Mortality Risk, Insurance, and the Value of Life[*]

Darius Lakdawalla
University of Southern California and NBER

Julian Reif
University of Illinois at Urbana-Champaign

Daniel Bauer
Georgia State University

Abstract. Economic models of risks to life and health are widely employed to assess the costs and benefits of public policies. Yet, these models fail to explain a number of common phenomena, such as the large fraction of spending that occurs near the end of life and a positive correlation between fatality risk and the willingness to pay for life-extension. These and other anomalies have led some authors to abandon the standard framework and posit a variety of ad hoc alternatives. However, we show that the standard framework accounts naturally for these phenomena, as soon as one relaxes its restrictive and unrealistic assumption of complete annuity markets. This generalized life-cycle model predicts that a fixed survival gain is worth more to individuals facing bleaker survival prospects, and vice-versa. This insight yields a number of novel policy implications. First, the commonly observed practice of spending disproportionately more resources on individuals facing limited life expectancies may be efficient, not problematic. Second, conventional methods currently used by public and private health insurers undervalue life-extension for severely ill patients compared to the moderately ill, because the value of a statistical life-year varies with the level of mortality risk. Third, public annuity programs like Social Security are strong complements for investments in retiree healthcare, because they increase the value of remaining life for the elderly. For example, annuitization roughly doubles the value of statistical life for a typical 85-year-old American. Finally, holding life-extension benefits constant, treatments are more valuable than preventive technology, because they extend life for people that value it the most.

I. INTRODUCTION

The economic analysis of risks to life and health has made enormous contributions to both academic discussions and public policy. Economists have used the standard tools of life-cycle consumption theory to propose a transparent framework that measures the value of mortality risk-reduction, the value of quality of life improvements, and the value of a statistical life-year (Arthur 1981, Rosen 1988, Murphy and Topel 2006). All these concepts and quantities now play central roles in public policy discussions surrounding investments in medical care, public safety, workplace safety, environmental hazards, and countless other arenas.

At the same time, a number of empirical anomalies have emerged that are difficult to reconcile with a key prediction of the standard framework, namely that – all else equal – the value of a statistical life-year does not vary with life expectancy. At first glance, this result appears contrary to the common economic intuition of diminishing returns – specifically that individuals with less life remaining ought to place more value on each year. For example, consider the plight of a forty-year-old father of two with an inoperable brain tumor. He may choose to spend significant sums of money or endure substantial suffering in pursuit of a one in one thousand increase in his probability of survival. On the other hand, a healthy forty-year-old might be willing to incur little if any cost to secure a similar increase in his survival. At the other extreme, the centenarian with prostate cancer may feel he is "playing with house money," and thus he refuses all aggressive treatment designed to extend life, preferring instead to live out his days without undergoing unpleasant chemotherapy or invasive procedures.

A body of empirical evidence suggests these anecdotes are not anomalies. A number of studies suggest diminishing returns to increases in survival probabilities – or, conversely, higher returns to survival for those who face shorter life expectancies. First, economists have found evidence that individuals are risk-averse over longevity gains (Delprat, Leroux and Michaud 2013). In addition, holding constant the cost per life-year gained, population surveys suggest that people prefer to allocate resources to treating severe diseases rather than milder ones (Nord et al. 1995, Green and Gerard 2009, Linley and Hughes 2013). Put differently, spending $X to gain one extra year of life is viewed as more valuable to a person who faces a more severe illness, contrary to the predictions of the standard model. Third, disproportionate medical resources appear to be spent on individuals near the end of life (Dartmouth Atlas Group 2011). That is, more money is spent per life-year gained on individuals with less time left to live. Fourth, Viscusi (2003) documents black-white disparities in the value of statistical life that remain even after controlling for income.

To be sure, various explanations have been proposed to address these anomalies. Some have proposed abandoning the traditional framework of life-cycle consumption theory by relaxing the assumption that preferences are additively separable over time (Bommier and Villeneuve 2010). Others have argued that healthcare market failures explain high spending at the end of life (Wennberg et al. 2007, Dartmouth Atlas Group 2011). There is also recent research arguing that high spending at the end of life is a rational response to the decline in the opportunity cost of spending money that occurs at the end of life (Philipson et al. 2010). While any or all of these have merit as explanations for particular facts, none is able to explain the full set of anomalies documented above, and none does so in the context of the conventional framework upon which so much economic theory has been built.

In this paper, we show that all these anomalies can be directly explained in the context of the traditional life-cycle framework by relaxing its unrealistic but long-held assumption of perfect annuity markets, which fully insure individuals against longevity risk. While this assumption provides a degree of

analytical convenience, incomplete annuitization is widely documented in empirical analyses of U.S. consumers (Davidoff, Brown and Diamond 2005). When annuity markets are incomplete, an individual's optimal consumption profile, and thus value of life, depend significantly on her mortality risk. For example, an incompletely annuitized person will find it optimal to "spend down" wealth at an accelerated rate when faced with a negative shock to longevity. This spend-down effect is costly to the individual, because it leads her to consume "too much" when hit with a negative survival shock and "too little" when hit with a positive one. Moreover, the size of the distortion in each period is larger when the reduction in survival is larger. For this reason, larger reductions in survival entail higher costs per life-year, and vice-versa.

A very simple example illustrates the intuition. Imagine a retiree with $120,000 in wealth and no bequest motive. Her life expectancy may be 3, 4, or 5 years. For additional simplicity, suppose the optimal consumption profile is flat. Taking a linear approximation to utility, her willingness to pay for a life-year is equal to annual consumption. If she did not annuitize her wealth, and if her life expectancy is 3 years, she will consume $40,000 per year and value a life-year at $40,000. If her life expectancy is extended to 4 years, she will derive just $30,000 of value from each year of life. If it is extended to 5 years, she derives even less, $24,000, and so on. Evidently, the value of a single life-year diminishes with the length of life remaining, as shown in Figure 1.[1] Conversely, the value of a life-year is larger when life expectancy is lower. However, if she were fully annuitized, the value of a life-year would be independent of life expectancy. We show in a general theoretical model that this basic insight holds in models with both deterministic and stochastic mortality.

The assumption of perfect annuity markets masks an important way in which life expectancy might influence the value of a life-year. Relaxing this assumption leads to a central implication: The value of reducing mortality risk is greater for diseases that occur closer to the end of life, and greater for individuals facing higher fatality risks – i.e., more severe diseases. This implication reconciles the standard life-cycle framework of mortality risk with the empirical evidence on medical spending and on the societal preference for allocating medical resources to the most severely ill.

Although the value of a statistical life-year rises immediately following a negative shock to longevity, this value then decays over time and eventually falls below the baseline (healthy) value. Since the value of statistical life – roughly speaking, the net present value of reductions in future mortality risk – integrates across all remaining life years, the overall effect of mortality risk on this value becomes theoretically ambiguous. Our empirical calibrations suggest, however, that increases in mortality risk raise the value of statistical life at older ages. This finding is notable because the elderly are most likely to suffer sudden changes in mortality risk and contribute disproportionately to healthcare spending.

From a policy perspective, a number of important corollaries follow. First, it is optimal to allocate disproportionately more resources to extending the lives of sick individuals closer to the end of life, even holding constant the cost per life-year gained. This provides a rational economic explanation for high

---

[1] Conceptually, this argument is most closely related to Philipson et al (2010), who implicitly assume that incompletely annuitized individuals will need to spend down their resources at the end of life. They correctly observe that this generates a relative increase in the demand for end-of-life medical care, compared to medical care earlier in life. We focus on the implications for the value of statistical life within a traditional life-cycle framework, and its subsidiary implication for social investments in life-extension.

end-of-life spending that is often decried as a market failure. Second, there is a need to develop disease-specific values of a life-year that vary with the fatality risk posed by particular diseases. Third, therapeutic treatments generally create more social value than preventive treatments that add the same amount of life-years, because therapeutic treatments extend the lives of people with higher mortality risks. This provides a rational economic explanation for the low level of observed investment in preventive treatments. Finally, the value of a statistical life-year varies less with income and wealth than conventionally believed, because the poor face higher fatality risk than the rich. This argues for more equity in healthcare allocation decisions than the standard framework suggests.

Another implication of our framework is that public annuity programs like Social Security divert the optimal distribution of healthcare resources away from the middle-aged and towards the elderly. For example, our calibrations suggest that annuitization roughly doubles the value of statistical life for a typical 85-year-old American. This generates the prediction that highly annuitized individuals should expend more resources on health and live longer than those with less generous annuities, which is consistent with existing empirical evidence (Virga 1996).

Our study connects the vast literature on the value of life (Arthur 1981, Rosen 1988, Murphy and Topel 2006, Hall and Jones 2007) with the literature on life-cycle consumption models that goes back to Yaari (1965). It is well known that annuitization provides substantial value by insulating individuals from consumption risk. We show that it also greatly increases the value of statistical life at older ages.[2] Our results suggest that more attention should be paid to the public finance interactions between pension and healthcare systems.

Section II reviews the predictions of the conventional model for the returns to life-extension and demonstrates how relaxing the perfect annuity assumption implies diminishing returns to life-extension. It also develops additional subsidiary implications of this framework. Section III presents empirical analysis that: (1) quantifies how health shocks change the value of statistical life when annuity markets are incomplete; (2) computes the effect of annuitization on the value of longevity; and (3) demonstrates why and to what extent treatment is more valuable than prevention. Section IV concludes.

II. THEORETICAL FRAMEWORK

Consider an individual who faces a mortality risk. We are interested in analyzing the value of a marginal reduction in this risk.[3] We first quantify the value of mortality risk-reduction for an individual who is fully annuitized. We then repeat the exercise for an individual who lacks access to annuity markets, and compare our findings. We first derive our results under the assumption that mortality is deterministic. This allows us to illustrate the basic insights of the paper in a very general setting. We then extend the model to accommodate stochastic mortality and show that our conclusions are unchanged.

---

[2] Reichling and Smetters (2015) show that stochastic mortality and correlated medical costs can explain the puzzling observation that many households do not sufficiently annuitize their wealth. They take the positive correlation between health shocks and medical spending as a given. Our study sheds light on *why* these two phenomena are positively correlated.

[3] We focus on improvements in longevity but allow for improvements in quality of life as well.

II.A. The fully insured value of life-extension

Let $c(t)$ be consumption at time t, W_0 be baseline wealth, $m(t)$ be exogenously determined income, ρ be the rate of time preference, and r be the rate of interest. Finally, define $q(t)$ as health-related quality of life at time t. Since it sacrifices little generality in our application, we take the life-cycle quality of life profile $q(t)$ as exogenous. The maximum lifespan of a consumer is T, and the unconditional probability of dying at any point in time is given by $\theta(t)$, where $0 \leq t \leq T$. The probability that a consumer will be alive at time t is:

$$S(t) = 1 - \int_0^t \theta(\tau)d\tau = \exp\left[-\int_0^t \mu(\tau)d\tau\right]$$

where $\mu(t) = \frac{\theta(t)}{S(t)}$ is the hazard rate at time t. At time period $t = 0$, the consumer fully annuitizes. We assume that annuitization is actuarially fair and based on the consumer's baseline survival, $S_0(t)$, so that the consumer's income is unaffected by unanticipated changes in mortality that might occur later in life.

The consumer's maximization problem is:

$$\max_{c(t)} \int_0^T e^{-\rho t} S(t) u(c(t), q(t)) dt$$

$$s.t. \int_0^T e^{-rt} S_0(t) c(t) dt \leq W_0 + \int_0^T e^{-rt} S_0(t) m(t) dt$$

The consumer's utility function, $u(c(t), q(t))$, depends on both consumption and health-related quality of life. Let $u_c(\cdot)$ denote the marginal utility of consumption. Associating the multiplier λ with the wealth constraint, optimal consumption is characterized by the first-order condition:

$$e^{(r-\rho)t} u_c(c(t), q(t)) = \lambda$$

The marginal utility of life-extension is given by:

$$\int_0^T e^{-\rho t} S(t) u(c(t), q(t)) dS(t) dt$$

Dividing this expression by the marginal utility of wealth, λ, yields the marginal value of life-extension, commonly called the value of a statistical life (VSL):

$$VSL \equiv \int_0^T e^{-rt} S(t) \frac{u(c(t), q(t))}{u_c(c(t), q(t))} dS(t) dt \tag{1}$$

VSL corresponds to the value that the individual places on a marginal reduction in the risk of death. For example, it is the amount that 1,000 people would be collectively willing to pay to eliminate a risk that is expected to kill one of them.

