The final publication is available in: International Journal of Business Continuity and Risk Management, 7, 3, 2017, 192-203

10.1504/IJBCRM.2017.088806

Does valuation research provide a credible basis for cost-benefit analysis of safety measures?

Rune Elvik

Institute of Transport Economics

Gaustadalleen 21, 0349 Oslo, Norway

ABSTRACT

Many studies have been made to obtain a monetary valuation of reduced risk of death, usually given as the value of a statistical life (VSL), which is a reduction in risk corresponding to the prevention of one fatality. This paper asks whether valuation research provides a credible basis for cost-benefit analysis of safety measures. A cost-benefit analysis is credible if its results cannot be criticised by reference to the valuation studies forming its basis. It is argued that a credible basis for cost-benefit analysis in this sense does not exist. The monetary valuations of a statistical life vary enormously. The enormous diversity in values is increasingly accepted by researchers working in the field as inevitable and consistent with individual utility maximisation, and thus not necessarily anomalous. Some recent contributions argue that the value of a statistical life ought to vary depending on, for example income. Such reformulations of the theory underlying valuation studies mean that the choice of a particular value of a statistical life within the huge range of such values is not necessarily more justified than the choice of a different value.

Key words: Value of a statistical life; diversity of values; choice of value; cost-benefit analysis; credibility

1 INTRODUCTION

Modern research on the valuation of safety, in particular reduced risk of accidental death, started around 1970. At that time, two widely quoted papers by Schelling (1968) and Mishan (1971) formulated the theoretical foundations for empirical research designed to obtain a monetary valuation of changes in the risk of death. Both papers argued that the only meaningful approach to the monetary valuation of changes in mortality risk is the willingness-to-pay approach. This point of view is virtually unanimously accepted by economists working in the field.

Several studies of willingness-to-pay (WTP) for changes, most often reductions, in mortality risk have been made. In most studies, risk is stated as the annual number of deaths from a specific cause per 100,000 inhabitants of a country. Willingness-to-pay is usually stated as the value of a change in risk that statistically corresponds to the prevention of one death, often referred to as the value of a statistical life (VSL) or the value of preventing a fatality. A recent meta-analysis of stated preference studies (Lindhjem et al. 2011) included 856 estimates of the value of a statistical life. These varied (in 2005 US dollars) between 4,450 and 197,000,000 – a factor of 44,000. There is, in other words, an enormous variation in estimates of the value of a statistical life. This, by itself, would not necessarily be a problem if a good explanation of the diversity of values could be given and, based on such an explanation, the best value selected for a given application. One would then define a context in terms of, for example, initial risk level, size of change in risk, mean income, etc. and select the most appropriate value of a statistical life for that context.

Using such a procedure would, however, only give meaningful results if: (1) The diversity of values of a statistical life was attributable mainly to theoretically relevant sources, and not, for example, to artefacts of the methods used in valuation studies, (2) A single value, or a narrow range of values, could be estimated, for example, by means of a meta-regression model.

Unfortunately, the procedure sketched above is unlikely to work. The chief reason for that is that recent developments in the theory of willingness-to-pay have made it increasingly difficult to distinguish between results of valuation studies that make sense from a theoretical point of view and results that do not make sense from a theoretical point of view. These theoretical developments have profound implications for the interpretation of diversity in estimates of the value of a statistical life and for analyses designed to identify theoretically meaningful patters of variation in estimates of the value of a statistical life. This paper argues that by trying to make sense of the wide diversity of values of a statistical life by developing new theoretical models, theorists have, perhaps unintentionally, undermined the basis for choosing the best-founded value of a statistical life in a particular context.

2 THEORETICAL MODELS OF VARIATION IN THE VALUE OF A STATISTICAL LIFE

The papers by Schelling (1968) and Mishan (1971) were quite general and did not put forward specific hypotheses about willingness-to-pay for changes in mortality risk. It did not take long, however, before more specific theoretical propositions were developed. An early paper contributing to the development of theory was written by Jones-Lee (1974).

