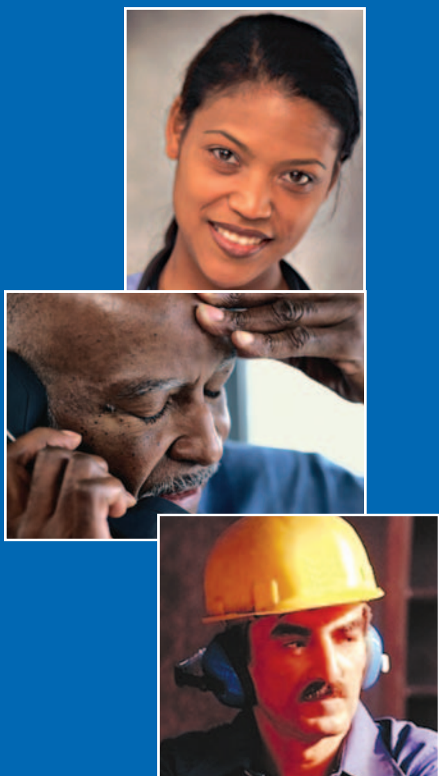

INSTITUTE FOR WORK & HEALTH
INSTITUT DE RECHERCHE SUR LE TRAVAIL ET LA SANTÉ

Lessons from the Literature on Valuing Reductions in Physical Risk

Working Paper #342

Address correspondence to: Richard Cookson, Ph.D.
Department of Social Policy and Social Work, University of York
YO10 5DD
York, England
Tel: +44 (0)1904 321248
Email: rc503@york.ac.uk

January 2008

WORKING PAPER # 342

Lessons from the Literature on Valuing Reductions in Physical Risk

Richard Cookson[1] and Peter Dorman[2]

[1]*Department of Social Policy and Social Work, University of York, England*
[2] *Evergreen State College, United States*

Lessons from the Literature on Valuing Reductions in Physical Risk

Richard Cookson[1] and Peter Dorman[2]

[1]Department of Social Policy and Social Work, University of York, England
[2] Evergreen State College, United States

Word count: 16, 017 (including all tables and references)

Correspondence to:
Richard Cookson, Department of Social Policy and Social Work, University of York, York YO10 5DD. rc503@york.ac.uk Tel: +44 (0)1904 321248.

Peter Dorman, Evergreen State College, Email: dormanp@evergreen.edu or dormanp@gmail.com

1. Introduction

The economic literature on valuing reductions in physical risk (or "valuing safety", for short) aims to value the prevention of fatal and non-fatal illness and injury in a manner suitable for incorporation into economic evaluation. There are a number of general literature reviews (Viscusi 1993, Beattie et al. 1998a) and textbooks (Jones-Lee 1989, Viscusi 1992, Dorman 1996).

At present, economic evaluations of workplace health and safety interventions typically focus on the financial cost savings associated with preventing worker illness and injury from an employer's perspective – for example, increased worker productivity, lower health care costs, lower worker compensation payments, and so on (Niven 2002). The valuing safety literature takes a rather different approach. It focuses on the "human" costs of illness and injury in terms of harm to individual well-being, rather than the financial costs in terms of lost profits to the firm or lost national income to the economy as a whole.

The valuing safety literature examines how individuals themselves value reductions in physical risk. It aims to quantify the value of intangible human outcomes such as pain, grief and suffering as well as more tangible human outcomes such as loss of life, limb and livelihood. It aims explicitly to value all such human outcomes in terms of a common index of value – i.e. money – that can be compared with all the other costs and consequences of the policy intervention under consideration. It is not concerned with *measuring* human outcomes – i.e. estimating the effectiveness of workplace interventions in terms of lives saved or injuries prevented or days of pain, grief and suffering avoided. Rather, it is concerned with *valuing* human outcomes of this kind in monetary terms suitable for economic evaluation.

The standard economic approach to valuing safety focuses on small changes in physical risks of death and ill-health. The economist analyses the trade-offs that individuals are willing to make (or that they might be imputed to make) between physical risk and wealth or other non-health benefits (e.g. time, convenience, consumption of goods and services, working conditions). Empirical estimates of the magnitude of such trade-offs can be made using either “revealed preference” data from real market transactions (in particular, imputed wage-risk trade-offs) or “expressed preference” data from questionnaires that ask people to make hypothetical trade-offs. The resulting monetary valuations can be thought of as values for *statistical* life. That is, they do not reflect the value of saving one identified person from imminent death. Rather, they ideally reflect the value of a small reduction in the probability of dying spread across a large number of individuals, which taken over the whole group will "on average" prevent one fatality – a scenario that more closely reflects the reality of most workplace health and safety interventions.

To see how a value of statistical life might be computed, consider the following highly simplified example. Imagine that a group of 100,000 workers benefit from an occupational health and safety regulation that reduces the probability of dying next year by 2 in 100,000 for each worker. The expected number of deaths next year will thus be reduced by two (100,000 workers * 2/100,000 probability of dying). Imagine that each worker is willing to pay (WTP) a maximum of $50 for this improvement, directly out of pocket or through lower wages. The group total maximum WTP is thus $50*100,000 = $5,000,000. Dividing this group total maximum WTP by the expected number of deaths prevented (i.e. two) yields a value of statistical life in this setting of $2,500,000.

If the total cost of the regulation were less than $5,000,000, then – according to the standard welfare economic approach – the regulation is considered worthwhile. This is because, hypothetically, the workers would be able and willing to pay for this regulation themselves, thus leaving themselves better off and no-one else worse off – a "no-brainer" (or, in the economics jargon, a "Pareto improvement"). For example, imagine the total cost were $4,000,000. Then the regulation could be funded if each worker were to pay $40. By hypothesis, each worker would then be better off – since their maximum WTP is $50. In the economics jargon, each would enjoy a "consumer surplus" of $10. Of course, in the real world things would probably not be so simple; nevertheless, the example serves to illustrate the basic idea.

Historically, empirical estimates of the value of statistical life started appearing in the literature in the 1970s, with most of the work originating from labour economics and transport economics. It is perhaps ironic that economic evaluations of workplace health and safety interventions have tended to ignore this literature – since so much of it originated from the study of workplace risk. The empirical literature on valuing safety mushroomed in the 1990s, with a substantial increase in both the volume and breadth of empirical studies, particularly in the area of expressed preferences – including studies in other sub-disciplines of economics (e.g. environmental economics and health economics), studies of the value of non-fatal injury and illness, and studies of variation in values by population subgroup (e.g. income and age) and hazard context. The literature on expressed preference has borrowed ideas and techniques from psychology and sociology, and is increasingly seen as a multi-disciplinary enterprise that overlaps with the fields of experimental and behavioural economics.

In this Chapter we will not review the health economic literature on valuing health outcomes in non-monetary terms such as quality-adjusted life-years (QALYs) or disability-adjusted life-years (DALYs) – the "valuing health" literature (Drummond et al. 2005). Health economists have developed and applied workable methods of valuing health for use in economic evaluations within the health sector. The QALY approach currently dominates in developed country work, and is increasingly becoming part of the formal apparatus of central government decision-making about the introduction of costly new pharmaceuticals and other health technologies in public health care systems (Drummond et al. 2005). The DALY approach currently dominates in developing country work, but has not (yet?) been taken up formally by any major decision-making body in the developing world (Disease Control Priorities Project 2006). Like the valuing safety literature, the valuing health literature also focuses on human outcomes rather than financial outcomes. It may therefore provide a helpful alternative approach to valuing life and other human outcomes within the workplace sector. Unlike the valuing safety literature, however, the valuing health literature does not attempt to express health outcomes in monetary terms that can be compared directly with non-health outcomes. This shortcoming is often acceptable in the context of allocating resources within a fixed health sector budget, insofar as (1) the primary policy goal of health care is to improve health and (2) implications for private consumption and for public sector budgets outside health can reasonably be ignored. However, neither of these conditions holds in the workplace setting, since health and safety are not the only (or even the primary) goals of work, and workplace interventions are likely to have important financial implications for firms (and, ultimately, for their customers, shareholders and/or employees) and for governments (and, ultimately, for spending on public services other than health and safety). The valuing health approach therefore only purports to evaluate the health-related outcomes of workplace interventions, rather than the full range of outcomes that may be relevant to policy-makers, firms and other stakeholders.

2. Why value "human" outcomes?

When faced with the task of valuing life in an economic evaluation, the practical question of whether to focus on financial or human outcomes is likely to turn primarily on the perspective of the evaluation – i.e. which individuals and/or organisations are likely to be the main users of the results of the evaluation? To caricature somewhat, financial outcomes may be of primary interest from the perspective of a firm – i.e. the "bottom line" of profitability. This is not to deny that firms may be interested in *measuring* human outcomes (as opposed to valuing them) insofar as these have important knock-on effects on financial outcomes – including long-term financial outcomes of intangible assets such as reputation and trust among employees and consumers.

From the perspective of the individuals affected by health and safety interventions, by contrast, human outcomes are likely to be of primary interest, quite apart from loss of future earnings and consumption. Similarly, from a public sector perspective – say, a health and safety regulator, or the government as a whole – human outcomes may be of considerable (or at least of some) interest in their own right. This is because saving lives and improving health and well-being are themselves important public policy goals, quite distinct from the goal of increasing national income. In some cases, saving lives and increasing national income may go hand in hand. But in other cases, there may be policy trade-offs between health and wealth. The purpose of valuing human outcomes is to help make those trade-offs explicit and to facilitate intelligent and transparent public decision-making.

To give a somewhat stark example, an exclusive focus on financial outcomes would imply that saving the lives of high-skill, high-wage workers is far more beneficial than saving the lives of low-skill, low-wage workers – since the latter make a far lower contribution to economic output when valued at market prices. Another example: a focus on financial outcomes implies that saving the lives of young workers is far more beneficial than saving the lives of elderly workers – since the latter are close to retirement and so have fewer years left to contribute productively to the economy. Indeed, strictly speaking, an exclusive financial focus implies that saving the lives of those likely to make a net negative long-term future contribution to economic output (e.g. the retired, the long-term unemployed, the chronically sick) ought to be discouraged! Faced with a series of somewhat unpalatable ethical implications of this kind, the valuation of life and limb in terms of financial outcomes – known generically in the economics literature as the "gross output" or “human capital” approach – has since the 1980s fallen out of favour among public policy-makers and economists in almost all areas of policy.

This shift to a "human cost" approach has had the practical implication of substantially increasing the monetary values of live typically used in economic evaluations. For example, the last gross output-based value for preventing a fatality used by the UK Department of Transport was £180,330 in 1985 prices. This is worth about £500,000 in 2004 prices, of which about 28% was a more or less arbitrary allowance for "pain grief and suffering" (Mason et al. 2006). By contrast, the current 2004 value based upon a willingness to pay approach to valuing safety is £1,312,260 of which about 66% is "human costs" and the remainder is lost output and medical and ambulance costs (Department for Transport 2004).

