5-2019

Utility Functions for Mild and Severe Health Risks

W. Kip Viscusi

Journal of Risk and Uncertainty (2019) 58:143–166
https://doi.org/10.1007/s11166-019-09301-9

Utility functions for mild and severe health risks

W. Kip Viscusi [1]

Published online: 11 May 2019

Abstract

This article reviews economic evidence on health-dependent utility functions and presents new estimates of utility functions for cancer. Estimates of health-dependent utility functions have found that mild adverse health impacts can be treated as monetary equivalents. Severe health consequences also reduce utility levels but have an additional effect of altering the structure of utility functions by reducing the marginal utility of income. The implications of past studies are often misleading when they fail to account for income losses and medical expenses associated with serious ailments. This article's estimates of the structure of utility functions for cancer indicate a substantially lower marginal utility of income at any given income level. This result is consistent with the welfare consequences of other severe health effects, which impose harms that are not tantamount to a monetary loss.

Keywords Value of statistical life · VSL · Cancer · Utility · Health risk · Willingness to pay

JEL classifications I10 · K13 · K32 · D8

1 Introduction

Risks of adverse health outcomes potentially impose losses that reduce individual utility. Recognition that adverse health effects lower individual welfare is not sufficient to fully characterize their economic implications. How one treats these losses and incorporates them in economic analyses has potentially profound effects on the welfare consequences of injuries and illnesses. The effect of adverse health impacts on the structure of utility functions also has ramifications for the appropriate policies for the prevention, insurance, and compensation of injuries and illnesses.

Paper prepared for Vanderbilt-JRU Symposium, Risk Guideposts for a Safer Society

✉ W. Kip Viscusi
kip.viscusi@vanderbilt.edu

[1] Vanderbilt Law School, 131 21st Ave. South, Nashville, TN 37203, USA

In the most general form, one can conceptualize the utility u of income y in any health state as being given by $u(y, health\ status)$. The standard monotonicity assumption is that more income is preferred to less. Similarly, for any given level of income, one would expect good health to be preferred to ill health, or $u(y, good\ health) > u(y, ill\ health)$. In some situations, knowing that a particular adverse health state is undesirable is all that is required. For example, in models of compensating differentials for mortality risk and job injury risks, it is possible to demonstrate that there should be a wage premium for risky jobs so long as being injury-free is preferable to being injured.

Much economic research has focused on the local rate of tradeoff between money and health risks, leading to estimates such as the value of a statistical life, or the VSL (Viscusi 2018). On a theoretical basis, these tradeoff rates are intrinsically tied to the utility function structure when healthy and the utility function in the ill health state. More specifically, the VSL and other risk tradeoff rates equal the difference in the utility when healthy and the utility in the ill health state, where this difference is divided by the expected marginal utility of income. For many purposes it is not essential to know how the adverse health event alters the shape of the utility function. Local rates of tradeoff are sufficient for valuing small changes in risk.

Some policy contexts require more information about the nature of individual preferences. Occasionally, policies may lead to substantial changes in the risk level. Valuation of such risk effects based on the marginal valuations leads to an understatement of the willingness-to-accept values for risk increases and an overstatement of the willingness-to-pay values for risk decreases. Similarly, understanding the effect of changes in the background risk is facilitated by knowledge of the structure of individual utility functions.

Detailed explorations with respect to matters such as optimal insurance or health care allocations often require that we know more than that good health is preferred to ill health or death and that we know how the adverse health status affects the form of utility functions. Does ill health simply lower the level of utility but otherwise leave the form of the utility function unaltered, which is the most straightforward situation? Or is there a more fundamental transformation of the utility function depending on the nature and severity of the health impact? Obtaining such information is essential to providing economic guidance in many areas. How should medical resources be allocated across different health conditions? If medical expenses can enable a person to transition to a healthy state or a less severe ill health state, how should these changes be valued? Even if the health state cannot be altered, knowing the structure of utility functions is essential to determining how valuable it is to transfer monetary resources to the ill health state. What is the optimal level of insurance and the optimal level of compensation after a personal injury that led to the health impact? Simply knowing that good health is preferable to ill health is insufficient to provide guidance in answering these questions, and equalizing the utility level in the ill health state and the healthy state will not generally be optimal when the utility function has been altered. Knowledge of how the adverse health effect influences the utility function and, in particular, the marginal utility of income resolves this broad set of issues.

There are two principal alternatives for characterizing the impact of adverse health effects on the structure of utility functions. The simplest approach analytically is to assume that the adverse health effect is tantamount to a monetary loss. That formulation has the advantage that the economic theory and empirical evidence used in economic analyses of monetary losses is well-developed. After converting the health loss into a monetary equivalent, adverse health impacts can be incorporated into the economic model in the same way as would a drop in income. All that is required is to ascertain the value of the monetary equivalent of the non-financial loss. However, unlike monetary losses, this amount is not observable. We may know that the person is totally disabled but are not able to ascertain what monetary loss would lead to an equivalent loss in welfare. Consequently, there may be empirical challenges even if the theory with respect to the treatment of monetary losses as components of utility functions is straightforward. A more sweeping characterization of the impact of adverse health effects is that there may be a transformation of the functional form of the utility function beyond simply having an effect on utility that is the same as a reduction in income. Particularly for severe adverse health effects, the health impact may also alter the marginal utility of income. Whereas a loss in income boosts the marginal utility of income, severe health impacts may impede individuals' ability to derive utility from additional consumption.

How adverse health effects influence utility functions is an empirical issue. Some analysts mistakenly assume that the effect of adverse health impacts either must always be equivalent to a monetary loss or must always alter the structure of utility functions, and that any deviation from a uniform effect is a sign of inconsistent empirical results. But the impact of different health effects may vary, as some may be tantamount to monetary equivalents and others may have more transformative effects. It is a mistake to assume that a one-size-fits-all approach correctly characterizes the influence of adverse health impacts on utility functions. Similarly, if there is empirical evidence indicating that utility functions for different ill health states are not the same, such a finding is not evidence of a deficiency of the analysis but a recognition that there are diverse impacts of injuries depending on their welfare consequences.

Section 2 presents a general discussion of how adverse health effects might be incorporated in the structure of utility functions. In Section 3, I outline different approaches that one might take to estimate utility functions conditional on health status, which typically involve treating health losses as equivalent to monetary losses or as impacts that also alter the marginal utility of income. Section 3 reviews the empirical evidence on utility functions conditional on health status. My previous estimates of health-dependent utility functions are quite robust, as mild injuries can be treated as tantamount to health losses and more severe injuries alter the marginal utility of income. Other studies in the literature often draw conclusions about the effects on the marginal utility of income, but the results are sometimes confounded by not also controlling for decreases in income levels and increases in medical expenses after adverse health effects. Section 4 estimates utility functions for cancer, a disease that produces a substantial reduction in the marginal utility of income. The implications of cancer for utility functions are consistent with other evidence for severe health effects. In Section 5, I discuss the fundamental ramifications of these estimates for policy and general principles for appropriate treatment of adverse health impacts.

2 Utility functions for mild and severe health risks

Mild health effects might plausibly be characterized as equivalent to a monetary loss. The adverse health effect makes the person worse off but does not affect the ability to derive utility from additional consumption expenditures. If $u(y)$ is the utility function in the healthy state, then after incurring a health impact that is equivalent to a monetary loss l, the utility level in the ill health state is given by $v(y) = u(y - l)$. As in the case of a loss of income, the loss lowers the utility level and boosts the marginal utility of income, each of which can be restored to their pre-injury level by receiving compensation l.