Jones-Lee (1974) proposed hypotheses about how an individual would value changes in mortality risk, based on the following assumptions:

1. The individual maximises expected utility (which is a probability-weighted utility of a lottery with life and death as potential outcomes).
2. The individual prefers more wealth to less and is financially risk averse (prefers an income received with certainty to an uncertain income).
3. The individual does not want descendants to be exposed to a greater financial risk than himself or herself.
4. At a given level of wealth, the individual prefers to be alive rather than dead.
5. The marginal utility of wealth is greater when the individual is alive than when the individual is dead.

Based on these assumptions, Jones-Lee could deduce that a positive willingness-to-pay for reduced risk of death will exist. He further deduced that willingness-to-pay will be positively related to income and positively related to the level of risk. This example shows how one can use theoretical predictions to assess whether empirical results make sense or not. If you find that willingness-to-pay varies systematically as predicted by theory, results make sense. This example shows the essential function of theory in willingness-to-pay research: It is to predict a systematic pattern of variation in willingness-to-pay that may serve as reference in assessing whether the results of empirical studies make sense or not.

It did not take long, however, before more complex models were developed and predictions became ambiguous. The extremely complex model proposed by Dehez and Drèze (1982) is an example. Here are the predictions of this model:

1. If an individual does not have life insurance or an annuity, and if the marginal utility of money is greater when alive than when dead, willingness to pay will increase when risk level increases.
2. If the individual has optimal life insurance and annuity at actuarially fair rates, willingness to pay is independent of the level of risk.
3. If the individual holds life insurance and annuity at less than actuarially fair rates, willingness to pay will increase as the level of risk goes down.
4. If the individual holds life insurance and annuity at more than actuarially fair rates, willingness to pay will increase as risk level increases.
5. If the individual has life insurance and annuity and the terms of the contracts are adjusted as risk level changes, willingness to pay will increase as risk level decreases.

One could say that they hedge their bets. Everything is possible; that willingness to pay does not depend on risk level, that it increases with risk level, or that it decreases with risk level. None of these findings is ruled out theoretically. No matter what you find, it has theoretical support – unless you can collect detailed data on the insurance coverage of respondents. But even if such data are available, it may be difficult to determine if insurance is actuarially fair or not.

As time went by, theorists expanded the range of topics they addressed. By now, the following topics have been discussed in theoretical contributions:

1. General characteristics of the individual valuation function for changes in risk,
2. The relationship between the level of the risk and willingness-to-pay,
3. The relationship between the size of the change in risk and willingness-to-pay,
4. The relationship between the direction of changes in risk and valuation of the changes,
5. The nature of the good producing changes in risk (private or public),
6. The relationship between individual characteristics and willingness-to-pay,
7. The effect of experiencing injury or a life-threatening event on willingness-to-pay,
8. The effects of income and insurance coverage on willingness-to-pay,
9. The relationship between human capital and willingness-to-pay,
10. The effects of the distribution of risk and wealth on willingness-to-pay,
11. Benevolence and altruism,
12. The degree of financial risk aversion,
13. The existence of background risks.

For many of these topics, alternative theoretical models have been developed. This means that when all theoretical models are viewed as a whole, few results of valuation studies can be ruled out on theoretical grounds. In other words: The theory of willingness to pay for changes in risk contains so many contradictory hypotheses that it can no longer be falsified and serve the essential function of distinguishing results that make sense from results that do not. Table 1 summarises the current state of affairs.

Table 1 about here

Consider, for example, the size of the change in risk. For a long time, the theoretical prediction was that willingness to pay would vary in proportion, or near proportion to the size of the change in risk. Many empirical studies did not find this, a phenomenon that was referred to as insensitivity to scope. Efforts were made to better explain changes in risk in order to make respondents more sensitive to the size of these changes. These efforts had moderate success. But then in 2003 and, more extensively, in 2010 Amiran and Hagen (2003, 2010) proposed a new model, directionally bounded utility functions, that predicted insensitivity to scope. They argued that conventional utility functions imply a property they call hypersubstitutability, which means that there are no limits to how much an individual will give up of a specific good in order to obtain another good. They argue that hypersubstitutabiliy is implausible (2010:294):

"It implies, for example, that consumers who prefer more birds to fewer birds would be willing to give up nearly all of their housing (and other material goods) in exchange for a sufficient number of additional Purple Martins. Consumers are unlikely to commit themselves to extreme material poverty in order to achieve an incremental gain in any environmental amenity."