In practice, of course, public sector decision-makers may only pretend to be interested in human outcomes and “ethical” approaches to policy making as a matter of political rhetoric. In reality, they may be more interested in their own “bottom line” in terms of administrative budgets or other political, professional or personal outcomes. Nevertheless, insofar as human outcomes are likely to be of direct interest to the general public and to key stakeholders in worker health and safety decisions (e.g. workers, local communities, customers, shareholders), there is a case to be made for encouraging public decision makers to engage in the explicit valuation of human outcomes, and to publish the results, so that the evidence and reasoning used to support public decisions can be subjected to public scrutiny. According to this line of argument, valuation of human outcomes is a mechanism for securing transparency in public decision-making and for preventing opportunistic public decision makers from using their discretion to feather their own nests and curry favour with special interests. The argument here is that human outcomes inevitably have to be valued *implicitly* by the decision-maker, by the very act of making a decision. The only question is whether those human outcomes are also valued *explicitly*.

There is also a case to be made against valuing human outcomes on precisely the opposite grounds. First, it may be argued that it is highly opaque activity that cannot possibly help to hold opportunistic decision to account. According to this counter-argument, the valuation of intangible human outcomes is an inherently subjective and opaque activity, unintelligible to the uninitiated and highly susceptible to manipulation. It therefore offers opportunistic decision makers the chance to hide their real motivations behind a screed of technocratic jargon about the public interest. Second, the reduction of health impact data to the one dimensionality of monetary value may result in the loss of important information. For instance, two health impacts that are equivalent from a monetary valuation perspective may differ ethically (for instance, in the degree of volition individuals exercise with respect to different risks) or in relation to other publicly determined priorities, such as concern for particular diseases, at-risk populations, etc. This second argument constitutes a call for *greater* (or preserved) discretion in policy-making, as against the case for reduced discretion implicit in formal valuation exercises.

These two opposing lines of argument represent an important fault-line in debates about how to value life and limb. The authors of the present chapter are evenly spread across this fault-line. It is probably fair to say, however, that economists tend to take the former view, more optimistic about valuation and more pessimistic about politics, whereas policy-makers and practitioners tend to adopt the reverse biases – unless they have come under the sway of economists!

3. Standard welfare economic theory

As pointed out above, until the 1970s the economic valuation of occupational injuries, illnesses and fatalities was concerned with what might be called the "financial" cost of such occurrences, using the "gross output" or "human capital" approach. There were calculations of lost wages and productivity, as well as the provision of health services to victims. This changed in the 1970s with the advent of new economic techniques to measure the "human" costs, or at least the *ex ante* expectations of loss of well-being forecast by individuals contemplating the risk of future health impairments. In this section we will examine the theoretical foundations for these methods as well as the unresolved doubts that have arisen in their wake.

3.1. Utility theory

The starting point for economic analysis in this field is welfare economics, which proposes that (1) individuals possess preferences that can be described mathematically by "utility functions" conforming to certain logical rules, (2) these functions are complete over all potential outcomes, (3) these functions uniquely determine individual behavior and are derivable from observations of behavior, and (4) these functions constitute the only acceptable basis for ascribing value to actions or outcomes. As we will see, all of these propositions are contested.

Assume for now that each individual possesses a utility function of the sort

(1) *utility = f(income, risk, x)*

where *utility* is the utility of a given individual, *income* is her (present value of lifetime) income, *risk* is the risk she faces of dying (or of some other physical impairment), and *x* is a set of other unspecified but utility-relevant circumstances or outcomes. The utility function *f* combines these elements in a manner that produces a calculation of her utility level. Such a function can be thought of as existing at a given moment in time; in principle it could vary from one moment to the next, but it is usual to suppose that it is stable over time.

Imagine now that we are interested in the effect of a reduction in risk. If the individual's income and other ("x") outcomes remain the same, it is reasonable to expect that her utility would go up. Similarly, if her income goes down and all else is as before, her utility would go down. So we might ask, how much reduction in income would exactly compensate the reduction in risk, leaving utility the same? For this particular individual and this particular risk, let's say that Δ*risk* (the change in risk) and Δ*income* (the change in income) are exactly these offsets. This would be equivalent to saying that it is worth Δ*income* in money to have risk reduced by Δ*risk*. In principle these two equivalents are both observable, if the individual we are considering is offered a sufficient range of choices or is asked the right questions (and responds accurately).

We have just begun, but already in this simple model we can identify several potential criticisms.

(a) Individuals may not possess utility functions that obey the properties of equation (1). A considerable body of evidence from psychology and behavioural economics suggests that individuals do not in fact behave in accordance with utility theory (Mullainathan and Thaler 2004). For instance, the laws of arithmetic require the transitive property (if the utility from outcomes A exceeds those from B, and B from C, then A must be greater than C). Yet there is considerable evidence that even this seemingly innocuous assumption is often violated in practice - known as the "preference reversal" phenomenon (Tversky and Thaler 1990). Further systematic "anomalies" have been documented in relation to relatively simple outcomes involving familiar market goods as well as relatively complex outcomes of the kind with which we are concerned in the health and safety sphere, such as those involving small probabilities, long time horizons and unfamiliar health states. Even more controversially, equation (1) encompasses the assumption that all the factors that influence individual well-being are commensurable, and this may be questioned (Anderson, 1993; Nussbaum, 2001).

(b) Even if behaviour were to conform to the canonical utility model, it is not evident that utility should be granted the interpretation of determining “value”. For instance, people may choose unwisely or (from some vantage point) irresponsibly - for example by smoking, drinking heavily or driving too fast. There is evidence, for instance, that individuals make systematic errors of judgment, such as unwarranted optimism about their performance relative to peers or mis-predictions of the satisfaction they will get from common life-experiences (Stutzer and Frey 2006, Kahneman 2002).

(c) A further skepticism regarding the identification of choice with value derives from role theory. Individuals, according to venerable traditions in philosophy and sociology, are not unitary creatures with a single utility-calculating engine. Rather, they are “multiple selves” capable of different sorts of value-creating mental processes in different contexts. Thus the political preferences exercised by individuals as citizens on a jury, a civic commission or in the voting booth are held to be different in kind from those they express as consumers (Sagoff 2004). This debate has empirical implications since the contexts in which risk preferences are estimated are typically ambiguous with respect to the role attributes that may govern observed responses.

3.2. Willingness to pay and accept

It is usual to classify the choices people make as representations of either willingness to pay (WTP) or willingness to accept (WTA). In the first case an individual wishes to acquire a benefit, and WTP is measured by the maximum amount he or she would pay for it. In the second an individual is asked to part with a benefit, and WTA is the minimum amount of compensation required if the benefit is to be relinquished voluntarily. According to the utility analysis above, the only difference in value between WTP and WTA would stem from the difference in initial utility: higher for WTA than WTP, all else being equal. Since utility is measured in monetary equivalents, and given the assumption of diminishing marginal utility of income (for a given individual), it is expected that WTA > WTP even though utility would be unchanged in either case. This is because there is higher utility associated with having the benefit initially (and being paid to give it up); hence a given amount of utility will correspond to a larger sum of money, and vice versa. Assuming the value of the benefit is small relative to total utility, however, this difference should also be small.

These assumptions play an important role in economic theory. First, they make it possible to apply values derived in WTP contexts and apply them to WTA contexts, and vice versa. Indeed, without this it would be difficult to speak of "the" value of a statistical life, even for a given individual and a given risk factor. This is particularly important to bear in mind when considering labor market studies, as we will see shortly. Second, these assumptions also underpin the formulation of economic efficiency as applied to the allocation of benefits such as reduced risk of injury or death. A potential Pareto improvement occurs whenever such a reduction can be achieved for a cost below the corresponding WTA/WTP measure of benefits. The idea here is that it will then hypothetically be possible for the "winners" from a policy change to compensate the "losers", still leaving at least one individual better off and no-one worse off. However, if WTA could diverge significantly from WTP then judgments of economic efficiency would be contingent on the initial allocation of the benefit. Indeed, there could be no determinate analysis of the "Pareto efficient" initial allocation of this benefit at all.

In fact, a lively debate rages around the empirical question of WTA/WTP divergence. One recent paper reviewed 45 studies that compared the two forms of valuation, and the average ratio of WTA to WTP was in excess of 7:1, far in excess of what could be explained via conventional utility theory (Horowitz and McConnell, 2002). Prospect theory, and its corollary, status quo bias, is often held to be responsible for thisfinding, but the evidence is not conclusive. (Kahneman and Tversky 1979).

3.3. Compensating wage differentials in the labor market

A particular application of utility-based risk analysis pertains to the labor market. Assume that workers have utility functions that take the form of equation (1); they constitute the supply side of the market. Assume that employers face profit functions of the form

(2) *Profit = Revenue (output * price) – Cost (wage * hours, risk)*

where *output* is itself a function of employment, and where *Cost* depends on both the wage bill (work hours times hourly wage) and the level of risk faced by workers. It is assumed that it is costly to provide safer working conditions, so that as *risk* goes up, so, all else being equal, so does *Profit*. Of course, at sufficiently high levels of *risk* the production process would unravel, but it is assumed that the initial level is well below this.

So put the supply and demand for labour together. Workers make offers based on their utility in *income* and *risk* which have the property that, for any given worker, greater *risk* must be offset by greater *wage* to maintain a particular level of utility. Employers make wage offers based on (2) with the property that any reduction in *risk* must be offset by a reduction in *wage* if a given level of profit is to be maintained, other things equal.

Equilibrium occurs where the wage/risk tradeoff is identical between workers and employers: Δ*risk* has exactly the same monetary value for both. To put it differently, workers are just willing to accept the wage tradeoff that employers also require. If workers were willing to pay more to reduce risk (by accepting a lower wage), then employers would accept this and thus move to a lower level of risk - at which worker WTP for further risk reductions will be less. And vice versa: if workers are not willing to accept a wage reduction sufficient to offset the employer's cost of providing an additional amount of safety, then the risk level will rise - and worker WTP for risk reduction will rise. Note that this logic depends on the assumption that increments of risk are increasingly disliked by workers (increasing marginal disutility of risk), while increments of safety are increasing costly to employers to provide (increasing marginal cost of safety). Of course, in a full specification we would take into account individual differences across both workers and employers, but the essential aspects of the model can be grasped from this simplification. The full model was first developed by Rosen (1974).