The monetary equivalent formulation has the advantage that the well-developed economic theory for financial losses carries over quite directly. In general, when offered actuarially fair insurance opportunities, the optimal level of insurance is to set the marginal utility of income equal in both states, or $u' = v'$. If adverse health impacts can be treated as financial losses, then the optimal level of insurance is full replacement of the loss. This full replacement also is equal to the "make whole" amount in legal contexts, where compensation at this level will fully restore the person's welfare to its former level in the good health state so that $v(y) = u(y) = u(y - l + compensation)$. The optimal level of *compensation* equals l when the functional forms of $u(y)$ and $v(y)$ are identical except that $v(y)$ equals $u(y - l)$. The theory with respect to the efficient level of accident avoidance behavior also follows the standard results for financial losses. Both strict liability and negligence rules potentially will lead to efficient levels of care by setting the damages amount paid by the injurer equal to the monetary loss.

More serious health impacts may not be readily converted into a financial loss so that it is not appropriate to treat $v(y)$ as being the same as $u(y - l)$. Serious health effects might lower the marginal utility of consumption so that $v' < u'$ for any given level of income. That relationship is likely to be viewed as compelling for extremely adverse health effects that lead to death, as the marginal utility of money in one's bequest is generally lower than the marginal utility of money when alive. Serious injuries might be viewed along a continuum from good health to death, at least within the context of a single-period model. Estimates of the marginal utility of bequests suggest that these values are very low. Kopczuk and Lupton (2007) and De Nardi and Yang (2014) consider the implications of retirement savings data and estimate that the marginal utility of bequests is close to zero. Moore and Viscusi (1990) estimate the value of bequests in a labor market model of jobs posing fatality risks and find that bequests have a value comparable to the utility of a 0.024 probability of a year of life.

A frequent consequence of disabilities and serious health effects is that there is a loss in income. Any reduction in income will tend to boost the marginal utility of income. However, the ill health event may have also changed the structure of the utility function. The marginal utility of income may not be higher than before when holding income levels constant. There may also have been a change in the utility function, but to assess whether there has been such an impact, it is essential to make the comparisons holding income levels constant.

Another possibility is that money may actually become more valuable after an injury by enabling the injured party to pay for medical and rehabilitation expenses. Such an observation is certainly on point. Ideally, analyses should set these expenditures aside and focus on the consumption-related expenditures rather than medical and rehabilitation costs. Unfortunately, many estimates in the literature fail to control for these monetary losses, leading to severe injuries boosting marginal utility of income. However, in personal injury contexts, the medical and rehabilitation expenditures are not properly viewed as consumption expenditures. Rather, they are treated as separate components of economic damages. Such expenditures also take us out of the simple two-state model that is under consideration. In particular, these medical and rehabilitation expenditures are often directed at enabling the individual to transition to a different health state that entails less of a welfare loss and presumably boosts the marginal utility of income. Such expenditures are a quite different matter than expenditures on personal consumption given the individual's current health state and the associated utility function.

A simple variant of the health state formulation that can be implemented empirically is to assume that the ill health state utility function is equivalent to some factor α multiplied by $u(y)$, or $v(y) = \alpha u(y)$, where $0 < \alpha < 1$. For any given income level, the health impact reduces the marginal utility of income but otherwise represents a rescaling of the utility function. There could be other variants of this approach as well, by having $v(y)$ also differ from $u(y)$ by an additive constant as well as a change in the marginal utility of income. Or one might characterize $v(y)$ as including both a monetary equivalent component as well as a change in the utility function, leading to forms such as $v(y) = \alpha u(y - l)$.

The economic implications of health risks that reduce the marginal utility of income are quite different than in the situation of monetary equivalents. Optimal insurance that equates the marginal utility of income in the good health and ill health states will generally provide less than full compensation for the loss and will not lead to equating u and v, as the post-insurance utility level in the ill health state will be less if the marginal utility levels are equalized. Given the diminished marginal utility of income, the "make whole" amount may be quite substantial. It may not even be feasible to restore the person's utility level even with extremely large income transfers, as in the case of catastrophic impacts such as quadriplegia.

The change in the utility function structure also undermines the possibility that efficient levels of insurance might also provide efficient levels of deterrence. The efficient level of insurance will not provide sufficient incentives for safety, and providing efficient levels of deterrence will lead to excessive levels of insurance. In the situation of wrongful death cases, the usual conceptualization of the appropriate level of compensation is the insurance amount that is typically needed to meet the economic losses incurred by the survivors. These monetary losses include medical expenses, lost income, and lost services. Optimal prevention of fatality risks requires that the price signals be provided at a level that is consistent with the value of a statistical life (VSL). These estimates, which are on the order of \$10 million, as reviewed in Viscusi (2018), may be an order of magnitude greater than the level of the economic loss to the family. Setting compensation levels equal to the economic loss leads to inadequate deterrence, while setting the compensation equal to the VSL provides excessive insurance. As a practical compromise, I have proposed that the VSL should be set equal to the damages amount in situations in which deterrence is a prominent concern, which is typically when punitive damages are warranted.

That there is an inevitable tradeoff between optimal insurance and efficient deterrence stems from the infeasibility of using one policy instrument to address both the optimal insurance and optimal deterrence objectives in a situation where the same penalty amount does not address both simultaneously. Spence (1977) presented an early economic analysis of this conflict in the product safety context. Although he did not note the critical importance of the VSL in his analysis because of the timing of the development of the VSL literature, it is interesting to note that the functional form of the VSL is the linchpin of the model for determining the efficient level of product safety. Consumers' VSL establishes the money-risk tradeoff that should serve as the guide for determining the cost-risk tradeoff firms make in setting the level of product safety.

The efficient level of disease prevention and medical interventions to ameliorate adverse health effects also hinges on the structure of utility functions (Crainich and Eeckhoudt 2017). If an illness has lowered the marginal utility of income, it may be preferable to allocate resources to medical expenditures that enable the individual to derive additional utility from a given level of consumption than to transfer funds that will be used for consumption purposes. Whether there is a rationale for such a reallocation, and the extent to which such a reallocation is desirable, hinges on how much the marginal utility of income is reduced by the illness and how successful medical expenditures can be in boosting the marginal utility.

3 Empirical evidence

3.1 Studies by Viscusi and coauthors

There have been numerous attempts to estimate how utility functions are influenced by adverse health effects. The efforts that I have undertaken either individually or with coauthors have used individual response to lotteries to construct the von Neumann-Morgenstern utility functions for different health states. Some other studies have followed a similar approach, while others have used a different procedure not based on attitudes toward lotteries and sometimes not with respect to estimating the overall structure of utility functions. I begin with a review of analyses in which I have been involved and then turn to review of other studies.

The overall objective of these empirical analyses is to assess how the ill health state has altered the utility function $u(y)$. Adverse health events may reduce income levels, which will not alter the functional form of $u(y)$ but will lead to a higher observed level of marginal utility. Similarly, the presence of additional medical expenses likewise reduces the funds available for consumption, which will affect the assessed level of marginal utility, but despite a change in marginal utility levels, there may not have been any shift in the functional form of the utility function. Thus, the task requires that one seek to assess whether there has been a change in the utility function by making a comparison with the utility function in the good health state at the same level of consumption.

All of my estimates that I discuss below are based on surveys of adult populations who are in situations in which the survey provided specific detailed information about the nature of the risk and its health consequences. In each case the survey sought to isolate the effect of the adverse health impact, controlling for possible income effects of ill health and medical expenses. In the case of the job risk studies, the workers' compensation system provides coverage of medical expenses as well as earnings replacement rates that are explicitly incorporated in the model. For the other consumer and patient studies, the survey informed the respondent that the medical costs associated with the injury or illness would be addressed by insurance so that income effects could be ignored. The studies reviewed here sought to isolate the utility consequences in situations in which these medical expense influences do not generate a confounding effect.