They illustrated the difference between a directionally bounded utility function and a standard utility function by means of a figure. Figure 1 is taken from their 2010-paper. Panel a shows a directionally bounded utility function, with the bounds indicated by dashed lines. Panel b shows a standard utility function. It is seen that sensitivity to scope is smaller according to the directionally bounded utility function than according to a standard utility function.

Figure 1 about here

Amiran and Hagen conclude as follows:

"Unlike other explanations for low sensitivity to scope, our results are shown to be consistent with a rational, self-interested consumer, whose preferences are consistent with all of the neoclassical axioms. Given the plausibility of directionally bounded utility functions, there is no simple a priori basis for the notion that the degree of sensitivity to scope should be large."

Their contribution is typical of many recent contributions to theory about the willingness to pay for changes in mortality risk. The predictions of the theoretical models are typically qualitative only; they predict the direction of an effect, but not its strength. Interpretation thus becomes difficult when hypotheses make contradictory predictions. If empirical findings are in a region of doubt, i.e. they can be consistent with more than one underlying utility model, attempting to ascertain which utility model best explains the results is likely to be inconclusive. This means that few results can be ruled out on theoretical grounds. Theory gives little guidance in the selection of the theoretically best-supported estimates of the value of a statistical life.

3 REFORMULATIONS OF THE NORMATIVE BASIS OF COST-BENEFIT ANALYSES

When modern research on the monetary valuation of changes in mortality risk started, an important element of the justification for it was that it was only by applying a single, uniform value of a statistical life that efficient priorities could be set. Hills and Jones-Lee (1983) put this argument in very clear terms and illustrated it by numerical examples. There is little doubt that the initial objective of empirical studies was to find a good estimate of the value of a statistical life that could be applied uniformly to all problems involving expenditures on safety programmes.

It did not take long, however, before the great diversity in estimates of the value of a statistical life cast doubt on the prospects of finding a single value that could be applied uniformly. Nevertheless, as explained above, if the variation in values could be explained theoretically, it might still be possible, guided by theory, to select the best-supported value of a statistical life. Today, this is impossible. Even findings reflecting insensitivity to scope are held, by some at least, to make sense from a theoretical point of view.

Some recent contributions that discuss the valuation of changes in risk from a prescriptive point of view have moved away from the traditional emphasis on using a uniform value to maximise efficiency. In two papers, Baker et al. (2008, 2009) discussed whether the widespread practice of using a uniform value of a statistical life, i.e. the arithmetic mean of willingness-to-pay (WTP) in a population, is consistent with the theoretical foundations of cost-benefit analysis, and, if not, if an acceptable normative foundation can be defined for applying a single value of a statistical life (VSL). They note that (2009:814, 815):

"In spite of the tendency to apply a uniform VSL in any given context, so far as we are aware of no satisfactory theoretical foundation has so far been provided that justifies the application of a common WTP-based VSL equal to, say, the overall population arithmetic mean of marginal rates of substitution (MRS) of wealth for risk of death by a given cause, other than under conditions which from a practical point of view appear somewhat implausible. … To the extent that the marginal

rates of substitution will typically depend upon the income, age and other personal characteristics of those affected by the safety improvement, the logic underpinning standard cost-benefit analysis would seem to require that the VSL employed in the evaluation of a safety improvement that affects a poorer (or older) group in society should be smaller than the value applied to a wealthier (or younger) group. … Clearly, though, if a normative rationale is to be provided for this "uniform valuation of safety" approach then this will have to rely upon value judgments that differ somewhat from those underpinning conventional cost-benefit analysis."

