of respondents whom we observe in all three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

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.

A.3 Environmental Exposures

We collected publicly available data on several measures of individuals' local environments that may affect their mortality risk. These include measures of access to health care, temperature extremes, air pollution, homicide mortality, and automobile mortality. In addition, we use measures for the local urbanization rate, median household income, and the high school graduation rate as proxy variables for unobserved features of the local environment that may lead to residential sorting in general, and sorting by income and education specifically.

After collecting data on each variable, mostly at the county level, we aggregated the data within each Hospital Referral Region (HRR) using geographic crosswalks available from the US Census Bureau and the Dartmouth Atlas of Health Care.[40] Specifically, we used Census Bureau crosswalks to map county data onto zip code tabulation areas, and then to map those areas onto zip codes. Finally, we used the Dartmouth Atlas crosswalk from zip codes to HRRs to aggregate to the HRR level.

Table A.1 reports summary statistics on the local environmental variables. We define each variable below.

A.3.1 Homicide mortality and automobile mortality

Annual county-level data on the homicide mortality rate and the automobile mortality rate were downloaded from the US Centers for Disease Control (CDC) database on underlying causes of death: https://wonder.cdc.gov/. The mortality rates are adjusted for age. Homicides are defined by ICD-10 codes X85-Y09 and Y87.1. Automobile deaths are defined by the following ICD-10 codes: V02 - V04, V09.0, V09.2, V12 - V14, V19.0 - V19.2, V19.4 - V19.6, V20 - V79, V80.3 - V80.5, V81.0 - V81.1, V82.0 - V82.1, V83 - V86, V87.0 - V87.8, V88.0 - V88.8, V89.0, V89.2.

The CDC suppresses age-adjusted mortality rates when the underlying mortality counts in a county-year cell are below 10, and it codes cells with between 10 and 20 deaths as "unreliable". These occurrences are most common in rural counties. We took a series of incremental steps to impute data in missing and unreliable cells. First, if possible, we replaced the age-adjusted mortality rate with a raw mortality rate (i.e. not adjusted for age) for the same county-year. If the data were still

[40] "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`.

Table A.1: Residential Environmental Variables

Meaures of Residential Environment	Mean	Standard Deviation
hospital compare index	0.79	0.04
acute care hospital beds per 1,000 residents	2.38	0.48
primary care physicians per 100,000 residents	71.48	11.22
total specialists per 100,000 residents	125.86	20.69
discharges for ambulatory care-sensitive conditions per 1,000 Medicare enrolles	65.96	14.53
auto mortality rate per 1,000 residents	14.55	5.59
homicide rate per 1,000 residents	5.99	3.01
annual average PM2.5 concentration mg/m^3	11.84	2.02
annual average daily maximum summer temperature	85.25	7.73
annual average daily minimum winter temperature	33.29	10.99
population share living in urban areas	0.78	0.16
median household income (thousand $2010)	52,812	9,606
high school graduation rate	0.86	0.04

Note: The table reports means and standard deviations for each measure of the residential environment. See the text for variable definitions.

missing or unreliable, then we substituted an age-adjusted mortality rate for the entire rural part of the state or, if necessary, a raw mortality rate.[41] Finally, in cases where the data were still missing, we replaced them with zeros. This seems likely to provide a reasonable approximation because if there are fewer than 10 deaths across all of the rural counties in a state then the number of deaths in any one county is likely to be near zero.

[41] We used CDC's rural and urban classifications for the year 2013.

A.3.2 Fine particulate matter

Annual county-level data on average outdoor concentrations of particulates smaller than 2.5 microns in diameter ($PM_{2.5}$) were downloaded from https://wonder.cdc.gov/. The data were originally developed by NASA using remote sensing and reported on a 10 - kilometer grid.

A.3.3 Average temperatures

Annual county-level data on average daily summer maximum temperature (June to August) and average daily minimum winter temperature (January to March) were downloaded from https://wonder.cdc.gov/. The data were originally developed by the North America Land Data Assimilation System.

A.3.4 Share urban

The share of the urban population in each county was downloaded from the US Census Bureau for the year 2010.

A.3.5 Median household income

The median household income by county-year was downloaded from the US Census Bureau and converted to constant year 2010 dollars using the consumer price index.

A.3.6 High school graduation rate

The high school graduation rate is defined as the proportion of people aged 25 and above who completed high school (at least). This variable was constructed from county-year data downloaded from the US Census Bureau on the number of residents aged 25+ by educational attainment.

A.3.7 Hospital compare index

The Hospital Compare Index is a measure of average hospital quality in a HRR reported by the US Centers for Medicare and Medicaid Services. CMS constructs the index from a weighted sum of measures describing: (1) whether hospitals follow best practices for care, (2) patient outcomes, (3) survey-based measures of patient experience, and (4) efficiency of imaging, emergency care, and other procedures. We use average measures for each HRR from 2005 – 2011.

A.3.8 Ambulatory Care Sensitive Conditions

The number of discharges for ambulatory care sensitive conditions per 1,000 Medicare enrollees in a HRR were downloaded from the Dartmouth Atlas of Health Care. We use average measures for each HRR from 2003-2007, 2008, 2009, 2010, and 2011. This is a measure of hospital admissions for conditions that can often be prevented by better outpatient management. Diagnoses in this group include diabetes, pneumonia, and congestive heart failure.