The implication of this analysis is that firms with greater difficulty in making work safe—roofing companies, for instance—would reach labour market equilibrium at higher levels of risk. Similarly, workers with a diminished tolerance for risk would, all things being equal, reach equilibrium at lower levels of risk. Furthermore, there are efficiency gains to be had by properly matching relatively risk-tolerant workers with employers who face relatively greater costs of occupational health and safety.

The preceding model of safety determination has three significant implications. First, it yields the property that, ceteris paribus, workers in dangerous jobs are no worse off than those in safer jobs. This is because workers are fully compensated for the risk they face; they receive an increment in their wage that leaves them at the same level of utility they would otherwise be on. If this were not the case, they would select another job that provided a sufficient tradeoff. Second, it identifies an efficient level of risk in the workplace: that which is attained in equilibrum. At this level, the marginal cost to the employer of reducing risk is exactly equal to the marginal benefit to the worker of having it reduced. Third, and for our purposes most important, it permits an estimation of the worker's implicit value of health or life. Assuming that workers are arrayed across a range of firms with different costs of safety, and assuming that they bring the same underlying utility functions to their choice of employment, the value of health impairments can be deduced from the estimated wage-safety tradeoffs arrived at in the market.

The specific techniques for doing this will be described later in this chapter.

This portrayal of labor markets is not broadly accepted by other disciplines, and it has attracted its share of critics among economists as well. Potential difficulties include the following:

a. The labor market does not equilibrate in the manner specified by the model. Excess supply is characteristic of most labor markets that have been studied (Solow, 1990), but the model under consideration reaches equilibrium where supply equals demand. One popular explanation for non-clearing labor markets is the efficiency wage hypothesis that wages above the market-clearing level are used to induce greater effort among workers. However, it can be shown that fully compensating wage differentials would not appear if efficiency wages are indeed used to induce effort (Dorman, 1998).

b. In practice, the most attractive employments also tend to be the best paid. Alternative employments available to a given worker (or two workers of a given degree of productivity endowment) do not typically provide a fixed level of utility. Rather, workers queue for desirable jobs offering greater pay or work amenities. Those who get such jobs are better off than those who do not . If dangerous jobs are, on average, less desirable than safe ones, the market wage-risk tradeoff will understate the utility tradeoff derived from equation (1). Indeed there may even be a negative tradeoff if dangerous jobs are so much worse that they also pay lower wages.

c. Workers may not "accept" the working conditions as provided by the employer at the point at which an employment agreement is struck. This may be because workers expect (correctly or otherwise) that they will be able to alter these conditions, or because they believe that the law (which has rejected the assumption of risk doctrine for more than one hundred years) will protect them. In any case, in the laws typically in force in industrialized countries, workers do not and cannot legally accept many of the risks they actually face at work.

d. There is little direct evidence that employers perceive a wage cost resulting from failure to provide safer working conditions. That is, the utility constraint embodied in equation (1) is not visibly binding on most firms. Textbooks used to train safety managers, for instance, go into great detail on productivity and insurance costs stemming from accidents, but they do not claim that employers will have to pay increased wage compensation (Dorman, 1996).

e. The model assumes full information on both sides. In practice, however, the extent of risk, particularly for factors that are either low probability or delayed in their effect, may not be known. In addition, there is typically information asymmetry: employers know the risks better than their workers do, and they have an incentive to conceal their knowledge.

f. If the model is correct, there is nothing to be gained by workers who would demand government regulation beyond the provision of accurate information. This can be seen by the efficiency properties of equilibrium risk: at any other level of risk the wage adjustment forced onto employers would leave workers worse off than before. Yet the history of occupational safety and health policy is one of constant agitation by workers for more assertive regulation.

g. This is implicitly a one-shot game played between a worker, or a set of workers, and an employer. Most economic theory prior to the 1980s was naive game theory in this sense. During the past two decades, however, self-conscious use of game-theoretic methods has flourished in economics. In this context, it is clear that a repeated game would be a better characterization of the economic factors surrounding the determination of wages and risk, since employees can be regarded as recontracting over time. Whether the properties of a correctly specified repeated game would be the same as of those in a one-shot game is currently unknown, since such a model has not been analysed.

h. The model is built upon the utility assumptions in section 3.1, but there is empirical evidence that belies this view.

A final point has to do with the emphasis placed on worker choice. The model presumes that workers face a range of jobs that are approximately equal to them in utility terms; thus the wage increment reflects the amount workers are willing to accept (WTA) to incur added risk. To the extent that this is correct, and given the substantial differences in the literature between WTA and WTP for identical risks, researchers should be cautioned against using labor market studies to derive values to be used in contexts in which WTP is the appropriate measure, such as investments in improved health care facilities.

4. Philosophical controversies and alternative theoretical approaches

The empirical literature using revealed preferences from labour market data remains closely tied to the standard welfare economic theory outlined above in section 3. Since the 1990s, however, the empirical literature using expressed preference data has borrowed heavily from ideas and techniques from psychology and sociology and in so doing has at times departed somewhat from this standard welfare economic theory in various ways – for example, by using focus group methods to help respondents arrive at more "considered" preferences and by asking hypothetical questions that invite respondents to adopt the perspective of a citizen or policy-maker making decisions about the health and wealth of groups of people, rather than the perspective of a consumer making decisions about his own health and wealth. Although no definitive alternative theoretical framework has been elaborated to justify these somewhat pragmatic departures, they may broadly be defended in terms of a "decision-making" or "extra-welfarist" approach to economic evaluation (Sugden and Williams 1978, Culyer 1989).

To appreciate the difference between standard or "welfarist" verses non-standard or "extra-welfarist" approaches, one must engage with some of the underlying philosophical controversies. The practice of valuing life brings into sharp relief a number of important general philosophical controversies about the theoretical underpinnings of economic evaluation. Historically, at least, there has been fairly widespread agreement among the economics profession on at least four philosophical principles that can cause controversy among member of other professions:

- *Consequentialism.* Economic evaluation should value interventions in terms of their likely consequences, rather than placing intrinsic value on processes.
- *Reductionism.* Economic evaluation should split consequences into different components, measure and value each component, and then aggregate to an overall figure.
- *Individualism.* Economic evaluation should value costs and benefits to society (or other groups, entities or organisations) in terms of costs and benefits to individual people.
- *Commensurability.* Economic evaluation should aim to explicitly to value all important consequences in relation to one another using a common currency.

Of course, all four principles can and have been questioned by economists who dissent from the mainstream view. For example, we have already noted that some economists are reluctant to value human outcomes in monetary terms: they question the assumption of commensurability.

However, it is on two rather different issues that there is a long tradition of widespread disagreement among economists:

- *Definition of individual benefit.* There are at least three main schools of thought about how to define and measure individual benefit or "utility": preference, happiness and capability schools.
- *Aggregation of individual benefit.* There are two main approaches to aggregating individual benefits to form an overall judgement of social benefit: the "potential Pareto improvement" approach and the "social welfare function" approach.

Although most economists define utility in terms of preferences (or choices), there is also a venerable tradition of defining individual benefit or "utility" in terms of happiness (or experienced pleasure and pain). The latter ("experienced utility") may depart from the former ("decision utility), since people can make mistakes about forecasting their own future experiences of happiness. The "happiness" tradition dates back to the classical utilitarian philosophers of the eighteenth and nineteenth centuries, such as Bentham and Sidgwick, and has seen a modern revival (Kahneman, Wakker and Sarin 1997, Layard 2005). There is also a venerable tradition, with roots going back to Aristotle, of defining individual benefit in terms of more "objective" values that may depart from the individual's own subjective preferences or experiences. This tradition includes the work of the philosopher Moore, writing at the turn of the twentieth century, who is sometimes referred to as an "ideal utilitarian". This Aristotelian tradition has also seen a modern revival in the form of Amartya Sen's "capability approach", which focuses on the individual's freedom to perform valuable activities or to achieve valuable states of being (Sen 1985, Sugden 1993).

Turning to the issue of aggregating benefits, the "potential Pareto improvement" approach attempts to avoid making explicit cardinal interpersonal comparisons of benefit i.e. value judgements of the kind that a benefit to one person or group (e.g. a safety improvement) is twice as valuable as the corresponding harm to another person or group (e.g. a loss of consumption). Instead, it relies on the apparently more innocuous ethical assumption that if a policy makes at least one person better off and no-one worse off - known as a "Pareto improvement - then that policy is a good idea. It does this by attempting to identify situations in which the sum total of individual WTP for a policy is greater than its cost, so that "winners" can hypothetically compensate "losers". This is known as a "potential" Pareto improvement.

By contrast, the social welfare function approach does make explicit value judgements about how cardinal interpersonal comparisons of benefit are to be made. These value judgements are embodied in a mathematical function – the "social welfare function" – that is supposed to represent the value judgements of a "social decision maker" about how much weight is to be attached to utilities (somehow measured) that accrue to different individuals. One simple approach, for example, is to measure individual utility in terms of willingness to pay (i.e. according to the principle "one dollar, one vote") but then to attach distributional weights that are proportional to each individual's marginal utility of income (i.e. according to the "widow's mite" principle, that two small coins from a poor widow are worth more than the sum total of all the money donated by many rich men; Luke 21:1-4). Distributional weights of this kind attempt to mitigate the in-built bias in favour of individuals with high ability to pay, by giving greater weight to individuals with low incomes who stand to gain more utility benefit from an extra dollar's worth of income. Another simple approach is to measure individual utility in terms of quality-adjusted life years (QALYs) and to weight each individual's utility equally - i.e. according to the principle "a QALY is a QALY" no matter to whom it accrues. This latter approach is often taken in the health care field, where in relation to publicly funded health care it is often possible to focus on health effects and to ignore effects on individual income or consumption - since the public health budget can often reasonably be taken as given and since care is often delivered free at the point of delivery. A complication arises in the workplace health and safety field, however, since wealth effects as well as health effects must typically be taken into account. In such cases, it is nevertheless possible to derive equity weights based on an ethical assumption similar to the assumption that "a QALY is a QALY" - for example, the assumption that the social value of a marginal change in the probability of dying is the same for all individuals, regardless of wealth or other characteristics (Baker et al. 2006). This assumption does not necessarily lead to a single common monetary value of statistical life applicable in all cases. This is because, although by assumption the monetary value does not depend on the distribution of health effects between individuals, it does depend on the distribution of wealth effects which in turn will depend on how the safety intervention is financed (Baker et al. 2006).

The combination of defining individual benefit as preference and adopting the potential Pareto improvement approach to aggregating individual benefits is sometimes referred to as “welfarism” or “standard welfare economic theory”; as outlined above in section 3. Anything else is known as the “decision-making approach” or “extra-welfarism”. However, among the thickets of this ideological jungle there is even controversy about how to draw the line between “welfarism” and “extra-welfarism” (Birch et al. 2005, Cookson 2005).