In short, the principles of cost-benefit analysis, as presented here by Baker et al. imply that VSL should vary according to the variation in WTP between different groups in society. A uniform value is inconsistent with cost-benefit analysis. Viscusi (2010) has also discussed the policy challenges associated with the huge variation in estimates of the value of a statistical life (VSL). He states that whether heterogeneity in estimates of VSL should be incorporated in policy evaluations (cost-benefit analyses) depends in part on the source of heterogeneity. However, as far as age and income are concerned, he clearly recommends taking account of these factors by applying VSL-estimates that vary by age and income. Regarding age, he states (2010:121):

"In many contexts, such as those involving regulations that affect people with very short remaining life expectancy, it is not appropriate to use the standard VSL measure. Rather, taking into account the difference in longevity often leads to the use of the VSLY (value of a statistical life year)."

He goes on to show that the value of a statistical life year depends on age. It has an inverted U-shape and reaches maximum at an age of about 50 years (based on empirical studies). Viscusi sums up whether income effects should matter in the following terms:

"The proper benefits measure should be grounded in the WTP of the beneficiaries of the policy. Whether these individual preferences indicate a positive income elasticity or other types of heterogeneity in preferences does not invalidate the importance of adhering to reliance on individuals' WTP for guidance in setting benefit levels. To impose constraints on income adjustments or any other aspects of the benefit assessment process in effect overrides individual preferences and the pivotal economic role of consumer sovereignty."

These points of view might seem to open a Pandora's box. If the value of a statistical life ought to vary according to age and income, how about other sources of variation, like gender, insurance coverage, or the nature of the good affecting risk. Which factors are relevant, which are not? And what is the plausible range of variation in the value of a statistical life to be permitted?

There are no obvious answers to these questions. By allowing the value of a statistical life to vary, theorists have not only abandoned the quest for a uniform value that initially motivated research, but implicitly approved of any set of values whose range is not outrageously wide. Yet, exactly how wide the permitted range of values can be remains unclear.

4 A RANGE OF VALUES BASED ON META-ANALYSIS

One option is to estimate a range of values based on meta-analysis. There have been many meta-analyses of the literature about the value of a statistical life; for a summary see Elvik (2016A). The most comprehensive of these analyses, in terms of the number of estimates of VSL included and the depth of statistical analysis, is the one reported by Lindhjem et al. (2011, 2012).

Lindhjem et al. (2011) developed several models to explain variation in the value of a statistical life. Models were fitted both to the full dataset and to subsets of it. Perhaps the most informative model as far as traffic risk is concerned, is Model V fitted to the subset that contained data about change in risk. The model explained 83 percent of the variation in the value of a statistical life. The strongest predictors of the value of a statistical life were GDP per capita and the size of the change in risk. To illustrate the range of predicted values, estimates have been developed for the following combinations of values:

GDP per capita: 40,000; 60,000; 80,000

Risk change: $5 \cdot 10^{-6}$; $10 \cdot 10^{-6}$; $15 \cdot 10^{-6}$

Type of good: private or public

GDP per capita is in the range of current GDP per capita in Norway (covering most individuals; values given in US dollars adjusted to purchasing power parity). The changes in risk represent realistic levels of change in the risk of a road accident fatality in Norway (current fatality rate is about 25 per million inhabitants; hence a reduction of 15 per million would imply reducing traffic fatalities by about 60 percent). Measures taken to reduce risk can either be private goods (buying a safer car) or public goods (installing road lighting). Results are given in Table 2. The values are in US dollars 2005-prices, adjusted to purchasing power parity.

Table 2 about here

It is seen that the estimated value of a statistical life varies by a factor of about 8 even for this limited range of combinations of values on the explanatory variables. All three variables included are seen to have a major effect on estimates of the value of a statistical life.