It is also useful to characterize the value of a statistical life-year, which is the value of a one-period change in survival from the perspective of current time:

$$v(t) \equiv \frac{u(c(t), q(t))}{u_c(c(t), q(t))}$$

Because annuity payments are unaffected by unanticipated changes in mortality, our expressions for the value of life do not include the net savings term found in other models (Shepard and Zeckhauser 1984, Rosen 1988, Murphy and Topel 2006). Arguably, this is more realistic than assuming that annuitization payments are updated on the basis of an individual's health profile.[4] It also facilitates comparison to the uninsured case by focusing on differences in the valuation of life that are attributable to differences in the individual's consumption profile. Nonetheless, it is straightforward to incorporate net savings into applications estimating society's willingness-to-pay for longevity.[5]

The value of statistical life is determined by consumption and quality of life. Define the elasticity of intertemporal substitution as:

$$\frac{1}{\sigma} \equiv -\frac{u_{cc}c}{u_c}$$

In addition, define the elasticity of quality of life with respect to the marginal utility of consumption as:

$$\eta \equiv \frac{u_{cq}q}{u_c}$$

When this term is positive, the marginal utility of consumption is higher in healthier states, and vice-versa.

Taking logarithms of the first-order condition for consumption and differentiating with respect to time yields the rate of change for consumption over the life cycle:

$$\frac{\dot{c}}{c} = \sigma(r - \rho) + \sigma\eta\frac{\dot{q}}{q} \tag{2}$$

If one assumes that $r > \rho$ and that the marginal utility of consumption is higher when health status is better, then life-cycle consumption will have the inverted U-shape observed in real-world data.[6]

Note the crucial feature of the conventional model that consumption growth over the life-cycle is independent of mortality risk, because the individual is fully insured against that risk. This feature in turn implies that the rate of change in the value of a life-year is also not a function of mortality risk:

$$\frac{\dot{v}}{v} = \left(\frac{1}{\sigma} + \frac{c}{v}\right)\frac{\dot{c}}{c} = \left(1 + \frac{\sigma c}{v}\right)\left((r - \rho) + \eta\frac{\dot{q}}{q}\right)$$

Although an unanticipated health shock has no effect on the value of a statistical life-year, inspection of equation (1) reveals that it lowers the value of a statistical life by reducing the expected net present

---

[4] Philipson and Becker (1998) argue this is particularly true for public annuity programs such as Social Security, which lack a profit motive to risk-adjust premiums.

[5] In our empirical application, incorporating net savings will decrease the estimated risk-reduction value of life-extension, thereby strengthening our conclusions. This is because our empirical analysis focuses on changes in mortality occurring towards the end of life, when most individuals consume more than they earn.

[6] The inverted U-shape for the age profile of consumption has been widely documented across different countries and goods (Carroll and Summers 1991; Banks, Blundell et al. 1998; Fernandez-Villaverde and Krueger 2007).

value of lifetime utility. Thus, the conventional model implies independence between mortality risk and VSLY. Moreover, the only effect of mortality risk on VSL operates by changing the probability that the individual will live to enjoy consumption in a particular year.

II.B. The uninsured value of life-extension

To characterize the model without annuitization, we employ the Yaari (1965) model of consumption behavior under mortality risk. The consumer's maximization problem is now:

$$\max_{c(t)} \int_0^T e^{-\rho t} S(t) u(c(t), q(t)) dt$$

$$s.t. W(0) = W_0$$

$$W(t) \geq 0$$

$$\dot{W} = rW(t) + m(t) - c(t)$$

If the non-negative wealth constraint binds, then the solution to the consumer's problem is simply to set $c(t) = m(t)$. Otherwise, the solution is to maximize subject to the constraint on the law of motion for wealth. We focus here on the latter, nontrivial case.

Associate the multiplier $\lambda(t)$ with the law of motion for wealth. Removing the annuity market causes the shadow value of wealth to vary over time. The consumer's first-order conditions for consumption and the multiplier are:

$$e^{(r-\rho)t} S(t) u_c(c(t), q(t)) = \lambda(0)$$

$$r\lambda(t) = -\dot{\lambda}$$

Unlike in the case of perfect markets, the survival function enters into the consumer's first-order condition for optimal consumption. Instead of setting the discounted marginal utility of consumption equal to the marginal utility of wealth, the consumer sets the *expected* discounted marginal utility of consumption at time t equal to the marginal utility of wealth. This effectively shifts consumption to earlier ages in the life-cycle. This is rational because consumption allocated to later time periods will not be enjoyed in the event of an early death.

The expression for the marginal utility of life extension is the same as in the case of perfect markets:

$$\int_0^T e^{-\rho t} S(t) u(c(t), q(t)) dS(t) dt$$

The expression for the willingness-to-pay for life-extension is different, however:[7]

$$VSL = \int_0^T e^{-\rho t} S(t) \frac{u(c(t), q(t))}{u_c(c(0), q(0))} dS(t) dt \qquad (3)$$

---

[7] The willingness-to-pay at time $t = 0$ depends on the marginal utility of wealth at time $t = 0$, which from the first-order condition for consumption is equal to $\lambda(0) = u_c(c(0), q(0))$.

As before, the value of statistical life is proportional to the expected discounted (lifetime) utility of consumption, and inversely proportional to the marginal utility of consumption. It is well known that removing annuity markets lowers lifetime utility (Yaari 1965). We also argued earlier, and will show more formally below, that removing these markets shifts consumption to earlier ages, thereby lowering the marginal utility of consumption, at least at those ages. Thus, the effect of annuity markets on VSL is in general ambiguous.

The expression for the value of a statistical life year, v, is the same under both perfect and imperfect markets because it is calculated from the perspective of current time and thus is unaffected by differential discounting. However, the actual values will in general differ because annuitization alters the time profile of consumption.

In particular, the life-cycle consumption profile of the non-annuitized individual depends explicitly on mortality risk. Taking logarithms of the first-order condition for consumption, differentiating with respect to time, and exploiting the relationship that $\frac{\dot{\lambda}}{\lambda} = -r$ yields:

$$\frac{\dot{c}}{c} = \sigma(r - \rho) + \sigma\eta\frac{\dot{q}}{q} - \sigma\mu(t) \tag{4}$$

Comparing this result to the standard case, given by equation (3), reveals both similarities and differences. As in the standard, fully annuitized model, the non-annuitized consumption profile described by equation (4) changes shape when the rate of time preference is above or below the rate of interest and when the quality of life changes. Unlike in the standard model, however, the consumption profile described by equation (4) depends explicitly on the mortality rate, $\mu(t)$. Higher rates of mortality depress the rate of consumption growth over the life-cycle. This rate of growth is always higher in the fully annuitized case, in which the last term drops out of the consumption growth equation (3). Put another way, removing the annuity market "pulls consumption earlier" in the life-cycle.

An appealing feature of the uninsured model is that it generates an inverted U-shape for the profile of consumption under quite natural assumptions. In particular, low income early in life and high mortality risk later in life are sufficient conditions. One need not impose the ad hoc assumptions on the signs of $r - \rho$ or η that are necessary in the fully annuitized model (Murphy and Topel 2006).

Because it depends on consumption, the life-cycle profile of the value of a statistical life-year also depends on mortality:

$$\frac{\dot{v}}{v} = \left(\frac{1}{\sigma} + \frac{c}{v}\right)\frac{\dot{c}}{c} = \left(1 + \frac{\sigma c}{v}\right)\left((r - \rho) + \eta\frac{\dot{q}}{q} - \mu(t)\right) \tag{5}$$

For example, according to equation (4) the current value of a statistical life-year, v, decreases in the event of an unexpected health shock that increases mortality, $\mu(t)$.[8] By contrast, v is unaffected by mortality in the standard, annuitized model. The effect on the value of *future* life-years is ambiguous,

---

[8] Because the current value of a statistical life-year, v, is always calculated from the perspective of current time, an unexpected (negative) health shock at t increases $\pi(t)$, the probability of dying at t, but does not affect $S(t)$, the probability of surviving until t.

however, because the effect of health shocks on the future hazard rate, $\mu(\tau) = \frac{\theta(\tau)}{S(\tau)}, \tau > t$, is ambiguous: all things equal, a shock that increases the probability of dying at age t necessarily reduces both $\theta(\tau)$, the unconditional probability of dying at age $\tau > t$, and $S(\tau)$, the probability of surviving to age $\tau > t$.

Pursuing this point further, we note that the absence of annuitization causes the net effect of a health shock on VSL to be ambiguous, rather than negative as in the traditional model. This result obtains from direct inspection of equation (2). VSL is equal to the expected present value of lifetime utility divided by the current marginal utility of consumption. A negative survival shock decreases lifetime utility, because survival is lower, but it also decreases the marginal utility of consumption, because such shocks shift consumption earlier. In other words, VSL decreases when the individual's lifetime is shorter, but the shock to longevity also makes consumption less valuable, so the individual is willing to trade more units of consumption for a given amount of life expectancy. The latter effect creates the ambiguity in the relationship between health shocks and VSL.

An important implication of this model is that willingness to pay for longevity depends critically on the life-cycle mortality profile. As societies become richer and live longer, the fraction of wealth spent on health will depend not just on the income elasticity of health, but also on the degree of survival uncertainty they face. We return to this point in our empirical exercise. Furthermore, our results imply that public programs such as Social Security that increase annuitization levels will affect society's willingness to pay for longevity, thereby creating a feedback loop that could dampen or increase program expenditures.[9] As a general matter, the model demonstrates that the degree of annuitization influences how people value gains in longevity.

II.C. The value of life-extension when mortality is stochastic

The analysis above demonstrates that mortality risk affects the value of a statistical life when annuity markets are incomplete. Earlier analyses of the value of life have overlooked this relationship by assuming complete annuitization. However, the conventional framework is ill-equipped to study the influence of mortality risk for another reason as well. Prior analysis, just like our model above, treats the mortality rate as a nonrandom parameter (cf, Murphy and Topel, 2006). Thus, shifts in mortality risk reflect preordained and anticipated changes in mortality. In the real world, however, neither the timing nor the size of shifts in mortality risk is known. Thus, we now extend our analysis to include random mortality risk and show that our results continue to hold.

Before we present the formal analysis, we illustrate the intuition behind this result using an heuristic approach to analyzing the deterministic model. In the context of the earlier "deterministic" model, imagine a wealth-compensated shift in mortality, or simply a shift in mortality that leaves the current stock of wealth unaffected. Since we are primarily interested in how individuals respond to mortality changes rather than in how they anticipate them, assume for simplicity that the health state is realized at time zero. Thus, define H as the realized health status, and assume $\theta(t; H)$ where $\pi_H < 0$. Since the health state is known, the mathematics of the analysis are largely unchanged from before, and consumption growth is now given by:

---

[9] Philipson and Becker (1998) make the important, but distinct, point that the moral hazard effects of public annuity programs also increase an individual's willingness to pay for longevity gains.

$$\frac{\dot{c}}{c} = \sigma(r - \rho) + \sigma\eta\frac{\dot{q}}{q} - \sigma\frac{\theta(t;H)}{S(t;H)}$$

It is evident that an increase in mortality risk, θ, decreases growth in consumption. Since W_0 is held fixed, this reduction in consumption growth must be matched by an increase in the initial level of consumption. Moreover, as shown in Figure 2, geometry implies that the new consumption profile will at some point have to cross below the old one. It is worth noting though that the individual may not survive to the point where the new consumption profile crosses below the original one. Highly fatal diseases may result in consumption profiles that are always above the original.

The value of a statistical life-year responds to the mortality shock just like consumption does. Recall that

$$v(t;H) \equiv \frac{u(c(t;H),q(t))}{u_c(c(t;H),q(t))}$$

The mortality shock "pulls consumption forward" in time. Therefore, VSLY rises for the immediate future, but ends up below the baseline level if the individual lives long enough.

