California Center for Population Research
University of California - Los Angeles

Valuing Health-Risk Reductions: Sick-Years, Lost Life-Years, and Latency

Trudy Ann Cameron
J.R. DeShazo

CCPR-053-05

October 2005

California Center for Population Research
On-Line Working Paper Series

Valuing Health-Risk Reductions: Sick-years, Lost Life-Years, and Latency

Trudy Ann Cameron and J.R. DeShazo[1,2]

October, 2005

Corresponding author:
Trudy Ann Cameron
Department of Economics, 435 PLC
1285 University of Oregon
Eugene, OR 97403-1285
cameron@uoregon.edu
(541) 346-1242 (ph)
(541) 346-1243 (fax)

JEL Classifications: I1, H43, Q51

Keywords: Value of a Statistical Life (VSL), mortality risk, morbidity risk, health, option price

[1]Senior authorship is not assigned. Nominal lead authorship will rotate though this paper series.

[2]Department of Economics, University of Oregon, and School of Public Affairs, UCLA, respectively. We thank Kip Viscusi, Vic Adamowicz, Richard Carson, Maureen Cropper, Baruch Fischhoff, Jim Hammitt, Alan Krupnick, and V. Kerry Smith for helpful comments. Rick Li implemented our survey very capably. Ryan Bosworth, Graham Crawford, Dan Burghart and Ian McConnaha have provided able assistance. This research has been supported by the US Environmental Protection Agency (R829485, Program Director Will Wheeler). It was was also encouraged by Paul De Civita and has been supported in part by Health Canada (Contract H5431-010041/001/SS). This work has not yet been formally reviewed by either agency. Any remaining errors are our own.

Valuing Health-Risk Reductions: Sick-years, Lost Life-Years, and Latency

Abstract

Many ex ante benefit-cost analyses of environmental, health, or safety interventions hinge upon the estimate used for the benefits of human health risk reductions. The standard approach to estimating these benefits is based on a single-period single-risk model, typically used to produce a single-valued estimate for the Value of a Statistical Life (VSL). We develop instead a utility-theoretic model in which individuals choose among alternative programs to reduce their risk of experiencing future years of illness and/or lost life-years. Our model is able to produce separate estimates of the marginal utilities of both avoided sick-years and avoided lost life-years as well as estimates of the willingness to pay to avoid wide range of arbitrary adverse health profiles over an individual's future life. Such benefits estimates are particularly where costs must be incurred now to reduce health risks that will not fully materialize until much later. We evaluate our model using data from an extensive nationally representative survey that contains a set of randomized choice experiments.

JEL Classifications: I1, H43, Q51

Keywords: Value of a Statistical Life (VSL), mortality risk, morbidity risk, health, option price

1 Introduction

Over their lifetimes, individuals face a portfolio of distinct health risks such as heart disease, accidents, cancers, strokes, respiratory disease and many others. Individuals and policymakers may mitigate these risks through expenditures on privately available preventative care and medical therapies and publicly provided environmental, safety and health programs. The consumer's problem is to optimally allocate expenditures to each risk-mitigating program for each future year of their life. An important dimension of this problem is that the severity of each health risk will vary over an individual's lifespan. Furthermore, the majority of these risk-mitigating programs involve multiple periods of costs and yield uncertain future benefits. How to accommodate long latencies in health risks has been a particular challenge in the benefit-cost analysis of many public programs.

In empirical analyses, researchers tend to simplify the consumer's problem to render it more tractable. It is common to estimate wage-risk (or wealth-risk) trade-offs by assuming the individual considers a single health risk in the current period (Dreze, 1962; Jones-Lee, 1974). These models do not incorporate multiple risks or inter-temporal decision-making under uncertainty. Traditional single-risk, single-period models have motivated hundreds of empirical demand analyses, including many of those currently used to evaluate the social benefits of life-saving public policies.

Much of this literature focuses on deriving a single one-size-fits-all measure of the Value of a Statistical Life (VSL), a construct that measures an assumed-constant marginal rate of substitution between mortality risk reductions and income, scaled up to a 1.00 change in risk. A VSL is intended to measure aggregate willingness to pay (WTP) for very small changes in mortality risks experienced by large numbers of people. The US Environmental Protection Agency, for example, currently relies upon a VSL of about \$6 million for its ex ante benefit-cost analyses for environmental regulations.[3] In contrast, the Department of

[3]This number is a central tendency among a set of 26 defensible studies. Five of these studies employ stated preference data (as does the present study) and the rest are based on revealed preference data (Viscusi, 1993).

Transportation uses a VSL on the order of \$1 million for road safety measures.

Our model generalizes the traditional single-risk, single-period VSL model in several ways. Parameter estimates for the marginal utilities of avoided degraded health-state years depend explicitly upon the latency of the program benefits, the stream of program costs, and the individual's discount rate and future income. We allow individuals to substitute mitgation expenditures across health risks with different time profiles, since omitting relevant substitutes from the individual's choice set may bias estimates of marginal utility (Rosen, 1988; Dow et al., 1999).[4] Furthermore, to assume that the individual's allocation of risk-mitigation expenditures is a one-period problem, when in fact it is a multi-period problem, has the potential to yield distorted estimates, so our model explicitly addresses the multi-period nature of the problem and the relevance of disease latencies and discounting.[5]

Most importantly, our model permits us to estimate both the marginal utility of avoiding a future year of morbidity and the marginal utility of avoiding a lost life-year. Most actual programs do not "save" lives; rather they extend life by deferring the future onset of morbidity or the event of death. We introduce the concept of the value of statistical illness profile ($VSIP$) which represents the individual's marginal rate of substitution between income, morbidity risks, and mortality risks. Both policymakers and scholars have long sought a tractable and theoretically consistent empirical model that describes how the marginal value of avoiding a year of morbidity or a lost life-year varies across an individual's remaining lifespan.[6] To the best of our knowledge, we provide the first such basic estimates in the literature.

Our theoretical choice model assumes that individuals face a portfolio of health risks, each

[4] Recently, scholars have sought to allow for substitution between pairs of risks (Liu and Hammitt, 2003).

[5] One might argue that hedonic wage studies are exempt from this critique, since wage contracts may be interpreted as one-period contracts. However, when choosing across occupations, individuals may, in effect, choose across time-paths of risk-wage premia that implicitly embody inter-temporal substitution of risk mitigation (Aldy and Viscusi, 2003).

[6] Several popular per-year estimates have been developed to meet this need, but none of these describes the marginal utility of avoiding a year of morbidity or premature death. For Quality-Adjusted Life Years (QALYs) see Gold et al., (1996) and for the Value of a Statistical Life Year (VSLY) , see Moore and Viscusi (1988).

with a different time profile of future health states. The consumer's problem is to choose the set of risk mitigation programs that maximizes the present discounted value of the expected utility that they anticipate deriving from their remaining lifespan, subject to the usual income constraint and stated program prices, and employing some rate of discount. This is explicitly a multi-period model of demand, based on the individual's indirect utility function in each future year of life. For consumers, the objects of choice in this study are patterns of health-states in future years. To accommodate the probabilistic nature of health states and program benefits, we recast the traditional model in an option price framework, where we use the term "option price" in the sense of Graham (1981), rather than the sense commonly used in the finance literature. In the context of benefit-cost analysis under uncertainty, an option price is the theoretically appropriate ex ante measure of WTP, before the uncertainty about possible future health states is resolved.