A.3.9 Access to Medical Care

Per capital measures of the numbers of acute care hospital beds, primary care physicians, and medical specialists were downloaded from the Dartmouth Atlas of Health Care. Each variable is reported at the HRR level for the years 2006 and 2011. We use an average over these two years.

A.4 HCC risk adjustment score

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 records.[42]

[42]Additional background information on the risk adjustment model can be found at http://www.nber.org/data/cms-risk-adjustment.html.

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).[43] We reduce the scope for such errors by adjusting HCC scores using the procedure from Finkelstein et al. (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 use the resulting predicted health index as a measure of objective health in the survival function.

A.5 References

Finkelstein, A., M. Gentzkow, and H. Williams. 2016. "Sources of geographic variation in health care: Evidence from patient migration." *Quarterly Journal of Economics*, 131(4): 1681-1726.

Song, Y. et al. 2010. "Regional variations in diagnostic practices." *New England Journal of Medicine*, 363(1): 45-53.

Welch, H.G. et al. 2011. "Geographic variation in diagnosis frequency and risk of death among Medicare beneficiaries". *Journal of the American Medical Association*, 305(11): 1113-1118.

[43]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.

Appendix B Additional Results

B.1 First Stage Results

Table B.1 reports the first-stage regression coefficients and Figure B.1 provides a graphical representation of the identifying variation. The histogram shows the density of the instrument constructed from the HRR fixed effects in (6). The solid line shows the conditional variation in medical spending predicted by the instrument. Specifically, it shows the fitted values from regressing residual medical spending on residual variation in the instrument, after controlling for all of the covariates in the main specification of our model. The dashed lines represent a 95% confidence interval on the prediction. Intuitively, the slope suggests that the HRR-based instrument is positively associated with individual medical spending.

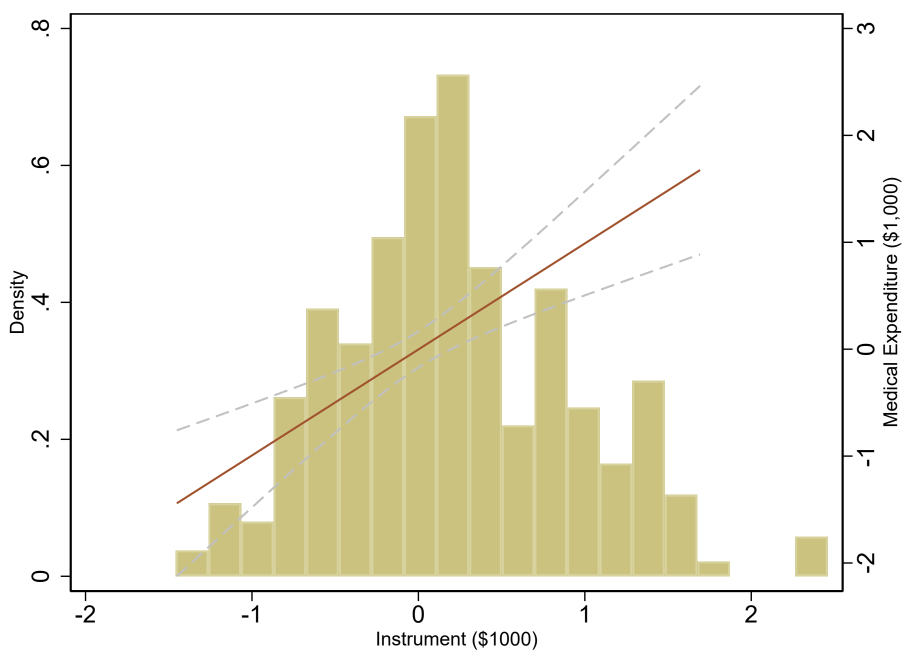

Note: The histogram shows the variation in medical spending due to place effects estimated for 306 hospital referral regions. The right vertical axis plots conditional variation in medical spending against conditional variation in the instrument after removing the variation in each that is explained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

Figure B.1: Identifying Variation in the Instrument

Table B.1: First Stage Results

	coefficient	standard error
IV	1.251	0.383
IV x HCC	2.413	0.656
IV x I{ADL}	1.241	0.342
IV x I{IADL}	-0.288	0.382
hcc index	21.239	0.527
ADL restrictions	3.435	0.243
IADL restrictions	2.222	0.293
health = poor	11.300	0.654
health = fair	3.562	0.346
health = very good	-2.350	0.203
health = excellent	-3.832	0.254
ever smoker	0.201	0.210
BMI = underweight	-0.965	0.517
male	-1.751	2.721
age x under90 x male	-0.561	0.032
age x under90 x female	0.011	0.035
age x over90 x male	-0.552	0.030
age x over90 x female	0.025	0.033
insurance = Medicare Advantage	-4.672	0.320
insurance = Medigap	2.265	0.264
insurance = Medicaid	-2.490	0.381
married	0.143	0.214
children	0.801	0.367
race = African-American	-1.541	0.462
race = Hispanic	-0.214	0.905
race = Other	-2.155	0.706
edu = less than high school	-1.313	0.296
edu = some college	0.790	0.279
edu = college	1.979	0.289
hrr hospital quality	-1.378	4.381
hrr hospital beds	-1.079	0.508
hrr primary care physicians	-0.040	0.023
hrr medical specialists	0.021	0.013
hrr physicians	0.044	0.017
hrr auto deaths	-0.262	0.039
hrr homicide deaths	-0.120	0.060
hrr PM2.5	-0.208	0.071
hrr average max summer temp	0.053	0.027
hrr average min winter temp	-0.024	0.037
hrr share urban	-1.956	1.930
hrr household income	0.000	0.000
hrr high school graduation rate	6.114	5.592