We now formalize this intuition using a model of stochastic mortality. In order to generate clean analytic expressions, we make the following simplifying assumptions:

1. Utility depends only on consumption and takes the CRRA form
    $$u(c) = \frac{c^{1-\gamma}}{1-\gamma}$$
    where the parameter γ is equal to the inverse of the elasticity of intertemporal substitution.
2. Labor income, $m(t)$, is equal to 0 in all periods.
3. There are two health states, sick (S) and healthy (H), with corresponding constant mortality rates μ_s and μ_h. The rate of transition from the healthy state to the sick state is given by the parameter π. The sick state is an absorbing state. Thus, the mortality rate evolves according to a two-state Markov chain with transition matrix
    $$\Delta = \begin{bmatrix} -\pi & \pi \\ 0 & 0 \end{bmatrix}$$

The individual's lifetime mortality risk is no longer preordained. Instead, it evolves randomly according to the transition matrix above. This allows for the random arrival of poor health and thus adverse shocks to mortality.

In this environment, the consumer's maximization problem is:

$$\max_{c(t)} E\left[\int_0^\infty e^{-\rho t}S(t)u\big(c(t)\big)dt\right]$$

$$s.t.\, W(0) = W_0$$

$$W(t) \geq 0$$

$$\dot{W} = rW(t) - c(t)$$

Proposition 1: Let $c^*(t|H)$ be the consumer's optimal consumption path conditional on being in the healthy state, H. If at any time t^* the consumer falls sick, her consumption will rise. That is, $c^*(t^*|S) > c^*(t^*|H)$.

Proof: See the appendix.

As a simple corollary of this result, the value of a statistical life-year also rises in response to illness. Observe that $v^*(t|i) = \frac{u(c^*(t|i))}{u'(c^*(t|i))}$, $i \in \{S, H\}$. Thus, whenever consumption rises, the value of a statistical life-year rises, and vice-versa. This analysis generalizes the finding that adverse shocks to mortality raise the value of a statistical life-year. Finally, since wealth does not change in the wake of a health shock, consumption following a health shock will eventually fall below its original level if the individual survives long enough. That is, there exists some $t' > t$ such that $c^*(t'|S) < c^*(t'|H)$ for all t.

Our analysis of health shocks reveals that the value of a treatment exceeds the value of prevention, even if they both increase life expectancy by the same amount. Imagine a treatment that adds $dS(t; S)$ additional survival at time t for patient in the sick state, S. Imagine a corresponding preventive health investment – e.g., exercise – that lowers the probability of the sick state at a point in time by $d\pi(t)$. Assume that the preventive and health investments add equal amounts of life expectancy so that

$$dS(t; S) = d\pi(t)[S(t; H) - S(t; S)]$$

The value of the treatment in the wake of the health shock is

$$value_{treat} = dS(t; S)\frac{u(c(t; S))}{u_c(c(t; S))}$$

The value of the preventive investment, however, is

$$value_{prevent} = d\pi(t)\left[S(t; H)\frac{u(c(t; H))}{u_c(c(t; H))} - S(t; S)\frac{u(c(t; S))}{u_c(c(t; S))}\right]$$

In the conventional model, the value of a statistical life-year is independent of health status ($\frac{u(c(t;H))}{u_c(c(t;H))} = \frac{u(c(t;S))}{u_c(c(t;S))}$). In this case, $value_{treat} = value_{prevent}$. However, in the incomplete annuities context, we have shown previously that $\frac{u(c(t;H))}{u_c(c(t;H))} < \frac{u(c(t;S))}{u_c(c(t;S))}$. This implies that the value of prevention is less than the value of treatment.

II.D. Health risks and the value of life-extension

Absent annuitization, shocks to survival affect consumption decisions and thus the value of a life-year. This contrasts with the standard fully annuitized model, in which consumption profiles are independent of mortality risk. A corollary of these results is that the presence of annuitization affects the willingness to pay for life-extension. Thus, in theory, social programs that increase the degree of annuitization also affect the aggregate willingness to pay for life-extension. We explore this possibility in our simulation analysis by performing calibration exercises that demonstrate how public annuity programs influence the aggregate willingness to pay for life-extension. In addition, we review evidence from the literature that demonstrates the extent to which better annuitized individuals are willing to invest more in life-extension.

We can characterize the implications of our theory in the form of four fairly intuitive sets of statements.

1. **The value of life-extension depends on baseline health status and mortality risk.** Any factor influencing mortality will alter the value of a statistical life-year. For example, countries or socioeconomic groups with different life expectancies will exhibit different values of life-extension. The arrival of different diseases will have systematically different implications for the value of life-extending treatments.
2. **Consumers exhibit diminishing returns to the value of life-extension in two respects.**
   a. *All else equal, a consumer suffering a sudden increase in mortality risk will initially place higher value on a life-year.* Life-years are more valuable to individuals facing bleaker survival prospects than their more robust counterparts. A corollary is that the value of a life-year is higher when the increase in mortality risk from disease is greater.
   b. *Consumers face costs associated with outliving their wealth.* In the wake of a mortality shock, the value of a statistical life-year will initially be higher than its baseline level, but must eventually fall below the baseline level. Patients who survive into old age for much longer than expected will place relatively low values on further life extension.
3. **Health shocks do not always lower the value of statistical life.** A negative survival shock decreases lifetime utility directly by lowering survival, but it also increases consumption in the short-term. Thus, negative survival shocks may raise the value of statistical life-years experienced in the short-term, even though this value will always fall in the long-term.
4. **Treatment and prevention are not equivalent.** In the conventional model, the willingness to pay for life-extension is the same before a health shock is realized as it is after. This equivalence breaks down in an incomplete annuities model, where the ex ante willingness to pay lies between the (lower) ex post healthy and the (higher) ex post sick valuations.

In our empirical analysis, we calibrate our model to quantify the empirical relevance of all these points.

III. EMPIRICAL ESTIMATES OF THE VALUE OF LIFE

Our theoretical framework highlights two important differences between an annuitized and unannuitized person's value of life. First, the value of a statistical life-year is unchanged when completely annuitized patients suffer a sudden shock to mortality because their annuities insulate their consumption from longevity risk. By contrast, incompletely annuitized patients will increase their consumption, and thus their value of a statistical life-year, immediately following a health shock, and this increase will be larger for more fatal shocks. Second, annuitization itself affects the willingness to pay to extend life, even absent an unanticipated health shock. Whether this effect is positive or negative depends on the individual and her environment.

How empirically relevant are these theoretical results? In this section we conduct calibration exercises in order to quantify the magnitude of these effects under reasonable parameterizations of our model. We also discuss evidence from the literature on public pension programs that are consistent with the hypotheses generated by our theoretical model and calibration exercises.

We begin by describing the estimation framework we employ for our calibrations. We then present our empirical results.

III.A. Calibration framework
We will work with the discrete time analogue of our model and abstract from the role of quality of life, since aggregate, nationally representative data on quality-of-life trends are not generally available. Let

c_t be consumption in period t, w_t (non-annuitized) wealth, ρ the utility discount rate, r the interest rate, and $1 - q_t$ the probability of surviving from period t to period $t + 1$. Assume that in each period the consumer receives an exogenously determined income, m_t, and that the maximum lifespan of a consumer is T (i.e., $q_T = 1$).[10] Our baseline model assumes there is no bequest motive, although we relax this assumption in a later exercise. The probability of surviving from period 1 to period t is equal to

$$S_t = \prod_{j=1}^{t-1}(1 - q_j)$$

The consumer's maximization problem is

$$\max_{\{c_t\}} E_t\left[\sum_{t=1}^{T}\frac{u(c_t)}{(1+\rho)^{t-1}}\right]$$

subject to

$$w_1\ given$$
$$w_t \geq 0$$

$$w_{t+1} = (w_t + m_t - c_t)(1 + r)$$

The expectation $E_t[\cdot]$ is taken over states of survival. We assume that utility takes a CRRA form:

$$u(c_t) = \frac{c_t^{1-\gamma} - \underline{c}^{1-\gamma}}{1 - \gamma}$$

We have normalized the utility of death at zero. The consumer receives positive utility if she consumes an amount greater than $\underline{c}$, which represents a subsistence level of consumption. Consuming an amount less than $\underline{c}$ generates utility that is worse than death.

The parameter γ is the inverse of the elasticity of intertemporal substitution, an important determinant of the value of life and the value of annuitization. We set $\gamma = 2$ in our analyses. As points of reference, Murphy and Topel (2006) argue that γ is approximately equal to 1, but Brown (2001) uses survey data to estimate a mean value of $\gamma = 3.95$.

We employ dynamic programming techniques to solve for the optimal consumption path. The value function is defined as:

$$V_t(w_t) = \max_{\{c_t\}} E_t\left[\sum_{t=1}^{T}\frac{u(c_t)}{(1+\rho)^{t-1}}\right]$$

We can use the value function to rewrite the optimization problem as a recursive Bellman equation:

---

[10] Hubbard, Skinner, and Zeldes (1995) show that failing to include a "welfare floor" in the budget constraint causes life-cycle models to overestimate savings for low-income households. Our calibration exercises model median-income individuals, however, for whom this issue is less important.

$$V_t(w_t) = \max_{\{c_t\}} u(c_t) + \frac{1 - q_t}{1 + \rho} V_{t+1}(w_{t+1})$$

The Bellman equation simplifies the maximization problem to a series of two-period problems, which can be solved using standard numerical methods. Details are available in the appendix.

Once we have solved for the optimal consumption path, we can use the analytical formulas derived in the previous section to calculate the value of life. The value of a statistical life-year from the perspective of current time is equal to

$$VSLY_t = c_t{}^\gamma \left(\frac{c_t{}^{1-\gamma} - \underline{c}^{1-\gamma}}{1 - \gamma} \right)$$

Under complete annuitization, the value of a statistical life at time t is equal to

$$VSL_t^* = \frac{1}{S_t} \sum_{\tau=t}^{T} \left(\frac{S_\tau}{(1+r)^{\tau-t}} \right) v_\tau$$

Under incomplete annuitization, this expression becomes

$$VSL_t = \frac{c_t{}^\gamma}{S_t} \sum_{\tau=t}^{T} \left(\frac{S_\tau}{(1+\rho)^{\tau-t}} \right) \left(\frac{c_\tau{}^{1-\gamma} - \underline{c}^{1-\gamma}}{1 - \gamma} \right)$$

We assume throughout that $r = \rho = 0.03$ (Siegel 1992).

We begin the model at age 20 and employ mortality data obtained from www.mortality.org. We assume that nobody survives past age 100.

We choose the individual's income profile, $\{m_t\}$, to fit data on life-cycle wages as reported by the Current Population Survey.[11] We assume income falls to \$15,528 beginning at age 65, which corresponds to the average annual Social Security benefit for retirees in 2013. The wage path is displayed as a dashed line in Figure 3.

We consider two scenarios. In the first, the consumer fully annuitizes and enjoys a constant annuity stream, $\overline{m}$, provided by an actuarially fair annuities market. In the second, she is partially annuitized (by virtue of her Social Security income) but she lacks access to financial markets and cannot borrow. The income streams in the two scenarios are related according to the following equation:

$$\sum_{t=1}^{T} \frac{m_t S_t}{(1+r)^{t-1}} = \overline{m} \sum_{t=1}^{T} \frac{S_t}{(1+r)^{t-1}}$$

Our assumed interest rate and our data on mortality and income imply an annuity value of $\overline{m} = \$35,293$.

---

[11] These data are available at http://data.bls.gov/pdq/querytool.jsp?survey=le. We smooth the data by fitting it to a quadratic polynomial in age.

We are unaware of any empirical evidence on the magnitude of $\underline{c}$, the subsistence level of consumption in the United States. We assume it is equal to $5,000. We set initial wealth, w_1, equal to 0.

The life-cycle profiles of consumption and the value of a statistical life-year for the two scenarios are displayed in Figure 3. Because the consumer discount rate is equal to the interest rate, the annuitized individual enjoys constant consumption over time.[12] Consumption for the non-annuitized individual, by contrast, is constrained by her low income in early life. She saves during middle age when income is high, and then consumes her savings during retirement. Her consumption is higher than the annuitized individual's consumption between the ages of 30 and 70. This is attributable to "shifting consumption forward" in response to mortality risk.

This difference in consumption causes a corresponding difference in the value of a statistical life-year (VSLY). The non-annuitized individual places a low value on VSLY in early and late ages, when consumption is low. Because the annuitized individual enjoys constant consumption, however, she values life-years at a constant rate throughout her life.