5. Empirical methodology

There are two general strategies for estimating trade-offs between income and safety: revealed and expressed preferences. Revealed preference techniques record actual market behavior in the presences of differing levels of risk and attempt to isolate the role that risk plays in altering prices. Expressed preference methods employ questionnaires directly to ask individuals how much they might pay (accept) for a given change in risk.

5.1. Revealed preference approaches

Using this approach, researchers utilize demand or supply data to estimate a relationship between changes in risk and offsetting monetary flows. Consider, for instance, the housing market. We might expressing the market value of a house as a function of its characteristics:

(3) $price = f(risk, x, y)$

where the price is a function of the health risk associated with the house's location, x is a set of house characteristics (like size and amenities), and y is a set of locational characteristics unrelated to the health risk (such as the quality of the local school system). Ideally we would want to compare houses with identical x and y in the manner of a controlled experiment, but this is not likely to be practicable. Instead, we can regress price on risk, x and y using a large enough sample, interpreting the coefficient on risk as the price effect of a unit change in that variable. If it happens that the price falls by $100 when the risk of a fatal cancer, associated with the proximity of the house to a source of toxic emissions, rises by 1 in 10,000, we can infer a value of statistical life of $1M: $100 / 0.0001.

A somewhat more tenuous connection can be drawn between the risk of a traffic fatality and the implicit monetary value of time. Suppose that this risk is a function of the driver's chosen velocity, among other factors:

(4) $risk = f(velocity, x)$

where the left-hand side measures the risk of fatality, which is a function of the velocity and a combination of other factors, *x*. A driver who chooses to drive faster is viewed as receiving added utility from reduced journey time of at least the same magnitude as the utility lost by the increased risk. To convert units of journey time into monetary equivalents, assume that the opportunity cost of time on the road is income that could be earned at the driver's hourly wage. Then compute the value of this income, divide (as above) by the change in risk associated with the change in velocity, and the result is a value of statistical life for this driver. What makes the traffic computation somewhat more problematic than the housing example is the addition of two assumptions required to convert highway speed into money, that the only attraction of greater speed is the saving of time, and that the value of time saved is the driver's hourly wage.

Both calculations, because they rely on multivariate regression rather than controlled experiments, are vulnerable to omitted variable or specification bias: if elements of *x* and *y* in (3) or *x* in (4) are not included in the regressions, or if the functional forms of the regressions do not capture the mechanisms of *f*, estimates of the coefficients on risk (3) and velocity (4) will contain error. There is an additional problem with the traffic example, emphasized by Ashenfelter (2006): while increased speed should result in more accidents all things being equal, all things are not equal. Specifically, when road conditions are safest (least congested, driest, etc.) drivers will respond by driving faster. Thus, increased safety can cause greater speed. This is a thorny endogeneity problem which calls for an instrumental variables approach, but such methods are normally less precise (they introduce more noise) than the single-equation technique they supercede.

Finally, both the examples given above share an imprecision in the context of workplace safety and health: they estimate VSL in other decision settings, and it is not evident that the values can be transferred from one setting to the other. In particular, both estimate WTP, whereas workplace valuation is or approximates WTA. There may also be social considerations in the workplace that are absent from housing or other markets, as discussed in Dorman (1996).

By far the largest number of studies examine the labor market, drawing on the analysis developed in section 3.3. In nearly every instance researchers employ the regression equation

$$(5) \qquad ln(wage) = constant + \beta_1 . risk + \beta_2 . x_1 + \beta_3 . x_2 + ... + \varepsilon$$

Here the natural logarithm of the wage is employed, since the principle of diminishing marginal utility of income implies that equal percentage changes in income are more likely to correspond to equal increments of utility than equal absolute changes in income. In this specification β_1 is the coefficient on the risk variable, the x's are various worker characteristics that might affect earnings, each with its own coefficient, and *e* is an error term. Samples are usually drawn from groups of workers concentrated in relatively more hazardous occupations; that is, they are likely to oversample manufacturing, construction and extractive industries relative to the population as a whole.

It is worth pausing for a moment to consider the theory behind this equation. Clearly the jobs that pay the highest wages (e.g. chemical engineering) are also likely to be relatively safe, while those that are the riskiest (e.g. meatcutting) pay less. Presumably the workers with the best jobs have qualifications that give them access to the higher rungs of the labour market, whereas those whose skills are less in demand must choose between less desirable options. Thus controlling for a range of worker characteristics is supposed to isolate the wage-risk relationship to more narrowly circumscribed market positions. If this statistical strategy is successful, the coefficient on risk would measure the tradeoff between jobs that compete for identically qualified workers. These would be much more likely to exhibit wage-compensation-for-risk results than a model which failed to distinguish between highly qualified and less-qualified workers. Also, note that *only* the job risk and the worker's characteristics are incorporated. This implies that other aspects of the employer or industry play no role in wage determination.

Several issues have emerged in the econometric practice:

a. In principle, workers face multiple risks on the job, and each has the potential to make its own contribution to wage determination. Very simply, one could sort them into fatal and nonfatal; this would yield two risk variables, $risk_f$ and $risk_{nf}$, and two coefficents, β_f and β_{nf}. In practice, however, most analysts include only one such variable, usually fatal risk. Assuming that fatal and nonfatal risk are correlated, this will have the effect of artificially inflating the estimate of β_f from which VSL calculations are drawn.

b. There is likely to be significant error in observations of risk at the individual level. We simply do not have data that tell us how much risk each worker faces; consequently the usual approach is to assign to each worker the average risk factor for his or her industry, occupation or industry/occupation cell. Because the error in this procedure is random, it reduces the potential magnitude and significance of the estimates. Moreover, by creating a variable *risk* with fewer true observations than the other variables in the equation, it alters the method by which statistical significance should be calculated. The antidote involves an application of the technique of "clustering", but in practice it is rarely employed.

c. Conventional economic theory holds that wages are determined primarily by the worker's productivity characteristics, but available observations on x are not likely to capture them fully. This will bias estimates of the risk coefficient in ways that are unpredictable *a priori*.

d. Nonfatal risks are costly to workers in part due to their monetary effects, but some of these are defrayed by workers' compensation. In the United States WC is a state-level program whose benefits and coverage vary. In principle, whenever $risk_{nf}$ is employed, a variable that estimates WC payments should be included as well.

e. For utility purposes workers may be assumed to care only about their after-tax income, but data sets used by researchers often report only pre-tax income. In that case, adjustments should be made based on marginal tax rates applicable to each individual.

f. Many of the x variables describe categorical differences between workers, such as age, race, sex and union membership. The specification presented in (5) implicitly assumes that the coefficients on all independent variables are the same across these categories, but this may be wrong. For instance, it is not clear that experience, a standard x variable, should play the same role in wage determination for men and women or for blacks and whites. Ideally each category should be estimated separately, although the data sets are not always large enough to permit this, particularly in light of the small sample error that would arise in particular industry/occupation cells. There may not be enough black workers, for instance, in each cell, even in the entire population, to generate a reliable observation on their level of risk or to reliably estimate its coefficient. As we will see later in this chapter, the evidence for group-specific VSL's is substantial but also highly controversial.

g. A different approach to age is to translate fatal risk into expected number of years lost, a quantity that falls with the worker's age and the extent of latency (if the risk is caused by illness). In principle this should yield a value per statistical life-year across individuals, although this may be age-related as well as a consequence of either changing preferences or life-cycle income effects. For instance, young workers may be less risk-averse than older ones; also, if workers save income for retirement, they may attach more value to their retirement years than their working years.

h. A more recent concern is that workers may vary not only by productivity and risk preferences, but also by capacity for self-protection in dangerous jobs (Shogren and Stamland 2002). If this is the case, the error stemming from assigning average group risk to individuals within that group (see b above) will not be random, since safer workers will presumably require less wage compensation. This exacerbates the problem of measurement error on *risk*.

i. In most instances researchers include only individual-level variables in x. This is in accordance with conventional economic theory, which holds that workers of equal productivity should receive equal utility packages in employment—the law of one price. Nevertheless, there is widespread evidence that employer characteristics play an important role in wage determination as well. This has been demonstrated in the "interindustry wage differential" literature, where it has been shown that persistent wage premia attach to particular industries, even across disparate occupations such as secretaries and machine operators. Economists who have studied this phenomenon are likely to attribute it to rent-sharing: firms in some industries are able to make above-average profits, and they are under pressure from their workforce to share some of this in the form of above-average wages. To the extent that this is the case, estimates of wage compensation for risk will be biased if they fail to include variables that capture these employer- and industry-level effects. The simplest form that this can take is industry dummies, but labor market research would endorse other variables as well, such as firm size, the firm's capital-labor ratio and union density and gender composition at the industry level (Dorman and Hagstrom 1998.)

From the standpoint of OSH policy, one irony implicit in the wage compensation literature is this: researchers are estimating VSL's on the basis of the assumption that all risk is fully compensated, yet these same values are often proposed as benchmarks to be used for evaluating public policies whose rationale depends on the *failure* of wage compensation.

5.2 Expressed preference approaches

The expressed preference approach that is most closely wedded to "standard welfare economic theory" is the Contingent Valuation (CV) approach. This focuses on hypothetical "wealth-health" trade-offs made by individuals (or households) between their own wealth and specified risks to their own life and health. In essence, it asks individuals how much they are willing to pay for a specified risk reduction (or, less often, how much they are willing to accept for a specified increase in risk). Unfortunately, at least when applied to the relatively small changes in individual health risk delivered by most health and safety interventions, this approach runs into some serious biases and methodological challenges (Beattie et al. 1998a, b; see section 6.2 below).

As a result, investigators have turned to alternative "Relative Valuation" approaches that ask respondents to make hypothetical trade-offs that are easier for respondents to handle from a cognitive point of view. These approaches may be somewhat less firmly grounded in standard welfare economic theory, but the advantage is that they place less cognitive strain on respondents and thus help to reduce some of the more serious biases and anomalies that have afflicted the standard CV approach.

One such approach is the "chained" CV/SG approach, which breaks the CV valuation task down into more manageable steps (Carthy et al. 1999). The first step involves standard CV questions involving trade-offs between wealth and a non-fatal injury – typically, both a question about WTP to be cured of the non-fatal injury and a question about WTA to remain injured. The second step is to ask a "Standard Gamble" (SG) question involving trade-offs between the certainty of a non-fatal injury and a risk of death. In this question, the respondent is asked to imagine a treatment for the non-fatal injury that carries a small risk of immediate death. They are then essentially asked what risk of immediate death they would accept in order to be cured of the non-fatal injury, thus giving a non-monetary "relative value" of dying compared with the non-fatal injury. The monetary value of a non-fatal injury from the first step can then be "chained" to the relative value of dying compared with the non-fatal injury, to yield a monetary value of dying. This is the approach that has been adopted as the basis for the value for preventing a fatality currently used by the UK Department of Transport (Carthy et al. 1999, Department of Transport 2003).