The mild and severe health impacts in the surveys I describe below differ not only in their severity but also in their duration, potentially leading to some confounding of the severity of the harm and the periods over which utility functions are being measured. The mild health risks produce levels of pain that would be rated as mild or moderate, and these risks do not typically have long-lasting effects. The severe health risks tend to produce more severe levels of pain and to affect a person's health for a longer period. For example, severe health impacts such as multiple sclerosis are permanent chronic conditions, implying a long time frame for analysis. The cancer risk study also dealt with health impacts that extended over a long treatment period. Given the severity of the illness, even with a long time perspective for utility functions there is no potential for shifting consumption within the period of interest to ameliorate the adverse health impact. The other studies had more limited time frames. The survey dealing with the mild consumer injuries used a survey time frame of up to a year in discussing the period of the consumer's usage of the product. Similarly, the survey dealing with job risks used an annual salary time frame in eliciting the respondent's willingness-to-accept amount for job risks.

Table 1 summarizes my previous studies and the current study. These analyses consider several different health outcomes involving jobs, products, and health care contexts. Even within the job and product groupings there are multiple health outcomes of differing severity that are involved so that it is not correct to interpret the estimates as pertaining to a homogeneous set of health impacts within a particular job or product category. The parameter for characterizing the adverse health effects on the structure of utility functions is denoted by the values of α, whereas monetary loss equivalents of injuries are indicated in Table 1 by the l values. Although the table includes seven sets of estimates, they are based on results from four different surveys. Some surveys appear more than once with somewhat different results, but the differences can be traced to sometimes substantial differences in the empirical specifications. For example, the study might be estimating the average value of utility function parameters or might estimate these values as functions of a series of different personal characteristics, such as income and education. In addition, while treating the probabilities that are stated in the survey or the probabilities that are reported by respondents at face value is the standard approach in stated preference studies, if respondents act as Bayesians and incorporate the survey information in conjunction with their prior beliefs, the probabilities implicit in their behavior may differ from these values. As a result, some studies jointly estimate both the parameters of the state-dependent utility functions as well as the risk perception functions that are implicit in the choices that the survey respondents express. Because von Neumann-Morgenstern utility functions are only defined up to a positive linear transformation, it is sometimes useful in empirical studies to set the marginal utility of income in the good health state equal to 1.0 so that the estimated utility in the ill health state is relative to this value.

There are two studies by Viscusi and Evans (1990, 2006) using a survey of chemical workers from four chemical plants by Viscusi and O'Connor (1984). The survey elicited a worker's baseline risk and level of annual earnings at the firm. Each worker then considered a hazard warning and was told that the chemical would replace the chemical with which the worker currently worked. The text of the warning described the nature of the risk and the health consequences. The worker then assessed the risk that the chemical would pose and the earnings increase that the worker required to work

Table 1 Summary of Viscusi and coauthors' utility function estimates

Article	Health outcome	Utility function estimation	Results
Panel A: Mild health impacts			
Evans and Viscusi (1991)	Insecticide: Skin poisoning Inhalation Inhalation (children) Toilet bowl cleaner: Eye burns Gassing Gassing (children) Child poisoning	Logarithmic utility, all α terms not significantly different from 1.0 except as noted	Insecticide: $l = 619$ skin poisoning $l = 849$ inhalation $l = 1433$, $\alpha = 0.998$ inhalation (children) $l = 2538$ child poisoning Toilet bowl cleaner: $l = 515$ eye burns $l = 486$ gassing $l = 582$ gassing (children) $l = 923$ child poisoning
Evans and Viscusi (1993)	Insecticide: Skin poisoning Inhalation Toilet bowl cleaner: Eye burns Gassing	1st order Taylor Series Logarithmic utility CARA, all nonlinear with income effects Results are for logarithmic evaluated at mean income	Insecticide: $l = 1290$ skin poisoning $l = 1546$ inhalation Toilet bowl cleaner: $l = 555$ eye burns $l = 683$ gassing
Evans and Viscusi (1998)	Insecticide Toilet bowl cleaner injuries	1st order Taylor Series with nonlinearities evaluated at mean income and risk perception function	Insecticide: $l = 1294$ skin poisoning $l = 1559$ inhalation Toilet bowl cleaner: $l = 564$ eye burns $l = 699$ gassing
Panel B: Severe health impacts			
Viscusi and Evans (1990)	Job injury	1st order Taylor Series 2nd order Taylor Series Logarithmic utility	1st order Taylor series: $\alpha = 0.773$ average $\alpha = 0.818$ TNT $\alpha = 0.415$ asbestos $\alpha = 2.522$ chloroacetophenone 2nd order Taylor series $\alpha = 0.701$ average Logarithmic: $\alpha = 0.928$ average $\alpha = 0.914$ TNT $\alpha = 0.939$ asbestos $\alpha = 0.959$ chloroacetophenone
Sloan et al. (1998)	Multiple sclerosis	Logarithmic utility with risk perception function	$\alpha = 0.668$ if have MS $\alpha = 0.084$ general population
Viscusi and Evans (2006)	Job injury	Logarithmic utility with risk perception function and functional dependence evaluated at mean tenure and education	$\alpha = 0.883$ average $\alpha = 0.897$ TNT $\alpha = 0.882$ asbestos $\alpha = 1.407$ chloroacetophenone
Viscusi, this article	Bladder cancer	1st order Taylor Series with Logarithmic utility	$\alpha = 0.545$

with the chemical. The survey consequently elicited the job risk lottery in the post-warning situation that the worker viewed as equally attractive as the job risk lottery in the pre-warning situation. As a result, the survey provided information with respect to two different points on the worker's constant expected utility locus. Using this information on two risk-earnings combinations, it is then possible to estimate the worker's state-dependent utility functions implicit in these stated preferences. The matter of particular interest is the utility function in the good health state and the utility function in the ill health state.

Figure 1 illustrates the mechanics of the analysis. There is a baseline probability p that the ill health state v will occur and a probability $1-p$ that the good health state u will occur. After being given the warning information, the assessed probability of an injury rises from p to q. To maintain the equivalence of the expected utility before and after receiving the warning, the respondent indicates the percentage earnings increase δ that is required to make the expected utility level the same in the pre- and post-warning situations. The survey consequently generates two points b and c on the constant expected utility locus EU, providing information on a pair of risk and income levels that the worker considers to be equally attractive. The empirical framework requires that one solve for the earnings increase value as a function of the other components of the model since the required pay increase is the dependent variable in the model. In order to estimate these utility functions, it is necessary to impose additional structure such as using Taylor series approximations and/or specific functional forms.

Each worker considered one of three chemicals: TNT, which poses the risk of explosion; asbestos, which poses the risk of cancer; and chloroacetophenone, which causes temporary eye irritation and tearing. The chloroacetophenone risk is the only health effect that is a minor transitory impact. One would expect the other hazards to lead to health consequences that would impede the person's ability to derive additional utility from consumption, thus lowering the marginal utility of income. The empirical estimates for several specifications of the model show a marked impact of job injury risks on the marginal utility of income v' for exposures to risks from TNT and asbestos. The only exception is the temporary health impact of the chemical chloroacetophenone for which the marginal utility of income is raised, not lowered, in two of the three

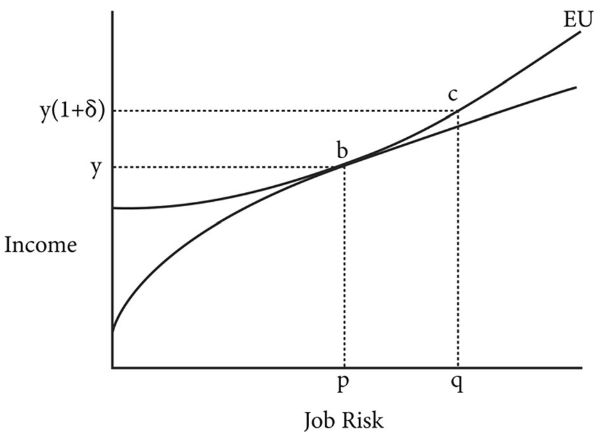

Fig. 1 Determining willingness-to-accept amount for increased job risk

estimates. An increase in the marginal utility of income would be expected if the injury were tantamount to a monetary loss. As expected, the more severe health impacts reduce the marginal utility of income and are not equivalent to adverse health impacts. These reductions in the marginal utility of income are remarkably robust, as they hold for general utility functions based on a first-order Taylor series approximation, a second-order Taylor series approximation, a logarithmic utility function, and a logarithmic utility function in a model in which the perceived risk beliefs implicit in people's choices may not necessarily coincide with the risk beliefs that the respondents report.