Figure 4 displays the value of statistical life (VSL) for these two different scenarios. Discounting and future mortality cause VSL to decline monotonically with age for a fully insured individual. By contrast, the rising value of VSLY early in life generates an inverted U-shape for uninsured individuals. At age 40, VSL is equal to $6.3 million and $4.9 million for fully insured and uninsured individuals, respectively. Both of these values are within the ranges estimated by empirical studies of VSL for working-age individuals (Viscusi and Aldy 2003).

III.B. The effect of health shocks on the value of life

In this section we estimate how the value of life differs across diseases with different mortality profiles. The effect of disease on survival is estimated using the Future Elderly Model (FEM), a widely published microsimulation model that employs nationally representative data from the Health and Retirement Study (Goldman et al. 2005, Lakdawalla, Goldman and Shang 2005, Goldman et al. 2009, Lakdawalla et al. 2009, Goldman et al. 2010, Michaud et al. 2011, Michaud et al. 2012, Goldman et al. 2013). The FEM uses real-world risks of disease incidence, and mortality rates by disease state, in order to estimate longevity for people over the age of 50 with different comorbid conditions.[13] This is quite useful for our current purposes, because it provides us with an empirically relevant set of estimates for what mortality risk looks like under different disease states.

Figure 5 shows survival curves for the average person diagnosed at age 50 with a particular disease. The "healthy" group includes everybody and thus corresponds to average survival for the cohort. The other groups correspond to individuals diagnosed with cancer, diabetes, hypertension, and stroke,

---

[12] Murphy and Topel (2006) generate an inverse U-shaped profile for consumption for fully annuitized consumers by assuming that (1) $r > \rho$; (2) health and the marginal utility of consumption are complements; and (3) quality of life declines with age. While the third assumption is not controversial, the empirical evidence on the first two assumptions is mixed.

[13] A description of its methodology is available at healthpolicy.app.box.com/FEMTechdoc.

respectively. As seen in the figure, stroke has the largest negative effect on survival, while hypertension has the smallest.

We use these survival data to calculate how the value of a statistical life-year (VSLY) varies with age for non-annuitized individuals experiencing these different health shocks at age 50. Those results are depicted in Figure 6. This figure illustrates our predictions that mortality shocks initially raise the value of a life-year above the average value in the population, but eventually lower them below average. It also demonstrates how the value of a life-year varies across diseases. As expected, stroke results in the biggest increase in the value of a life-year, because it also involves the greatest risk to mortality. Over time, this effect decays, as stroke survivors outlive their wealth. Eventually, the value of a life-year ends up being among the lowest for long-lived stroke survivors.

Figure 7 shows the corresponding values for the value of statistical life (VSL). This demonstrates that mortality shocks have ambiguous effects on VSL, in the absence of annuities, but not in the fully annuitized case. The left panel of the figure shows VSL for fully annuitized individuals. As predicted by the standard model, VSL is always higher for diseases with the highest survival probabilities – e.g., hypertension – and lower for those with lower survival probabilities – e.g., stroke. And, it is highest overall for the solid blue curve representing VSL for individuals not experiencing any health shock. The right panel of the figure illustrates the effect of incomplete annuitization. In the absence of annuities, mortality shocks can actually increase the overall VSL in the short-term, because they make each remaining year of life more valuable in the short-term. This is evident, as VSL is initially higher for all the disease groups than it is for the baseline group. Eventually, VSL for the sick groups decays to levels below the baseline, as the long-term survivors outlive their wealth.

Figure 8 illustrates that the rate of decay in the value of a life-year is highest for diseases that involve the greatest mortality risk. This is mechanically true, because these diseases also involve the sharpest increase in the value of a life-year at onset. For example, stroke, which is the deadliest disease in our list, initially raises the value of a life-year by more than 50% above its baseline (healthy) level, but within five years this value has fallen almost halfway back to baseline. In contrast, hypertension raises the value of a life-year by less than 25%, and sees about half that increase disappear after five years. These results imply that first-line medical treatments that extend life at the onset of disease are more valuable than second- and later-line treatments. It also implies that treatments for severe illnesses are more valuable than mild illnesses, holding constant the level of mortality reduction. This differs from the standard result of the annuitized model, which places equal value on all treatments regardless of disease severity.

Figure 9 contrasts the fully annuitized model predictions with those of the non-annuitized model. With annuities, diseases that impose greater mortality risk steadily lower the value of a statistical life at disease onset. In contrast, the relationship can possibly reverse, although not always, when dealing with non-annuitized individuals. From a policy standpoint, this figure sheds light on the question of how much sick individuals are willing to pay in order to reduce their risk of dying. This question is important when considering the reimbursement of new medical technologies as well as policies to encourage investments in medical innovation. On the one hand, a fully annuitized patient who suffers a stroke has a VSL equal to $3.5 million, about $1.5 million less than a healthy fully annuitized person. In other words, a group of 100 stroke patients would be willing to be pay $3.5 million for a 1 percentage point reduction in stroke mortality. Under full annuitization, mortality risk always lowers VSL, because it reduces life expectancy. On the other hand, a non-annuitized stroke victim would be willing to pay $5.2 million for a marginal reduction in mortality, which is actually higher than VSL for the healthy. The

relationship between mortality risk and VSL becomes more complicated without annuities, because in this case, higher mortality risk reduces the individual's time-horizon but also increases the value of a life-year. Given our parameter assumptions, in this particular case the absence of annuity markets makes medical technology more valuable on balance for patients suffering the onset of a stroke.

III.C. The effect of annuitization on value of life

Another implication of our framework is that public annuity programs should influence the overall social willingness to pay for life-extension. Our theoretical results suggest that there may be important complementarities between annuity programs like Social Security and retiree health care programs like Medicare. Here, we illustrate this relationship in a calibration exercise, and review the literature for evidence that tests our predictions.

III.C.1. Calibration estimates

The United States and many other countries have private and public insurance programs that provide annuities to their citizens. For example, about 60% of elderly income comes from annuities in the U.S.[14] How does this affect society's willingness to pay for health? Our two scenarios – partial annuitization and full annuitization – can shed light on this question.

Figure 4 displays VSL by age for these two different scenarios. VSL declines monotonically with age for a fully insured individual. By contrast, it exhibits an inverted U-shape for uninsured individuals. This has implications for both the optimal level and distribution of healthcare spending.

First, the optimal level of healthcare spending could go up or down, depending on the population distribution. The presence of a public annuity program in Iran – where more than 60 percent of the population is under 30 – likely increases the optimal level of healthcare spending. The opposite is likely true for Japan, where only about 25 percent of the population is under 30.

Second, the presence of an annuity program will divert the optimal distribution of healthcare resources away from the middle-aged and towards the young and the elderly. Given that most public annuity programs only apply to the elderly, we focus on that age group. Figure 4 shows that the VSL for, say, an 85-year-old doubles in the presence of an annuity. Thus, a program like Social Security drastically increases the VSL for the older elderly. This dovetails with the point, made by Philipson and Becker (1998), that the moral hazard effects of these programs also increase the willingness to pay for longevity. It is therefore not surprising that public spending on healthcare – particularly for the elderly – has grown enormously in developed countries. Indeed, public annuity programs are complementary with retiree healthcare programs, but substitutable with universal healthcare access for the middle-aged.

III.C.2. Evidence from aggregate data

Our calibration estimates predict that annuitization raises VSL for the elderly. This should cause them to spend more on healthcare and invest more in healthy behaviors, which in turn should ultimately manifest in increased life expectancy. Philipson and Becker (1998) analyze data from Virga (1996) to show that people with more generous annuities live longer than those with less generous annuities. They interpret this as the effect of endogenous longevity investments, which are encouraged among

---

[14] 40 percent comes from Social Security and 20 percent comes from private pensions and annuities. See http://www.ebri.org/pdf/notespdf/EBRI_Notes_06-June10.Inc-Eld.pdf

highly annuitized individuals. An additional explanation is that annuitization increases the value of statistical life, as we emphasize here. These are compatible and consistent interpretations.

Another implication of our model is that, all things equal, a society with a public annuity program will have greater demand for healthcare than a society without a public annuity program. While this is difficult to test directly, it suggests the importance of large complementarities in public spending on old-age annuities and old-age healthcare.

III.D. The effect of bequest motives on the value of longevity

A bequest motive encourages an individual to delay consumption because money saved for consumption in old age also has the added benefit of increasing bequests in the event of death. Its effects on consumption and the value of longevity are therefore similar to that of increased annuitization. All else equal, a stronger bequest motive dampens the effect of health shocks on the value of longevity.

These points are illustrated in Figure 10, which replicates the exercise from Figure 8 but also incorporates a bequest motive.[15] Comparing the two figures reveals that individuals with a bequest motive have a lower value of a life-year at age 50. This happens because the bequest motive causes individuals to postpone consumption at this age. Notice also that health shocks have a smaller effect on values when bequest motives are present. A stroke patient in the baseline model values a life-year by $130,000 more than a healthy individual at the onset of disease (Figure 8), but by only $50,000 more when a bequest motive is present (Figure 10). In both cases the value of a life-year decays with age following the onset of illness, but the drop-off is less significant in the presence of a bequest motive.

This implies that the absence of annuity markets impacts behavior the most when bequest motives are absent. Since bequests are much more common among the wealthiest consumers (Hurd and Smith 2002), incomplete annuity markets are expected to make the biggest impact on the value of life for less wealthy groups.

III.E. Comparing the benefits of prevention to the benefits of treatment

While therapeutic medical treatments such as kidney transplants improve health by treating the sick, preventive medical treatments such as vaccines work by preventing the onset of disease among the healthy. Individuals can also prevent illness through non-medical means by avoiding risky behaviors like smoking. In a standard value of life model that assumes complete annuity markets, there is no reason to prefer preventive treatments to equally effective therapeutic treatments. This equivalence breaks down when annuity markets are incomplete. Heuristically, since prevention is used by healthy people facing lower mortality risk, it produces less value than therapeutic technology used by sick patients facing higher risk. In this section, we sharpen this intuition.

Prevention is necessarily an ex ante concept, but therapeutic treatments can be valued ex ante or ex post. Thus it is important to be precise when comparing the two. From an ex ante point of view, equally effective preventive and therapeutic treatments are always equivalent, regardless of whether or not annuity markets are complete, because they are both valued from the perspective of a healthy person.

---

[15] Figure 10 omits the "Annuitized (baseline)" case from Figure 8 because a bequest motive has no effect on the behavior of a fully annuitized individual with no wealth.

For example, it makes little difference to an individual whether she elects to avoid a disease by getting vaccinated when healthy or by purchasing an option to take a drug that completely cures her in the event that she falls ill.[16] More formally, it does not matter whether the $dS(t)$ term in equation (2) corresponds to a marginal reduction in the probability of dying (prevention) or a marginal increase in the probability of surviving (treatment). Another way of putting this is that there is no difference between prevention and therapy in the long run.

But as Keynes dryly noted, "in the long run, we are all dead." In the short-run, society includes adults who already suffer from diseases that currently lack adequate treatment, and these individuals necessarily value new medical treatments from an ex post perspective. Medical research policy decisions made on behalf of society should account for the current realized health status of its population rather than just a single, representative consumer who is just beginning the cycle of life.

Consider the case of stroke. Is it better to invest in preventive treatments (e.g., anti-smoking and exercise programs, or hypertension diagnosis and treatment) for the healthy or to invest in the development of new stroke treatments that increase longevity among the sick? For simplicity, consider the population of individuals over the age of 50 who are at risk of suffering, or have already suffered, a stroke at age 65. The relevant survival curves are obtained from the Future Elderly Model and displayed in Figure 11. The corresponding values for VSLY are given in Figure 12.

