Economics and Philosophy, **23** (2007) 163–183
doi:10.1017/S0266267107001356

PHILOSOPHICAL PROBLEMS IN COST–BENEFIT ANALYSIS

Sven Ove Hansson
Royal Institute of Technology, Stockholm

Cost–benefit analysis (CBA) is much more philosophically interesting than has in general been recognized. Since it is the only well-developed form of applied consequentialism, it is a testing-ground for consequentialism and for the counterfactual analysis that it requires. Ten classes of philosophical problems that affect the practical performance of cost–benefit analysis are investigated: topic selection, dependence on the decision perspective, dangers of super synopticism and undue centralization, prediction problems, the indeterminateness of our control over future decisions, the need to exclude certain consequences for moral reasons, bias in the delimitation of consequences, incommensurability of consequences, difficulties in defending the essential requirement of transferability across contexts, and the normatively questionable but equally essential assumption of interpersonal compensability.

1. INTRODUCTION

Cost–benefit analysis (CBA) is a collection of decision-aiding techniques that have in common the numerical weighing of advantages against disadvantages. In a typical CBA, two or more options in a public decision are compared to each other by careful calculation of their respective consequences. These consequences can be different in nature, e.g. economic costs, risk of disease and death, environmental damage etc. In the final analysis, all such consequences are assigned a monetary value, and the option with the highest value of benefits minus costs is recommended or chosen.

Cost–benefit analysis is built on a very sound fundamental principle: advantages should be weighed against disadvantages, costs against benefits. However, many steps are needed to take us from this basic

principle to any of the forms of cost–benefit analysis that are currently in use. Some of these steps do not share the immediate intuitive appeal of the fundamental principle.

Cost–benefit analysis is controversial, and has repeatedly been subject to severe criticism not least from philosophers (Anderson 1988; Sagoff 1988). Most of this criticism has focused on two practices. One of these is the assignment of a monetary price to (the loss of) a human life. The other is contingent valuation, in which the price of non-market goods such as environmental assets are determined by asking people what they are willing to pay for them. It should be noted that neither of these practices is an essential component of cost–benefit analysis. It is possible, and indeed quite common, to perform cost–benefit analysis without applying either of them.[1]

However, there are other reasons why cost–benefit analysis is highly interesting from a philosophical point of view: it is associated with two interconnected philosophical issues that combine to make it (unintentionally) into a true testing ground for important philosophical principles.

The first of these issues is consequentialism, or consequential evaluation. Cost–benefit analysis is the only well-developed form of applied consequentialism.[2] The other issue is counterfactual analysis. In cost–benefit analysis, an alternative is not evaluated by itself but in comparison to other alternatives (or, at least, in comparison to not choosing that alternative). Therefore, cost–benefit analysis always involves an appraisal of what will happen if various options are chosen.[3] This feature of cost–benefit analysis is closely related to its consequentialism,

[1] An otherwise insightful article by Cass R. Sunstein exemplifies how restrictedly the problems of CBA are usually conceived. Sunstein suggests that 'there are two serious problems with CBA' (Sunstein 2005: 355). The first is that willingness to pay is sometimes an inappropriate basis for decisions and the second that probabilities cannot always be assigned. On the list of ten problems to be presented here, Sunstein's first problem is a variant of the incommensurability problem, and as we have just noted it is a variant that does not affect all CBAs. His second problem is an aspect of the prediction problem.

[2] Deontological requirements can be included in a CBA, for instance by assigning negative weights to violations of prima facie rights and duties, and including them on the cost side in the analysis. However, when this is done, a consequentialist structure is imposed on these requirements.

[3] The counterfactual nature of CBA has been recognized by several of its practitioners. A few examples are: 'Ideally, we would compare the future world with and without the regulation in terms of fatalities and other implications such as the quality of life' (Keeney 1990: 147). 'To evaluate regulations, we should compare the future world as it would be if a proposed regulation were passed to the future world as it would be if the regulation were not passed' (Keeney 1994: 97). 'I conclude, therefore, that if hazardous waste workers were *not* working on hazardous waste sites they would be performing other jobs with comparable risk of premature fatality from accidents. Thus the comparison of worker versus residential risk is not a simple one' (Gochfeld 2004).

which is much more closely connected with counterfactual analysis than seems to be generally recognized.

The philosophical issues in cost–benefit analysis come out clearly in three major phases in the process leading up to a finished CBA: (1) the *framing* of the decision, that determines which decision alternatives are included in the analysis; (2) the *option characterization*, i.e. the delimitation of those consequences of these alternatives that will be subject to valuation; and (3) the *valuation* of these consequences. In what follows philosophical problems in cost–benefit analysis will be discussed under these three headings.