Three studies that analyzed the valuation of mild health risks using a sample of adult consumers considering potentially dangerous products are Evans and Viscusi (1991, 1993, and 1998). The survey described the risks for the two hypothetical products considered in the survey, insecticide and toilet bowl cleaner. Each of these products was patterned after commercially available household chemicals, but were given different names. The survey indicated the baseline risk and the product price. It then informed respondents that the products were reformulated and would pose a greater risk, whose magnitude was also communicated to respondents. The survey elicited information about price discounts that the consumer required to leave the consumer indifferent to the change in the risk level. As in the case of the job risk study, the expected utility before and after receiving the product warning was equalized, in this case by eliciting the price discount that was sufficient to maintain the equality. The survey consequently elicited information about two price-product risk points on a constant expected utility locus, which is the product market counterpart of the wage-job risk equivalent lottery pairings. The dependent variable in the regressions was the change in the consumer expenditure on the product that the consumer was willing to make to reduce the increase in the health risk.

To ensure that the respondents understood the health impacts that were at risk for each chemical, for each product the survey described the health effects and their welfare consequences. The survey indicated that the risks posed by the products involved only temporary adverse health impacts such as skin irritation and headaches from skin poisoning, tearing and coughing from inhalation or gassing by the chemical, and nausea and stomach pains from child poisonings.

Given the temporary and minor nature of the health impacts, one would expect that the monetary loss equivalent model would be appropriate to capture these health impacts. The results in Table 1 for these three studies indicate that these injuries are tantamount to monetary losses ranging from \$486 to \$1546 depending on the health outcome, the utility function model, and whether the estimation also accounted for the possible difference between stated and perceived probabilities. The estimates are consistent with theoretical predictions in that mild health effects don't impede the ability to derive additional utility from consumption but instead can be treated as monetary loss equivalents.

It is also possible to use the equivalent lottery approach involving two points on a constant expected utility locus to assess utility functions in health care contexts. The study by Sloan et al. (1998) administered a survey dealing with multiple sclerosis to two samples, a sample of adult patients who have multiple sclerosis and a sample of otherwise healthy adults. For the healthy respondents, the survey described the implications of multiple sclerosis, which one would expect to have a negative impact on their marginal utility of income given the severity of the health effects. The lottery prospect

Table 2 Other estimates of utility functions

Article	Health outcome	Dependent variable	Utility Function estimation	Results
Lillard and Weiss (1997)	Poor health after retirement and dying, leaving surviving spouse	Compare consumption paths across individuals who vary in their predicted probability of entering poor health	Longitudinal Retirement History Survey; Iterative maximum likelihood; sample is retirees	Marginal utility of consumption higher in poor health than good health
Smith et al. (2005)	Disability	Average subjective well-being responses to happy, enjoy life, sad (reverse-scores), and lonely (reverse scored)	HRS; Higher marginal utility for disability captures income and medical cost effects; OLS	Smaller decline in post disability condition at higher net worth
Finkelstein et al. (2013)	Number of diseases from a set of seven	Subjective well-being 0–1 happiness score	HRS; Age 50 and over, outside the labor force, with insurance; Linear probability model	Diseases reduce the marginal utility of income
Levy and Rizansky Nir (2012)	Cancer and diabetes versus good health	Patient percentage of consumption willing to pay to cure illness	Interviews with 180 cancer patients and 132 diabetes patients; Willingness to pay based on actual and hypothetical household consumption levels; Logarithmic utility. Generalized Power, Negative exponential	Logarithmic function $U(h, w) = h \cdot \log(aw)$, with $h = 0.83$ for cancer and 0.97 for diabetes
Tengstam (2014)	Future grandchild paralyzed in both legs	Income lottery for hypothetical grandchild	Author's survey of 292 undergraduates; "Society pays all extra economic costs (e.g. for special transportation service and for adjusting her house) that arise due to being mobility impaired."; Choose between income lotteries if impaired and not impaired.	63.7% of respondents had higher marginal utility of income when disabled.

Table 2 (continued)

Article	Health outcome	Dependent variable	Utility Function estimation	Results
Ameriks et al. (2017)	Needing long term (LTC) care for activities of daily living and death	Main survey: how to allocate $100,000 to provide insurance in long-term care state	Vanguard Research Initiative and survey of 9000; Insurance allocation for long-term care assuming no public support; Method of simulated moments with joint wealth and moments, and separately	Higher marginal utility for long-term care than bequests.
Brown et al. (2016)	Physical or mental disabilities preventing work or needing long-term care	$10,000 allocation between sick and healthy states	Hypothetical survey using American Life Panel (ALP); Basic nursing home costs fully covered; Physical health and lifespan unaffected; Allocation ratio in division of money between hypothetical sick and healthy states	Long-term care states reduce marginal utility of income, but higher values for work disability.
Gyrd-Hansen (2017)	Problems in walking, washing or dressing, performing your usual activities, moderate pain or discomfort, moderately anxious or depressed	Difference in income across two years	Author survey of 2000 Danish citizens; Income allocation after operation; Ask subjects to allocate 1000 DKK/month consumption between this year and next year when they will not be able to work	No state dependence on average, but find higher marginal utility for "intermediate" health states compared to higher and lower health states.

that the respondents considered was how great of a risk of death they would be willing to incur for a procedure that would return them to the good health state if they had multiple sclerosis. This information was used to estimate the marginal utility of income in the multiple sclerosis health state.

As one would expect for a severe ailment such as multiple sclerosis, the estimated marginal utility of income in the ill health state was below that in the good health state, which was normalized to have a value of 1.0. For the general population, the effect of multiple sclerosis is to reduce the marginal utility of income to a level of 0.084, or in effect reducing the marginal utility of income to close to zero. In contrast, the population with multiple sclerosis was much less willing to incur a risk of death for the prospect of a cure, leading to the utility function factor α of 0.668, so that the level of marginal utility is about one-third less than in the good health state.

This discrepancy in the impact of catastrophic health impacts on the marginal utility of income provides striking and compelling economic evidence that people may adapt to severe health impacts. Studies by psychologists have found that reported happiness levels are often not greatly different for people who suffer from health impairments and those who are healthy (Kahneman 2011). However, reported happiness scores rated on scales such as from 1 to 10 or some other ordinal metric may not be conclusive. Is the person rating the current level of happiness conditional on permanent aspects of the physical condition? Given that they currently have multiple sclerosis, are they relatively happy today? Or is the thought experiment reference point what the happiness level would be if the person could be returned to perfect health? There is no such problem with the lottery-based responses, as these are based on specified risk levels that establish indifference between the multiple sclerosis state and a lottery involving a risk of death and a complementary probability of being in good health.

The final study listed in Table 1 is the assessment of utility functions for cancer that will be presented in Section 4. As in the case of the serious injuries examined above, cancer reduces the marginal utility of income. Because the empirical framework is also based on an equilibrating lottery approach, it provides another example of evidence based on the kinds of models used in the studies in Table 1.