Figure 12 indicates that the value of a statistical life-year is large immediately following a stroke, indeed larger than the value of a statistical life-year for healthy individuals at any age. Because of diminishing marginal returns to the value of life-extension, this value eventually falls below the healthy value for the small fraction of stroke sufferers that survive long enough. This implies that therapeutic technologies used soon after stroke incidence will be worth the most to the patient – more than prevention, and also more than treatments with mortality benefits that accrue long after stroke incidence.[17]

Figure 13 illustrates this point with a simple example drawn from the data displayed in Figure 11 and Figure 12. The figure illustrates the value of five different medical technologies delivering equal discounted life expectancy gains, but at different points in the disease course. The first column depicts the value of a preventive technology that reduces mortality at all ages beginning at age 52, before the stroke occurs. The subsequent columns depict the value of four different new therapeutic technologies that treat patients who suffered a stroke at age 65. The first one only treats 65-year-olds, the second one treats 65-69-year-olds, the third one treats 65-74-year-olds, and the last one treats 65-100-year-olds. The benefit of each treatment is calibrated so that it increases discounted life expectancy by 0.5 life-years from the perspective of a 52-year-old, after accounting for the interest rate and expected survival. In all cases the therapeutic treatments are more valuable than the preventive treatment, although the difference is greatest for a therapy that is targeted at 65-year-olds who were just

---

[16] The therapeutic treatment is perhaps slightly preferable to preventative treatment because then action only needs to be taken in the event of illness, but this difference in cost is generally insignificant when compared to the costs of a major illness.

[17] Preventive technologies are also worth more at earlier ages, reflecting the usual profile of life-cycle consumption. But the reduction in value is much smaller than the corresponding reduction in the value of a therapeutic technology, as illustrated in Figure 13.

diagnosed. There is little difference in value between the last two therapeutic treatments because so few stroke patients live past age 74.

IV. CONCLUSION

The economic theory surrounding the value of life has been put to many important uses. Yet, like most theories, it suffers from a few anomalies that appear at odds with intuition, common sense, or empirical facts. We have demonstrated that several of these anomalies are easily explained without abandoning the standard framework, simply by relaxing its strong assumptions around the completeness of annuity markets. Moreover, relaxing this assumption generates a number of new predictions with implications for health policy and behavior. In particular, we show diminishing returns to life-extension. A given gain in longevity is more valuable to a consumer that has less life remaining, and vice-versa. In addition, we demonstrate an interaction between annuity policy and health policy: Completing the annuity market may significantly increase the value of life-extension, especially for the elderly.

Diminishing returns to life-extension yield a number of subsidiary predictions. First, the onset of fatality risk creates a short-term spike in the value of statistical life, but the lucky ones that survive such risks over the long-term value life-years by less. This explains the anecdotal perception of desperation among the newly diagnosed cancer patient, along with complacency among long-term survivors who may feel they are "playing with house money." Second, the value of a statistical life-year will tend to vary across types of risk, not just across types of people. It is more valuable to add one month of life for a patient facing a highly fatal disease than for one facing a much milder ailment. Third, contrary to the old saying, treatment might be more valuable than prevention, at least when the expected gain in longevity is held fixed. Many healthcare researchers decry the unwillingness of policymakers and patients to invest in preventive activities (Dranove 1998, Finkelstein and Brown 2006). Our findings suggest there may be an economic basis for this unwillingness, other than market failures, time-discounting, or myopia. Finally, public programs that expand the market for annuities might simultaneously boost the demand for life-extending technologies. Intuitively, annuities calm consumer fears about outliving their wealth and thus enable more aggressive investments in life-extension. Viewed differently, our results also show that market failures in annuities affect the value of statistical life, and thus the socially optimal level of health care spending.

Our analysis raises a number of important questions for further research. First, how does the value of longevity vary with endogenous demand for quality of life? Elsewhere, we have studied how incomplete health insurance enhances the value of medical technology that improves quality of life, because such technology acts as insurance by compressing the difference in utility between the sick and healthy states (Lakdawalla, Malani and Reif 2015). Less clear is how demands for the quantity and quality of life interact with financial market incompleteness of various kinds. Second, what does the generalized value of life model mean for the value of different kinds of medical technologies? For instance, the model suggests that short-term survival gains for high-risk diseases are more valuable than previously believed, but very long-term survival gains might actually be less valuable than previously believed. Finally, what are the implications for the empirical literature on the value of statistical life? Empirical analysis has typically proceeded under the assumption that different kinds of mortality risk are all valued the same way, as long as they imply similar changes in the probability of dying (Hirth et al. 2000, Mrozek and Taylor 2002, Viscusi and Aldy 2003). Our framework casts doubt on this assumption and suggests the

need for a more nuanced empirical approach. This missing insight may be one reason for the widely disparate estimates in the empirical literature on the value of a statistical life.

V. REFERENCES

Arthur, W. B. (1981). "The economics of risks to life." The American Economic Review: 54-64.

Banks, J., R. Blundell and S. Tanner (1998). "Is there a retirement-savings puzzle?" American Economic Review: 769-788.

Bommier, A. and B. Villeneuve (2010). Risk Aversion and the Value of Risk to Life. Zurich, CER-ETH -- Center of Economic Research at ETH Zurich.

Brown, J. R. (2001). "Private pensions, mortality risk, and the decision to annuitize." Journal of Public Economics **82**(1): 29-62.

Carroll, C. D. and L. H. Summers (1991). Consumption growth parallels income growth: Some new evidence. National saving and economic performance, University of Chicago Press**:** 305-348.

Dartmouth Atlas Group (2011). Trends and variation in end-of-life care for Medicare beneficiaries with severe chronic illness. Hanover, NH, Dartmouth Institute for Health Policy and Clinical Practice.

Davidoff, T., J. R. Brown and P. A. Diamond (2005). "Annuities and individual welfare." American Economic Review **95**(5): 1573-1590.

Delprat, G., M.-L. Leroux and P.-C. Michaud (2013). Evidence on Individual Preferences for Longevity Risk. Quebec, CIRPEE.

Dranove, D. (1998). "Is there underinvestment in R & D about prevention?" Journal of Health Economics **17**(1): 117-127.

Fernandez-Villaverde, J. and D. Krueger (2007). "Consumption over the life cycle: Facts from consumer expenditure survey data." The Review of Economics and Statistics **89**(3): 552-565.

Finkelstein, E. A. and D. S. Brown (2006). "Why does the private sector underinvest in obesity prevention and treatment." NC Med J **67**(4): 310-312.

Goldman, D. P., D. Cutler, J. W. Rowe, P.-C. Michaud, J. Sullivan, D. Peneva and S. J. Olshansky (2013). "Substantial health and economic returns from delayed aging may warrant a new focus for medical research." Health Affairs **32**(10): 1698-1705.

Goldman, D. P., P.-C. Michaud, D. Lakdawalla, Y. Zheng, A. Gailey and I. Vaynman (2010). "The fiscal consequences of trends in population health." National Tax Journal **63**(2): 307-330.

Goldman, D. P., B. Shang, J. Bhattacharya, A. M. Garber, M. Hurd, G. F. Joyce, D. N. Lakdawalla, C. Panis and P. G. Shekelle (2005). "Consequences of health trends and medical innovation for the future elderly." Health Affairs **24**: W5.

Goldman, D. P., Y. Zheng, F. Girosi, P.-C. Michaud, S. J. Olshansky, D. Cutler and J. W. Rowe (2009). "The benefits of risk factor prevention in Americans aged 51 years and older." American Journal of Public Health **99**(11): 2096.

Green, C. and K. Gerard (2009). "Exploring the social value of health-care interventions: a stated preference discrete choice experiment." Health Econ **18**(8): 951-976.

Hall, R. E. and C. I. Jones (2007). "The Value of Life and the Rise in Health Spending." The Quarterly Journal of Economics **122**(1): 39-72.

Hirth, R. A., M. E. Chernew, E. Miller, A. M. Fendrick and W. G. Weissert (2000). "Willingness to pay for a quality-adjusted life year in search of a standard." Medical Decision Making **20**(3): 332-342.

Hubbard, R. G., J. Skinner and S. P. Zeldes (1995). "Precautionary Saving and Social Insurance." Journal of Political Economy **103**(21).

Hurd, M. and J. P. Smith (2002). Expected bequests and their distribution, National Bureau of Economic Research.

Kopczuk, W. and J. P. Lupton (2007). "To leave or not to leave: The distribution of bequest motives." The Review of Economic Studies **74**(1): 207-235.

Lakdawalla, D., A. Malani and J. Reif (2015). The Insurance Value of Medical Innovation, National Bureau of Economic Research.

Lakdawalla, D. N., D. P. Goldman, P.-C. Michaud, N. Sood, R. Lempert, Z. Cong, H. de Vries and I. Gutierrez (2009). "US pharmaceutical policy in a global marketplace." Health Affairs **28**(1): w138-w150.

Lakdawalla, D. N., D. P. Goldman and B. Shang (2005). "The health and cost consequences of obesity among the future elderly." Health Affairs **24**: W5R30.

Linley, W. G. and D. A. Hughes (2013). "Societal views on NICE, Cancer Drugs Fund, and Value Based Pricing Criteria for Prioritising Medicines: A Cross-Sectionalk Survey of 4118 Adults in Great Britain." Health Economics **22**(8): 948-964.

Michaud, P.-C., D. Goldman, D. Lakdawalla, A. Gailey and Y. Zheng (2011). "Differences in health between Americans and Western Europeans: Effects on longevity and public finance." Social science & medicine **73**(2): 254-263.

Michaud, P.-C., D. P. Goldman, D. N. Lakdawalla, Y. Zheng and A. H. Gailey (2012). "The value of medical and pharmaceutical interventions for reducing obesity." Journal of health economics **31**(4): 630-643.

Mrozek, J. R. and L. O. Taylor (2002). "What determines the value of life? a meta‐analysis." Journal of Policy analysis and Management **21**(2): 253-270.

Murphy, K. M. and R. H. Topel (2006). "The Value of Health and Longevity." Journal of Political Economy **114**(5): 871-904.

Nord, E., J. Richardson, A. Street, H. Kuhse and P. Singer (1995). "Maximizing health benefits vs egalitarianism: an Australian survey of health issues." Soc Sci Med **41**(10): 1429-1437.

Philipson, T. J., G. Becker, D. Goldman and K. M. Murphy (2010). Terminal care and the value of life near its end, National Bureau of Economic Research.

Philipson, T. J. and G. S. Becker (1998). "Old-age longevity and mortality-contingent claims." Journal of Political Economy **106**(3): 551-573.

Reichling, F. and K. Smetters (2015). "Optimal Annuitization with Stochastic Mortality and Correlated Medical Costs." The American Economic Review **105**(11): 3273-3320.

Rosen, S. (1988). "The value of changes in life expectancy." Journal of Risk and Uncertainty **1**(3): 285-304.

Shepard, D. S. and R. J. Zeckhauser (1984). "Survival versus consumption." Management Science **30**(4): 423-439.

Siegel, J. J. (1992). "The real rate of interest from 1800–1990: A study of the US and the UK." Journal of Monetary Economics **29**(2): 227-252.

Virga, M. (1996). "Earn More, Live Longer-Variation in Mortality by Income Level." Pension Section News(28): 1-7.

Viscusi, W. K. (2003). "Racial differences in labor market values of a statistical life." Journal of Risk and Uncertainty **27**(3): 239-256.

Viscusi, W. K. and J. E. Aldy (2003). "The value of a statistical life: a critical review of market estimates throughout the world." Journal of risk and uncertainty **27**(1): 5-76.

Wennberg, J. E., E. S. Fisher, J. S. Skinner and K. K. Bronner (2007). "Extending the P4P agenda, part 2: how Medicare can reduce waste and improve the care of the chronically ill." Health Affairs **26**(6): 1575-1585.

Yaari, M. E. (1965). "Uncertain lifetime, life insurance, and the theory of the consumer." The Review of Economic Studies: 137-150.

VI. APPENDIX

VI.A. Theory with stochastic mortality

Proposition 1: Let $c^*(t|H)$ be the consumer's optimal consumption path conditional on being in the healthy state, H. If at any time t^* the consumer falls sick, her consumption will rise. That is, $c^*(t^*|S) > c^*(t^*|H)$.