More generally, the “Relative Valuation” or “relativities” approach focuses on valuation tasks that ask respondents to compare one health risk against another – for example, fatal versus non-fatal injury, different types of non-fatal injury, and different types of risk in different contexts. This allows one to develop a system of relative valuations for different health risks. In principle, this entire system of relative values can then be converted into a system of monetary values by attaching a monetary value to just one of the health risks – rather like hanging the system on a single monetary “peg”. Although developed in other policy sectors, one preliminary study suggests that this “relativities” approach may be a fruitful avenue to pursue in the workplace health and safety field (Karnon, Tsuchiya and Dolan 2005).

These relative valuation questions may involve risks to the individual’s own health, as in the SG approach described above and as in the “Time-Trade-Off” approach that involves trade-offs between length of time spent in different health states. Alternatively, relative valuation may involve risks to population health. For example, the “Person Trade-Off” value elicitation technique asks individuals to make trade-offs between numbers of lives saved (or illnesses prevented) from different policy interventions. This involves an important change of perspective from that of consumer (i.e. thinking about risks to one’s own life and health) to that of citizen (i.e. thinking about risks to other people’s life and health) – and a further departure from standard welfare economic theory, which adopts the consumer perspective

These and other relative valuation techniques allow the investigator to establish "relativities" between values for preventing fatality and values for preventing non-fatal injury (e.g. Jones-Lee Loomes and Philips 1995), between values for different hazard contexts (e.g. Jones-Lee and Loomes 1995) and also, in principle, between different population subgroups. These relative values can then be converted into absolute monetary values by using one or more standard "peg" monetary values – for example, the standard Department for Transport monetary value for preventing a road accident fatality (described below). Such methods go a considerable distance toward responding to the second criticism of utility theory described in Section 3 above.

6. Characteristic findings

In this section we will summarize the large empirical literatures that have developed on the basis of the theory and methodology in sections 2 and 3.

6.1. Revealed preference findings

In recent years several literature reviews, some incorporating meta-analysis, have appeared, and we will make use of them in this section.

Table 1 is drawn from de Blaeij et al. (2003); it includes only the studies based on revealed preference. All utilize data related to auto safety, either through the purchase of safer vehicles, related safety equipment or driving speeds. Monetary estimates were converted to US dollars by the authors according to PPP; the GDP deflator was used to convert 1997 into 2000 dollars. Dropping one outlier at each end, the estimates of VSL range from approximately $1.5M – $4.8M, with most clustering at the low end. (The meta-analysis of de Blaeij et al. is not reported here, because most of the studies on which it was based employed other methodologies, expressed preference and implicit VSL from regulatory choice.)

Table 1: Estimates of the value of statistical life in road safety, in millions of $US 2000

Authors	Country	Data year	Point estimate	Lowest estimate	Highest estimate
Atkinson and Halvorson (1990)	US	1986	4.8		
Blomquist (1979)	US	1988	1.6		
Blomquist and Miller (1992)*	US	1987		1.5	5.9
Dreyfus and Viscusi (1995)	US	1987	4.3		
Ghosh et al. (1975)	UK	1973	1.8		
Jondrow et al. (1983)	US	1988	2.0		
Melinek (1974)	UK	1974	0.8		
Miller and Guria (1991)	New Zealand	1990	1.5		
Winston and Mannering (1984)	US	1988	2.0		

*Three estimates.
Source: de Blaeij et al. (2003)

Table 2 is a selection from Blomquist (2004). It encompasses a few studies not included in de Blaeij et al., as well as a pair of studies on non-automobile-related behavior. Leaving aside motorcyclists, the range of VSL for adults is $2.8M – $7.2M, with no particular bias toward either end. This points to a higher VSL than is found in Table 1.

Table 2: Additional estimates of the value of statistical life from risk avoidance, in millions of $US 2000

Authors	Study	Data year	VSL
Jenkins, Owens, and Wiggins (2001)	Bicycle helmet use with fatality risk reductions and costs	1997	$4.3 adult $2.9 child 5-9 $2.8 child 10-14 users of helmets
Blomquist, Miller and Levy (1996)	Car seat belt use with fatality risk reductions and time and disutility costs	1983	$2.8 - 4.6 adult $3.7 - 6.0 child under 5 $1.7 - 2.8 motorcyclist typical driver or rider

Carlin and Sandy (1991)	Child safety seat use with fatality risk reductions with time and money costs	1985	\$0.8 child under 5
Mount, Weng, Schulze, and Chestnut (2001)	Hedonic analysis of motor vehicle prices with fatality risks	1995	\$ 7.2 adult \$ 7.3 child \$ 5.2 elderly typical vehicle occupant
Gayer, Hamilton, and Viscusi (2000)	Hedonic analysis of housing prices with fatality risk near Superfund sites	1988-93	\$4.7 (4.3 - 5.0) typical resident

Source: Blomquist (2004)

Despite the overall range encompassed in these two sets of studies, is might be considered a sign of success that they lie within a single order of magnitude.

Labor market studies have been comprehensively summarized in Viscusi and Aldy. Table 3 presents the majority of studies they cite that estimate a VSL using US labor market data. (We have excluded studies that employ occupational and locational measures of risk, which are less accurate than industrial attribution.) The table indicates whether the studies employed a variable for nonfatal risk, since failure to do so should result in the overestimation of wage compensation for fatal risk. It also records the sample mean income, since the use of ln(w) as the dependent variable signifies that the coefficient on risk applies to the percentage of the worker’s income.

Table 3: Estimates of VSL from US labor market data, in millions of 2000 US$

Author	Nonfatal risk	Sample income mean	VSL
Smith (1974)	Yes	$29,029	$9.2
Smith (1976)	yes	$31,027	$5.9
Viscusi (1978, 1979)	yes	$31,842	$5.3
Viscusi (1981)	yes	$22,618	$8.3
Olson (1981)	yes	$36,151	$6.7
Butler (1983)	no	$22,713	$1.3
Dorsey and Walzer (1983)	yes	$21,636	$11.8, 12.3
Leigh and Folsom (1984)	yes	$29,038 and $36,946	$10.1 – 13.3
Smith and Gilbert (1984, 1985)	no		$0.9
Dillingham and Smith (1984)	yes	$29,707	$4.1 – 8.3
Dillingham (1985)	no	$26,781	$1.2, 3.2 – 6.8
Leigh (1987)	no		$13.3
Moore and Viscusi (1988a)	no	$24,931	$3.2, 9.4
Moore and Viscusi (1988b)	no	$31,092	$9.7
Garen (1988)	yes	$29,865	$17.3
Viscusi and Moore (1989)	no	$24,611	$10.0
Herzog and Schlottman (1990)	no	$48,364	$11.7
Moore and Viscusi (1990a)	no	$24,611	$20.8
Moore and Viscusi (1990b)	yes	$24,611	$20.8
Kniesner and Leeth (1991)	yes	$33,627	$0.7
Gegax, Gerking and Schulze (1991)	no	$41,391	$2.1
Leigh (1991)	no	$32,961	$7.1 – 15.3
Berger and Gabriel	no	$46,865 and $48,029	$8.6, 10.9
Leigh (1995)	no	$29,587	$8.1 – 16.8
Dorman and Hagstrom (1998)	yes	$32,243	$8.7 – 20.3
Lott and Manning	no	$30,245	$1.5, 3.0

Source: Viscusi and Aldy (2003)

It is difficult to spot trends in the data; there is no overall tendency for later studies to generate higher VSL's than earlier ones, nor do the inclusion of a variable for nonfatal risk or differences in average income have the predicted effects. Heterogeneity in the estimates is driven primarily by the choice of samples, the use of different sources for measures of risk and choice of specification. It should be noted that, where researchers try multiple specifications, Viscusi and Aldy select those which, in their view, produce results most supportive of compensating differentials theory. We will return to the problem of publication bias at the conclusion of this section.

Table 3 reports their summary of non-US labor market studies of fatal risk.

The range is even greater for these studies, which employ a much wider variety of risk measures. In particular, there is only the slightest relationship between average worker income and average VSL.

Viscusi and Aldy also summarize 39 studies of nonfatal (injury) risk in the US. There would be little value in reproducing this list, since the type/severity of injury differs across studies, and since many of them did not control for the extent of income replacement by workers compensation. Nearly all report statistically significant coefficients on risk, with the value of an injury lying primarily within the range of $10,000 – $60,000 (in 2000 US$).

An important consideration is that there has yet to be an exploration of injury reporting bias in this literature. In most countries data on fatal and nonfatal injuries are collected from employers, who have an incentive to underreport. This incentive presumably plays a larger role in nonfatal than fatal injuries, and there is evidence to suggest that underreporting is most severe among small firms. (Dorman, 2000) If this latter hypothesis is valid, there would be a spurious size-risk relationship across firms which would interact with the known size-wage relationship to generate spuriously positive coefficients on risk in wage equations. This would be a particular problem for those studies that fail to control for firm size, which is to say nearly all of them.