Ideally, we would estimate our model of demand for risk mitigating programs using market data. However, revealed-preference data that would cleanly identify the subjective tradeoffs that we wish to quantify do not exist. Thus, we chose to administer a nationally representative survey that elicits individuals' choices over alternative risk-mitigation programs in a stated-preference experiment. Stated-preference methods are widely used in the transportation, marketing, and environmental economics literatures.[7] Each health risk in our study is presented as an illness profile that describes a probabilistic time pattern of health states that the individual could experience. Based upon each individual's gender and current age, each health profile consists of randomly assigned values for the individual's future age at the time of onset, the severity and duration of treatments and morbidity, the age at recovery (if recovery occurs), and the number of lost life-years (if any).

We present respondents with an illness-specific health-risk reduction program that involves diagnostic screening, remedial medications and life-style changes that would reduce their probability of experiencing that illness profile. Individuals must pay an annual fee to

[7] Manski (2004) encounters a similar paucity of data concerning consumer expectation and likewise resorts to eliciting this critical information by using a consumer survey.

participate in each risk-reducing program. They are asked to choose one of two risk reducing programs (each associated with a different illness profile) or to reject both programs. An advantage of this choice setting is that the individual faces a portfolio of health risks that resemble those they actually face. Through their choices, individuals reveal trade-offs across specific illnesses and a full continuum of health states of different durations. We also observe them strategically allocating expenditures for risk mitigating programs across the current year and future years of their remaining lifespan.

To analyze individuals' program choices, we estimate a modified translog indirect utility function using over 7,500 choices made by respondents to a representative national survey of over 2,400 U.S. citizens. Our estimated model permits estimates of the marginal utility of avoiding a year spent in each of three health states: morbidity, a post-morbidity recovered state, and premature mortality. Two important innovations in our specification are that it (a) accomodates dependence of the marginal utility of years of morbidity upon years of premature mortality, and (b) controls for the correlation between an individuals' age and the set of health outcomes they potentially face over their remaining lifespan.

To illustrate the implications of our model, we estimate individuals' willingness to pay to avoid five archetypical illness profiles. This exercise demonstrates how our model generalizes the concept of a VSL. Rather than capturing only the special case of mortality in the current year, we show how our model can value the discounted benefits of avoiding a wide variety of statistical illness profiles across an individual's remaining lifespan. Finally, by way of sensitivity analysis, we illustrate how different conditions–with respect to (1) discount rates, (2) income levels, (3) the respondent's age, and (4) the latency of program benefits–affect the marginal value of risk mitigation and, in turn, the value of avoiding different types of statistical illness profiles. The structural model in Section 2 shows that the conventional concept of a VSL is essentially a special case of our more-general Value of Statistical Illness Profile ($VSIP$). We outline our survey methods in Section 3, and our model's parameter estimates in Section 4, along with some sensitivity analyses. Section 5 concludes.

2 Survey Methods and Data

Market data that adequately illustrate how individuals allocate risk mitigation expenditures across competing risks and across their remaining years of life are not available.[8] Therefore, we have conducted a survey of over 2,400 randomly chosen adults in the United States. The centerpiece of the survey is a conjoint choice experiment that presents individuals with specific illness profiles and programs to mitigate these illness risks.

The development of this survey instrument involved 36 cognitive interviews, three pretests (n=100 each) and an unusually large pilot study (n=1,400).[9] Knowledge Networks administered the final version of the demand survey and the health-profile survey to a sample of 2,439 of their panelists.[10,11] Our response rate for those panelists contacted was 79 percent. (See our discussion of sample selection correction techniques below.)

We designed this survey to ameliorate several limitations of existing risk valuation methods. First, many studies have focused on non-representative sub-populations (e.g., working age men) while our sample is of the general population of men and women, including a wide range of ethnicities, age groups, and income groups. Second, many studies focus upon only mortality risks from one source, often ignoring indivduals' marginal rates of substitution between morbdity and mortality states. Furthermore, many of these stated preferences studies focus on only one, or at most two, risk reduction(s). To enhance representativeness of the $VSIP$, we assess the most common health risks over a wide range of risk reductions.

[8] Most market data characterize at best only one source of risk (e.g. hedonic wage data) and are often missing essential variables such as the baseline risk, risk reduction, the latency of the programs or the costs of programs. For example, using the Health and Retirement Survey, Picone, Sloan and Taylor (2004) explored how time preferences, expected longevity and other demand shifters affect women's propensities to get mammograms or pap-smears and to conduct regular breast self-exams. However, missing data on program costs, baseline risks, and latency of program benefits prevented a fuller demand analysis.

[9] We thank Vic Adamowicz, Richard Carson, Maureen Cropper, Baruch Fischhoff, Jim Hammitt, Alan Krupnick, and V. Kerry Smith for their careful reviews of the second of four versions of this instrument.

[10] Panelists are recruited in the Knowledge Network sample using standard RDD techniques. Recruits without home computers are equipped with WebTV technology that enables them also to receive and answer our web-based surveys. More information about Knowledge Networks is available from their website: www.knowlegdenetworks.com.

[11] Respondents were paid 10 dollars for completing our survey, in addition to the usual benefits of Knowledge Networks panel membership.

Third, the results of many revealed and stated preference studies may be subject to a biases because they omit relevant substitute risks and mitigating programs from the individual's choice set. In contrast, we strive to establish in the individual's mind a more complete health risk decision environment before valuing the reduction of any one given risk.

Here, we review the structure of the survey only briefly.[12] The first module evaluates the individual's subjective risk assessments for the major illnesses they face, their familiarity with each illness, and any current mitigating and averting behavior they may undertake. The second module consists of an extensive tutorial that introduces individuals to the idea of an illness profile and programs that may manage these illness-specific risks and prepares them for the information to be presented in the upcoming choice sets, attribute by attribute. The attribute levels in the tutorial are identical to those used in the first choice set. As indicated in Table 1, these illnesses include prostate cancer, breast cancer, colon cancer, skin cancer, lung cancer, heart disease (i.e., heart attack, angina), stroke (e.g., blood clot, aneurysm), respiratory diseases (i.e., asthma, bronchitis, emphysema) as well as diabetes and Alzheimer's.[13]

Each illness profile is a description of a time sequence of health states associated with a major illness that the individual is described as facing with some probability over the course of his or her lifetime. The attributes of the illness profiles are randomly varied, subject to a few plausibility constraints for each illness type.[14] We summarize the key attribute levels employed in our choice set in Table 1, including the frequencies with which each of the twelve randomly assigned illness names appeared in the choice sets. Up to eleven attributes characterize each illness profile and program, although we concentrate on just the main

[12] An annotated version of the survey is available at
http://darkwing.uoregon.edu/~cameron/vsl/Annotated_survey_DeShazo_Cameron.pdf

[13] There is also an adaptation for traffic accidents.

[14] Each illness was randomly assigned a particular name, although we then took great care to avoid having individuals reject the scenario because it was completely implausible (e.g., one does not recover from Alzheimer's or die suddenly from diabetes). In this paper, we rely on the extensive randomization of this assignment to minimize omitted variables bias in the specifications we consider here.. Controlling for illness names would of course reduce the error variances in the model. We explore the systematic effects of illness names in a separate paper.

attributes in this paper.[15] In other work, we explore heterogeneity in the marginal utilities associated with future health states according to the illness names and specific symptoms and treatments described in the choice scenarios. Here, we assume homogeneity in preferences along these dimensions and focus primarily on the timing and duration in each health state. Given that these other attributes were randomly assigned, the preference parameters we estimate here can be viewed as the averages across the types of major health threats covered by our study. In terms of the number and type of attributes, our design is comparable to existing state-of-the-art health valuation studies (Viscusi et al., 1991; O'Connor and Blomquist, 1997; Sloan et al., 1998; Johnson, et al., 2000). It should also be made clear that we seek to estimate demand conditional on the individual's ex ante information set about each health risk.[16] Appendix A provides one example of a choice set from the primary survey instrument.