It may perhaps strike readers as counterintuitive that the value of a statistical life declines as a function of the size of the change in risk. Thus, when reduction in risk is a private good, the value of a statistical life at a risk reduction of 15 per million is only about half the value at a risk reduction of 5 per million. Hence, all else equal, a small risk reduction seems to be more attractive than a large risk reduction. This, however, is a result of insensitivity to scope. It is easy to calculate backwards and find the mean willingness-to-pay (WTP) for the stated changes in risk. Thus, at an income of 60,000, one finds that mean WTP increases from 200 for a risk reduction of 5 per million, to 268 for a risk reduction of 10 per million and 318 for a risk reduction of 15 per million. In other words, individuals prefer the largest risk reduction, for which the value of a statistical life is lowest. This is the consequence of accepting insensitivity to scope as an expression of rational trade-offs made by individuals. There is perhaps no alternative to doing so. Were one to reject studies showing insensitivity to scope, virtually all valuation studies would have to be rejected.

Acceptance of insensitivity to scope does, however, lead to a preference reversal when mean individual willingness to pay is aggregated to the value of a statistical life. In practice, no government would use a range of values like those given in Table 2 in cost-benefit analysis. A single value would be selected and used for all projects. Yet all values presented in Table 8 are equally well justified. Choice of a specific value could therefore be based on strategic considerations. One could, for example, argue that incomes are growing, that improvements in safety will come in the form of safer cars (an individual good), and that we can at present only estimate the effects of road safety measures (including safer cars) that would reduce risk by 5 per million to justify a choice of the highest value. Equally plausible considerations might lead to the choice of a lower value.

5 DISCUSSION

Cost-benefit analyses are controversial for many reasons, but one of them is that not everybody trusts the monetary valuations of non-market goods used in cost-benefit analyses. Indeed, as far as the valuation of human life is concerned, the values are, as one critic (Hauer 2011) puts it "all over the place". It is impossible to disagree with this description, given the fact that a comprehensive meta-analysis (Lindhjem et al. 2011) included values of a statistical life (in 2005 US dollars) ranging from 4,450 to 197,000,000. Many would say that this fact alone discredits any cost-benefit analysis based on any value of human life within the huge range of such values found in the literature.

Yet, proponents of cost-benefit analysis might argue that not all values in the huge range are to be trusted. Some of these values must be rejected and the remaining lie within a narrower range. This is correct, but even if the presumably "best" values are selected, the range remains huge (Lindhjem et al. 2012). Even within a single study, based on state-of-the-art methods, the range of values is enormous. Thus, based on a

Norwegian valuation study (Veisten et al. 2010, 2013), Veisten (2016) extracted 66 estimates of the value of a statistical life ranging from 15.8 to 362.7 million NOK. Which of these estimates is the best one? It is difficult to select a single one of the 66 estimates as clearly superior to the others. However, it is plausible to argue that some of the estimates are more credible, i.e. more trustworthy, than the others in the sense that they are less likely to be influenced by methodological weaknesses or answers that are inconsistent with economic theory. A subset of 22 estimates was defined, by excluding those who answered lexicographically in stated choice experiments (to answer lexicographically is to always prefer the alternative that is best with respect to a specific attribute, for example safety, regardless of the values of other attributes) and including only respondents who indicated that they were certain about their answers (using a certainty scale).

The range of values of a statistical life in the subset of 22 estimates was from 15.8 to 151.5 million NOK. This still leaves considerable room for choice. A weighted mean, estimated by means of meta-analysis, was 30.3 million NOK according to a fixed-effects model and 52.8 million NOK according to a random-effects model. According to the rules of the game for meta-analysis (Elvik 2016B), one would normally prefer the random-effects estimate. The data show, however, that there is a positive relationship between the estimate of the value of a statistical life and its standard error, which means that higher estimates of the value of a statistical life contribute proportionately more to the random-effects statistical weights than to the fixed-effects statistical weights, thereby pulling the weighted mean estimate upwards. If one accounts for this relationship by means of the precision-regression method proposed by Stanley and Doucouliagos (2013, 2014), the mean estimate of the value of a statistical life becomes 23.2 million NOK. The results of the meta-analysis are thus open to discussion.