Proof: Define the value function as

$$V_t(w_t) = \max_{c(t)} E_t \left[\int_0^\infty e^{-\rho t} S(t) u\big(c(t)\big) dt \right]$$

Because the sick state, S, is an absorbing state, we can rewrite the optimization problem as

$$V(w_0) = \begin{bmatrix} V(w_0; H) \\ V(w_0; S) \end{bmatrix} = \max_{c(t)} \begin{bmatrix} E\left[\int_0^\infty e^{-\rho t} e^{-\int_0^t \mu_\tau d\tau} u\big(c(t)\big) dt \,|\mu_0 = \mu_h \right] \\ \int_0^\infty e^{-(\rho + \mu_s)t} u\big(c(t)\big) dt \end{bmatrix}$$

where $\mu_\tau \in \{\mu_s, \mu_h\}$ is the (constant) mortality rate in state $\tau \in \{S, H\}$. Assuming CRRA utility allows one to solve for a closed-form solution for the value function in the sick state:

$$V(w; S) = \left(r + \frac{(\rho + \mu_s) - r}{\gamma} \right)^{-\gamma} \frac{w^{1-\gamma}}{1 - \gamma}$$

$$c^*(t;S) = w\left(r + \frac{(\rho + \mu_s) - r}{\gamma}\right)$$

The value function for the healthy state satisfies the following Hamilton-Jacobi-Bellman equation:

$$V(w,H) = \max_{c(t)}\{u(c(t)) - (\rho + \mu_h)V(w,H) + V_w(w,H)(rw - c(t)) \qquad (6)$$
$$+ \pi(V(w,S) - V(w,H))\}$$

The first-order condition together with the CRRA assumption implies that

$$c^*(t,H) = V_w(w,H)^{-1/\gamma}$$

The CRRA utility assumption suggests that the form of the solution for the value function is

$$V(w,H) = B\frac{w^{1-\gamma}}{1-\gamma}$$

for some constant B. Plugging these results into the Hamilton-Jacobi-Bellman equation implies that

$$0 = \frac{\gamma}{1-\gamma}B^{1-1/\gamma}w^{1-\gamma} + rBw^{1-\gamma} - (\rho + \mu_h)B\frac{1}{1-\gamma}w^{1-\gamma} + \pi\left[\left(r + \frac{(\rho + \mu_s) - r}{\gamma}\right)^{-\gamma} - B\right]\frac{w^{1-\gamma}}{1-\gamma}$$

Rearranging shows that B must satisfy the expression

$$\gamma B^{1-1/\gamma} + rB - rB\gamma - (\rho + \mu_h)B + \pi\left[\left(r + \frac{(\rho + \mu_s) - r}{\gamma}\right)^{-\gamma} - B\right] = 0 \qquad (7)$$

From the first-order condition, we know that

$$c^*(t,H) = wB^{-\frac{1}{\gamma}} = \frac{w\left(r + \frac{(\rho + \mu_s) - r}{\gamma}\right)}{\left(r + \frac{(\rho + \mu_s) - r}{\gamma}B\right)^{1/\gamma}} = \frac{c^*(t,S)}{\left(r + \frac{(\rho + \mu_s) - r}{\gamma}B\right)^{1/\gamma}}$$

All that remains is to show that $G \equiv \left(r + \frac{(\rho + \mu_s) - r}{\gamma}B\right) > 1$. Rearranging equation (7) yields

$$-\gamma B\left(r + \frac{(\rho + \mu_h) - r}{\gamma}\right)\left(r + \frac{(\rho + \mu_s) - r}{\gamma}\right)^{\gamma} + \gamma B^{1-\frac{1}{\gamma}}\left(r + \frac{(\rho + \mu_s) - r}{\gamma}\right)^{\gamma}\frac{\left(r + \frac{(\rho + \mu_s) - r}{\gamma}\right)}{\left(r + \frac{(\rho + \mu_s) - r}{\gamma}\right)} + \pi$$
$$- \pi B\left(r + \frac{(\rho + \mu_s) - r}{\gamma}\right)^{\gamma} = 0$$

Plugging in the definition of G then implies:

$$-G\left(r + \frac{(\rho + \mu_h) - r}{\gamma}\right) + G^{1-\frac{1}{\gamma}}\left(r + \frac{(\rho + \mu_s) - r}{\gamma}\right) + \frac{\pi}{\gamma}[1 - G] = 0$$

Rearranging then yields:

$$\frac{\pi}{\gamma}[G-1] = G^{1-\frac{1}{\gamma}}\left(r + \frac{(\rho + \mu_s) - r}{\gamma}\right) - G\left(r + \frac{(\rho + \mu_h) - r}{\gamma}\right)$$

Suppose by way of contradiction that $G < 1$. Then:

$$G^{1-\frac{1}{\gamma}}\left(r + \frac{(\rho + \mu_s) - r}{\gamma}\right) < G\left(r + \frac{(\rho + \mu_h) - r}{\gamma}\right)$$

Since $\mu_s > \mu_h$, this implies

$$1 \leq \frac{\left(r + \frac{(\rho + \mu_s) - r}{\gamma}\right)}{\left(r + \frac{(\rho + \mu_h) - r}{\gamma}\right)} \leq G^{\frac{1}{\gamma}}$$

But then $G > 1$, which is a contradiction.

VI.B. Empirical approach

Our calibration exercises use numerical methods to solve the following Bellman equation:

$$V_t(w_t) = \max_{\{c_t\}} u(c_t) + \frac{1 - q_t}{1 + \rho} V_{t+1}(w_{t+1})$$

Because the problem is finite, we can work backwards from the final period. We discretize the state space into $N_w = 2{,}000$ points evenly distributed across the interval $[0, w_{max}]$. Let that set of values be $\{w_n\}$. Define $g_t(w_t) = w_{t+1}$ as a mapping from the current wealth state, w_t, to the optimal wealth state in the following period, w_{t+1}.

It is clear that the consumer should consume all her wealth in the final period, i.e., $g_T(w_T) = 0$ for all $w_T \in \{w_n\}$. This implies that $V_T(w_T) = u(w_T + m_T)$ for all $w_T \in \{w_n\}$.

Next, we calculate $V_{T-1}(w_{T-1}) = \max_{g(w_{T-1}) = w_T} u(w_{T-1} + m_{T-1} - w_T/(1 + r)) + \frac{1 - q_{t+1}}{1 + r} V_T(w_T)$. In other words, for each $w_{T-1} \in \{w_n\}$, we calculate the optimal $V_{T-1}(w_{T-1})$ by determining which choice of $g_{T-1}(w_{T-1}) = w_T \in \{w_n\}$ will maximize utility. This algorithm is then repeated for $t = T - 2, T - 3, \ldots 1$.

Given the initial condition, w_1, we can then employ our results to calculate $w_2 = g_1(w_1)$, $w_3 = g_2(w_2), \ldots w_T$. Period consumption, c_t, is then calculated using the equation for the budget constraint.

When accounting for a bequest motive, we follow Kopczuk and Lupton (2007) and assume the utility from leaving a bequest is linear in wealth:

$$V_t(w_t) = \max_{\{c_t\}} u(c_t) + \frac{1}{1 + \rho}[(1 - q_t)V_{t+1}(w_{t+1}) + q_t \alpha w_{t+1}]$$

Kopczuk and Lupton (2007) estimate that the constant $\alpha^{-\gamma}$ is approximately equal to \$50,000, where γ is the coefficient of relative risk aversion from a CRRA utility function. We adopt their estimate in our calibration exercise when accounting for a bequest motive. This parameterization implies that the marginal utility of consumption exceeds (is less than) the marginal utility of leaving a bequest when consumption in the last year of life is less than (more than) \$50,000.

VII. TABLES AND FIGURES

Figure 1. Annual consumption for consumer with $120,000 in wealth and a life expectancy of 3, 4, or 5 years


Notes: This illustrative example assumes there is no uncertainty in mortality and the consumer discount rate is equal to the interest rate, so that the optimal consumption profile is flat. Increasing life expectancy from 3 to 4 years reduces annual consumption by $10,000. Increasing it from 4 to 5 years reduces annual consumption by $6,000.

Figure 2. Effect of a negative health shock on life-cycle consumption when the consumer does not have an annuity
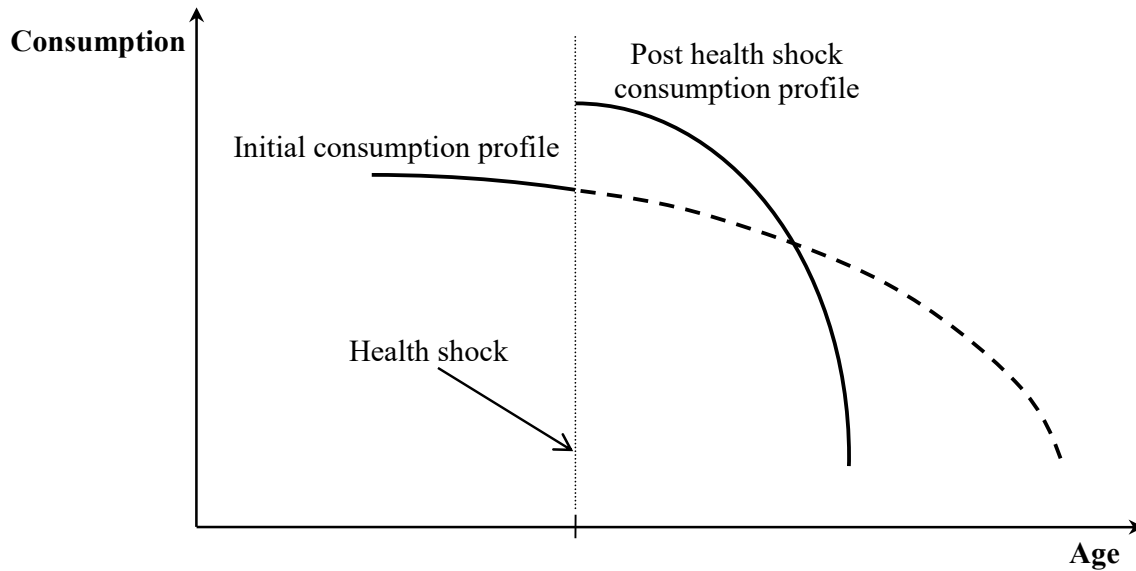
Notes: The health shock increases annual mortality or, equivalently, reduces expected survival. This causes the consumer to "spend down" her wealth.

Figure 3. Life-cycle profiles of consumption and value of statistical life-year, by annuity status
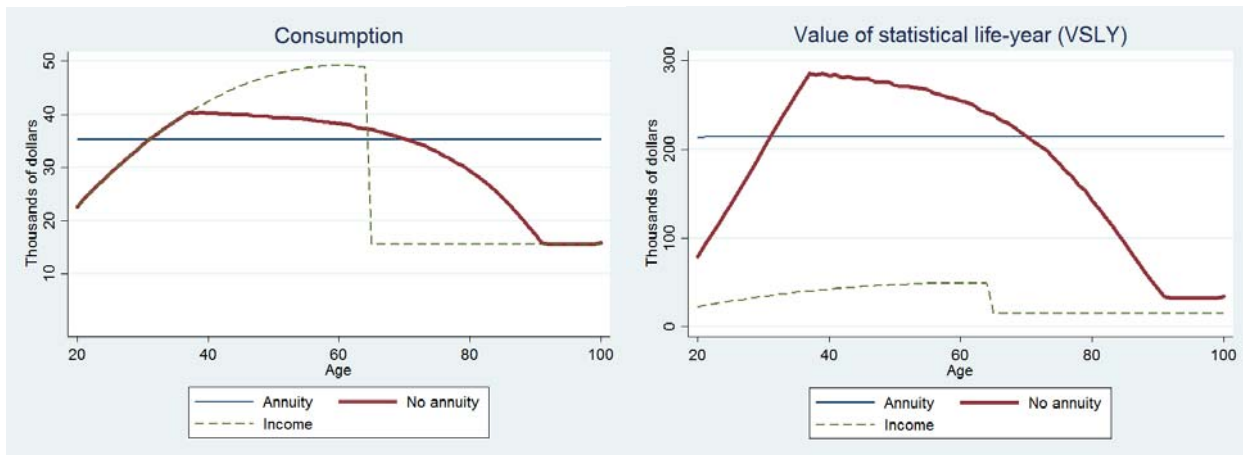
Notes: Figures plot results from a life-cycle calibration exercise. See Section III.A for parameters and additional details.

Figure 4. Value of statistical life (VSL), by age and annuity status

Notes: Figure plots results from a life-cycle calibration exercise. See Section III.A for parameters and additional details.