After presenting an illness profile, we next explain to individuals that they could purchase a new program that would be coming on the market that would reduce their risk of experiencing specific illnesses over current and future periods of their life. These programs are described as involving annual diagnostic testing and, if needed, associated drug therapies and recommended life-style changes. We choose this class of interventions because pretesting showed that individuals view this combination of programs (diagnostic tests, followed by drug therapies) as feasible, potentially effective and familiar for a wide range of illnesses.[17] The effectiveness of these programs is described in four ways: 1) graphically, with a risk grid, 2) in terms of risk probabilities, 3) in terms of measures of relative risk reduction across the two illness profiles and 4) as a qualitative textual description of the risk reductions (Corso

[15] These illness profiles included the illness name, the age of onset, medical treatments, duration and level of pain and disability, and a description of the outcome of the illness. Our selection of these attributes was guided by a focus on those attributes that 1) most affected the utility of individuals and 2) spanned all the illnesses that individuals evaluated (Moxey et al. 2003).

[16] Prior to the choice experiments, we ask individuals questions about their subjective assessment of: 1) various background environmental risks, 2) their risk of each illness, 3) their personal experience with illness, and 4) the experience of friends and family with each illness.

[17] Depending upon their gender and age, individuals were familiar with comparable diagnostic tests such as mammograms, pap smears and prostate exams, or the new C-reactive protein tests for heart disease.

et al., 1999; Krupnick et al., 2000). The payment vehicle for each program is presented as a co-payment that would have to be paid by the respondent for as long as the diagnostic testing and medication are needed.[18] For the sake of concreteness we ask respondents to assume that to reap the health risk reduction offered, these payments would be needed each year for the remainder of their lifespan unless they actually experienced that illness.

The third module contains the five main choice sets, each offering the individual two programs that reduced the risk of two distinct illness profiles. We carefully explain to individuals that they can choose neither program. We also point out several possible explanations why reasonable people might choose neither program in some cases.[19] If individuals choose "neither program," we assume that they prefer their status quo illness profile to either of the two costly illness-reducing programs in each choice set.

The fourth module contains various debriefing questions that are used to document the individual's status quo health profile and to cross-check the validity of the responses (Baron and Ubel, 2002). Module five was administered separately from the choice experiment. It collects a detailed medical history of the individual, as well as household socioeconomic information.

Data gathered using any stated preference survey that is hastily developed, poorly designed, or insufficiently validated is rightfully suspect. For this study, we conducted ### in-person one-on-one interviews with subjects during the design phase, and we mounted a pre-test survey that gathered over 1000 responses. The instrument was approximately ### months in development, went through several major revisions and benefitted from expert external review. We subjected individuals' responses to an extensive set of robustness and validity checks. Due to space limitations, we merely summarize our quality assurance efforts here.

[18] Costs were expressed in both monthly and annual terms. The interventions (diagnosis and treatment regimens) were selected to be as minimally invasive (or onerous) as possible, while still remaining credible.

[19] These reasons include that they 1) cannot afford either program, 2) did not believe they faced these illness risks, 3) would rather spend the money on other things, 4) believed they would be affected by another illness first. If the individual chooses neither program, we ask them why they did so in a follow-up question.

Risk Comprehension Verification. After administering an extensive risk tutorial and presenting the risk changes in three forms (textually, graphically and mathematically), we tested individuals' risk comprehensions. This comprehension test required individuals to rank the sizes of the risk reductions associated with two risk mitigation programs. Approximately eighty percent of the individuals demonstrated adequate comprehension of the relative risk size reductions of the programs, which is a rate consistent with risk comprehension levels documented in other surveys (Alberini, et al., 2004 and Krupnick et al., 2003).[20]

Mitigating Bracketing Biases Associated with Omitted Substitutes. In contrast with many valuation studies that focus on just one risk and just one risk-mitigating program, we endeavored to reduce biases associated with bracketing (Read, et al, 1999) by ensuring that nearly all major substitute risks (and specific programs to reduce them) were included in at least one of each individuals' choice sets. Presenting a broad spectrum major health threats and mortality risks also increases the representativeness of our estimates and makes the motivation of a fuller range of illness profiles plausible, and thus, possible. Of course, a potential disadvantage of this approach is the cognitive complexity associated with the choice task, which we seek to minimize through careful survey design, and evaluate ex post.[21]

Mitigating Hypothetical Bias. At the beginning of the valuation module, we include a "cheap talk" reminder to ensure that respondents carefully consider their budget constraint and to discourage them from overstating their willingness to pay (Cummings and Taylor, 1999; List, 2001). Individuals are instructed, "In surveys like this one, people sometimes do not fully consider their future expenses. Please think about what you would have to give up to purchase one of these programs. If you choose a program with too high a price, you may not be able to afford the program when it is offered. . . ." (See the online Annotated Survey for a complete description.)

[20] We discuss the effects on the estimated parameters of including and excluding individuals from the sample based on their risk comprehension below.

[21] We assess this concern directly in the survey. After each choice set we ask individuals how difficult each choice was. On a scale of 1 to 5 (very easy to very difficult), the average response for the first choice set was 3.2. This rating fell with each subsequent choice set, suggesting that the choice task became easier with increasing familiarity.

Mitigating Bias from Provision Rules and Order effects. In order to clarify provision rules for each choice set (Taylor, et al, 2004) and to avoid potential choice set order effects (Ubel et al., 2002; de Bruin and Keren, 2003), we instructed individuals to assume that every choice is binding and to evaluate each choice set independently of the other choice sets. Our empirical analyses showed that the first four choice sets appeared largely free of order effects. Individuals did exhibit a slightly higher propensity to select a program from the last choice set, an effect that has also been demonstrated in similar settings (Bateman, et al, 2004).

Testing for the Effects of Scope on Willingness to Pay. We explore whether individual choices are sensitive to the scope of the illness profile and the scope of the risk mitigating program (Hammitt and Graham, 1999; Yeung et al., 2003). We show, using simple ad hoc conjoint choice analyses, that individuals were highly sensitive to changes in the scope or level of our central attributes. (See models 1 and 2 in Table 2.) These models evaluate the two most crucial attributes of the program, its cost and the size of the risk reduction, as well as the two most important dimensions of the illness profiles, the number of years spent in a morbid condition and the number of lost life-years.

Other Validity Checks on Willingness to Pay. We also show that individuals' willingness to pay for these programs varies with several factors as economic theory would predict it should. It rises with income, as shown in the analysis in this paper. For any given age, it rises with the expected incidence of health risks in future years (DeShazo and Cameron, 2004a). It also varies systematically as predicted by economic theory with same-illness and other-illness morbidity (DeShazo and Cameron, 2004b)

Validating the Representativeness of Our Estimating Sample. Our estimating sample is very close to being representative of the U.S. population in terms of standard demographic characteristics. Appendix B illustrates this by comparing the individuals in our estimating sample with corresponding population characteristics (e.g., age, income, and gender) from the 2000 Decennial Census. Our final estimating sample consists of 7,520 choices involving 22,560 alternatives. We arrived at this sample after cleaning the data based on two primary quality control criteria. We exclude individuals if they failed to answer correctly the simple risk comprehension question at the end of the survey's risk tutorial. We also exclude individuals if they explicitly rejected the choice scenario.[22] Sensitivity analyses with respect to these conservative data exclusion criteria are presented in Appendix C.