2. FRAMING PROBLEMS

The framing of a decision for a CBA has several components. The general topic of the analysis has to be selected (section 2.1). Depending on the purpose of the particular CBA, a choice has to be made between the different decision perspectives in which it can be performed (section 2.2). As an important aspect of this, a choice has to be made whether decisions should be made piecemeal or combined into large units (section 2.3).

2.1 Topic selection

For obvious reasons of resource limitations, CBAs cannot be performed for every decision; a selection of topics has to be made. This selection is in practice not done in a coordinated fashion but depends entirely on the willingness of particular decision-makers to pay for a CBA. It may nevertheless be the most important factor that determines the social effects of cost–benefit analysis.

To perform a CBA means to make an economic analysis that is concerned with the economy and the society as a whole, not only those parts of it that affect a particular decision-maker's own economy. Hence, a CBA for a decision to build a manufacturing plant should include externalities such as the (economic and non-economic) costs of its emissions to air and water, the disposal costs for its products, the healthcare costs and human suffering of workplace accidents etc. Similarly, a CBA for a regulatory measure undertaken by an environmental agency should include the costs to the private sector of that measure, including investments costs, the future maintenance costs for new purification equipment, lost business opportunities due to decreased competitiveness etc.

CBAs are much more common in the public than in the private sector. Probably, the major reason for this is that governments have more responsibility for the total social and economical effects of the activities of the public sector than business managers have for those of their companies. Furthermore, there are political leaders who want to lighten the burden on

business that is imposed by regulations and other measures taken in the public sector. Since CBAs for proposed regulations highlight the predicted costs to business, such CBAs can often be used as tools to hold back regulations that would be costly for the affected companies. In contrast, it would be hard to find business leaders who, analogously, use CBAs as a means to alleviate the burden on others of the external costs created by their companies.

In the current system, business decisions that give rise to increased pollution are not required to pass a test showing that their total effects, including externalities, have a positive balance. On the other hand, regulatory decisions that reduce pollution are required to pass such a test. Given the incentives structure of our political and economic systems, this imbalance in the topic selection for CBAs may be unavoidable. It can nevertheless be problematic. Suppose, for instance, that decisions on public infrastructure investments are based on CBAs that take into account the negative effects of these investments on all sectors in society, whereas private infrastructure investments are based on business calculations in which no externalities are taken into account. Then different investments will be assessed according to different criteria. The difference may be seen as a form of bias against the public sector investments that are required to balance out more indirect costs than private investments have to do. Even if the professionals who perform cost–benefit analyses succeed in avoiding bias or partisanship, so that each individual CBA is a paragon of impartiality and objectivity, this does not guarantee that the total effect of all CBAs is unbiased.

A possible defence of current practices is that CBAs are needed in the public sector as a substitute for market mechanisms, whereas business managers do not need CBAs since they operate on real markets. This argument can be used to justify the use of market substitutes in public sector decisions, but it does not justify the practice of including aspects such as social and environmental effects in these market substitutes although they are external to the 'substituted' markets in the private sector.

2.2 Perspectives

A decision-guiding evaluation of an object (a project, say) can be performed with respect to different decisions about this object. Therefore, a CBA that is adequate for one decision-maker need not be adequate for another decision-maker who is considering another decision concerning the same project. We can see this from the following example:

Example 1: A CBA is made for a large road construction project. One of its (statistically) expected effects is the death of two construction workers

https://doi.org/10.1017/S0266267107001356 Published online by Cambridge University Press

in workplace accidents. This is clearly a negative term in the CBA for the project. But how should this be treated in a comparison with the alternative course of action, namely to reject the project? The standard approach is to conclude that no workplace deaths will follow from not carrying on the project. However, reasonable arguments can be made for other approaches. Perhaps, if this project is rejected the funds intended for it will instead be used for some other construction project. In that case it could be argued that the correct comparison is not with zero deaths but with the expected number of deaths from a project of that size. Or perhaps these workers will instead be out of work, or absorbed into the general labour market. Then the corresponding expected number of deaths can be entered into the analysis.

First, consider this example in the perspective of a government agency that is going to decide whether or not to fund the construction project. Such an agency has typically been assigned a certain amount of funds that are intended for the construction and maintenance of roads. Presumably, the agency has alternative plans for how to use its resources if the project is rejected, and in accordance with its instructions these alternative plans are all road construction projects.[4] A CBA that guides the agency's decisions can therefore compare alternative uses of the resources. From the agency's point of view, the relevant comparison is between this project and its alternatives. The expected death tolls of these projects can therefore be entered into the analysis.

