Reliability Engineering and System Safety 94 (2009) 1954–1961

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress

The role of social cost–benefit analysis in societal decision-making under large uncertainties with application to robbery at a cash depot

M. Jones-Lee [a], T. Aven [b,*]

[a] University of Newcastle, UK
[b] University of Stavanger, Norway

ARTICLE INFO

Article history:
Received 17 February 2009
Received in revised form
15 June 2009
Accepted 28 June 2009
Available online 1 July 2009

Keywords:
Social cost–benefit analysis
Risk assessment
Uncertainties

ABSTRACT

Social cost–benefit analysis is a well-established method for guiding decisions about safety investments, particularly in situations in which it is possible to make accurate predictions of future performance. However, its direct applicability to situations involving large degrees of uncertainty is less obvious and this raises the question of the extent to which social cost–benefit analysis can provide a useful input to the decision framework that has been explicitly developed to deal with safety decisions in which uncertainty is a major factor, namely risk analysis. This is the main focus of the arguments developed in this paper. In particular, we provide new insights by examining the fundamentals of both approaches and our principal conclusion is that social cost–benefit analysis and risk analysis represent complementary input bases to the decision-making process, and even in the case of large uncertainties social cost–benefit analysis may provide very useful decision support. What is required is the establishment of a proper contextual framework which structures and gives adequate weight to the uncertainties. An application to the possibility of a robbery at a cash depot is examined as a practical example.

1. Introduction

Over the past few decades there has been a growing tendency to rely on some form of social cost–benefit analysis (CBA) as a means of informing decisions concerning investment and regulation related to the safety of members of public, as well as workplace safety. More specifically, significant research efforts have been directed towards the definition and estimation of monetary values of safety, which can then be used to quantify the benefits of a proposed safety improvement in terms that are directly comparable with its costs of implementation and hence facilitate a decision as to whether or not the benefits are sufficient to justify the costs.

However, quite apart from the conceptual and practical difficulties posed by the task of defining and estimating monetary values of safety, there are, in addition, several other substantial problems that must be dealt with in the decision-making process. For example, how is one to obtain reliable estimates of the outcomes of any given safety improvement in terms of lives saved, injuries avoided and the prevention of damage to buildings and property? And how is one to take account of the fact that such estimates will almost inevitably be subject to a greater or lesser degree of uncertainty? And how are safety effects – even when estimated and quantified in monetary terms – to be weighed against or aggregated with other effects, such as the impact on public confidence or goodwill, or social stability, which may not be so readily amenable to quantification in monetary terms? And so on.

Arguably, it is at this point that risk analysis comes to the rescue. Essentially, risk analysis provides a carefully and clearly specified framework within which difficult decisions with uncertain outcomes can be structured, analysed and undertaken in a systematic and balanced manner, with due allowance being made for the risk and uncertainty associated with each aspect of the decision concerned. Thus, rather than consisting of a largely mechanical procedure in which pre-determined values of safety are applied to unique estimates of fatal and non-fatal injuries avoided – as well as physical damage to plant and equipment – in order to establish whether the benefits of a safety improvement will exceed its costs of implementation, risk analysis facilitates a very much more circumspect and reflective assessment of potential gains and losses as well as their likelihood of occurrence. In addition, considerations that do not lend themselves naturally to monetary quantification – such as social or political consequences – can also be weighed in the decision-making process. This having been said, there can be little doubt that clearly specified values of safety and the estimated costs of safety improvement will still constitute vitally important inputs to the

* Corresponding author.
E-mail address: terje.aven@uis.no (T. Aven).

0951-8320/$ - see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.ress.2009.06.011

decision-making process. The key point is that within the framework of risk analysis, these values and costs will not constitute the *only* input, nor will they, on their own, provide the final answer. In short, risk analysis can perhaps most fruitfully be viewed as providing a decision-making framework within which cost–benefit analysis plays a key role, but is not the only spanner in the analytical toolbox.

In the next two sections of the paper the key features of safety-related cost–benefit analysis and risk analysis will therefore be outlined. The application of these techniques – in their role as essentially complementary rather than competing analytical devices – will then be illustrated via their use in assessment of the appropriateness (or otherwise) of the location of a cash depot in a Norwegian residential area in the vicinity of a kindergarten, with the obvious risks to residents and/or children in the event of a robbery at the depot.

2. Cost–benefit analysis and the valuation of safety

The basic purpose of social cost–benefit analysis is to provide a mechanism by which decisions concerning the allocation of society's scarce resources can take due account of the preferences – and more particularly the *strength of preference* – of those members of society who will be affected by the decision concerned. Naturally, it is important to ensure that these preferences are, so far as possible, adequately informed and carefully considered.