3 A Utility-Theoretic Choice Model

Our structural model interprets individuals' choices as revealing their option prices, in the sense of Graham (1981), for programs that mitigate the risks of adverse future health states.[23] While program choices have inter-temporal consequences, our model is one of static decision-making, with future costs and benefits converted into the appropriate present values. We focus on four distinct health states: 1) a existing pre-illness healthy state (pre), 2) an illness state (ill), 3) a post-illness "recovered" state (if the illness is non-fatal) (rcv), and 4) a lost life-year (lyl). Let i index individuals and let t index time periods.[24] Let $1\,(pre_{it})$, $1\,(ill_{it})$, $1\,(rcv_{it})$, and $1\,(lyl_{it})$ be a set of mutually exclusive and exhaustive 0,1-variables that indicate individual i's health state in time period t.[25] Let α_0, α_1, α_2, and α_3 be the undiscounted marginal utilities from one period in each health state.[26] In its simplest form, the individual's indirect utility function in period t might be specified as:

$$V_{it} = \beta f(Y_{it}) + \alpha_0 \; 1\,(pre_{it}) + \alpha_1 \; 1\,(ill_{it}) + \alpha_2 \; 1\,(rcv_{it}) + \alpha_3 \; 1\,(lyl_{it}) + \eta_{it} \tag{1}$$

[22] We excluded 2,236 choices because the respondent selected "Neither Program" and indicated as the *only* explanation, "I did not believe the programs would work." If any other (economic) reason was given, we retained the choice.

[23] Cameron (2005) employs a less-elaborate model in a similar vein to the problem of climate change mitigation programs, where costs must be incurred starting now to reduce the chance of adverse consequences many years into the future.

[24] Time is measured in years or months, as needed.

[25] Algebraically, the indicators for each health state will play a role that is equivalent to adjusting the limits of the summations used in calculating the present value of future continued good health, future intervals of illness, post-illness time, and life-years lost.

[26] We interpret the disutility of each adverse health state as equivalent to the utility associated with avoiding it.

where β is the undiscounted marginal utility of some function of current income, $f(Y_{it})$.

In our data, individuals will face choices that involve three alternatives: Program A, Program B, or neither program (labeled A, B, and N). In developing our estimating specification, however, we will describe our model in terms of just two choices: Program A versus no program (just A and N). The three-alternative case is completely analogous.

Let undiscounted indirect utility be V_{it}^{jk} for the i^{th} individual in period t, where $j = A$ if Program A is chosen and $j = N$ if the program is not chosen. The superscript k will be S (denoting "sick") if the individual suffers the illness and H (denoting "healthy") if the individual does not suffer the illness. From the perspective of a program choice made today, individuals will discount the streams of utility derived from each future health state.[27] Let the discount factor be $\delta^t = (1+r)^{-t}$, and employ it to calculate the present discounted indirect utility from these profiles of future health states, which we will denote $PDV(V_i^{jk})$.

Given the ex ante uncertainty about future health states, we need to calculate *expected utilities* to derive the individual's option price for any given program. In this case, the expectation is taken across the binary uncertain outcome of getting sick, S, or remaining healthy, H. The probability of illness or injury differs according to whether the respondent participates in the risk-reducing intervention program. Let the baseline probability of illness be Π_i^{NS} if the individual opts out of the program, and let the reduced probability be Π_i^{AS} if the individual opts to participate in the program.

Expected utility (taken across the uncertain sick (S) and healthy (H) states) differs according to whether the individual selects Program A or "no program" (N):

$$\begin{aligned} E_{S,H}\left[PDV\left(V_i^A\right)\right] &= \Pi_i^{AS} \times PDV(V_i^{AS}) + \left(1-\Pi_i^{AS}\right) \times PDV(V_i^{AH}) \qquad (2) \\ E_{S,H}\left[PDV\left(V_i^N\right)\right] &= \Pi_i^{NS} \times PDV(V_i^{NS}) + \left(1-\Pi_i^{NS}\right) \times PDV(V_i^{NH}) \end{aligned}$$

In presenting the expected utility difference formula, $\Delta E_{S,H}[PDV(V_i^A)] = E_{S,H}\left[PDV\left(V_i^A\right)\right] -$

[27] When discounting, we assume the individual uses the same discount rate, r, to discount both future money costs and health states.

$E_{S,H}\left[PDV\left(V_i^N\right)\right]$, to be discussed next, we will make use of a number of notational abbreviations. The basic discounting term to be applied to any constant ("c") stream of payments between now and the individual's nominal life expectancy, T_i, is $pdvc_i^A = \sum_{t=1}^{T_i} \delta^t$. Other discounted terms, also summed from $t = 1$ to $t = T_i$ include $pdve_i^A = \sum \delta^t 1\left(pre_{it}^A\right)$, $pdvi_i^A = \sum \delta^t 1\left(ill_{it}^A\right)$, $pdvr_i^A = \sum \delta^t 1\left(rcv_{it}^A\right)$, and $pdvl_i^A = \sum \delta^t 1\left(lyl_{it}^A\right)$. Since these four different health states are mutually exclusive and exhaustive, the constant-payment discounting term is $pdvc_i = pdve_i + pdvi_i + pdvr_i + pdvl_i$. Finally, to accommodate the time profiles for program costs ("p") over the individual's remaining lifespan, it is convenient to define an additional term, $pdvp_i^A = pdve_i^A + pdvr_i^A$.

The expected utility difference that drives the individual's choice between Program A and the "No Program" alternative can then be expressed in terms of the quantities defined above to produce the most basic version of our estimating specification. (There will be an analogous utility difference for Program B versus the "Neither Program" alternative in the three-alternative case.) Let $cterm_i^A = \left[\left(1-\Pi_i^{AS}\right) pdvc_i^A + \Pi_i^{AS} pdvp_i^A\right]$ and $yterm_i^A = \left[-pdvc_i^A + \Pi_i^{AS} pdvs_i^A + \Pi_i^{NS} pdvl_i^A\right]$. Then

$$\begin{aligned}\Delta E_{S,H}\left[PDV\left(V_i^A\right)\right] &= \beta\left\{f(Y_i - c_i^A)\ cterm_i^A + f(Y_i)\ yterm_i^A\right\} \\ &\quad + \alpha_1\left\{\Delta\Pi_i^{AS} pdvi_i^A\right\} + \alpha_2\left\{\Delta\Pi_i^{AS} pdvr_i^A\right\} + \alpha_3\left\{\Delta\Pi_i^{AS} pdvl_i^A\right\} + \varepsilon_i^A\end{aligned} \tag{3}$$

The four terms in braces in equation (3) can be constructed from the data, given specific assumptions about the discount rate and about respondents' perceptions of the time profiles of future income and program payments.[28] The basic utility parameters β, α_1, α_2, and α_3, which are the same marginal utilities appearing in equation (1), are the focus of our empirical illustration. In what follows, however, it is convenient also to abbreviate the set of terms in equation (3) that involve the discounted health states in the illness profile:

[28] The underlying complexity of the first term in braces is an artifact of the need to acknowledge different time profiles of income and program costs in the sick and healthy states. We will assume that individuals anticipate being able to sustain their current real income while they are sick, but not if they are dead. We also assume that they expect not to have to pay the costs of the program if they are either sick or dead.