3.2 Other studies of utility functions

Economists also have used approaches other than the equivalent lotteries procedure that I have used to explore the characteristics of utility functions. Table 2 summarizes several analyses along these lines. While some studies have analyzed how ill health states affect the structure of utility functions, most have had the narrower objective of simply assessing whether the marginal utility of income is greater or lower in ill health states. Note that this marginal utility issue is different than examining how ill health affects the marginal utility of income for any given income level. If the ill health event leads to a loss in income or additional medical expenses, for a given utility function the marginal utility of income will have increased. But that result does not imply that the utility function itself has changed.

One study that seeks to estimate the overall form of the utility function is that by Levy and Rizansky Nir (2012). Their study used a survey approach for 180 cancer patients and 132 diabetes patients. The survey asked the patients what percentage of consumption they would give up to be fully cured. Based on a generalization of the standard logarithmic utility function, the authors concluded that these two ill health states had a lower marginal utility of income than if the patients were fully cured, particularly for cancer, which had a much more substantial impact on utility. Whereas Sloan et al. (1998) used a lottery approach in which the multiple sclerosis patients specified a risk of death that the person was willing to incur in order to be cured, this study asked respondents to equate the utility level in the current ill health state with a utility level in good health after incurring an equilibrating drop in consumption. Based on their estimates of a logarithmic utility function, they concluded that the findings were consistent with Viscusi and Evans (1990) since these adverse health outcomes reduced the marginal utility of income.

Other surveys also have sought to present subjects with hypothetical resource allocations that could then be used to assess the marginal utility of income. Brown et al. (2016) used the American Life Panel sample in which they included hypothetical questions regarding the division of money between a healthy and an unhealthy state. The survey included two variants, a work disability risk and a long-term care risk. For work disability, there presumably would be a loss in income, while for the long-term care risk, the survey informed respondents that medical expenses would be covered. There was no evident health state dependence on average for physical work-related disability, but there was for long-term nursing home care. One would expect a reduction in income in the disability state to boost the marginal utility of income, thus ameliorating any drop in marginal utility attributable to the health state effect. The reduced marginal utility of income in the long-term nursing home care state is consistent with severe health impacts reducing the marginal utility of income.

Gyrd-Hansen (2017) asked a sample of 2000 Danish citizens to consider a hypothetical situation in which they would have an operation next year. The costs of the operation would be fully covered. In the year following the operation, they would have to take a year off of work. The survey asked them what level of funds that they would like to shift to the future health state during their year of recovery. The future health states included five possible health conditions, such as moderate pain or some problems walking. The study found that these temporary health effects often had a neutral effect on marginal utility, but that there was a positive increase in marginal utility for the intermediate health states.

Analyses of consumption, saving, and bequest behavior serve as another source of evidence regarding changes in the marginal utility of income with ill health. A promising study that permits a comparison of ill health states and good health, as opposed to the value of bequests, is that of Tengstam (2014). The author ran a survey on a student sample comprised of 292 consistent responses to analyze hypothetical transfers of wealth to possible future grandchildren. The survey stated that there was the risk that both legs of the respondent's grandchild would be paralyzed. There was no further description of the health impacts other than to note that mobility-related expenses would be covered. The majority of respondents believed that income would be more valuable if the grandchild was paralyzed than if the grandchild was in good health, so that being paralyzed increases the marginal utility of income. For studies such as this in which there is a preference to allocate more funds to the ill health state, it would be useful to include additional probing to explore the purposes that respondents envisioned for these funds. Was the objective that respondents were seeking to achieve that of equating the utility in the two health states or the marginal utility?[1]

Ameriks et al. (2017) ran a survey on 9000 participants in the Vanguard Research Initiative to explore the comparative value of long-term care insurance and bequests. The survey asked respondents to divide $100,000 at age 80 between a healthy state and a long-term care state, where the amounts would reflect actuarially fair insurance opportunities given the stated probabilities of those states. The authors concluded that the marginal utility of income was higher in the long-term care situation than for bequests, which is what one would expect.

This study yielded evidence similar to that in the analysis of patterns of actual consumption behavior in the Health and Retirement Survey (HRS) by Lillard and Weiss (1997). Ameriks et al. (2017) considered bequests but also found that the marginal utility of consumption is higher for people in poor health. Poor health states have diverse impacts on financial resources and medical expenses so that an increase in marginal utility levels would be expected even if the utility function itself was unaltered.

In lieu of von Neumann-Morgenstern utility functions, there have also been efforts to use reported subjective well-being measures as an index of cardinal utility levels. Using the four subjective well-being questions in the HRS data pertaining to whether respondents were happy, enjoyed life, were sad, or were lonely, Smith et al. (2005) found that higher wealth levels buffered the drop in well-being with disabilities. This result can be related to the marginal utility of income in the disability states but in a manner that is much less direct than in the other studies. Suppose that a healthy affluent person has a utility score of 5 that drops to 3 when disabled, and that a healthy middle income person has a utility score of 4 that drops to 1 when disabled. By buffering the utility loss with disability, with extra wealth the utility loss from disability is 2 points rather than the 3-point loss for middle income people. In the healthy state, there is a utility difference of 1 point between the wealthy and the middle income person, while in the ill health state the utility gap is 2 points. Transforming a person from middle income to wealthy has a utility benefit of 2 points in ill health and 1 point in good health. While the marginal utility of income is greater when in poor health, such a result might be expected because of the income loss due to disability. Thus, the utility functions in the two health states are being evaluated at different income levels.

Finkelstein et al. (2013) used an HRS subsample consisting of those aged 50 and over who were outside the labor force and who also had health insurance. Overall, 97% of the sample had Medicare. The approach to assessing the utility impacts was to assess the effect of ill health on a 0–1 subjective well-being measure of utility based on whether the person was mostly happy in the past week: “Much of the time during the past week I was happy. (Would you say yes or no?)” A high 87% of the person-year responses indicated that the person was happy, which received a utility score of 1. The empirical approach involved estimating the effect of the number of diseases on the estimated probability that the individual was happy. The authors used the following set of seven diseases, which received equal weight: hypertension, diabetes, cancer, heart disease, chronic lung disease, stroke, and arthritis. The regression of the happiness score on the number of

[1] The author is indebted to James K. Hammitt for the observation that respondents might have been seeking to equate the utility levels rather than the marginal utilities in the two health states.

reported diseases found that each additional disease from the set of seven possible ailments reduced the probability that the respondent was happy by −0.011. Each additional disease consequently reduces the marginal utility of income by 1.1%.

The diversity of the findings and the evidence of some marked differences these other health state utility function studies have with the results in Table 1 is not an indication that there is considerable uncertainty about how ill health affects the marginal utility of income. Most of the differences arise because there is substantial heterogeneity in what is being measured. In some instances, there are income losses that will boost the marginal utility of income for any given utility function. In other cases, there are medical expense requirements that make the financial resources no longer comparable in the two health states. Ill health events that reduce income levels or impose medical costs often increase the marginal utility of income. For the most part, the results indicate that severe health impacts usually decrease the marginal utility of income but less severe effects do not. The analyses often differ from those in Table 1 not only in terms of the research approach but also in terms of the economic questions being addressed and what is being measured. In most of the analyses, the focus was on the issue of whether ill health events raise or lower the marginal utility of income. That is a quite different issue than whether ill health alters the marginal utility of income for any given income level.