This case illustrates the virtually endless complexities and choices one lands in when trying to extract a single estimate of the value of a statistical life from even a single study. What, then, to make of the vast diversity of estimates found in the literature? It is perhaps only a slight exaggeration to say that recent theoretical contributions to the value of life literature tell us not to worry about it. Do studies find insensitivity to scope? Well, do not worry about it. This is entirely what a directionally bounded utility function predicts. Is there, on the contrary, sensitivity to scope? Well, that is what a standard utility function would predict.

Theory, in other words, no longer gives much guidance in selecting values of a statistical life. The general public is no doubt unaware of this state of affairs. It nevertheless is clear that there is a large element of arbitrariness in the choice of a single value of a statistical life for use in cost-benefit analyses.

5 CONCLUSIONS

The main conclusions of the discussion presented in this paper are:

1. Estimates of the value of preventing a fatality, the value of a statistical life, vary enormously.
2. Over time, there has been a tendency towards reformulating the theory of willingness to pay for non-market goods so that findings that were initially regarded as anomalous, i.e. inconsistent with theory, are no longer regarded

as so. This means that the great diversity of estimates of the value of a statistical life is no longer regarded as inconsistent with economic theory.

3. This development has broadened the range of values of a statistical life that are regarded as theoretically plausible. However, theory cannot specify how wide the theoretically plausible range of values is, as its predictions are qualitative only (e.g. it predicts insensitivity to scope, but not how strong the insensitivity will be).
4. Even if one restricts permissible values of a statistical life by screening studies according to methodological criteria, a huge range of values remains. There are many ways of extracting a best estimate from such a restricted sample, and they do not necessarily give the same results.
5. Choice of a particular value within any range of values considered has a profound element of arbitrariness.

REFERENCES

Amiran, E. Y., Hagen, D. A. 2003. Willingness-to-pay and willingness to accept: How much can they differ? Comment. American Economic Review, 93, 458-463.

Amiran, E. Y., Hagen, D. A. 2010. The scope trials: Variation in sensitivity to scope and WTP with directionally bounded utility functions. Journal of Environmental Economics and Management, 59, 293-301.

Baker, R., Chilton, S., Jones-Lee, M., Metcalf, H. 2008. Valuing lives equally: Defensible premise or unwarranted compromise? Journal of Risk and Uncertainty, 36, 125-138.

Baker, R., Chilton, S. M., Jones-Lee, M. W., Metcalf, H. R. T. 2009. Valuing lives equally in a benefit cost analysis of safety projects: A method to reconcile theory and practice. Safety Science, 47, 813-816.

Dehez, P., Drèze, J. H. 1982. State-dependent utility, the demand for insurance and the value of safety. In Jones-Lee, M. W. (Ed): The value of life and safety, 41-65. Amsterdam, North-Holland Publishing Company.

Elvik, R. 2016A. The value of life. The rise and fall of a scientific research programme. Report 1531. Oslo, Institute of Transport Economics.

Elvik, R. 2016B. Meta-analytic methods. Chapter 19 in Lord, D., Washington, S. (Eds) Safe mobility: challenges, methodology and solutions (Volume VIII) (forthcoming).

Hauer, E. 2011. Computing what the public wants: some issues in road safety cost-benefit analysis. Accident Analysis and Prevention, 43, 151-164.

Hills, P. J., Jones-Lee, M. W. 1983. The role of safety in highway investment appraisal for developing countries. Accident Analysis and Prevention, 15, 355-369.

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

Lindhjem, H., Navrud, S., Braathen, N. A., Biausque, V. 2011. Valuing mortality risk reductions from environmental, transport and health policies: a global meta-analysis of stated preference studies. Risk Analysis, 31, 1381-1407.

Lindhjem, H., Navrud, S., Biausque, V., Braathen, N. A. 2012. Mortality risk valuation in environment, health and transport policies. Report. Paris, OECD Publishing.

Mishan, E. J. 1971. Evaluation of life and limb: a theoretical approach. Journal of Political Economy, 79, 687-705.

Schelling, T. C. 1968. The life you save may be your own. In Chase, S. B. (Ed), Problems in public expenditure analysis, 127-176. Washington D. C., The Brookings Institution.