$pterm_i^A = \Delta\Pi_i^{AS}\left[\alpha_1 pdvi_i^A + \alpha_2 pdvr_i^A + \alpha_3 pdvl_i^A\right].$

The Graham-type option price for the program is the maximum common certain payment that makes the individual just indifferent between paying for the program and enjoying the risk reduction, or not paying for the program and not enjoying the risk reduction. The annual option price that will make $\Delta E_{S,H}\left[PDV(V_i^A)\right]$ exactly zero, $\widehat{c_i^A}$, can be calculated as:

$$\widehat{c_i^A} = Y_i - f^{-1}\left(\frac{\beta f(Y_i) yterm_i^A + pterm_i^A + \varepsilon_i^A}{-\beta cterm_i^A}\right) \tag{4}$$

While the payment $\widehat{c_i^A}$ is the maximum *annual* payment the individual is willing to make, these payments are necessary for the rest of the individual's life, so the present value of these payments must be calculated. In this context, however, there is uncertainty over just what will constitute "the rest of the individual's life," since this may differ according to whether the individual suffers the illness. We use the expected present value of this time profile of costs:

$$E_{S,H}\left[PV(\widehat{c_i^A})\right] = cterm_i^A\left[Y_i - f^{-1}\left(\frac{\beta f(Y_i) yterm_i^A + pterm_i^A + \varepsilon_i^A}{-\beta cterm_i^A}\right)\right]$$

To convert our expected present-value option price to the "value of a statistical illness profile" ($VSIP$), we normalize arbitrarily on a 1.00 risk change by dividing this WTP by the absolute size of the risk reduction to produce $VSIP = E_{S,H}\left[PV(\widehat{c_i^A})\right] / \left|\Delta\Pi_i^A\right|$. In the special case where indirect utility is merely linear in net income (i.e. $f(Y_i) = Y_i$), the $VSIP$ is just:

$$VSIP = Y_i pdvl_i^A - \frac{\alpha_1}{\beta} pdvi_i^A - \frac{\alpha_2}{\beta} pdvr_i^A - \frac{\alpha_3}{\beta} pdvl_i^A - \frac{\varepsilon_i}{\beta\left|\Delta\Pi_i^{AS}\right|} \tag{5}$$

This linear case illustrates clearly how the $VSIP$ depends on the different marginal utilities of avoided periods of illness, post-illness status, and premature death, and on the time profiles for each of these states as embedded in the terms $pdvi_i^A$, $pdvr_i^A$, and $pdvl_i^A$, and upon the individual's own discount rate (implicit in the *pdv* terms).[29] Heterogeneity in

[29] Subsequent work will preserve individual discount rates as systematically varying parameters that depend upon respondent characteristics. In a separate subsample for our survey, we elicited choices that allow us to infer individual specific discount rates. Here, however, discount rates are presumed to be exogenous and

preferences can be accommodated by making the indirect utility parameters α_1, α_2, and α_3, and even β, depend upon other individual characteristics, notably age.[30,31]

To calculate a measure closest to the conventional VSL, one would assume death in the current year, with no period of illness or post-illness status. The terms in $pdvi_i^A$ and $pdvr_i^A$ will be zero. The remainder of the individual's nominal life expectancy would be experienced as lost life-years. If we assume that $E[\varepsilon_i] = 0$, our analog to the conventional VSL formula in the linear case will be $VSIP = (Y_i - \alpha_3/\beta)\, pdvl_i^A$. Note that the summation in the $pdvl_i^A$ term, in this case, runs from now until the end of the person's nominal life expectancy, and this interval depends upon the individual's current age. Thus, the $VSIP$ will vary with age even in a model with homogeneous preferences. The term α_3/β is the monetized disutility of a lost life-year, so the overall value of avoiding a lost life-year can be equated to this monetized avoided disutility plus the chance to enjoy future income (other consumption) during that period.

The linear-in-net-income form is simple and convenient, but in this paper we use a model that is more general in terms of the permitted relationships between indirect utility and net income. An indirect utility function that is quadratic in net income is not necessarily monotonic, but can be very flexible in terms of capturing arbitrary degrees of risk aversion with respect to income without necessitating generalized nonlinear optimization. In this case, we generalize the coefficient β to depend linearly on the value of net income, either $\beta f(Y_i) = (\beta_0 + \beta_1 Y_i)(Y_i)$ or $\beta f(Y_i - c_i^A) = (\beta_0 + \beta_1(Y_i - c_i^A))(Y_i - c_i^A)$. (This subsumes the linear model described above as a special case when $\beta_1 = 0$.) While the quadratic-

constant across individuals although our empirical analyses explores the senstivity of our results to different discount rates.

[30] For example, illness characteristics can be expected to shift the value of α_1, the marginal (dis)utility of a sick-year, and possibly the marginal utility of each period in the post-illness state, α_2, since the type of illness may connote the degree of "health" that nominal recovery from that illness actually implies. Also, the marginal utility of a lost life-year may depend upon the health state prior to death. Many of these dimensions of heterogeneity will be explored in detail in subsequent papers.

[31] The error term ε is assumed to be identically distributed across observations in a manner appropriate for conditional logit estimation. Given the transformation needed to solve for the $VSIP$, however, the error term in the $VSIP$ formula will be heteroscedastic, with smaller error variances corresponding to cases with larger absolute risk reductions, $\left|\Delta\Pi_i^{AS}\right|$.

in-income specification allows for fitted marginal utilities of income that become negative at high enough levels of income, the empirical estimates of these marginal utilities remain positive within the range of our sample, and the flexibility of this form is an appealing property.

The fitted $VSIP$s from our estimating specification, however, correspond to sets of illness attributes designed into our choice experiments, rather than those empirically associated with specific illnesses. Two things are needed for construction of predicted $VSIP$s using our present results. For the illness in question, one must have an approximate joint distribution for the illness profile (possible ages of onset, possible reductions in lifespans, and possible outcomes (recovery, sudden death, limited morbidity, chronic morbidity). Also, for the population affected by this health threat, one must have an approximate joint distribution of age, gender, and income level. This distribution may be based on expert judgment combined with exposure and epidemiological data. With these two joint distributions in hand, one needs to make a large number of random draws from this pair of joint distributions and combine the illness profiles and individual characteristics for each draw using our VSIP formulas. Across a large number of random draws, one then builds up a sampling distribution for the implied $VSIP$s. The mean of this distribution can be interpreted as our model's prediction about the average of $VSIP$s for this type of health threat affecting this particular population. The overall $VSIP$ distribution, estimated in this fashion and calculated for a given policy by simulation methods, allows the researcher to more fully capture the policy choice context for the risk in question.

4 Empirical Analysis

We begin this analysis by estimating marginal utilities of years spent in each of three health states: morbidity, a post-illness state, and a lost life-year. We also explore interactions between years of morbidity and lost life years in order to assess the assumption of additive separability that characterizes our most basic model. Using the implied marginal rates of substitution between illness profiles and money, we then construct individual measures of willingness to pay to avoid five archetypical illness profiles to be introduced in the next section. Our underlying structural model requires (for now) that we make assumptions about individuals' time preferences and income expectations if they get sick. Thus, in Section 4, we explore how our implied $VSIP$s vary systematically with these two assumptions.