Note: The table reports coefficients from the first-stage regression for the survival function in Table 2 column (5). State fixed effects are suppressed for brevity. Standard errors are clustered by HRR.

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. Conditional on the covariates in our main specification, a standard deviation increase in the instrument is associated with a 0.035 standard deviation increase in expenditures not processed by Medicare compared to a 0.04 standard deviation increase in expenditures processed by Medicare.

B.2 Second Stage Results

Table B.2 reports average marginal effects from the survival function in Table 2 column (5). They are expressed as percentage point changes in the one-year probability of death.

The marginal effects of health measures are intuitive and quantitatively important. For example, a standard deviation increase in the HCC morbidity index of observable chronic illnesses is associated with an 8 percentage point increase in the one-year probability of death. 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 5.6 percentage point increase in the probability of death, whereas moving from "good" to "excellent" is associated with a reduction of 3.3 percentage points.

B.3 Model Fit

Figure B.2 compares model-based 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 does not reproduce as much of the idiosyncratic year-to-year variation around the trend.

Table B.2: Average Marginal Effects on Mortality

	Average marginal effect (pp)	Bootstrapped standard error	95% confidence interval	
$1,000 in medical spending	-0.189	0.091	-0.366	-0.013
1st stage residual morbidity	0.238	0.090	0.064	0.420
HCC index	8.501	1.967	4.819	12.409
one or more ADL restrictions	2.630	0.400	1.889	3.405
one or more IADL restrictions	1.048	0.299	0.513	1.655
health = poor	5.573	1.103	3.558	7.848
health = fair	2.265	0.435	1.400	3.147
health = very good	-2.063	0.360	-2.804	-1.366
health = excellent	-3.252	0.547	-4.319	-2.212
ever smoked	1.255	0.239	0.778	1.709
underweight BMI	2.368	0.413	1.596	3.179
male	6.853	3.274	0.931	13.629
age x {male} x {under 90}	0.085	0.059	-0.033	0.196
age x {female} x {under 90}	0.068	0.041	-0.008	0.152
age x {male} x {over 90}	0.108	0.056	-0.007	0.212
age x {female} x {over 90}	0.063	0.037	-0.005	0.142
married	-0.509	0.208	-0.938	-0.125
has living children	0.451	0.385	-0.300	1.236
African-American	-0.236	0.430	-1.137	0.596
Hispanic	-1.261	0.599	-2.512	-0.244
race = other	-2.147	0.948	-4.272	-0.657
eduction = less than high school	-0.324	0.302	-0.859	0.343
education = some college	0.105	0.285	-0.477	0.679
education = college	-0.492	0.363	-1.195	0.183
Medicare advantage coverage	-0.406	0.553	-1.521	0.741
Medigap coverage	-0.408	0.355	-1.037	0.301
Medicaid coverage	-1.248	0.412	-2.104	-0.442
hospital compare index	-2.161	6.136	-12.846	11.648
hospital beds / capita	-0.035	0.618	-1.119	1.298
primary care physicians / capita	0.005	0.026	-0.055	0.048
medical care specialists / capita	-0.001	0.018	-0.027	0.041
ambulatory discharges / capita	0.007	0.022	-0.037	0.050
automobile mortality	0.025	0.047	-0.070	0.111
homicide mortality	-0.074	0.060	-0.174	0.067
fine particulate matter	-0.030	0.077	-0.180	0.123
mean summer high temperature	0.023	0.036	-0.044	0.093
mean winter low temperature	-0.032	0.038	-0.106	0.046
share urban	3.342	2.091	-1.693	6.596
median household income	0.000	0.000	0.000	0.000
high school graduation rate	5.769	5.594	-5.076	17.033
number of person-years	44,697	44,697	44,697	44,697
number of people	22,206	22,206	22,206	22,206

Note: The table reports average marginal effects from the survival function in Table 2 column (5). State fixed effects are suppressed for brevity. Standard errors and confidence intervals are based on 1,000 bootstrap repetitions.