Stanley, T. D., Doucouliagos, H. 2013. Neither fixed nor random: weighted least squares meta-analysis. Statistics in Medicine, 34, 2116-2127.

Stanley, T. D., Doucouliagos, H. 2014. Meta-regression approximations to reduce publication selection bias. Research Synthesis Methods, 5, 60-78.

Veisten, K. 2016. Et konkret forslag til en liten nedjustering av ex-ante verdsetting av "lettere skade" basert på Verdsettingsstudien. Arbeidsdokument 51015. Oslo, Transportøkonomisk institutt.

Veisten, K., Flügel, S. Elvik, R. 2010. Den norske verdsettingsstudien. Ulykker – verdien av statistiske liv og beregning av ulykkenes samfunnskostnader. Rapport 1053C. Oslo, Transportøkonomisk institutt.

Veisten, K. Flügel, S. Rizzi, L. I., Ortuzar, J. de D., Elvik, R. 2013. Valuing casualty risk reductions from estimated baseline risk. Research in Transportation Economics, 43, 50-61.

Viscusi, W. K. 2010. Policy challenges of the heterogeneity of the value of statistical life. Foundations and Trends in Microeconomics, 6, 99-172.

LIST OF FIGURES AND TABLES

Figure 1:

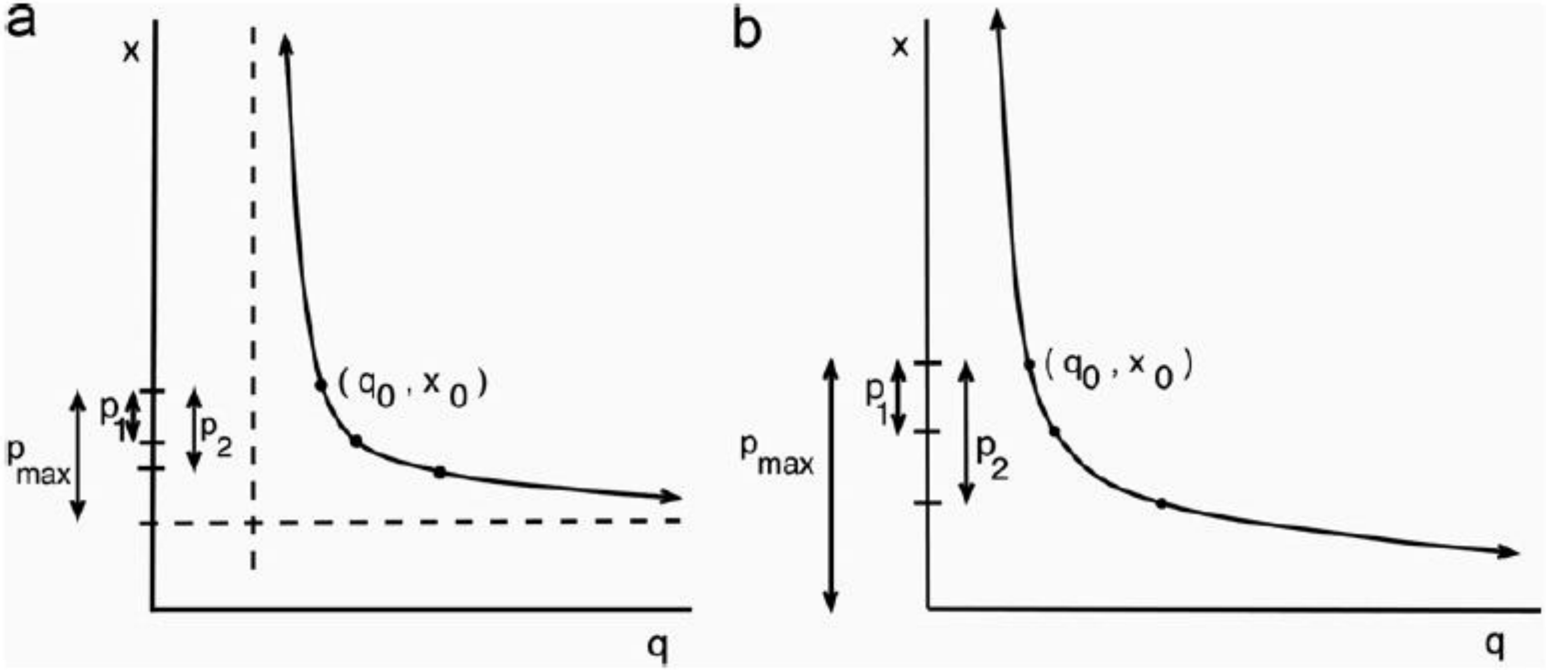

Table 1:

Topic addressed	Hypotheses proposed	Falsification possible
Shape of valuation function	Demand curve bounded at both ends	Yes, a curve with a different shape may be found
Level of risk	Relationship can be positive, negative or independent	No, any relationship found, or absence of a relationship, is consistent with at least one theoretical model
Life-threatening event	Relationship is indeterminate if individual is risk averse	No, any relationship found is consistent with at least one utility model
Size of change in risk	There can be sensitivity or insensitivity to the size of change in risk	No, both finding sensitivity and not finding it is consistent with a utility model based on hard core assumptions
Direction of change in risk	There can be both a small and a very large difference between compensation demanded for an increase in risk and willingness-to-pay for a reduction in risk	No, both finding a small difference between WTA and WTP and finding a large difference between WTA and WTP is consistent with utility models based on hard core assumptions
Nature of good	Safety as a public good can be valued both lower than, equal to or higher than safety as a private good	No, any differences between valuations of safety as a public or private good are consistent with some utility model based on hard core assumptions
Age of individual	The valuation of safety can have any relationship or no relationship to age	No, any finding is consistent with some model of the optimal path of lifetime consumption