Our basic quadratic-in-net-income structural model, which assumes homogeneous preferences, produces the five parameter estimates shown as Model 3 in Table 2. These homogenous-preferences specifications are estimated without sign restrictions and show robust significance and the expected signs on all five primary parameters.[32] The marginal utility of income is positive, but declines with the level of income (yet does not go negative within the range of incomes in our sample). The marginal utilities of sick-years, post-illness-years, and lost life-years are all negative and very strongly significantly different from zero.[33] While simple intuition might suggest that death should be "worse" than illness and recovery, it is important to keep in mind that the units involved are *years* in each health state. The relatively large (dis)utility associated with recovered state probably reflects the general seriousness of the illnesses our survey describes. Respondents seem not to interpret being recovered from any of this list of major illnesses as being fully equivalent to the pre-illness state.

We now relax the maintained hypothesis in Model 3 that the marginal utilities from each state are independent of the duration of that state and the durations of other health states that characterize the profile in question. Our original model was developed in terms of the individual's undiscounted *per-period* indirect utility, where current-period health status is captured only by a set of mutually exclusive and exhaustive dummy variables. At the moment of the individual's program choice, however, each alternative is likely to be perceived in terms

[32]Not surprisingly, the additional structure in Model 3, as opposed to Models 1 and 2, produces a lower maximized value of the log-likelihood function. This is a common tradeoff. The structure is required for a rigorous utility-theoretic interpretation of the results, but the ad hoc model provides a better fit to the data.

[33]A positive marginal utility associated with a lost life-year might be expected only when the illness is question constitutes a "fate worse than death."

of the present value of the sequence of future health states it represents. These present values reflect the mix of future health states in each illness profile. If they capture the relevant attributes of each alternative in the individual's choice set, we can consider richer models that allow for diminishing, rather than constant, marginal utilities from present discounted health-state years, and for interactions between the numbers of present discounted years in different health states. In contrast, Model 3 constrains the marginal utility of each health state to be constant and imposes a constant marginal rate of substitution between different health-state-years.

The final line in the estimating specification in equation (3), $\alpha_1 \left\{\Delta\Pi_i^{AS} pdvi_i\right\} + \alpha_2 \left\{\Delta\Pi_i^{AS} pdvr_i\right\} + \alpha_3 \left\{\Delta\Pi_i^{AS} pdvl_i\right\}$, can easily be adapted to be non-linear in $pdvi_i^A$, $pdvr_i^A$, and $pdvl_i^A$. We first factor out the common $\Delta\Pi_i^{AS}$ term. Then the original form of the term involving the present discounted health states is $\Delta\Pi_i^{AS} \left\{\alpha_1 pdvi_i^A + \alpha_2 pdvr_i^A + \alpha_3 pdvl_i^A\right\}$. To accommodate zeros, we shift the data by one unit, then we take logarithms. The resulting alternative logarithmic form for the final substantive term in the estimating equation becomes $\Delta\Pi_i^{AS}\{\alpha_1 \log\left(pdvi_i^A + 1\right) + \alpha_2 \log\left(pdvr_i^A + 1\right) + \alpha_3 \log\left(pdvl_i^A + 1\right)\}$.Estimates for this form are presented as Model 4 in Table 3, which produces an improvement in the log-likelihood function compared to the linear and additively separable structural specification in Model 3. This suggests diminishing marginal utility in avoided present discounted degraded health-state years.

Model 5 then illustrates the consequences of allowing the parameters of the model to vary according to the fitted probability that each respondent appears in our estimating sample. Full-fledged selectivity correction models in multiple-choice conditional logit models are challenging, so we do not attempt them in this paper, although we do estimate a response/non-response model that produces fitted response probabilities for each individual in our sample.[34] We have allowed each basic parameter of our model to vary systematically

[34] Our selection model takes the over 525,000 original random-digit dialed recruiting contacts for the Knowledge Networks panel and fits a probit model to explain the presence or absence of each household in our final estimating sample. As explanatory variables, we use a set of 15 orthogonal factors derived from a factor analysis of over 100 census tract characteristics, county voting records, county mortality from each

with the deviation of that individual's fitted response propensity from the median response propensity among all 500,000-plus members of the random-digit-dialed initial Knowledge-Networks recruiting sample. Only the coefficient on the lost life-years term differs significantly with the fitted probability that the respondent shows up in our estimating sample. The greater the probability of being in our sample, relative to the median probability, the lesser the disutility the individual appears to experience from a percentage increase in discounted lost life-years. While the shift is statistically significant, comparison of Model 5 and Model 4 reveals that the minimal actual difference in the magnitude of this key lost life-years coefficient across individuals with different response propensities.[35]

Whenever a linear-in-logs form is a better predictor of consumer choices than a linear form, the researcher is typically inclined to explore even more general logarithmic forms. In particular, the translog form represents a second-order local approximation to any arbitrary functional relationship. This form is fully quadratic in all of the log terms and their pairwise interactions. We have explored the inclusion of all three squared terms and all three interaction terms. Only the squared term in $pdvl_i^A$ and the interaction term between $pdvi_i^A$ and $pdvl_i^A$ are robustly significant. This more general specification is presented as Model 6. Again, it produces a substantial improvement in the log-likelihood. The estimates suggest that the disutility of an additional discounted lost life-year shrinks as the number of discounted lost life-years increases. They also suggest that the disutility of an additional discounted lost life-year is reduced by increases in the number of discounted illness-years that precede it (e.g., as the number of years of morbidity preceding death increases, dying earlier becomes less bad).

In this application, however, there is a further complication. The illness profiles that were

major disease over the previous decade as a fraction of 2000 census population, and the number of hospitals in the same census tract(s) as the address (or telephone exchange) of the contacted household. Discussion of this response/nonresponse model constitutes a separate manuscript, currently under preparation.

[35] We employ differences from the median response probability so that the estimated utility parameters correspond to the simulated case where all response probabilities are exactly equal to the median in the population. We employ the median because the distribution is skewed, with a number of large positive outliers.

eligible to be considered by each respondent were constrained by the respondent's current age. No respondent considered illnesses that could strike at an age younger than their current age, so current age defines the maximum duration of any illness profile. The result is a degree of multicollinearity between the respondent's remaining nominal life expectancy and the range of sick-years, post-illness years, and lost life-years they were eligible to consider. In particular, when including interactions between the $pdvi_i^A$ terms and the $pdvl_i^A$ terms, occasional large values of these interaction terms were closely associated with the youth of the respondent.

It is not possible to include current age as a factor that might have an additively separable effect on the individual's level of utility, since terms such as these drop out of the utility-difference calculation across alternatives. To control for the effect of current age on the apparent marginal utility of each health state, we need to allow current age, age_{i0}, to shift the marginal utility parameters. An intermediate model (not shown in Table 3) assessed the consequences of allowing age_{i0} to shift only the coefficients on each of the linear terms in the logs of discounted years in each adverse health state. Each of the additional coefficients, α_{11}, α_{21}, and α_{31}, was statistically significant. Older respondents appear to anticipate lesser disutility from discounted sick-years and discounted lost life-years, but greater disutility from discounted post-illness years.