Figure 5. Survival rates for five different health conditions
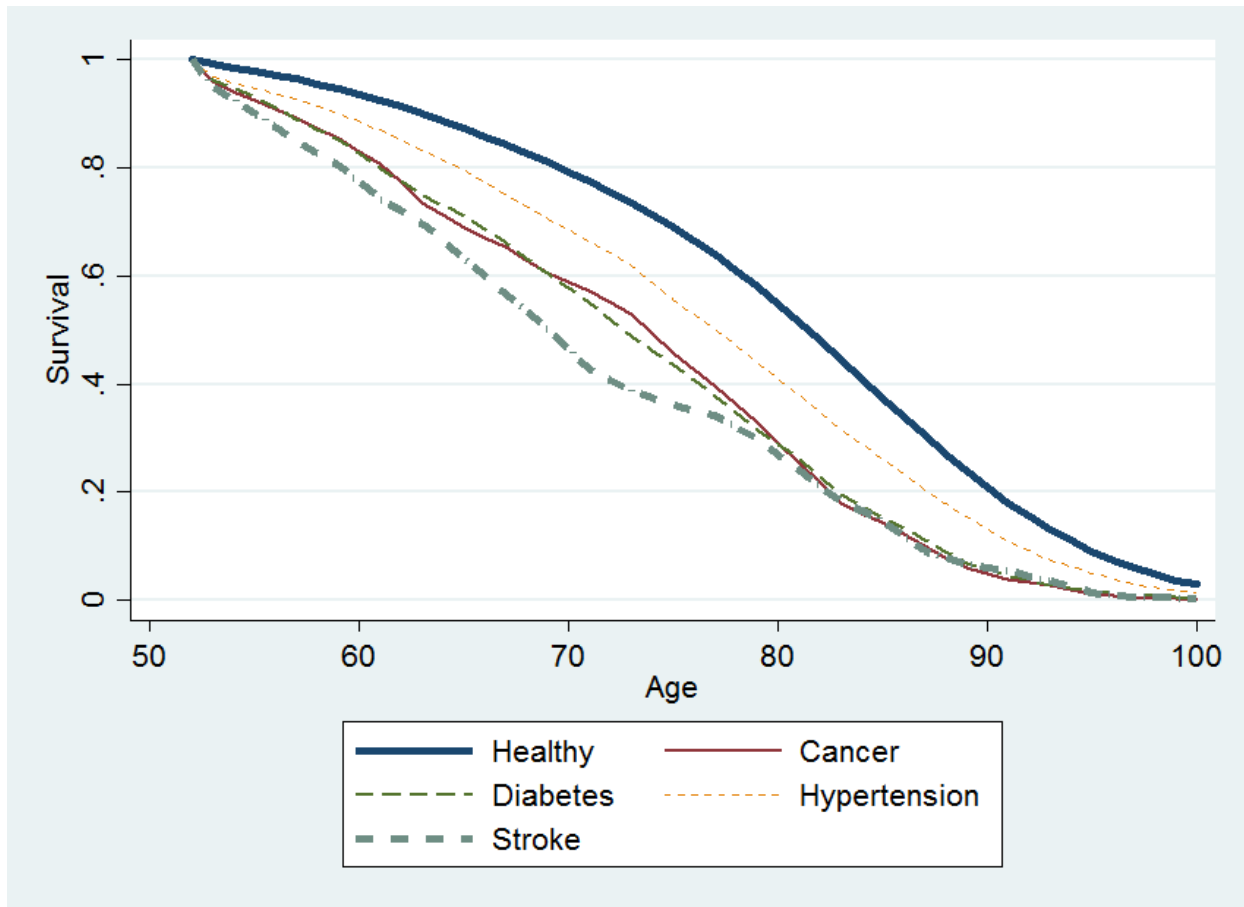
Notes: Survival rate estimates are derived from the Future Elderly Model (see main text for cites). They are conditional on surviving to age 50 and assume that disease incidence occurs at age 50.
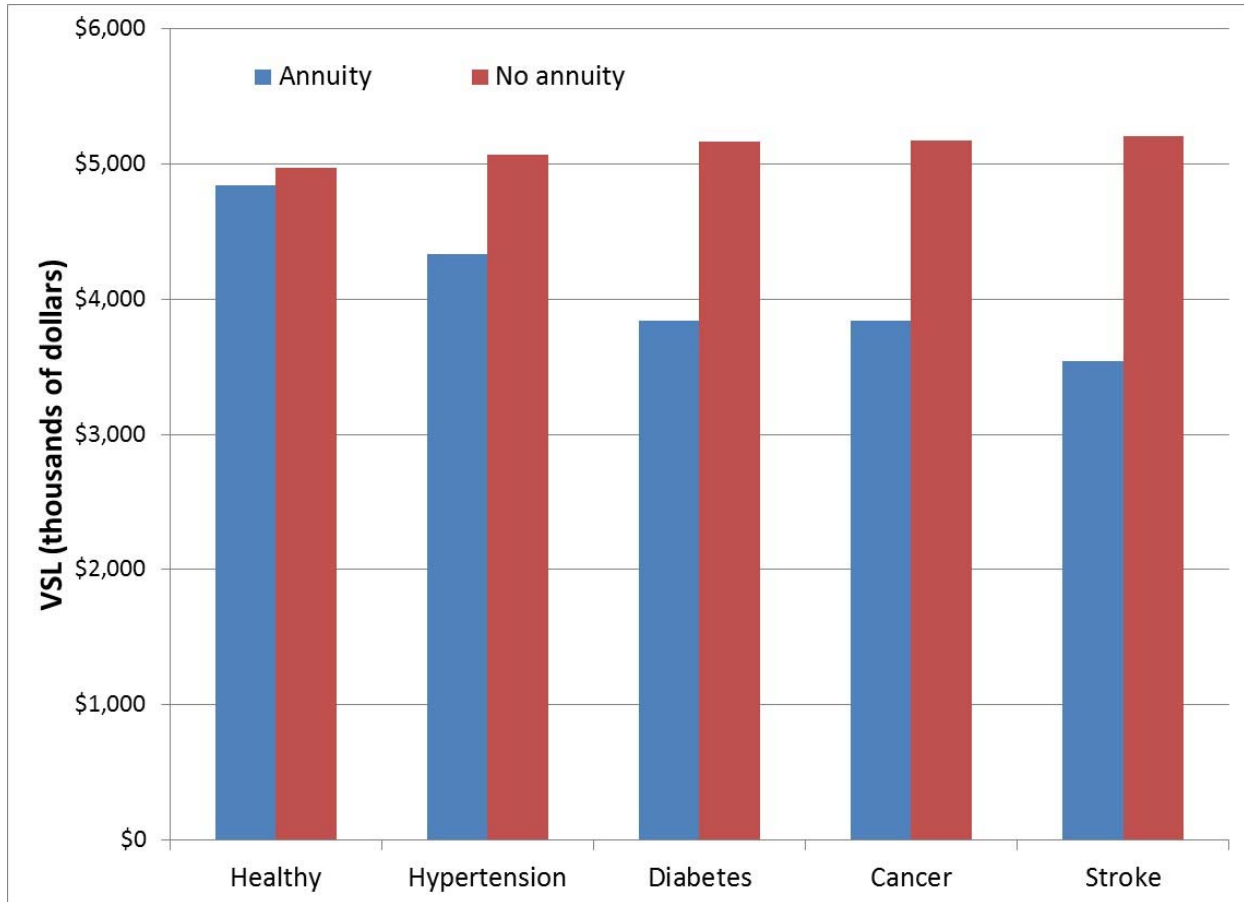
Figure 6. Value of statistical life-year (VSLY), by disease and age

Notes: Figure plots results from a life-cycle calibration exercise using the survival data displayed in Figure 5. The model assumes consumers lack access to annuity markets, except in the baseline "Healthy (annuity)" case.

Figure 7. Value of statistical life (VSL), by disease and age
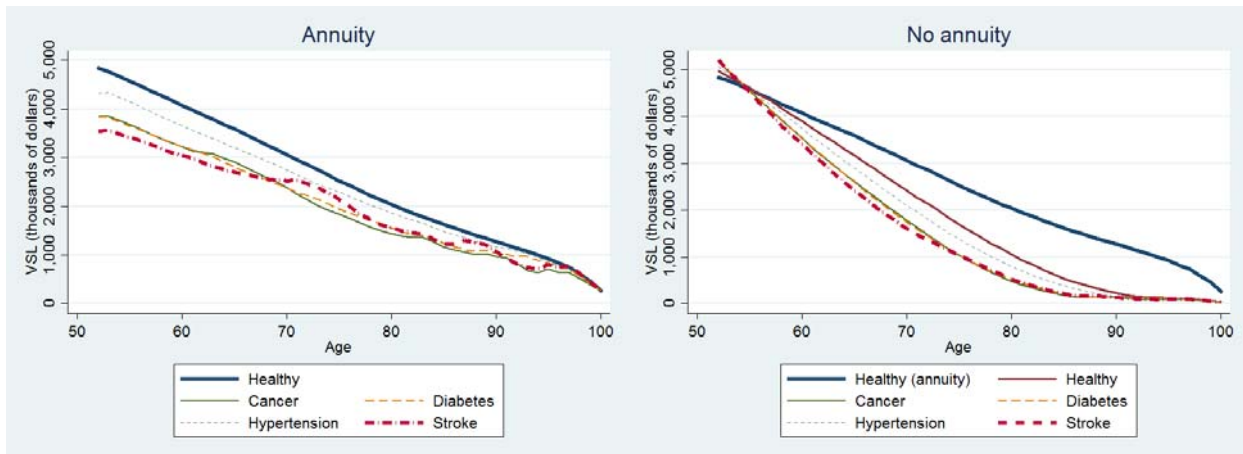
Notes: Figures plot results from a life-cycle calibration exercise using the survival data displayed in Figure 5. In the left panel, consumers have access to annuity markets. In the right panel, consumers do not have access to annuity markets, except in the baseline "Healthy (annuity)" case.

**Figure 8. Value of statistical life-year (VSLY) at onset of disease and 5 years later, when annuity markets are absent**


Notes: Onset of disease occurs at age 50. Values are calculated using the life-cycle model described in the main text and the survival data displayed in Figure 5.

32
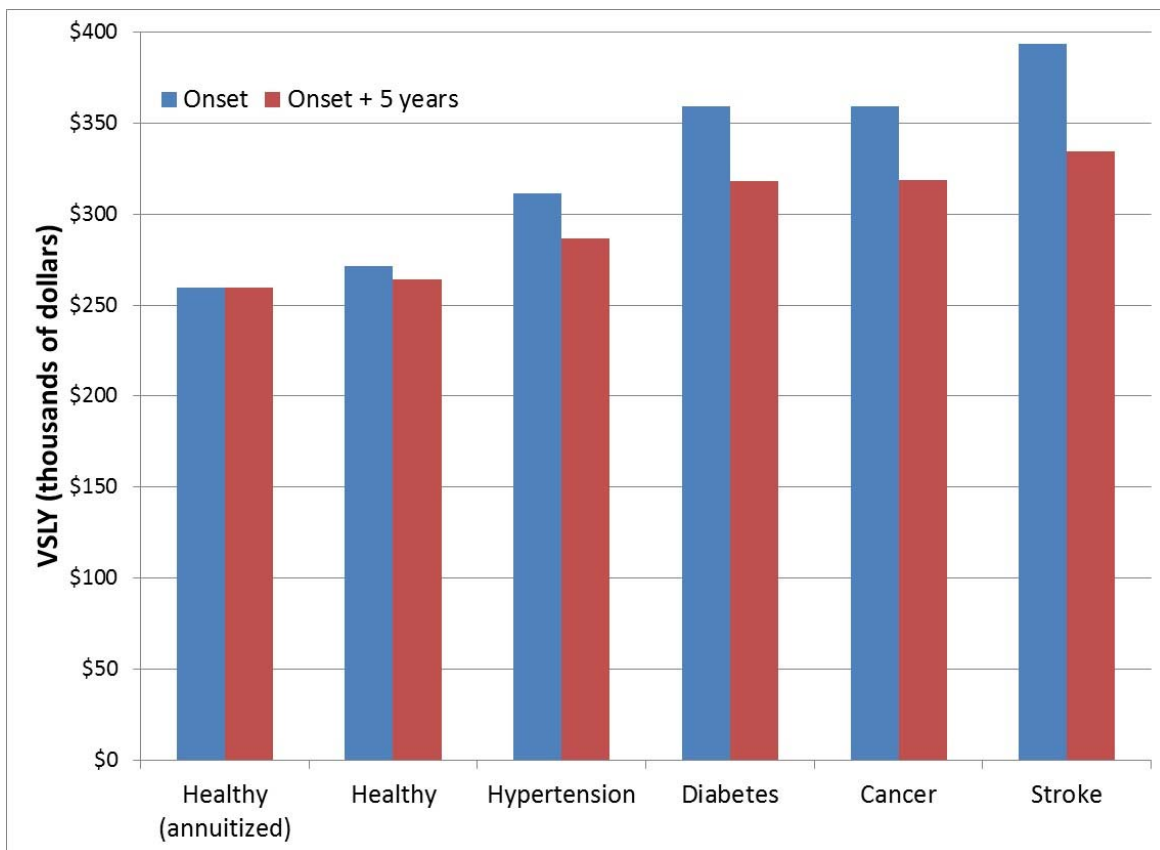
Figure 9. Value of statistical life (VSL) at age 50, by disease and annuitization status

Notes: Onset of disease occurs at age 50. Values are calculated using the life-cycle model described in the main text and the survival data displayed in Figure 5.