Insurance coverage interacting with level of risk	Depending on the type of insurance coverage, valuation may be negatively, positively or unrelated to the level of risk	No, any sign and strength of the relationship between level of risk and willingness-to-pay is consistent with some utility model based on hard core assumptions
Income and wealth	A standard utility model predicts that willingness-to-pay is positively related both to income and wealth; a model based on prospect theory shows that a negative relationship to wealth is possible	Yes, for income there is a clear prediction of a positive relationship; the hypothesis is falsified if this is not found No, for wealth models predict both a positive and a negative relationship, making falsification impossible
Human capital	Under weak assumptions, VSL as estimated from willingness-to-pay will exceed human capital	Yes, finding that a WTP value is lower than a human capital estimate of VSL would falsify the hypothesis
Unequal distribution of risk	Willingness-to-pay may increase, decrease or first increase, then decrease when the distribution of risk becomes more egalitarian	No, the relationship can have different directions depending on the size of the risk reduction – the theory does not address the issue of promoting an egalitarian distribution of risk by reducing high risks more than low risks
Altruism	Willingness-to-pay should be higher for a paternalistic altruist than for an egoist	Yes, provided the motivations for willingness-to-pay can be truthfully revealed; this, however, is highly doubtful
Financial risk aversion	Different versions of utility models cannot predict the relationship between financial risk aversion and willingness-to-pay	No, the theoretical results are ambiguous – any relationship would therefore be consistent with theory
Background risks	A competing mortality risk may decrease willingness-to-pay for a target risk or have no effect on it	Partly, if a competing risk is found to increase willingness-to-pay for the target risk that would be inconsistent with theory
	A background financial may both decrease and increase willingness-to-pay for a target risk	No; while the conditions for a decrease or increase in willingness-to-pay are different, they are in practice unobservable

Table 2:

Risk reduction (per million)	GDP per capita (US dollars, 2005)		
	40,000	60,000	80,000
	Safety as a private good		
5	29,094,768	39,966,337	50,063,520
10	19,503,861	26,791,686	33,560,396
15	15,435,330	21,202,906	26,559,654
	Safety as a public good		
5	11,676,267	16,039,229	20,091,415
10	7,827,259	10,751,998	13,468,406
15	6,194,483	8,509,118	10,658,880