4 Estimation of utility functions for cancer

4.1 The cancer survey data

Cancer risks have figured prominently in the stated preference literature given their pivotal role in the assessment of the benefits of government regulations. Studies directed at eliciting the willingness to pay to reduce cancer risks include: Magat et al. (1996), Hammitt and Liu (2004), Van Houtven et al. (2008), Hammitt and Haninger (2010), Viscusi et al. (2014), and McDonald et al. (2016). The principal focus has been on the VSL and the morbidity-risk counterpart for cancer-related risks. From a policy standpoint, a principal concern has been whether valuations of cancer risks merit a premium relative to other risks of death. The debate over whether there should be a premium relative to the VSL has been a policy issue in the United Kingdom, where cancer receives a 50% premium, and the United States, where there has been consideration of a premium but no formal adoption of an increased value for cancer risks. The impetus for higher valuation of cancer risks stems from the associated morbidity effects. These impacts for nonfatal cases of cancer are the focus of the survey discussed below. If cancer does have substantial morbidity consequences, these presumably should be reflected in the structure of utility functions in the ill health cancer state whereby cancer reduces the marginal utility of income.

The stated preference cancer survey dataset to examine the effect of cancer is that used in Viscusi et al. (2014). The cancer risk context presented in the survey is the risk of bladder cancer from arsenic exposures in drinking water, which had been the subject of a contentious environmental regulation. The survey was designed by the authors and administered by Knowledge Networks (now the GfK Knowledge Panel) to an online panel based on a probability sampling approach. The characteristics of the sample closely followed the distribution of the adult U.S. population. The sample size was 3430 adult respondents.

The survey included two training components before introducing respondents to the valuation task. First, because the survey dealt with drinking water risks, the survey engaged respondents in that context by exploring their drinking water practices and their familiarity with water-based illnesses. Second, to train respondents to deal with probabilistic information, the survey included a tutorial including risk ladders and a grid with 1000 colored squares to use in conceptualizing risk levels and changes in risk levels.

For characterization of the health consequences of cancer, the survey sought to strike a middle ground between an excessively terse description of the illness, such as simply noting that it involved a case of bladder cancer, and an overly detailed discussion that would be difficult to process. The survey indicated that arsenic exposures in drinking water posed the risk of bladder cancer that might occur after a 10-year latency period. A survey slide then summarized eight possible health consequences of bladder cancer and asked respondents to rate their severity of each possible consequence to ensure that they processed the symptom information and were aware of the stated preference commodity that they were valuing.

The survey design elicited the stated preference willingness-to-pay value for a reduction in the cancer risk. The payment mechanism was an increase in the respondent's water bill or, in the case of those households on well water, the cost of a water filter. For respondents to make the assessment of the cost amount they were willing to pay, they had to know both the initial risk and the post-water treatment risk probability. To enable respondents to conceptualize risks of 100,000, the survey presented the information in reference to population denominators such as the size of a large college football stadium and the population of cities such as Green Bay, Wisconsin and Berkeley, California. The starting risk levels p, which were consistent with the scientific literature on the levels of arsenic-related cancer risks in drinking water, were either 4/100,000 or 2/100,000. The post-treatment risk levels q were 2/100,000, 1/100,000, or 0/100,000. After presenting the risk information, respondents then considered an iterative series of valuation tasks based on a survey decision tree. For example, would they be willing to pay $200 to reduce the risk from 4/100,000 to 2/100,000? If they indicated "yes," they would be given a larger cost number. If they indicated "no," they were given a choice involving a smaller cost number. If they expressed indifference, the procedure ended as their rate of tradeoff was determined and set equal to the midpoint of the two values. The initial cost levels differed across respondents and ranged from $50 to $300 per year. The iterations continued for up to three rounds of choices.[2] In several previous studies, we found that respondents could process these iterative binary choices much more readily than when using other stated preference formats.

As for other stated preference studies, it is essential that respondents process the information provided and give thoughtful answers. The survey included across subject scope tests whereby larger risk decreases did in fact receive greater willingness-to-pay values. In addition, the survey itself incorporated rationality tests so that the small group of respondents indicating preference for a dominated alternative received a training module and did not continue the survey until their responses met usual rationality criteria.

[2] Subjects who reached the tips of the iteration tree before indifference were assigned the values at that tip for which the results closely paralleled the estimates based on interval regressions that accounted for this aspect of the design.

The overall structure of the survey is to elicit the willingness-to-pay value of the cost c that establishes the equivalence between the expected utility in the baseline condition and the expected utility in the post-treatment condition after incurring the treatment cost c. Let us assume that the individual incurs the cost c in both health states in the post-treatment situation. Then the value of c satisfies:

$$(1-p)u(y) + pv(y) = (1-q)u(y-c) + qv(y-c) \quad (1)$$

After imposing some structure on the utility functions, we will examine how the utility function for the ill health state differs from that in the good health state.

Table 3 summarizes the characteristics of the panel. On average, respondents were willing to pay $219 for water treatment that reduced the risk level from 3.49×10^{-5} to 7.53×10^{-6}, or a risk change of 2.74×10^{-5}. The average money-risk tradeoff implied by these estimates is a value of $8.0 million per expected cancer case, or $9.4 million in 2017 dollars, which is similar to the regression-based tradeoff estimate of $10.9 million for this dataset in Viscusi et al. (2014), which also takes into account other aspects of the survey design and sample characteristics. The cost value c, the baseline risk value p, and the post-treatment value q are pivotal variables in the estimation of utility functions, as indicated in Eq. 1 above. Utility functions are typically defined with respect to income levels, and respondents had an average income of $63,111.[3] Utility functions are likely to differ across the population. The personal characteristic variables used in the analysis to explore the heterogeneity of risk preferences are respondent *age*, whether the respondent considers herself to be an *environmentalist*, and whether the respondent is *female*. On average, the sample was 49 years old. The effect of *age* on willingness to pay for risk reduction is unclear since respondent age may be correlated with familiarity with the health consequences of cancer risks, remaining life expectancy, and available economic resources. Environmentalists, who comprise 42% of the sample, should be likely to place a greater value on water treatment to reduce cancer risks given their stated policy orientation, which reflects an avowed commitment to reducing environmental risks generally. Whether respondents consider themselves to be environmentalists is strongly correlated with membership in environmental organizations and has a positive effect on willingness to pay for environmental benefits in a variety of contexts. To the extent that the 52% of the sample who are women have a greater willingness to pay for reducing health risks, as some studies suggest is likely to be the case, their utility function should reflect their relatively greater utility loss from cancer risks.

Although the survey provided numerical risk information for the situations before and after water treatment, the perceived risk levels could depend on the respondent's prior beliefs and on whether the treatment has reduced the risk to zero. To account for the 5% of the sample who view their risks as particularly high, the regressions include a variable for whether the person *considers own cancer risk high.* If people place a premium on achieving a zero risk level, that will tend to boost valuations. A certainty premium could arise from the decreased worry about risks once they have been reduced to zero or the possibility that small nonzero risks are overestimated, leading to a discontinuity in valuations once risks are eliminated. Survey options in which the risk was reduced to zero are indicated by *end risk*

[3] Only 2.7% of respondents had top coded income values of $175,000. To adjust for the influence of top coding, the top coded income levels were multiplied by 1.5.

Table 3 Sample characteristics

Variable	Mean	Standard deviation
Cost	\$218.77	208.74
Income	\$63,110.67	49,880.26
Initial cancer risk p	3.49×10^{-5}	8.71×10^{-6}
Final cancer risk q	7.53×10^{-6}	8.32×10^{-6}
End risk zero	0.50	0.50
Considers own cancer risk high	0.052	0.222
Age	48.7	16.1
Environmentalist	0.42	0.49
Female	0.52	0.50
Northeast	0.18	0.39
Midwest	0.24	0.43
South	0.35	0.48
West	0.22	0.42
Lives in an SMSA	0.84	0.37
N	3430	

zero. These variables for whether respondents believed that they faced high risks of arsenic-related bladder cancer or received a survey option in which the risk was reduced to zero will be included in the regression models to account for whether the cancer valuations may be capturing perceived risk levels other than those stated in the survey. Along similar lines, some of the analyses will include indicator variables for whether the respondent lives in a standard metropolitan statistical area (SMSA) as well as for three of the four major regions of the country. These variables will also capture differences in risk exposures and water treatment facilities.