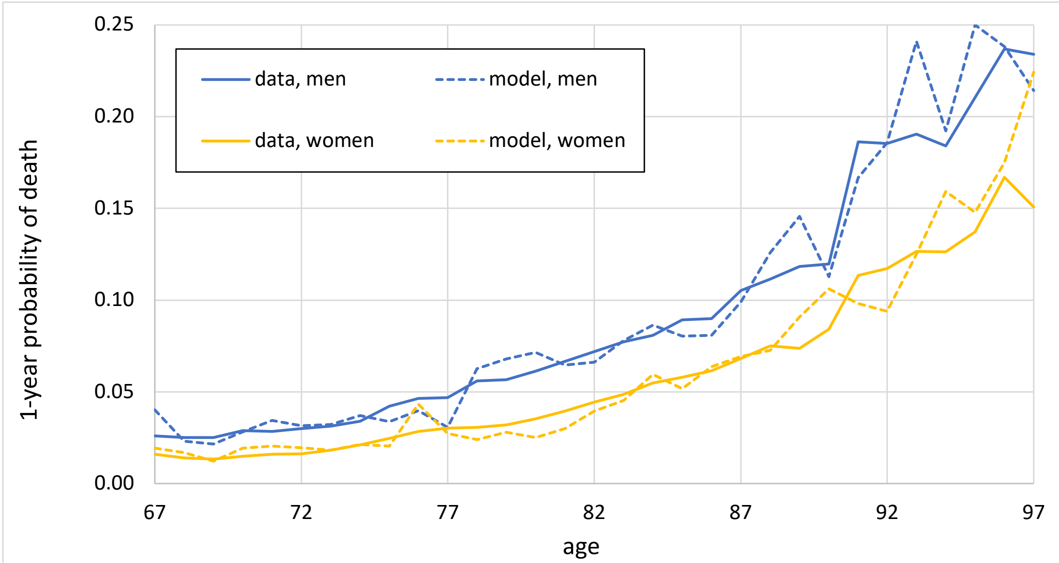

Note: The dashed lines show one-year mortality rates by age and sex in the data. The solid lines show model-based predictions.

Figure B.2: Predicted and Actual One-Year Mortality Rates for Males and Females

B.4 Heterogeneity in the Return to Medical Spending

Figure B.3 shows how our estimates for the return to medical spending vary with age. The solid line shows the estimated average marginal effects of spending on the one-year probability of survival for seniors at each age from 67 to 97. The shaded area defines 95% confidence bands based on 1,000 bootstrap repetitions of the survival function summarized in column (5) of Table 2.

Figure B.4 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 0.35 percentage points for the average person who reports their health as "poor" compared to 0.04 for the average person who reports their health as "excellent". Panel (B) shows the same qualitative pattern such

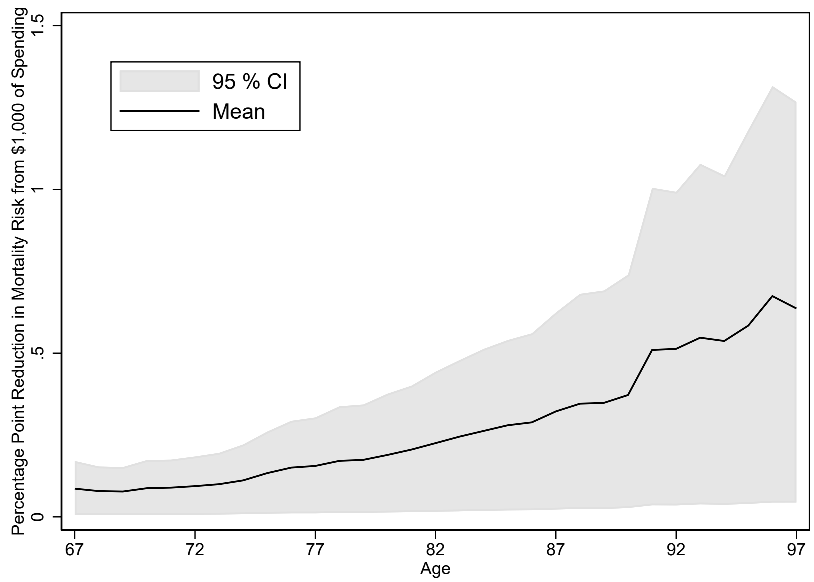

Note: The solid line shows the age-specific average marginal effect of medical spending on the one-year probability of survival. The shaded area defines 95% confidence bands based on 1,000 bootstrap repetitions.

Figure B.3: Average Marginal Effects of Medical Spending on Survival, by Age

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 instrumental 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.

Panels (C) and (D) show that the 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 in Figure B.4 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 (2019), Doyle et al. (2015), and Doyle (2011) collectively suggest a range of