We next allow all of the translog coefficients to vary systematically with both age_{i0} and age_{i0}^2 since earlier empirical research has suggested the presence of quadratic age effects in VSLs.[36] The age shifters on the sick-years and post-illness years terms ($pdvi_i^A$ and $pdvr_i^A$) become statistically insignificant. However, the presence of significant quadratic-in-age shifters on the linear and quadratic lost life-years terms ($pdvl_i^A$) and on the interaction between the $pdvi_i^A$ term and the $pdvl_i^A$ term, prevents counter-intuitive negative fitted $VSIP$ point

[36] See for example Jones-Lee et al. (993) or Krupnick et al. (2002). The specification with just linear age effects on the linear-in-logarithms terms in discounted health-state years produces a substantial improvement in the log-likelihood function, but leads to some outliers in the simulation results when we use the parameter estimates to predict VSIs for specific illness profiles. Quadratic forms in age for each of the systematically varying parameters appear necessary to accommodate nonlinearities in these relationships.

estimates for some illness profiles for young respondents. Therefore, we prefer the specification presented as Model 7 in Table 3, even though two coefficients (on the lower-order "level" and "linear age effects" on the interaction between the $pdvi_i^A$ and $pdvl_i^A$ terms) are not individually statistically significant. To our knowledge, these are the first attempts to estimate, within a common model, the age-varying marginal utilities of avoiding a present discounted year of morbidity and a present discounted lost life-year. We assess the validity of our estimates by exploring whether they vary systematically in a manner that economic theory or simple intuition would predict.[37]

It is relevant to examine how our estimates vary with assumptions about average time preferences, as well as with the data concerning each individual's income, and with current age and disease latency. We employ the estimated parameters reported for Model 7 in Table 3 to characterize the $VSIP$ associated with selected combinations of years of morbidity, years in post-illness status, and years of premature mortality. A vast range of different illness profiles could be selected, but for illustrative purposes, we examine five representative illness profiles: 1) a period of shorter-term morbidity followed by recovery, 2) a period of longer-term morbidity followed by recovery, 3) a combination of shorter-term morbidity followed by premature mortality, 4) a combination of longer-term morbidity followed by premature mortality, and 5) sudden death. Recall that a $VSIP$ reflects willingness to pay for typically a very small risk reduction, scaled up (proportionately) to the amount that would correspond to a risk reduction of 1.00. As for VSLs, there is of course no presumption that the resulting number would ever be used to value the life of one individual who would die with certainty. While the individual's budget will constrain the underlying willingness-to-pay, the $VSIP$ is not so limited, and the benchmarks against which we compare our $VSIP$ figures are the roughly \$6 million VSL used by the U.S. EPA and the roughly \$1 million used by the U.S. Department of Transportation.

[37] The only other ordinal utility measure expressed per year is the concept of the value of a statistical life year. However, this is not a measure of marginal utility, rather it is constructed by dividing a VSL estimate by the remaining number of expected life-years.

Our models currently require that the researcher specify each individual's time preferences. In Table 4, we consider an individual who is now 45 years old with an income of \$42,000 and calculate the fitted $VSIP$ (in millions of dollars) for each of our five illness profiles to illustrate the sensitivity of our models to our choice of discount rate. Results for Models 3 through 7 were all derived under the assumption that $r = 0.05$. The middle column of Table 4 shows the medians and 90% ranges of simulated point estimates of the $VSIP$ for our five different illness profiles assuming a current age of 45 and immediate onset.[38] The first and third columns of results in Table 4 are produced by re-estimating Model 7, having first constructed the present discounted value terms using two alternative discounting assumptions: $r = 0.03$, and $r = 0.07$.[39] Table 4 shows that fitted $VSIP$ estimates vary inversely with the assumed discount rate. For our 45-year-old, the case of sudden death (most common in the conventional VSL context) the 5% discount rate produces a $VSIP$ of roughly \$4.5 million, whereas the median estimates for 3% and 7% rates and about \$5.5 million and \$3.8 million.

Responsiveness of $VSIP$s to income level can also be assessed. Table 5 reverts to a discount rate of $r = 0.05$ and again reports in the center column the results of simulating $VSIP$s for an individual who is now 45 years old, with an income of \$42,000, and faces each of our five representative illness profiles with immediate onset. The first and third columns illustrate $VSIP$ estimates for arbitrarily selected alternative income levels of \$25,000 and \$67,500.[40] As expected, $VSIP$ is larger when income is greater. For our 45-year-old and the case of sudden death, the fitted median $VSIP$ at \$25,000 income is only about \$3.5 million, whereas the fitted median $VSIP$ at \$67,500 income is about \$7.0 million. The bottom panel

[38] These simulations are taken across 1000 draws from the joint distribution of the estimated parameters. We acknowledge that the mean of the theoretical distribution of a ratio of asymptotically normal quantities is undefined. However, we present finite-sample medians and 90% ranges to convey a sense of the precision of the parameter estimates and the implications of this precision for fitted VSIs.

[39] Table C2, in Appendix C (available from the authors) details the consequences of these different discount rate assumptions for the estimated parameters of Model 7.

[40] These corresponding roughly to the 25^{th} percentile and median of the household income distribution according to the 2000 Census (\$25,000 and \$42,000), as well as for the 75^{th} percentile of individual income for our sample (\$65,000).

of Table 5 provides the corresponding implied arc elasticities of $VSIP$ with respect to income in these two income intervals.

Table 6 explores the effect of illness latency on willingness to pay to avoid health risks, for a subject with an assumed 5% discount rate and an income of $42,000. In this table, we array our five examples of different illness profiles across the top of the table. In the body of the table, we display fitted median $VSIP$ estimates and 90% ranges for one respondent aged 35 now and another aged 65 now. The age at onset of each illness is varied to include immediate onset, as well as onset at decade intervals starting five years from now. Considerable variability is present. Focusing again on the sudden death scenario, our model suggests that the 65-year-old is willing to pay less to avoid sudden death than the 35-year-old. Looking forward, however, both individuals are willing to pay less to avoid the same illness profile when it commences at a later age. Our model allows $VSIP$s to reflect the duration of each type of health state. The numbers of prospective sick-years and life-years lost can be expected to have a substantial effect on willingness-to-pay to avoid each illness profile.

As a visual summary of the effect of the respondent's age now on the $VSIP$ for sudden death, we offer Figure 1, which shows the simulated median and 90% confidence interval for this fitted $VSIP$ as a function of age now. As the algebraic form of the term in (**??**) indicates, age has an nonlinear effect on several of the parameters of the model. The combined influence of these three different types of quadratic age effects on the fitted $VSIP$ for this particular illness profile is captured by Figure 1. Figure 2 illustrates one other possible illness profile. In this case, it is an illness that lasts five years, ending in death, but with ten years of latency prior to onset. Willingness to pay to avoid this illness profile also differs systematically with age.[41]

When evaluating the social benefits of a policy change that alters the incidence of a

[41] For some age levels, the simulated 90% confidence interval for the predicted Value of a Statistical Illness includes zero. There is no opportunity for respondents to express a negative willingness to pay, so it may be appropriate to adopt a Tobit-type interpretation of the implied fitted distribution and to move any probability in the negative domain to zero. Alternatively, by resorting to a fundamentally nonlinear estimating model, one could estimate the logarithm of willingness-to-pay, instead of its level, thereby constraining fitted willingness-to-pay to be strictly positive.

particular illness, there are great advantages to being able to estimate the continuum of statistical illness profiles associated with that particular illness. Our approach offers the flexibility to evaluate changes in the type, future timing, and duration of heterogeneous illness profiles. Additionally, it does so within a consistent theoretical and empirical model, rather than requiring researchers to cobble together estimates for current period morbidity and mortality from separate valuation methods and studies.