4.2 The empirical model

The model to be estimated will involve two principal procedures. After taking a first-order Taylor series approximation to Eq. 1, I adopt a logarithmic utility function similar to that used in most of the studies listed in Table 1. In large part because of the small probabilities for p and q that were involved, it was not feasible to estimate some more general frameworks. Subsequent studies might explore preferences at larger risk levels as well as the robustness of the results to additional functional forms.

The first-order Taylor series approximation to $u(y-c)$ is $u(y)-cu'(y)$, and the counterpart value for $v(y-c)$ is given by $v(y)-cv'(y)$. Thus, we can rewrite Eq. 1 as

$$(1-p)u(y) + pv(y) = (1-q)(u(y)-cu'(y)) + q(v-cv'(y)). \quad (2)$$

After some simplification, this equation can be rewritten as

$$c = ((p-q)[u(y)-v(y)])/\left[(1-q)u'(y) + qv'(y)\right]. \quad (3)$$

This expression does have an economic interpretation. If both sides of the equation are divided by $(p - q)$, then the result is that $c/(p - q)$ equals the standard formula for the VSL in the situation in which p and q pertain to mortality risks.

For logarithmic utility functions, let $u(y) = \ln y$ so that $u'(y)$ equals $1/y$, and $v(y) = \alpha \ln y$ so that $v'(y) = \alpha/y$. Substituting these values into Eq. 3 leads to

$$c = (p-q)(\ln y)(1-\alpha)y/[(1-q) + q\alpha]. \tag{4}$$

The key matter of interest is the value of α, which we expect to satisfy $0 < \alpha < 1$ if cancer reduces the marginal utility of income.

To account for average influence on willingness-to-pay amounts of other variables that affect the value of c, the estimating equation will be of the form

$$c = \gamma + [((p-q)(\ln y)(1-\alpha)y)/((1-q) + q\alpha)]. \tag{5}$$

In the most extensive version of the model, both γ and α are linear functions of other variables. Thus, the intercept term γ is a function of *considers own cancer risk high* and *end risk zero*, or $\gamma = \gamma_0 + \gamma_1 \times$ *considers own cancer risk high* $+ \gamma_2 \times$ *end risk zero*. Some equations will also include three regional variables and one SMSA variable as part of the set of intercept terms. Similarly, the utility function parameter α is a linear function of three personal characteristic variables, leading to $\alpha = \alpha_0 + \alpha_1 age + \alpha_2 female + \alpha_3 environmentalist$. Negative values for α_1, α_2, and α_3 indicate greater adverse impacts of these characteristics on the utility function in the post-cancer state.

4.3 Cancer utility function estimates

Table 4 presents a series of three equations estimating different versions of Eq. 5. The first equation includes no covariates and reports on the average values that are estimated for γ and α. The second equation permits the value of γ to be dependent on *considers own cancer risk high*, *end risk zero*, and a set of four locational characteristic variables. The third equation also permits the utility function parameter α to be a function of *age*, *environmentalist*, and *female*.

The results in column 1 indicate an average utility function parameter α value of 0.55, with a 95% confidence interval of (0.18, 0.91). The estimate is consequently consistent with cancer diminishing the marginal utility of income. Although the point estimate for the effect of cancer on the marginal utility of income is not estimated precisely, the confidence interval does not include a value for utility function parameter of 1.0. This is the expected result for severe health effects and is consistent with the description of the extensive effects of cancer that were described in the survey.

The addition in Eq. 2 of Table 4 of the set of variables that are likely to influence the level of risk beliefs does not alter the overall result. Accounting for the four locational characteristics and for whether the respondent considers *own cancer risk high* and took a survey with *end risk zero* has effects that do not alter the overall utility function parameter estimate. The average value remains largely unchanged and is 0.546. The small group of respondents who believe that they face a particularly high cancer risk are willing to pay just over $100 more for water treatment that reduces the arsenic-related cancer risk. Despite the evidence in the literature that people sometimes place

Table 4 Nonlinear least squares regression of cost respondents are willing to pay

	1 Baseline model	2 Model with intercept interactions	3 Model with parameter interactions
Constant terms			
γ_0	209.8***	222.4***	217.5***
	(5.1)	(12.7)	(12.7)
γ_1 Considers own risk high	–	115.7***	117.8***
		(15.9)	(15.9)
γ_2 End risk zero	–	12.25	12.83
		(7.11)	(7.07)
Utility Function			
α_0	0.545**	0.546**	2.637***
	(0.185)	(0.186)	(0.445)
α_1 Age	–	–	−0.030**
			(0.009)
α_2 Environmentalist	–	–	−1.139***
			(0.265)
α_3 Female	–	–	−0.552*
			(0.257)
R^2	0.02	0.02	0.03

Standard errors are in parentheses. Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Eqs. 2 and 3 include intercept interactions with SMSA, Midwest, South, and West

irrationally high values on policies that reduce the risks to zero, the *end risk zero* variable has a small coefficient and is not statistically significant. The general result that cancer reduces the marginal utility of income continues to hold, with negligible effect on the utility function estimates.

The final equation in Table 4 also includes the three personal characteristic variables. The results indicate the presence of substantial heterogeneity in utility functions based on these personal characteristics. All three variables have negative coefficients and are statistically significant at the 0.05 level or better, indicating a lower marginal utility in the ill health state. The negative signs on these variables indicate that each of these groups suffers a greater utility loss from cancer risks and should be willing to pay more to reduce these risks. These results are quite reasonable. Older respondents may place a higher value on cancer risk due to their greater familiarity with the morbidity effects of cancer. Since the risks in the survey arose from arsenic contamination in the water, it is not surprising that respondents who consider themselves to be environmentalists place a greater value on arsenic risk reductions. Consistent with evidence that females often display a greater willingness to pay to avoid health risks, female respondents had lower utility function parameters indicating a lower valuation of ill health states.

A useful corroboration of these results is to examine whether they are consistent with the patterns of preferences displayed in more conventional stated preference values. The results reported in Viscusi et al. (2014) indicate significantly greater willingness to pay for cancer risk reductions for the *age*, *environmentalist*, and *female* variables. The lower estimated values of γ for these groups are consistent with the direct evidence on the differences in willingness to pay with respect to these personal characteristics.

5 Implications

The effect of adverse health effects on well-being is not tantamount to a monetary loss in the case of serious health impacts, but mild health effects can be treated in that manner. If the health outcome was equivalent to a monetary loss l, the economic implications would be simple. Compensation equal to the amount of the loss l would restore the injured person to the pre-injury level of utility, thus serving as the "make whole" amount in legal contexts. This amount would also provide efficient levels of deterrence in personal injury contexts. Risk-averse individuals purchasing insurance would select insurance providing a coverage amount of l for such losses if they faced actuarially fair insurance rates. The value of l would also serve as the value of a statistical injury, or the amount per unit risk that they would be willing to pay to reduce the risk.

Once the adverse health event alters the utility function structure by reducing the marginal utility of income for any given income level, matters become quite different. Just as the marginal utility of income in a person's bequest function is typically lower than the marginal utility of income when the person is alive, there is a similar drop in the marginal utility of income after severe health events. There will be parallel influences of severe health impacts on three measures of possible interest that one could construct—the risk-money tradeoff for cancer (i.e., the analog of the VSL), the amount that is required to restore the person's welfare after getting cancer, and the optimal insurance amount that a person would select to provide income to the post-cancer state.