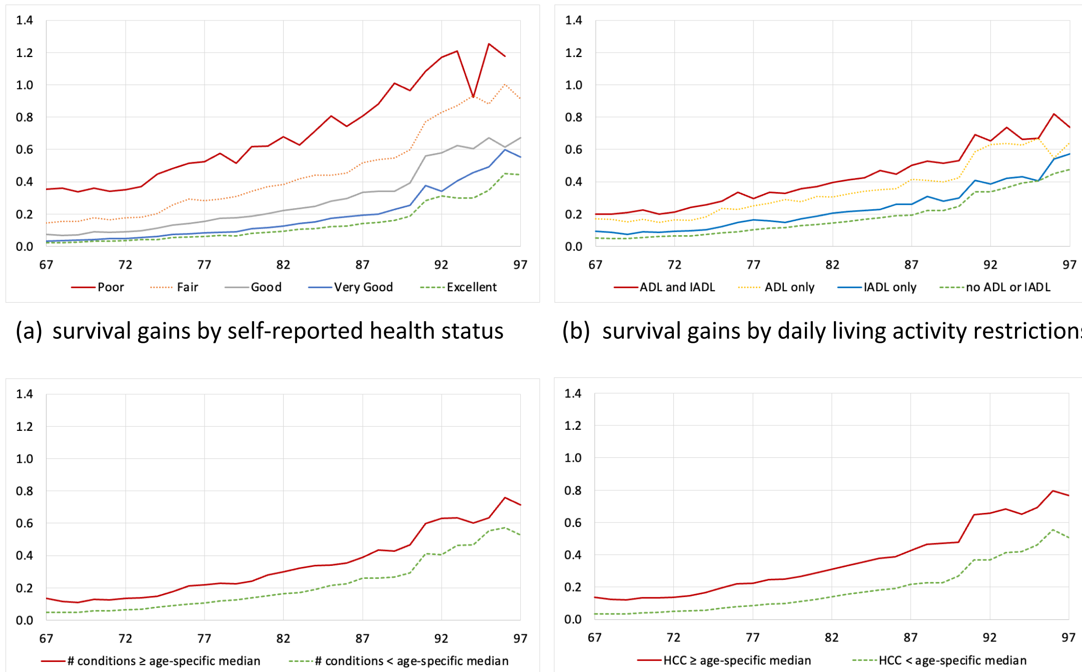

(a) survival gains by self-reported health status

(b) survival gains by daily living activity restrictions

(c) survival gains by number of chronic conditions

(d) survival gains by HCC score

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 (5) of Table 2.

Figure B.4: Survival Gains from Marginal Increase of $1,000 in Medical Spending

marginal returns to $1,000 of medical spending from about 0.1 to 1.5, with relatively higher returns among sicker cohorts.[44]

Figure B.5 helps to illustrate why the VSL declines with age in Figure 3. The average coinsurance rate decreases very slightly with age whereas the returns to

[44]We summarize prior results here, converting to 2024 dollars. At the lower end, Huh and Reif (2017) find that spending $1,000 more on prescription drugs due to the implementation of Medicare Part D reduced mortality by 0.1 percentage points. Among the younger, poorer Medicaid population, however, an additional $1,000 spending on prescription drugs led to a 1.5 percentage point reduction in mortality (Clayton, 2019). 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 reduced annual mortality of 0.14 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 reduced annual mortality by about 1.3 percentage points.

survival from medical spending increase steadily with age. Thus, dividing the former by the later yields a mean VSL that declines with age.

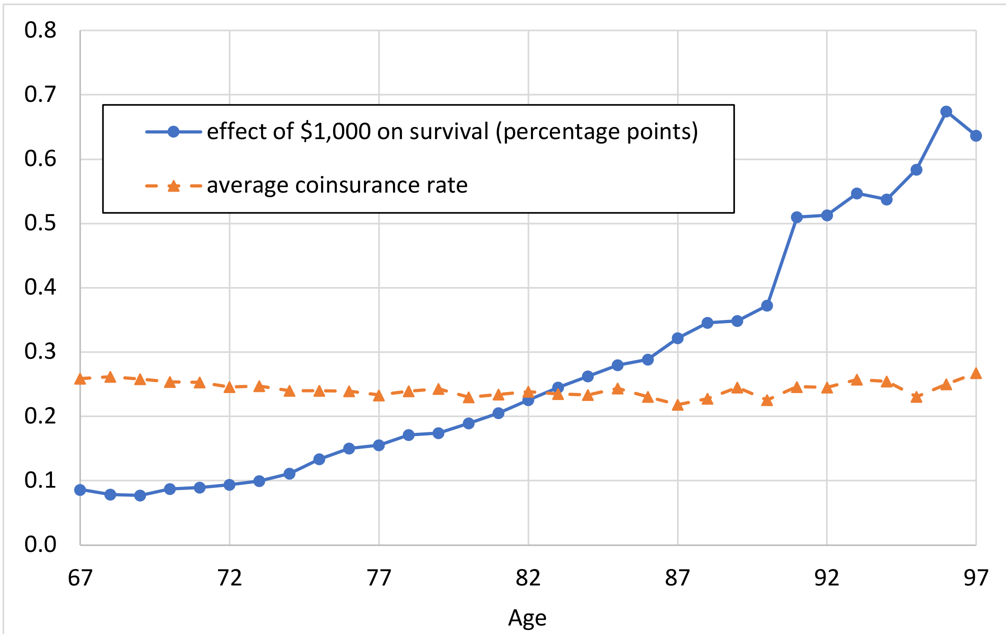

Note: The dashed line shows the average coinsurance rate from the data, i.e., the ratio of out-of-pocket to total medical expenditures. The solid line shows the average marginal effect of a $1,000 increase in medical spending on the probability of surviving to the end of the following year measured in percentage points and calculated from the model shown in col (5) of Table 2.