Table 3: Estimates of VSL from non-US labor market data, in 2000 US$

Author	Country	Nonfatal risk variable	Mean sample income	VSL (M)
Marin and Psacharopoulos (1982)	UK	no	$14,472	$4.2
Weiss, Maier and Gerking (1986)	Austria	yes	$12,011	$3.9, 6.5
Meng (1989)	Canada	no	$43,840	$3.9 – 4.7
Meng and Smith (1990)	Canada	no	$29,646	$6.5 – 10.3
Kniesner and Leeth (1991)	Japan	yes	$44,863	$9.7
Kniesner and Leeth (1991)	Australia	yes	$23,307	$4.2
Cousineau, Lacroix and Girard (1992)	Canada/Quebec	yes	$29,665	$4.6
Martinello and Meng (1992)	Canada	yes	$25,387	$2.2 – 6.8
Kim and Fishback (1993)	S. Korea	yes	$8,125	$0.8
Siebert and Wei (1994)	UK	yes	$12,810	$9.4 – 11.5
Lanoie, Pedro and Latour (1995)	Canada/Quebec	yes	$40,739	$19.6 – 21.7
Sandy and Elliott (1996)	UK	no	$16,143	$5.2 – 69.4
Shanmugam (1996/7)	India	no	$778	$1.2, 1.5
Liu, Hammitt and Liu (1997)	Taiwan	no	$5,007– 6,088	$0.2 – 0.9
Miller, Mulvey and Norris (1997)	Australia	no	$27,177	$11.3 – 19.1
Siebert and Wei (1998)	Hong Kong	no	$11,668	$1.7
Liu and Hammitt (1999)	Taiwan	yes	$18,483	$0.7
Meng and Smith (1999)	Canada	yes	$19,962	$5.1 – 5.3
Arabsheibani and Marin (2000)	UK	yes	$20,163	$19.9
Shanmugam (2000)	India	yes	$778	$1.0, 1.4
Baranzini and Ferro Luzzi (2001)	Switzerland	no	$47,400	$6.3, 8.6
Sandy et al. (2001)	UK	no	$16,143	$5.7, 74.1
Shanmugam (2001)	India	yes	$778	$4.1

Source: Viscusi and Aldy (2003)

A final point has to do with differences in estimated values of fatal and nonfatal risk across different sub-populations. This has proved to be a flashpoint in the debate over monetization of risk and its role in public policy. We will consider differences by income, age, sex, race and union status.

a. Income. Within any given study, there is a direct dependence of estimated VSL (or value of nonfatal injury) on income, provided the dependent variable is ln(w). To identify this relationship across studies, researchers have treated individual studies as data points and regressed their VSL estimates on mean sample risk, the log of mean sample income and other variables. Coefficients on income in such specifications can be used to generate estimated income elasticities of VSL. Viscusi and Aldy describe four such efforts, and they re-estimate them using their own methods, arriving at elasticities in the neighborhood of .5. This would be consistent with the view that health is a "necessity" rather than a "luxury", in the sense that WTA measures of its value do not change as readily as income. Of course, any relationship between income and VSL suggests that, if WTA/WTP measures are used as the basis for public policy, interventions that promote the safety of the rich will be preferred to those that promote the safety of the poor, ceteris paribus. For a partial defense of this procedure, see Sunstein (2004).

b. Age. As discussed previously, there is no theoretical reason to presume one relationship between age and VSL over another. Viscusi and Aldy review eleven studies that attempted to differentiate VSL by age; in no instance did any find that VSL was greater for older workers, and in eight it was significantly less. Since their article was published, however, new evidence has appeared. A careful study by Smith et al. (2004) indicates that VSL may indeed be greater for those nearing retirement. This finding was corroborated by Kniesner et al. (2004), whose results display a U-curve relating age and VSL: high VSL for young and old workers, lower VSL for those in between. Whether the new studies supercede the old ones is unclear at this time.

c. Sex. It has been common to estimate wage/risk equations for men only under the assumption that their jobs are more hazardous, and it would be a mistake to attribute group risk averages to women in the same industry or occupational cells. On the basis of new BLS data, however, Hersch (1998) has shown that women face an average level of injury risk equal to 71% of the male rate. Her study and that of Viscusi (2004) attempt to calculate values of fatal and nonfatal risk for men and women separately. Hersch finds that the value per injury of her full sample of women workers is approximately equal to that of *blue-collar* men. The reason she conducts this comparison is that white-collar men receive significantly negative compensation: they have a negative value of health according to her methods. She eliminates this result from her published tables.

Table 4 presents a summary of Viscusi's results, derived from risk measures by industry-occupation cells. These are presumably fine-grained enough to capture most sex-specific differences in risk exposure. Only those results derived from a semi-log specification and pooled 1992-97 risk measures are reproduced.

Table 4: Value per Injury and Statistical Life in 1997 US$

Sample	VSL (M)	Value of injury (K)
male	4.9	13.4
female	-1.7	10.90
blue-collar male	7.0	12.2
blue-collar female	8.5	29.6

Source: Viscusi (2004)

These results are anomalous. Overall, blue-collar workers, who would be viewed as having higher tolerance for physical risk, are assigned higher values for injury and life. Overall, men have higher values than women, but the comparison is reversed among the blue-collar subsample. Women, and particularly white-collar women, have a *negative* VSL. It is difficult to escape the impression that disaggregation of wage compensation by sex has not yet yielded convincing results.

d. Race. The most extensive treatment of this is in Viscusi (2003). By partitioning his sample of workers into nested black and white subsamples, he derived two separate measures for VSL. In the largest such paired regressions, encompassing men and women and all occupations, he found that white workers have a VSL approximately twice that of blacks. This discrepancy is much too large to be explained by income differences (it appears in the risk coefficients themselves), nor is it the result of large differences in risk tolerance, since experienced (and therefore, according to this theoretical perspective, chosen) risk levels are only slightly unequal. He attributes the result to a flatter wage/risk offer curve on the part of employers, although it is not clear why this gradient, which in theory depends only on the technical cost of making particular jobs safer, should be so different across categories of workers. In any case, the difference in VSL's poses troublesome issues for the interpretation of wage compensation for risk.

e. Union status. Viscusi and Aldy describe eleven studies based on US data that allow for separate union and non-union risk premia. (It is important to note that, in all but one case, the studies measured the effect of the worker's own union status only and not industry-level union density, as economic theory might suggest.) It is a general pattern that they find much higher values for injury and life among union workers, and they sometimes find significantly negative values for nonunion workers. The international pattern is more diverse, with unionized workers sometimes receiving less compensation than their nonunion counterparts. Why unionization might matter depends on one's theoretical outlook. For those who believe that labor markets in general function to generate utility-equalizing risk compensation, unions may enhance the flow of information to the workforce or facilitate collective action over the public good aspects of safety. (Whether unions address the problem of marginal workers setting compensation levels for those with inframarginal preferences depends on the political model believed to characterize the setting of union policy.) For those who believe that labor markets are non-equilibrating, it may be only unionized workers who have the tenure and bargaining power to negotiate compensation.

A second question has to do with the role that union-nonunion VSL differentials should play in public policy. If nonunion workers receive less (or negative) compensation because of a lack of information, for example, should the implicitly lower (negative) value of their lives be discounted? What if bargaining power is at the root of the difference? And how might higher VSL's for nonunion workers in some non-US contexts be interpreted?

Overall, the aspect of revealed preference analysis that causes the most difficulty is the tendency for different social groups to be assigned different values of life and health. Straightforward interpretations of VSL as valid measurements of intangible value are vexed by disturbing group-specific patterns in the empirical literature; selective use of these results, on the other hand, requires justifications that undermine the case for revealed preference analysis in the labor market altogether. This is most clearly seen in the union-nonunion distinction, but applies to race and perhaps income as well.

A final concern which emerges from any close inspection of the literature is publication bias. Most researchers in the field of revealed preference analysis believe that properly-conducted studies *should* find statistically significant values for life and health that roughly correspond to those already reported. Any paper submitted for publication would have to explain and validate non-conforming outcomes beyond what would be required for corroborative results. As a result, the sample of published papers—and the subsample of results that appears in influential literature reviews—is not representative of the full breadth of research. Ashenfelter and Greenstone (2004) illustrate this point in their study of state-level decisions to alter highway speed limits. The main result is based on the pooled response of 28 states, but they entertain a thought experiment: what if each state was submitted as a separate study, and only those for states with significant results of the appropriate sign were accepted for publication. In that case, they show, the implicit value of life would be twice as high. Given that coefficients on risk in wage equations, for instance, are highly sensitive to the specification of the regression model, it is likely that many theoretically plausible specifications end up on the cutting room floor, discarded lest they interfere with reviewer acceptance. One can only speculate on how serious this problem may be.

6.2 Expressed preference findings

The largest CV studies, with nationally representative samples, have been conducted in relation to road safety. There have been a number of such studies in OECD countries, producing broadly similar results. Overall, large CV studies of this kind tend to give similar results to well-designed revealed preference studies (Beattie at al 1998a). However, early labour market studies tended to give lower estimates, since they failed to account for the fact that wealthy people tend to choose safer jobs – an important endogeneity issue that well designed studies address using simultaneous equation regression modelling (Jones-Lee 1989, Beattie et al. 1998a).

However, CV estimates of the value of statistical life vary considerably with 700-fold variation between highest and lowest estimates (Beattie et al. 1998a). Covey et al. (1995) estimated a food poisoning value of life of £48,840,000 in 1995 prices, whereas Acton (1976) estimated a heart attacks value of life of £69,800 in 1994 prices. One of the main drivers of this variation is the size of the risk reduction: smaller risk reductions yield larger values for statistical life. Covey et al. (1995) used an exceptionally small reduction in risk of death of 1.67 in 10 million, whereas Acton (1976) used an exceptionally large risk reduction of 2 in 1,000. This phenomenon has been linked with to a serious bias in CV – an under-sensitivity of WTP/WTA responses to the magnitude of risk reduction. Well-conducted experimental studies have found this bias to be particularly serious in relation to small health risks of the kind used in CV studies of the value of statistical life (Beattie et al. 1998b).

Other drivers of variation include the reference point (willingness to pay for risk reduction yields lower values than willingness to accept for risk increases), the sequencing of questions, the population subgroup (in particular, wealth and age), the hazard context (e.g. the degree of "voluntariness" of the risk), and the time horizon (e.g. whether reductions in current health risks from classic occupational health and safety programmes are valued more highly than equivalent reductions in future risks from disease management or health promotion programmes).

Studies that directly compare familiar hazard contexts – such as road safety, rail safety, food safety and air pollution – using CV and RV techniques have generally found rather small variations in relative values for statistical life, of around plus or minus 50%. Loomes and Jones-Lee (1995) found a premium of 51% for underground fatalities relative to road accident fatalities. Cookson (2000) found variation of plus or minus 50% between six different contexts, with similar values for road, rail and food safety, the value for air pollution about 50% higher and the value for birth control pills and medical radiation about 50% lower. The main factor underlying such variations appears to be the degree of perceived "voluntariness" of the risk.

Anomalies in expressed preference studies

A number of systematic and persistent "anomalies" have been documented in expressed preference studies, both in the valuing safety literature and the wider economic literature on expressed preference valuation of non-market goods (Beattie et al. 1998a, b, Dubourg, Jones-Lee and Loomes 1997).

Some of the well-documented anomalies are as follows.

- Protest zero responses
- Right skew in responses, resulting in the *median* response being used as the basis for estimated values for statistical life rather than the theoretically appropriate mean response

- Under-sensitivity to the magnitude of risk reduction. This and related phenomena are known variously in the wider literature on valuing non-market goods as the embedding effect, part-whole bias, and scope/scale effects.
- Over-sensitivity to theoretically irrelevant "cues" such as cost, baseline risk, payment card starting points and ranges, and the order in which items are valued

These anomalies also accord with a substantial body of evidence from psychology and experimental economics, which suggest that individual preferences may not confirm to the classical model of rational choice (Loomes 1999). Anomalies are perhaps particularly striking in relation to non-market goods that involve unfamiliar outcomes, small probabilities and long time horizons. However, they have also been documented even in relation to familiar household goods sold on the market place - such as mugs and chocolate bars.