5 Discussion and Conclusions

Unlike many previous empirical efforts to measure willingness to pay to reduce mortality risks, our model does not produce just a single best estimate for the Value of a Statistical Life (VSL) for use in all policy contexts. Instead, our model is best understood as a generalization of the standard single-period, single-risk valuation model. It explicitly allows the individual to allocate risk-reduction expenditures across health risks that come to bear across different future time periods. Our model allows for substitution across health risks with different time profiles that more-completely characterize the duration of morbidity and the eventual health outcomes that result from those risks. Rather than focusing on only a single risk of death in the current period, the model considers entire future illness profiles as its objects of choice. The most significant advantages of this simple generalization are that it allows us to accommodate (a) varying latencies in health risks, (b) morbidity in addition to mortality, and (c) non-fatal as well as fatal risks. Along these three dimensions, our model represents a major departure from previous empirical specifications.

Although the model is a generalization, it nonetheless produces a new and important type of economic information: distinct estimates of the marginal utilities of avoiding a year of morbidity and a lost life year within a single model. It also appears that these marginal utilities are *not* simple constants. From these heterogeneous marginal values, which seem to depend upon the current age of the respondent (and therefore possibly upon other factors which are correlated with age) and the mix of health states in an illness profile, we have illustrated how to construct average values for a wide range of illness profiles, for individuals of different ages.

To further enhance program and policy evaluation, we organize our analysis around the task of estimating the value of a statistical illness profile ($VSIP$), although we allow for the identification of a simpler concept that is similar to the more-traditional value of statistical life (VSL). The $VSIP$ evaluates the set of heterogeneous health outcomes associated with a given illness risk. Policy changes that affect the prevalence and severity of that illness will shift the joint distribution of the duration of morbidity and premature mortality, for specified populations, and our model is capable of assessing the benefits of such broad shifts.

Our analyses illustrate some initial results concerning how marginal utility of risk mitigation varies systematically across individuals. Specifically, we evaluate how the demand for mortality risk reduction varies with the individual's current age and the disease latencies that dictate the future ages at which degraded health states would be experienced. Our results suggest that, however convenient it may be, the presumption that there should be a single-valued VSL is probably misguided. While the use of a single number may continue to be dictated by political concerns, the willingness to pay to reduce health risks should be viewed as a function (rather than a scalar that is merely proportional to the magnitude of the risk reduction). A VSL is ultimately derived from a type of inverse demand function, so the prospect of systematic variation in willingness to pay, according to the qualities of the good and with indicators of individual preferences, should not be surprising.

6 References

Aldy, J.E. and W.K. Viscusi. 2003. "Age variation in worker's value of statistical life." Mimeograph

Baron J., and Ubel P.A. 2002. "Types of inconsistency in health-state utility judgments." *Organizational Behavior and Human Decision Processes* 89 (2): 1100-1118.

Cameron, T.A. 2005. "Individual option prices for climate change mitigation." *Journal of Public Economics*, forthcoming.

Corso, P.S., Hammitt, J.K., and Graham, J.D. 2001. "Valuing mortality-risk reduction: Using visual aids to improve the validity of contingent valuation." *Journal of Risk and Uncertainty* 23 (2): 165-184.

Cummings, R.G., and Taylor, L.O. 1999. "Unbiased value estimates for environmental goods: A cheap talk design for the contingent valuation method." *American Economic Review.* 89 (3): 649-665.

de Bruin, W.B., and Keren, G. 2003. "Order effects in sequentially judged options due to the direction of comparison." *Organizational Behavior and Human Decision Process.* 92 (1-2): 91-101.

Dow, W.H., Philipson, T.J., and Sala-I-Martin X. 1999. "Longevity complementarities under competing risks." *American Economic Review* 89 (5): 1358-1371

Dreze, J. 1962. "L'Utilite Sociale d'une Vie Humaine." *Revue Francaise de Recherche Operationalle* 6: 93 - 118.

Gold, Marthe R., Russel, Lousie B., Siegel, Joanna E., Weinstein, Milton C. (eds). .1996. *Effectiveness in Health and Medicine.* New York: Oxford University Press.

Graham, D.A. 1981. "Cost-benefit analysis under uncertainty." *American Economic Review* 71 (4): 715-725 1981

Hammitt, J.K., and Graham, J.D. 1999 "Willingness to pay for health protection: Inadequate sensitivity to probability?" *Journal of Risk and Uncertainty* 18 (1): 33-62.

Johnson, F.R., Banzhaf, M.R., and Desvousges, W.H. 2000. "Willingness to pay for improved respiratory and cardiovascular health: A multiple-format, stated-preference approach." *Health Economics* 9 (4): 295-317.

Jones-Lee, M.W. 1974. "The value of changes in the probability of death or injury." Journal of Political Economy 82(4): 835 - 849.

Jones-Lee, M.W., G. Loomes, D. O'Reilly, and P.R. Phillips .1993. "The value of preventing non-fatal road injuries: findings of a willingness-to-pay national sample survey." TRY Working Paper, WP SRC2.

Krupnick A., Alberini A., Cropper M., Simon N., O'Brien B., Goeree R., and Heintzelman M. 2002. "Age, health and the willingness to pay for mortality risk reductions: A contingent valuation survey of Ontario residents" *Journal of Risk and Uncertainty.* 24 (2): 161-186.

List J.A. 2001. "Do explicit warnings eliminate the hypothetical bias in elicitation procedures? Evidence from field auctions for sportscards" *American Economic Review.* 91 (5): 1498-1507.

Liu, J.-T. and Hammitt,J.K. 2003. "Effects of disease type and latency on the value of mortality risk," NBER Working Paper no. w10012.

Manski, C. 2004. "Measuring expectations." Econometrica. 72 (5): 1329-1376.

Moore, M., and Viscusi, W.K. 1988. "The quantity-adjusted value of life." *Economic Inquiry* 26 (3): 369-388.

Moxey, A., O'Connell, D., McGettigan, P., and Henry, D. 2003. "Describing treatment effects to patients - How they are expressed makes a difference." *Journal of General Internal Medicine* 18 (11): 948-U8.

O'Connor, R.M. and Blomquist, G.C. 1997. "Measurement of consumer-patient preferences using a hybrid contingent valuation method." *Journal of Health Economics* 16 (6): 667-683.

Picone, G., Sloan, F., and Taylor, D. 2004. "Effects of risk and time preference and expected longevity on demand for medical tests." *Journal of Risk and Uncertainty* 28 (1): 39-53.

Rosen, S. 1988. “The value of changes in life expectancy.” Journal of Risk and Uncertainty 1: 187-203.

Sloan F.A., Viscusi, W.K., Chesson, H.W., Conover, C.J., and Whetten-Goldstein K. 1998. “Alternative approaches to valuing intangible health losses: the evidence for multiple sclerosis” *Journal of Health Economics* 17 (4): 475-497.

Ubel P.A., Richardson J., and Baron J. 2002. “Exploring the role of order effects in person trade-off elicitations.” *Health Policy* (2): 189-199.

Viscusi, W.K. 1993 “The value of risks to life and health.” *Journal of Economic Literature* 31: 1912 - 1946.

Viscusi, W.K., Magat, W.A., and Huber, J. 1991. “Pricing environmental health risks: survey assessments of risk-risk and risk-dollar trade-offs for chronic bronchitis.” *Journal of Environmental Economics and Management* 21 (1): 32-51.

Yeung, R.Y.T., Smith, R.D., and McGhee, S.M. 2003. “Willingness to pay and size of health benefit: an integrated model to test for ’sensitivity to scale.” *Health Economics* 12 (9): 791-796.