Figure B.5: Coinsurance Rate and Return to Spending: Age 67 to Age 97

B.5 Heterogeneity in the Value of Statistical Life

Figure B.6 illustrates heterogeneity in the VSL at age 70. To construct the histogram, we first used our model to calculate VSL measures at each unique combination of covariates found among the 2,698 people who we observe in the data at age 70. The figure reports the resulting distribution of type-specific measures. While the distribution is mostly concentrated below $2 million, the right tail extends above $9 million.

Figure B.7 summarizes how the VSL-age profile varies by measures of sex, health, and health behaviors. Figure B.7a shows a large VSL gap between ever-smokers and never-smokers. At age 67 the VSL among never-smokers is approximately twice

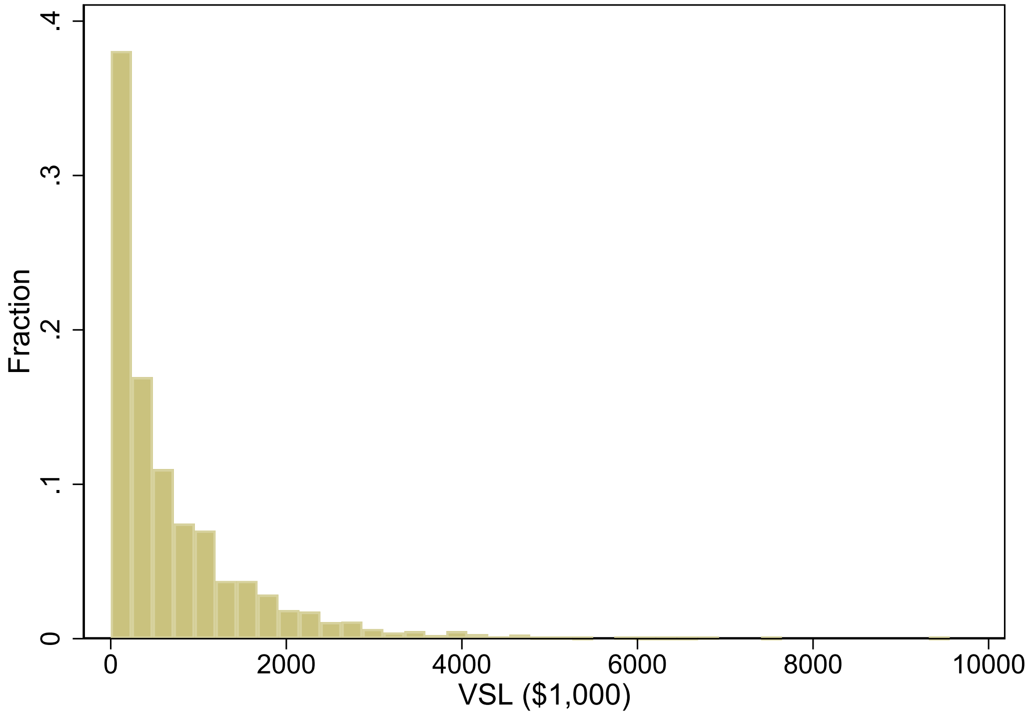

Note: The histogram shows the variation in VSL estimates based on 2,698 people who we observe at age 70. Conditional on age, the VSL differs across person-types due to differences in their health and demographics.

Figure B.6: Heterogeneity in the VSL at Age 70

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. 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.