One response to such anomalies is to urge improved survey design (Carson and Mitchell 1995). However, the pervasive nature of such anomalies suggests that they may reflect fundamental features of individual psychology and preference, and cannot exclusively be due to shortcomings in survey design. A second response is to use "deliberative opinion polling" techniques that allow respondents opportunities for deliberation and reflection. Together with a battery of inconsistency checks, opportunities to revise preferences in the light of feedback, and qualitative research to investigate people's reasons for their responses, the hope is that this might uncover a more "considered" set of preferences. A third response is to re-design value elicitation survey instruments so as to place less cognitive burden on the respondent and to focus on the central valuation issues at hand. For example, "Relative Valuation" questions about health-health trade-offs may be easier for respondents to tackle than "Contingent Valuation" questions about wealth-health trade-offs. A combination of all three responses seems appropriate.

7. Using off-the-shelf value of safety estimates in practice

Two main practical questions face the researcher wishing to incorporate a monetary value of life or injury into an applied economic evaluation of a workplace health and safety intervention:

1. What base case and range of values should I select from the thousands of "off-the-shelf" values reported in the valuing safety literature?
2. Once the human outcomes of illness and injury have been valued in monetary terms, should I also simply add in financial outcomes such as loss of earnings or is there a risk of "double counting"?

There is no general algorithm for selecting the most appropriate base case monetary value of life or injury for economic evaluation. In general, as we have seen, the appropriate base case value will depend on the policy context – e.g. the cause of risk and the characteristics of those at risk. It is known, for example, that the degree of "voluntariness" of physical risk typically contributes up to 3-fold variation in values for statistical life; other factors may appropriately contribute similar or even larger variations. There may therefore sometimes be a trade-off between the policy relevance of the study setting and the scientific quality of the study (e.g. sample size and design). Whichever estimate is chosen for the base case, however, the high degree of variability in estimates (up to 700-fold variation), and the serious methodological concerns about all methods in this field, means that it will always be important to conduct sensitivity analysis using alternative estimates. Sensible ranges of estimates for sensitivity analysis might include: (1) the typical range suggested by current reviews of wage-risk studies, (2) the typical range suggested by current reviews of expressed preference studies, and (3) the typical range currently used by other relevant national or international decision making bodies. For example, a recent review of wage-risk estimates of the value of preventing a single fatality from US labour market studies suggested a reasonable range of between $US4-9million in 2004 prices (Viscusi and Aldy 2004). We give two examples below of estimates used by major UK government bodies: the Department of Transport and the National Institute for Health and Clinical Excellence (NICE).

It is also hard to provide definitive practical guidance on the question of "double counting", since human and financial outcomes are hard to distinguish – both conceptually and empirically. A sensible guiding principle is that a worker's lost earnings should either be counted as a "human" outcome (i.e. as part of individual WTP for safety) or as a financial outcome (i.e. as part of lost output) but not both. Ideally, one would like to separate out the value a worker attaches to her own earnings (a "human" outcome) from the value to the economy as a whole from the worker's net contribution to national output (a "financial" outcome) – i.e. net of the consumption of national output enjoyed by the worker by spending her own earnings. This is not entirely satisfactory from a conceptual point of view, since the worker's net contribution to national output ultimately has human consequences in terms of the value other individuals place on their own consumption: output and consumption are two sides of the same coin. Nevertheless, it does seem a reasonable approach – there seems little to be gained in practice from attempting to convert a monetary figure for lost output into a series of estimates of reduced consumption for individuals, and then re-converting back into a monetary value. The difficulty, however, lies in separating out the two components empirically. In particular, it is hard to disentangle the component of individual WTP for safety attributable to loss of livelihood from components attributable to other factors such as the "pure" loss of life and the pain, grief and suffering of relatives and friends.

UK Department for Transport Values for Preventing Fatalities and Injuries

The UK Department for Transport value for preventing a single fatality is currently 1,312,260 in £UKSterling in 2003 prices (Department for Transport 2004). This is made up of three components: (1) lost output of £451,110, representing the present value of expected loss of earnings plus employer non-wage payments such as pension contributions, (2) ambulance and health care costs of £770, and (3) human costs of £860.380, representing the "pain, grief and suffering to the casualty, relatives and friends, and, for fatal casualties, the intrinsic loss of enjoyment of life over and above the consumption of goods and services." (Department for Transport 2004). This third element is derived from a series of large national expressed preference studies conducted in the late 1990s and subsequently uprated for inflation. This element is calculated as a preference-based estimate of the value of statistical life, based on the CV/SG approach (Jones-Lee, Loomes and Philips 1995, Carthy et al. 1999), less an estimate of lost lifetime consumption. The corresponding value for preventing a serious injury is £147,460 and for preventing a slight injury £8,750.

The NICE Value of a QALY

The UK National Institute for Health and Clinical Excellence (NICE) draws heavily on economic techniques to guide the UK National Health Service (NHS) in making "hard choices" about the allocation of scarce resources to costly new pharmaceuticals and other health technologies (Williams 2004). It values health improvements using a quality-adjusted life-year approach rather than the WTP-based approach adopted in the valuing safety literature. NICE's own official estimate of the monetary value of a QALY that it adopts in practice is between £25,000 to £35,000 (Rawlins and Culyer 2004). One analysis of historical decisions made by

NICE suggests that the monetary value (or "threshold") actually adopted is somewhat higher, at around 30,000-45,000 in £UK Sterling in 2005 prices (Parkin and Devlin 2005).

Note that this monetary value is a reflection of historical custom and practice, and does not rest explicitly on theoretical foundations or empirical evidence. However, the theoretical foundations currently favoured by NICE are quite distinct from those of standard welfare theory (Rawlins and Culyer 2004). The QALY threshold adopted by NICE does not represent the sum of individual willingness to pay for reductions in health risk. Rather, it represents NICE's estimate of the amount of money the UK National Health Service typically has to spend at the margin to improve health (i.e. to purchase the corresponding number of QALYs). This concept – known in the literature as the "shadow price of the budget constraint" – is based on linear programming theory, and can be thought of as one rather special form of the more general "decision-making" or "extra-welfarist" approach.

It is possible to convert monetary values for a statistical life into monetary values for a QALY (and vice versa), by making various modelling assumptions about the remaining length and quality of life of the relevant at risk population (Hirsh et al. 2000, Mason et al. 2006). The most recent and sophisticated exercise of this kind was based on the 2003 UK Department of Transport value for preventing a fatality of £1,311,490 (i.e. the full 2003 value less medical and ambulance costs). This generated values for a QALY ranging from £6,000 to £63,000 with a preferred range between £45,000 to £60,000. The threshold value of a QALY currently adopted by NICE (i.e. ranging from £25,000 to £45,000) could thus be expected to translate back into a value of statistical life somewhat lower than the UK Department of Transport value.

One possible interpretation of this discrepancy is that the NICE value is too low – and perhaps that the NHS budget should be expanded to allow government spending on health improvement in the health care field to match government spending on health improvement in the transport field. Another possible interpretation, of course, is that government spending on health improvement in the transport field is too high; or that there are good reasons for spending more on health improvement in the transport field (for instance, that health risks generally relate to younger individuals). A third possible interpretation, of course, is that the methodological challenges of valuing life and health are so great, and the theory and evidence used in the two different policy sectors are so different and so fraught with uncertainty and potential bias, that this apparent discrepancy is in fact remarkably small given the bounds of error that might reasonably be expected.

8. Methodological caveats and challenges

In concluding, we would like to emphasise some of the methodological challenges and caveats that beset this area, in the hope of encouraging further methodological work specifically related to the workplace health and safety arena.

First, although literature reviews often cite apparently narrow "preferred" ranges of values, it is important to note that the full range of values in the literature is much wider. Preferred ranges of values cited by literature reviews are always somewhat selective and, inevitably, those selections are based on the value judgements of the investigators.

Second, when interpreting the results of a particular valuation study, the devil is in the detail. That is to say, the findings of both revealed preference and expressed preference studies are highly dependent – often in systematic ways – on the details of the study design and the methods of statistical analysis. So use of any particular value in practice should be accompanied by extensive sensitivity analysis.

Third, much of the existing theoretical and empirical literature on valuing safety draws heavily on standard preference-based welfare economic theory, which has come under a cloud. It would be helpful to have alternative approaches based on competing approaches, such as happiness approaches and capability approaches, for purposes of comparison.

Fourth, revealed preference analysis should ideally be accompanied by research indicating the degree to which risk perceptions of the target population are in accordance with objective measures of risk, to the extent that the latter can be determined. There is little merit to valuations based on error.

Fifth, given serious methodological problems with attempting directly to set an absolute monetary value on safety using "Contingent Valuation", the alternative "Relative Valuation" approach warrants further investigation and development in the workplace health and safety area. Such studies should ideally be accompanied by experimental checks on internal consistency of responses and by other efforts to identify and minimise bias – potentially including qualitative research to identify the reasons underlying quantitative responses, and opportunities for respondents to arrive at more "considered" responses.

Sixth, more attention should be given to the problem of group-specific and context-specific values for statistical life and injury's. The idea that there should be a single, universally applicable monetary value of statistical life for use in all settings is not supported by economic theory or evidence. Rather, the appropriate value is likely to depend – at least to some extent – upon the characteristics of the population group and the nature of the hazard context. But how large are such differences in practice? And to what extent should these differences be recognized and incorporated into policy analysis?

Finally, we return to the issue of whether or not the "human" (as opposed to "financial") outcomes of death and illness need to be quantified in monetary terms at all when conducting economic evaluations of workplace health and safety interventions. An alternative approach would be to measure and report the full range of "human" outcomes in non-monetary units (e.g. lives saved, illnesses and injuries of different kinds avoided, QALYs gained; possibly broken down by different population groups), but without attempting to value these different outcomes in monetary terms in order to reduce them to a single figure representing overall benefit. This alternative approach is sometimes known as "cost consequence analysis" as opposed to "cost-benefit analysis" (Drummond et al. 2005). The choice of approach raises important philosophical issues about the role of economic evidence in the policy-making process, which we have discussed in sections 2 and 4. Cost-consequence analysis implicitly requires decision-makers to use their own judgement and discretion in weighing up the value of the different "human" outcomes, whereas cost-benefit analysis performs that valuation task explicitly. Ultimately, this methodological choice is as much a political question as a technical one: how much discretion should be left to the decision-making authority?

References

Acton, J P. (1976). Evaluating public programs to save lives: the case of heart attacks. RAND Corporation, Santa Monica, Report No. R950RC.

Anderson, Elizabeth. 1993. *Value in Ethics and in Economics*. Cambridge: Harvard University Press.

Ashenfelter, Orley and Michael Greenstone. 2004. Estimating the Value of a Statistical Life: The Importance of Omitted Variables and Publication Bias. AEI Brookings Joint Center Working Paper No. 04-06.