Figure 10. Value of statistical life-year (VSLY) at onset of disease and 5 years later, when annuity markets are absent but a bequest motive is present


Notes: This figure employs the same data used to create Figure 8, but estimates the value of a life-year using a utility function that incorporates a bequest motive (see the mathematical appendix for details). Onset of disease occurs at age 50. Values are calculated using the life-cycle model described in the main text and the survival data displayed in Figure 5.

Figure 11. Survival as a function of whether or not individual suffers a stroke at age 65
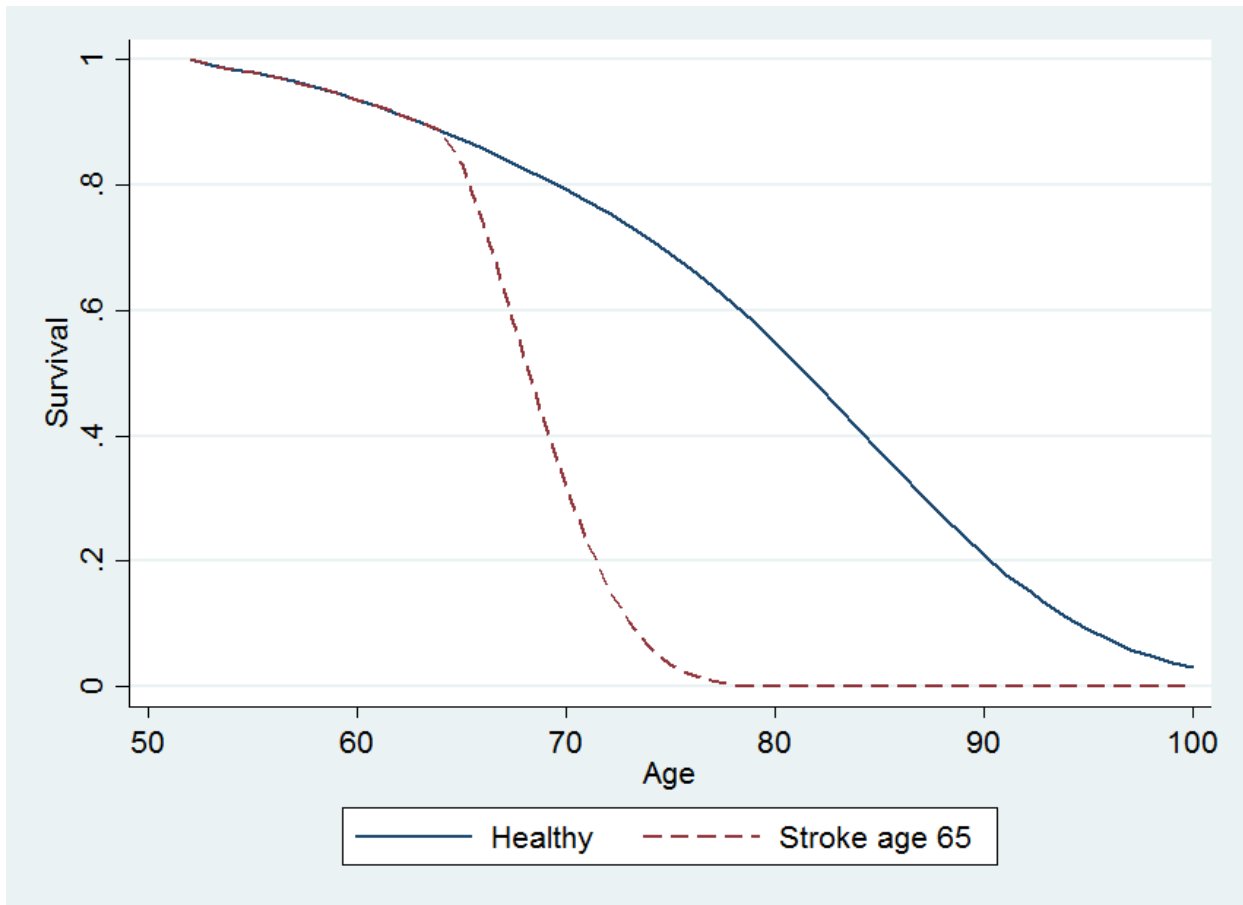

Notes: Survival rate data come from the Future Elderly Model (see main text for cites). They are estimated conditional on surviving to age 50. Data for stroke patients are estimated assuming stroke occurs at age 65.

Figure 12. Value of statistical life-year (VSLY), by disease and age
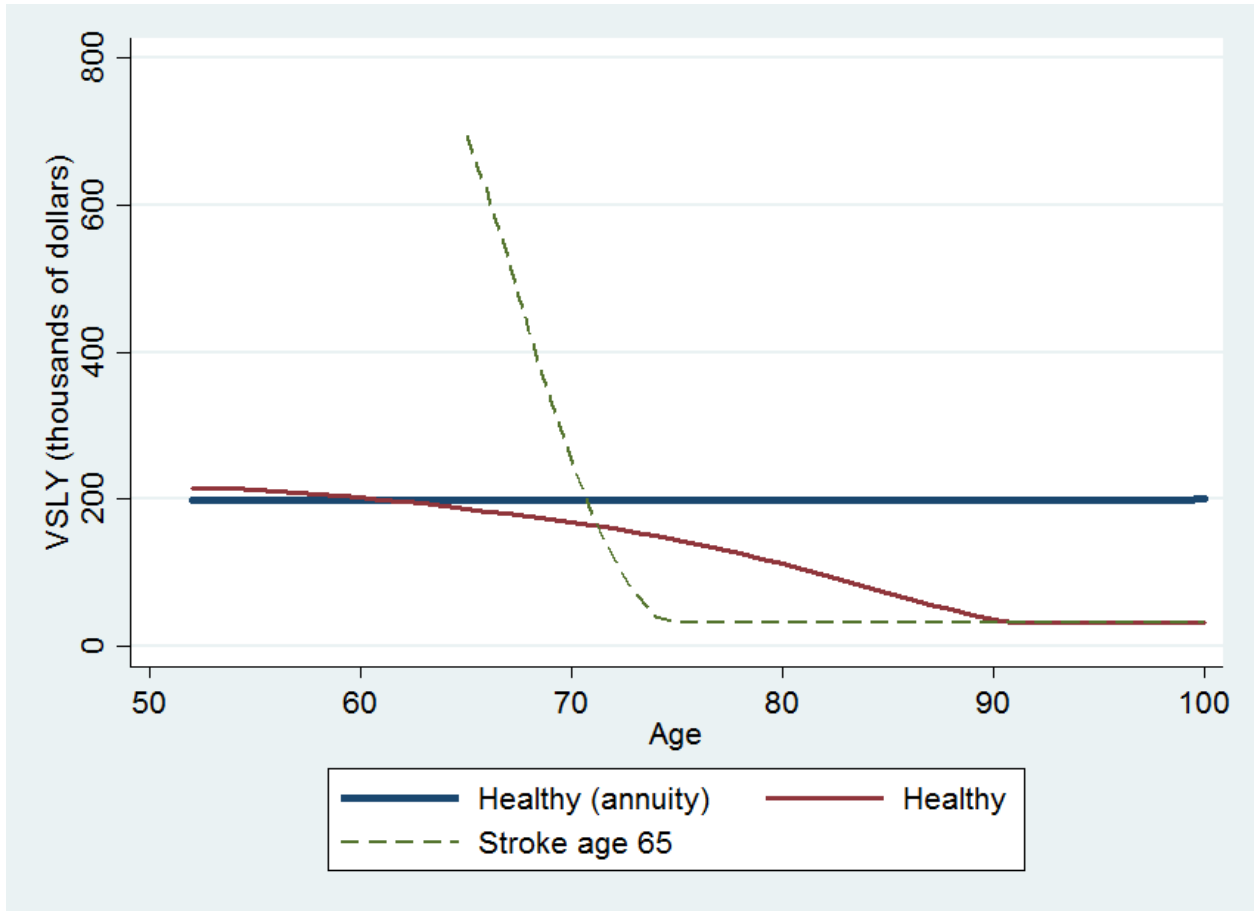
Notes: Figure plots results from a life-cycle calibration exercise using the survival data displayed in Figure 11. The model assumes consumers lack access to annuity markets, except in the baseline "Healthy (annuity)" case.

Figure 13. Value of different medical technologies that deliver 0.5 discounted life-years to the individual
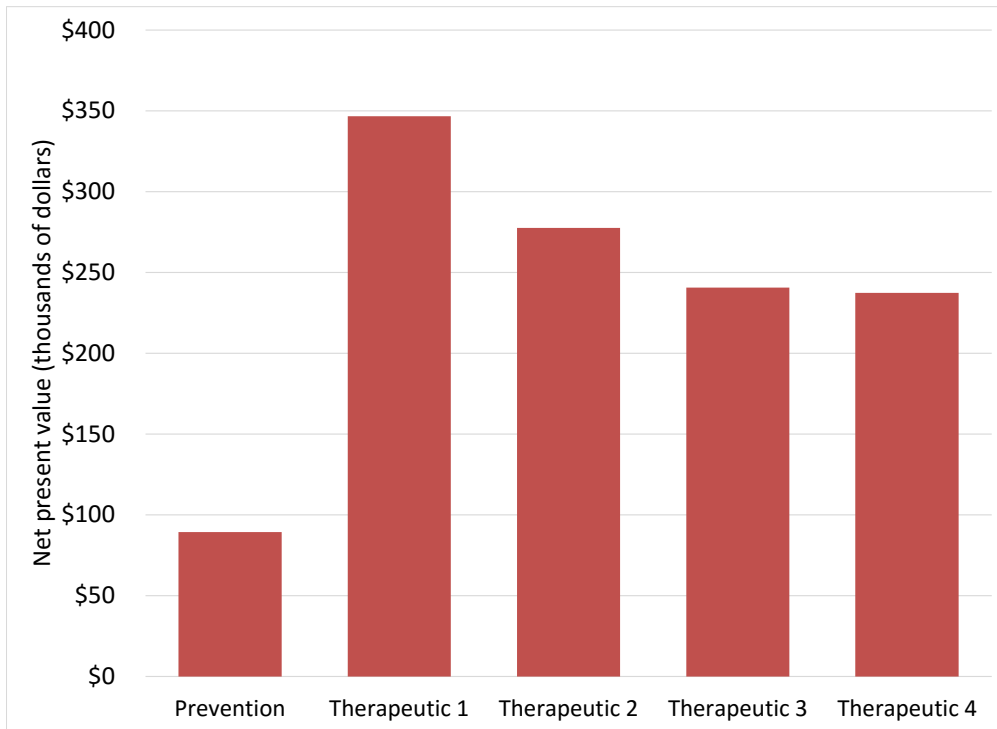
Notes: Values are calculated from the perspective of a 50-year-old. Preventive treatment reduces mortality by a constant amount at all ages. Therapeutic treatments 1-4 reduce mortality by a constant amount following a stroke at age 65 for one year, five years, ten years, and thirty-five years, respectively.