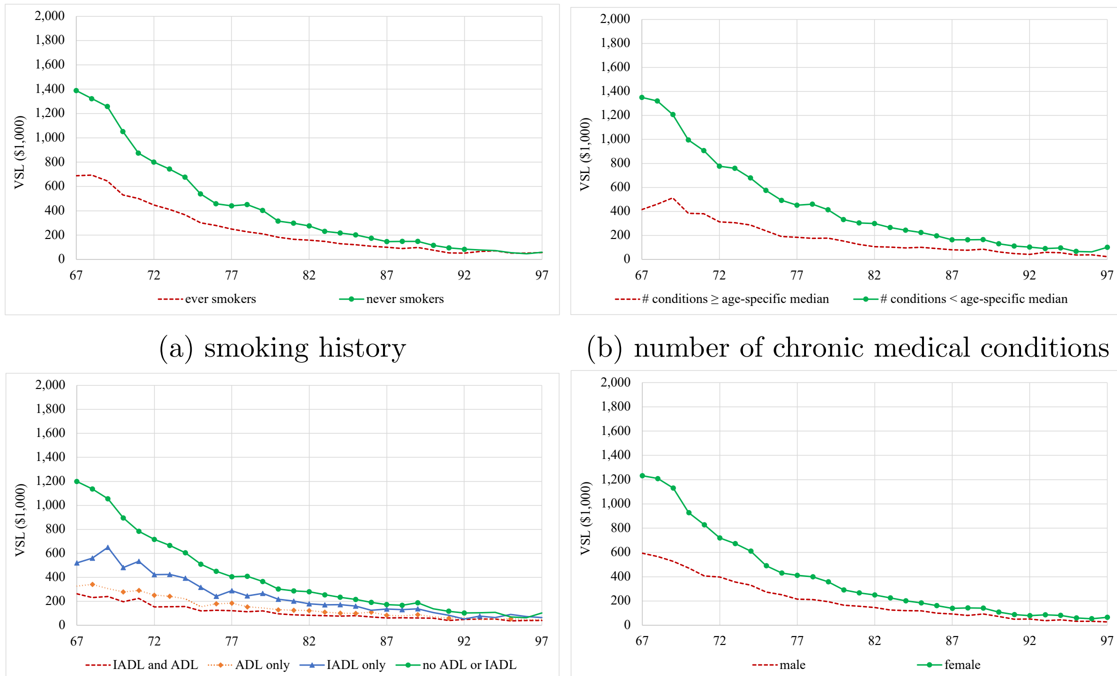

(a) smoking history (b) number of chronic medical conditions

(c) restrictions on activities of daily living (d) sex

Figure B.7: VSL Measures by Age, Health, Sex and Smoking History

Note: Each panel shows the mean age-specific VSL in \$1,000 (2024) dollars stratified by demographics. Markers along each line denote ages at which the VSL measure exceeds the VSL measure for the lower adjacent line in at least 95% of 1,000 bootstrap repetitions of the model in col (5) of Table 2.

Figures B.7b and B.7c provide additional evidence that morbidity is associated with lower VSL measures. Conditional on age, Figures B.7b shows that the VSL is lower for people who are diagnosed with more than the median number of chronic medical conditions. Similarly, Figure B.7c shows that the VSL is lower for people who face restrictions on basic and/or instrumental activities of daily living.

Figure B.7d 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 is consistent with life-cycle model-based predictions in Aldy and Smyth (2014) and Murphy and Topel (2006).

Table B.3 reports the VSL and VSLY point estimates shown in Figure 3.

B.6 References

Clayton, D.H. 2019. "The effect of prescription drug coverage on mortality: evidence from Medicaid implementation." *Journal of Health Economics*, 63: 100-113.

Doyle, J.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 J.J., Graves, J.A, Gruber, J., and Kleiner, S.A. 2015. "Measuring returns to hospital care: Evidence from ambulance referral patterns". *Journal of Political Economy*, 123(1): 170-214.

Huh, J. and Reif, J. 2017. "Did Medicare Part D reduce mortality?" *Journal of Health Economics*, 53: 17-37.

Jha, P. et al. 2013. "21st-century hazards of smoking and benefits of cessation in the United States." *New England Journal of Medicine*, 368(4): 341-350.

Viscusi, W.K. and Hersch, J. 2008. "The mortality costs to smokers" *Journal of Health Economics*, 27(4): 943-958.

Table B.3: VSL and VSLY, by age

age	mean VSL	mean VSLY
67	930	69
68	915	70
69	863	69
70	725	60
71	647	55
72	578	51
73	539	50
74	489	47
75	395	40
76	348	37
77	324	36
78	316	37
79	290	35
80	238	30
81	222	30
82	207	30
83	184	28
84	169	27
85	159	27
86	140	26
87	123	24
88	120	25
89	125	28
90	98	23
91	77	19
92	72	19
93	72	21
94	72	23
95	54	18
96	49	17
97	58	22

Note: The table reports the VSL and VSLY by age, based on results from the survival function in Table 2 column (5). Measures are reported in year 2024 dollars ($1,000).

Appendix C Additional 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 following Leamer (1983), Banzhaf and Smith (2007), and Greenstone et al. (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 decisions.

C.1 Modifiable features of the research design

C.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.

C.1.2 Alternative instruments for medical expenditures

Our main specification for the instrument in equation (6) followed Finkelstein et al. (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 sensitivity 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 (6) with dummies for integer age. Next we add additional granularity by using dummies for sex-by-integer-age. As a third alternative, we reconstruct the IV after extending the sample to include people who never moved. This increases statistical power and yields a more nationally representative sample of seniors. Fourth, we construct an alternative instrument from data on end-of-life spending based on evidence from 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.

C.1.3 Interacting the instrument with observed health

Our main specification of the first-stage regression in Equation (7) defines the vector of instruments, Z_{it} to include four elements: the instrumental variable defined above and interactions between that variable and the HCC morbidity index, an indicator for whether the individual has restrictions on basic activities of daily living, and an indicator for whether the individual has restrictions on instrumental activities of daily living. This set of instruments is designed to capture between-person variation in medical spending induced by geographic variation treatment options, as well as within-person variation along a given menu induced by changes in individual health. As a sensitivity check on our main specification, we repeat estimation after excluding the three interactions.

C.1.4 Interacting medical spending with health, age, or education

Our main specification of the survival function in Equation (8) is a nonlinear function of a linear index of medical spending and covariates. As a sensitivity check on our main specification, we repeat estimation after adding interactions between medical spending and measures of health, age, or education. For health we use the HCC morbidity index, an indicator for whether the individual has restrictions on basic activities of daily living, and an indicator for whether the individual has restrictions on instrumental activities of daily living. For age we use indicators for whether the individual is in their 60's, 70's, 80's, or over 90. For education we use indicators for whether the individual did not complete high school, has a high school degree, has some college education but no degree, or has a college degree.

C.1.5 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. The Gompertz form assumes that the log of the mortality rate is linear in covariates. It is used to model mortality rates in Chetty et al. (2016), Finkelstein et al. (2019), and Bishop et al. (2024).