Ashenfelter, Orley. 2006. Measuring the Value of a Statistical Life: Problems and Prospects. Institute for the Study of Labor Discussion Paper No. 1911.

Baker, S, Chilton, A, Jones-Lee, M, Metcalf, H. (2006). Valuing Lives Equally – A Theoretical Justification. Newcastle Discussion Papers in Economics 2006/03.

Beattie, J., Chilton, S., Cookson, R., Covey, J., Hopkins, L., Jones-Lee, M., Loomes, G., Pidgeon, N., Robinson, A., Spencer, A., 1998a. Valuing health and safety controls: A literature review. Contract Research Report 171/1998. Health and Safety Executive, HSE Books, London. Available via www.hse-databases.co.uk/ research/crr_pdf/1998/CRR98171.pdf (accessed 20 March 2006).

Beattie, J, Covey, J, Dolan, P, Hopkins, L, Jones-Lee, M, Loomes, G, Pidgeon, N, Robinson, A and Spencer, A. (1998b). On the contingent valuation of safety and the safety of contingent valuation: part 1 – caveat investigator. Journal of Risk and Uncertainty 17: 5-25.

Birch, S and Donaldson, C. Valuing the bene?ts and costs of health care programmes: where's the 'extra' in extra-welfarism? Social Science & Medicine 2003; 56: 1121–1133.

Blomquist, Glenn C. 2004. Self-Protection and Averting Behavior, Values of Statistical Lives, and Benefit Cost Analysis of Environmental Policy. *Review of Economics of the Household.* 2(1): 89 - 110.

Carthy T, Chilton S, Covey J, Hopkins L, Jones-Lee M, Loomes G, Pidgeon N, Spencer A: On the contingent valuation of safety and the safety of contingent valuation: Part 2 - the CV/SG 'chained' approach. Journal of Risk and Uncertainty 1999, 17(3):187-213.

Cookson, R. (2000) "Incorporating psycho-social considerations into health valuations: An Experimental Study" Journal of Health Economics 19(3): 369-401.

Cookson, R. (2005). "QALYs and the capability approach." Health Economics 14: 817-829.

Covey, J, Jones-Lee, M.W, Loomes, G and Robinson, A. (1995). The exploratory empirical study. In: Ives, D., Soby, B., Goats, G., Ball, D. J., Covey, J., Jones-Lee, M.W., Loomes, G., Robinson, A . Eds. , Exploratory Study of Consumers' Willingness to Pay for Food Risk Reduction. University of East Anglia, Report to MAFF.

Culyer, A J. (1989). The normative economics of health care finance and provision. Oxford Review of Economic Policy 5(1): 34-58.

Department for Transport. (2004). Highways economics note number 1: 2003 Valuation of the Benefits of Prevention of Road Accidents and Casualties. London, HMSO. Available via www.dft.gov.uk (accessed 8 September 2006).

Dorman, P. (1996). Markets and Mortality: Economics, Dangerous Work, and the Value of Human Life. (Cambridge: Cambridge University Press).

Drummond M, O'Brien B, Stoddart GL, Sculpher, M Torrance GW. 2005. Methods for the economic evaluation of health care programmes (3rd Edition). Oxford, UK: Oxford University Press.

Dubourg, W R, Loomes, G, and Jones-Lee, M W. "Imprecise Preferences and Survey Design in Contingent Valuation," Economica, 1997, 64: 681-702.

de Blaeij., Arianne, Raymond J.G.M. Florax, Piet Rietveld and Erik Verhoef. 2003. The Value of Statistical Life in Road Safety: A Meta-Analysis. *Accident Analysis and Prevention.* 35(6): 973–986.

Laynard, R. (2005). Happiness: Lessons from a New Science. Allen Lane, London.

Loomes, G, Jones-Lee, M W and Philips, P. "Valuing the Prevention of Non-Fatal Road Injuries: Contingent Valuation vs Standard Gambles," Oxford Economic Papers, 1995, 47: 676-695.
Hirth RA, Chernew ME, Miller E, Fendrick MF and Weissert WG (2000) Willingness to Pay for a Quality-Adjusted Life Year: In Search of a Standard, Medical Decision Making, 20: 332-342.

Disease Control Priorities Project (2006). http://www.dcp2.org/main/Home.html (accessed 8 September 2006)

Guria, J., Leung, J., Jones-Lee, M. and Loomes, G. The Willingness to Accept Value of Statistical Life Relative to the Willingness to Pay Value: Evidence and Policy Implications. Environmental and Resource Economics 2005, 32, 113-27.

Hersch, Joni. 1998. Compensating Differentials for Gender-Specific Job Injury Risks. *American Economic Review*. 88(3): 598-607.

Horowitz, J. K. and K. E. McConnell. 2002. A Review of WTA/WTP Studies. *Journal of Environmental Economics and Management*. 44: 426-447.

Jones-Lee, M W. (1989). The economics of safety and physical risk. Basil Blackwell, Oxford.

Jones-Lee, M W, Loomes, G, and Philips, P R. (1995). Valuing the Prevention of Non-Fatal Road Injuries: Contingent Valuation vs. Standard Gambles. *Oxford Economic Papers* 47: 676-95.

Mullainathan, S and Thaler, R H. (2004). Behavioral Economics. In N. J. Smelser and P. B. Baltes (eds.) International Encyclopedia of the Social and Behavioral Sciences. (Elsevier Science) pp. 1094-1100

Jones-Lee, M W and Loomes, G. (1995). Scale and Context Effects in the Valuation of Transport Safety. *Journal of Risk and Uncertainty* 11: 183-203.

Loomes, G. Some Lessons from Past Experiments and Some Challenges for the Future," Economic Journal, 1999, 109: F35-45.

Loomes, G, J Beattie, J Covey, P Dolan, L Hopkins, M W Jones-Lee, N Pidgeon, A Robinson and Anne Spencer, "The Contingent Valuation of Safety and the Safety of Contingent Valuation, Part 1: Caveat Investigator," Journal of Risk and Uncertainty,1998, 17: 5-25.

Kahneman, D. (2002). Maps of bounded rationality: a perspective on intuitive judgment and choice. Nobel Prize Lecture, December 8, 2002. Available via http://nobelprize.org/nobel_prizes/economics/laureates/2002/kahneman-lecture.html (Accessed 8 September 2002).

Kahneman, Daniel and Amos Tversky. 1979. Prospect Theory: An Analysis of Decisions Under Risk. *Econometrica*. March, 47(2): 263-91.

Kahneman, D, Wakker, P and Sarin, R. (1997). Back to Bentham? Explorations of experienced utility. Quarterly Journal of Economics 112(2): 375-405.

Karnon, J, Tsuchiya, A and Dolan, P. (2005). "Developing a relativities approach to valuing the prevention of non-fatal work-related accidents and ill health" *Health Economics* 14: 1103-1115.

Kniesner, Thomas J., W. Kip Viscusi and James P. Ziliak. 2005. Life-Cycle Consumption and the Age-Adjusted Value of Life. NBER Working Paper 10266.

Kniesner, Thomas J., W. Kip Viscusi and James P. Ziliak. 2004. Life-Cycle Consumption and the Age-Adjusted Value of Life. NBER Working Paper 10266.

Mason, H, Marshall, A, Jones-Lee, M and Donaldson, C. (2006). Estimating a monetary value of a QALY from existing values of preventing fatalities and serious injuries. Report to UK NHS National Coordinating Centre for Research Methodology. Available via http://pcpoh.bham.ac.uk/publichealth/nccrm/publications.htm (accessed 8 September 2006).

Mrozek, J. R. and Taylor, L. O. 2002. What Determines the Value of Life? A Meta-Analysis. *Journal of Policy Analysis and Management.* 21(2): 253–70.

Niven, K J M (2002) "A review of the application of health economics to health and safety in healthcare", Health Policy, 61, 291-304

Nussbaum, Martha. 2001. The Costs of Tragedy: Some Moral Limits of Cost-Benefit Analysis. In Matthew D. Adler and Eric A. Posner (eds.), *Cost-Benefit Analysis: Economic, Philosophical, and Legal Perspectives*. Chicago: University of Chicago Press.
Sagoff, Mark. 2004. *Price, Principle and the Environment*. Cambridge: Cambridge University Press.

Rawlins, M D and Culyer, A J. (2004). National Institute for Clinical Excellence and its value judgments British Medical Journal 329: 224-227.

Sen, A. Commodities and capabilities. Amsterdam: North Holland, 1985. Reprinted in 1999 by Oxford University Press, New Delhi.

Shogren, Jason F. and Tommy Stamland. 2002. Skill and the Value of Life. *Journal of Political Economy*. 110(5): 1168-73

Smith, V. Kerry, Mary F. Evans, Hyun Kim and Donald H. Taylor Jr. 2004. Do the Near-Elderly Value Mortality Risks Differently? *Review of Economics and Statistics*. 86(1): 423-29.

Stutzer, Alois and Bruno Frey. 2006. What Happiness Research Can Tell Us About Self-Control Problems and Utility Misprediction. Institute for the Study of Labor (IZA) Discussion Paper No. 1952.

Sugden, R. (1993). Welfare, Resources, and Capabilities - a Review of Inequality Reexamined by Amartya Sen. Journal of Economic Literature 31(4): 1947-1962.

Sugden, R and Williams, A. (1978). The principles of practical cost-benefit analysis. Oxford University Press, Oxford.

Sunstein, Cass R. 2004. Are Poor People Worth Less Than Rich People? Disaggregating the Value of Statistical Lives. AEI-Brookings Joint Center for Regulatory Studies Working Paper No. 04-05.

Tversy, A and Thaler, R. (1990). Anomalies: preference reversal. Journal of Economic Perspectives 4(2): 201-211.

Viscusi, W K and Aldy, J E. "The value of a statistical life: a critical review of market estimates throughout the world" Working Paper 9487 Available via http://www.nber.org/papers/w9487 (accessed 30 November 2005).

Viscusi, W.K. (1992). Fatal Tradeoffs: Public and Private Responsibilities for Risk. New York: Oxford University Press.

Viscusi, W.K. (1993). "The Value of Risks to Life and Health," Journal of Economic Literature 31, 1912-1946.

Viscusi, W. Kip. 2003. Racial Differences in Labor Market Values of a Statistical Life. *Journal of Risk and Uncertainty*. 27(3): 239-56.

Viscusi, W. Kip and Joseph Aldy. 2003. The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World. *Journal of Risk and Uncertainty*. 27(1): 5-76.

Viscusi, W. Kip. 2004. The Value of Life: Estimates with Risk by Occupation and Industry. *Economic Inquiry.* 42(1): 29-48.