Understanding the shape of the utility function also enables one to assess the willingness to pay for non-marginal changes. A common concern of government agencies is whether non-marginal changes in risk reductions should be valued more or less than much smaller reductions in risk for which the usual risk-money tradeoffs for small risks are applicable. Using utility functions such as those discussed here demonstrates that the willingness to pay for these risk reductions is reduced if the decrease in risk levels is large. However, even large risk reductions for government policies are typically quite small so that the usual willingness-to-pay values remain applicable. For example, even a large reduction in the arsenic-related cancer risk of 4/100,000 is about the same as the annual U.S. workplace fatality rate used in estimates of the VSL (Viscusi 2018). Should risk changes be far outside of that range, one can use the insight provided by the functional form of the utility function to assess the extent of the modification in the benefit value that is warranted.

There has been substantial confusion in the literature to date on the effect of adverse health impacts on the structure of utility functions for two principal reasons. First, the impact on the utility function structure depends on the nature of the health impact. Mild health losses will not influence a person's subsequent marginal utility of income, whereas more catastrophic health losses will. Second, there are substantial differences in what the studies are measuring. Ideally, we would like to obtain an assessment of how the pre-injury utility function $u(y)$ differs from the post-injury utility function $v(y)$. Such a comparison requires that one equate the income levels y in each state. If, for example, there are medical expenses in the ill health state that are not taken into account, that will tend to boost the observed marginal utility of income when in ill health. Assessing such an effect on the marginal utility of income is a matter of legitimate academic interest, but it does not address the question of whether the utility function has been transformed by the ill health event.

The effect of severe health impacts also indicates why it is not appropriate to award values linked to the VSL as routine measures of compensation in wrongful death cases. As discussed in Viscusi (2018), in some jurisdictions the courts have permitted the application of the VSL to establish appropriate levels of compensation. The implication of the estimates of the utility functions for severe health effects is that there is a substantial reduction in the marginal utility of income, making full insurance less attractive. Similarly, high levels of compensation implied by the VSL are also unwarranted. In the case of the utility functions estimated for job risks in Table 1, the optimal level of income replacement after job injuries is 85% not 100%. Less than full insurance is the optimal level of insurance based on workers' own preferences. Reduction in the marginal utility of income diminishes the desirability of shifting resources to the ill health state for purposes of consumption.

The potentially substantial effect of severe adverse health effects on the marginal utility of income has fundamental ramifications for policy. In addition to implying that the optimal level of insurance is reduced, the shift in utility function structure highlights the powerful welfare-enhancing value of medical expenditures that can alter the health state. Allocations that involve either rehabilitation efforts or other interventions that improve the underlying health status may be especially valuable to the extent that they enable people to derive greater welfare benefits from their consumption expenditures. As in the situation of people in good health, money matters. But for those who suffer from severe health impairments, allocation of funds to enhance their health may be preferable to comparable expenditures on personal consumption.

Acknowledgements The author is indebted to James K. Hammitt for extremely helpful comments. Rachel Dalafave and Scott Jeffrey provided excellent research assistance.

References

Ameriks, J., Briggs, J. S., Caplin, A., Shapiro, M. D., & Tonetti, C. (2017). *Long-term care utility and late-in-life saving.* NBER Working Paper No. 20973.

Brown, J. R., Goda, G. S., & McGarry, K. (2016). Heterogeneity in state-dependent utility: Evidence from strategic surveys. *Economic Inquiry, 54*(2), 847–861.

Crainich, D., & Eeckhoudt, L. (2017). Average willingness to pay for disease prevention with personalized health information. *Journal of Risk and Uncertainty, 55*(1), 29–39.

De Nardi, M., & Yang, F. (2014). Bequests and heterogeneity in retirement wealth. *European Economic Review, 72*, 182–196.

Evans, W. N., & Viscusi, W. K. (1991). Estimation of state-dependent utility functions using survey data. *Review of Economics and Statistics, 73*(1), 94–104.

Evans, W. N., & Viscusi, W. K. (1993). Income effects and the value of health. *Journal of Human Resources, 28*(3), 497–518.

Evans, W. N., & Viscusi, W. K. (1998). Estimation of revealed probabilities and utility functions for product safety decisions. *Review of Economics and Statistics, 80*(1), 28–33.
Finkelstein, A., Luttmer, E. F. P., & Notowidigdo, M. J. (2013). What good is wealth without health? The effect of health on the marginal utility of consumption. *Journal of the European Economic Association, 11*(S1), 221–258.
Gyrd-Hansen, D. (2017). A stated preference approach to assess whether health status impacts on marginal utility of consumption. *Health Economics, 26*(10), 1224–1233.
Hammitt, J. K., & Haninger, K. (2010). Valuing fatal risks to children and adults: Effects of disease, latency, and risk aversion. *Journal of Risk and Uncertainty, 40*(1), 57–83.
Hammitt, J. K., & Liu, J.-T. (2004). Effects of disease type and latency on the value of mortality risk. *Journal of Risk and Uncertainty, 28*(1), 73–95.
Kahneman, D. (2011). *Thinking, fast and slow.* New York: Farrar, Straus, and Giroux.
Kopczuk, W., & Lupton, J. P. (2007). To leave or not to leave: The distribution of bequest motives. *The Review of Economic Studies, 74*(1), 207–235.
Levy, M., & Rizansky Nir, A. (2012). The utility of health and wealth. *Journal of Health Economics, 31*(2), 379–392.
Lillard, L. A., & Weiss, Y. (1997). Uncertain health and survival: Effects on end-of-life consumption. *Journal of Business and Economic Statistics, 15*(2), 254–268.
Magat, W. A., Viscusi, W. K., & Huber, J. (1996). A reference lottery metric for valuing health. *Management Science, 42*(8), 1118–1130.
McDonald, R. L., Chilton, S. M., Jones-Lee, M. W., & Metcalf, H. R. T. (2016). Dread and latency impacts on a VSL for cancer risk reductions. *Journal of Risk and Uncertainty, 52*(2), 137–161.
Moore, M. J., & Viscusi, W. K. (1990). Models for estimating discount rates for long-term health risks using labor market data. *Journal of Risk and Uncertainty, 3*(4), 381–401.
Sloan, F. A., Viscusi, W. K., Chesson, H. W., Conover, C. J., & Whetten-Goldstein, K. (1998). Alternative approaches to valuing intangible health losses: The evidence for multiple sclerosis. *Journal of Health Economics, 17*(4), 475–497.
Smith, D. M., Langa, K. M., Kabeto, M. U., & Ubel, P. A. (2005). Health, wealth, and happiness: Financial resources buffer subjective well-being after the onset of a disability. *Psychological Science, 16*(9), 663–666.
Spence, M. (1977). Consumer misperceptions, product failure, and producer liability. *Review of Economic Studies, 44*(3), 561–572.
Tengstam, S. (2014). Disability and marginal utility of income: Evidence from hypothetical choices. *Health Economics, 23*(3), 268–282.
Van Houtven, G., Sullivan, M. B., & Dockins, C. (2008). Cancer premiums and latency effects: A risk tradeoff approach for valuing reductions in fatal cancer risks. *Journal of Risk and Uncertainty, 36*(2), 179–199.
Viscusi, W. K. (2018). *Pricing lives: Guideposts for a safer society.* Princeton: Princeton University Press.
Viscusi, W. K., & Evans, W. N. (1990). Utility functions that depend on health status: Estimates and economic implications. *American Economic Review, 80*(3), 353–374.
Viscusi, W. K., & Evans, W. N. (2006). Behavioral probabilities. *Journal of Risk and Uncertainty, 32*(1), 5–15.
Viscusi, W. K., & O'Connor, C. J. (1984). Adaptive responses to chemical labeling: Are workers Bayesian decision makers? *American Economic Review, 74*(5), 942–956.
Viscusi, W. K., Huber, J., & Bell, J. (2014). Assessing whether there is a cancer premium for the value of a statistical life. *Health Economics, 23*(4), 384–396.