Since an individual's maximum willingness to pay (WTP) for a good or service is a clear indication of what that good or service is worth to the individual relative to other potential objects of expenditure – and given that willingness to pay is ultimately constrained by ability to pay (i.e. income), which reflects the underlying scarcity of resources – individual willingness to pay constitutes a natural measure of strength of preference and is therefore used as the fundamental measure of value or 'benefit' in CBA.

Under what has naturally come to be known as the 'willingness to pay' approach to the valuation of safety, one therefore attempts to determine the maximum amounts that people would individually be willing to pay for (typically small) improvements in their own and (possible others') safety. These amounts are then simply summed (possibly with distributional weights to take account of considerations of 'fairness') across all affected individuals to obtain an overall value for the safety improvement concerned (see, for example, [1,2] or [3]). The resultant value is thus a clear indication of what the safety improvement is 'worth' to the affected group, relative to the alternative ways in which each individual might have spent his or her limited income. Furthermore, defining values of safety in this way effectively 'mirrors' the operation of market forces – in circumstances in which markets do not exist – given that such forces can be seen as vehicles for allowing individual preferences to interact with relative scarcities and production possibilities in determining the allocation of a society's scarce resources.

In order to standardise values of safety derived under the WTP approach, the concept of the prevention of a 'statistical' fatality or injury is employed. Thus, suppose that a group of 100,000 people enjoy a safety improvement that reduces the probability of premature death during a forthcoming period by, on average, 1 in 100,000 for each individual in the group. While the safety improvement could turn out to prevent no deaths, or one death (in fact, the most likely outcome) or two deaths (with a lower probability) and so on, the arithmetic mean (or statistical expectation) of the number of deaths prevented is precisely one and the safety improvement is therefore described as involving the prevention of one 'statistical' fatality.

Now suppose that individuals within this group are, on average, each willing to pay £x for the 1 in 100,000 reduction in the probability of death afforded by the safety improvement. Aggregate willingness to pay will then be given by £x times 100,000. This figure is naturally referred to as the WTP-based 'value of preventing one statistical fatality' (VPF). An alternative term often used is the 'value of statistical life' (VSL). Thus, if on average, the members of the population were willing to pay £15 per year to reduce their risks of death to this extent, the VPF (or VSL) would be £1.5 million.

In order to avoid possible confusion it is very important to appreciate that, as defined above, the VPF is *not* a 'value or (price) of life' in the sense of a sum that any individual would accept in compensation for the certainty of his or her own death—for most of us no sum, however large, would suffice for this purpose so that in this sense life is literally priceless. Nor is the VPF a measure of the amount that the typical individual would be willing to pay to avoid the certainty of death—even if an individual was prepared to pay all that he or she could afford, the resultant sum would tend to be relatively modest given the impact of income and wealth constraints. Rather, the VPF is in fact aggregate willingness to pay for typically *very small* reductions in individual risk of death (which, realistically, is what most safety improvements actually offer at the individual level). The Treasury Green Book emphasises the point when it notes that:

> "The willingness of an individual to pay for small changes in their own or their household's risk of loss of life or injury can be used to infer the value of a prevented fatality (VPF). The changes in the probabilities of premature death or of serious injury used in such WTP studies are generally small" [1, p. 61].

But of course under the 'aggregate willingness to pay' definition, strictly speaking the VPF applicable to a safety improvement that will affect a relatively well-off group in society will almost certainly exceed the figure for a poorer group, simply because willingness to pay is ultimately constrained by ability to pay. It is essentially for this reason that most advocates of the WTP approach recommend VPFs which reflect the aggregate willingness to pay of a *representative sample* of the population as a whole. Using an overall average figure effectively amplifies the (typically lower) willingness to pay of the less well-off and somewhat deflates the (typically higher) willingness to pay of upper-income groups. In this respect, the VPF involves a modification that might be said to reflect a social or 'citizen's' value judgement being applied to individual 'consumer' values.

Before proceeding to consider the various ways in which researchers have attempted to obtain empirical estimates of values of safety using the WTP approach, two further refinements should be noted. First, so far there has only been brief mention of peoples' concern – and hence willingness to pay – for others', as well as their own safety. To the extent that many people do display such 'altruistic' concern then it might seem reasonable to expect that it would be appropriate to 'boost' the WTP-based VPF to reflect the amounts that such people would be willing to pay for an improvement in others' safety. However, it turns out that the appropriateness of including peoples' willingness to pay for others' safety in the definition of a WTP-based VPF depends crucially on the nature of their altruistic concern for other peoples' general wellbeing. If this concern is 'safety-focused' in the sense that, while person A cares about person B's safety, A is in fact quite indifferent about B's ability to pay for say, a holiday abroad, a meal out or a new carpet, then A's willingness to pay for B's safety should be included in the definition of VPF. If, by