C.1.6 Allowing heterogeneity in agency and information frictions

We repeat the estimation for two alternative sets of covariates that allow for heterogeneity in agency and information frictions. The first set interacts spending with indicators for whether the individual makes health insurance decisions on their own, gets help from someone else, or uses a proxy. In this case, we calculate VSL measures only for the subset of people who make their own decisions. The second set interacts spending with an indicator for whether the MCBS knowledge module provides reason to suspect the individual's decisions may be affected by information frictions. In this case, we calculate VSL measures only for the subset of people whose decisions are not obviously affected by information frictions.

C.2 Results

Altogether we consider five different instruments for medical spending, first-stage regressions with and without interactions between the instrument and observed health, survival functions with and without interactions between medical spending and health, age or education, survival functions with and without interactions between medical spending and indicators for agency and potential information frictions, and models including and excluding workers. Considering all permutations of these modeling decisions yields 240 different specifications. We estimate each one and calculate the mean VSL by age.

Figure C.1 shows VSL-age profiles from the 240 models. Each line corresponds to a distinct combination of modeling decisions. The heavy dashed line toward the

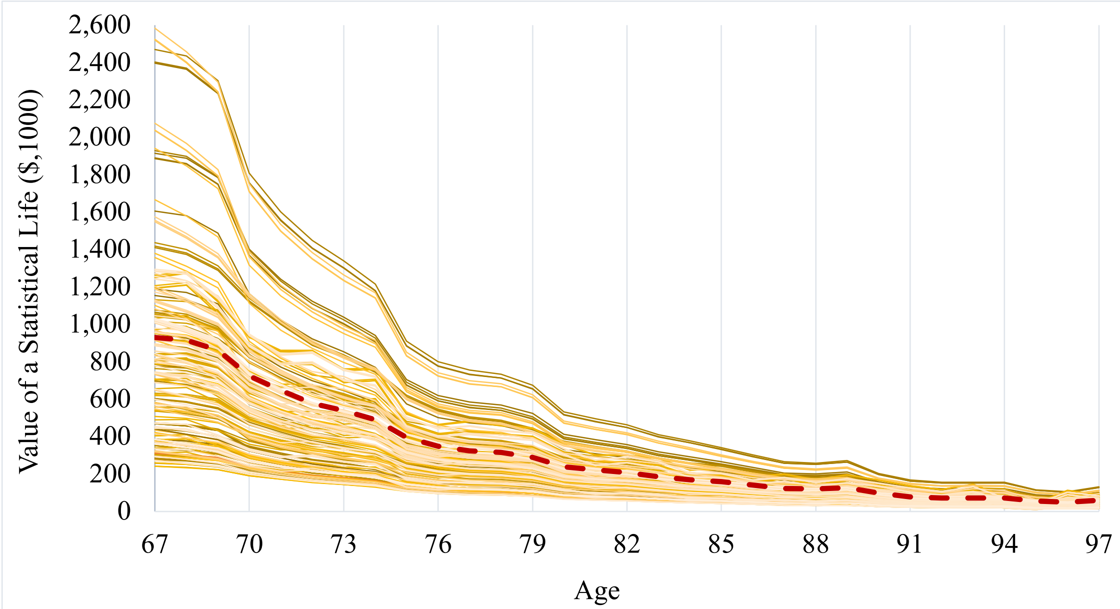

Figure C.1: Sensitivity of VSL Estimates to Model Features

Note: The figure shows the estimated mean VSL by age for 240 different specifications of the survival function described in the text.

middle of the range of estimates is our main specification from col (5) of Table 2. At age 67 our preferred estimate is $930,000. The 95th percentile is $1,891,000 and the maximum is $2,583,000. 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 et al. 2013). Notably, every specification yields mean VSL estimates that lie below $2.6 million at ages 67 and above. Overall, we find that the level and curvature of the VSL-age profile varies with modeling decisions, but two of its most important features are thoroughly robust. First, the VSL declines with age. Second, $2.6 million provides an upper bound on the VSL implied by seniors' medical expenditures.

C.3 References

Banzhaf, H.S. and V.K Smith. 2007. "Meta-analysis in model implementation: Choice sets and the valuation of air quality improvements." *Journal of Applied Econometrics*, 22: 1013-1031.

Bishop, K.C., N.V. Kuminoff, S.M. Mathes, and A.D. Murphy. 2024. "The marginal cost of mortality reduction: evidence from housing markets." *Forthcoming in Journal of Urban Economics*.

Chetty et al. 2016. "The association between income and life expectancy in the United States, 2001-2014." *Journal of the American Medical Association*, 315(16): 1750-1766.

Cutler, D., J.S. Skinner, A.D. Stern, and D. 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.

Finkelstein, A., M. Gentzkow, and H. Williams. 2016. "Sources of geographic variation in health care: Evidence from patient migration". *Quarterly Journal of Economics*, 113(4): 1681-1726.

Finkelstein, A., M. Gentzkow, and H. Williams. 2021. "Place-based drivers of mortality: Evidence from migration". *American Economic Review*, 118(8): 2697-2735.

Greenstone, M., E. Kopits, and A. 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.

Leamer, E.E. 1983. "Let's take the con out of econometrics". *American Economic Review*, 73(1): 31-43.