Second, consider the same example in the perspective of an organization campaigning against the project due to what it considers to be excessive occupational risks. The campaigners are presumably against all projects with this level of risk. Therefore, an analysis showing that alternative projects have the same prospects will not give them a reason to accept the project in question but rather to extend the campaign to the other projects. When protesting against the project they aim also to influence the projects that are, from the government agency's point of view, its alternatives. In a CBA that guides the organization's decisions, the relevant alternatives are not the same as for the government agency. A comparison with the alternative use of resources promoted by the organization would be more adequate. (In practice, however, organizations campaigning against a particular project often do not have a policy for the alternative use of resources. This makes it difficult if not impossible to perform a CBA from their perspective.)

[4] In principle, the agency may have the alternative not to use the allotted funding at all, but in the culture of government agencies this 'do nothing alternative' will typically not be treated as a serious option.

The problem with decision perspectives comes out even more clearly in the following example:

Example 2: A proposal has been put forward to reduce CO_2 emissions in a member country of the European Union. A CBA is performed, which includes the economic consequences for the country of the expected competitive disadvantages in relation to other European countries that do not reduce their emissions. From the viewpoint of the country's industrialists and its government this seems to be adequate. However, from a European viewpoint it would be more reasonable to perform a CBA in a wider perspective that includes, on the positive side, the benefits in the other countries from their ensuing competitive advantages. These may very well cancel out the disadvantages in the country in which the reform is undertaken.

Clearly, a national perspective is as legitimate as a European one, and the same can be said about a global perspective or a regional or local one. Such differences in perspective often lead to a divergence of practical conclusions, since neither environmental problems nor the consequences of measures taken to abate them are restricted by national or regional boundaries.

More generally, Example 2 shows that a change in perspective can lead to a change not only in the identity of the alternatives that have to be considered (as in Example 1), but also in the range of possible consequences that should be evaluated for each of the alternatives. What seems to be one and the same decision may give rise to different decision problems when treated in the perspectives of different decision-makers. A CBA produced for one decision perspective may require careful adjustment in order to be relevant in another decision perspective.

Perspective problems are seldom recognized or openly dealt with in cost–benefit analysis. The best way to deal with them is to recognize that a CBA is always framed relative to a particular decision (or choice between alternatives). What is usually called a CBA for a project or a policy is strictly speaking always a CBA for that project or policy in the perspective of a particular decision. There are often other equally legitimate perspectives, related to other decisions concerning the same project or policy, for which the CBA is inadequate.

2.3 Synopticism

CBAs are intended to cover in principle all decision-relevant aspects of the alternatives under evaluation. The most obvious way to achieve this is to make CBAs as all-encompassing as possible. Therefore, cost–benefit

analysts often strive to frame decisions in large coordinated units that cover as many aspects as possible. Attempts to combine aspects and policy areas into large unified decision agendas have been called 'super synopticism' (Hornstein 1993: 387). If a perfect optimization that covers all aspects is feasible, then it would seem obvious that decisions should be made in as big pieces as possible, so that all aspects are combined. However, such large-scale optimization does not always work, for at least three reasons: (1) it may be impossible to collect and process massive amounts of information with sufficient efficiency; (2) due to cognitive limitations, we may be less competent decision-makers when the task has high complexity; (3) procedures that subdivide decisions may be preferable from a democratic point of view. Implementation may be facilitated if decisions have the legitimacy obtained in decentralized, participative procedures.

The dangers of aggregating and centralizing decisions too much are well known and much discussed in contexts where the alternative to the synoptic approach is a market. One of the major reasons to prefer a market to centralized planning is the impracticability of collecting and handling all the information needed, for instance for production decisions, in a central agency. The same problem applies to decisions that cannot be transferred to the market. However, these problems have been less recognized in contexts where the alternative to centralization is not a market but instead a combination of decisions made in different agencies with partly conflicting agendas. Two examples can clarify this:

Example 3: In Sweden, public decisions on traffic policies are made by a number of agencies with potentially conflicting responsibilities. There are four national agencies responsible respectively for railroad traffic, road traffic, aviation and maritime traffic. In each of the 21 counties, there is a regional administration responsible for coordinating traffic and transportation in that area. In addition, other agencies are responsible for coordinating policy goals that concern traffic and transportation, such as accessibility for handicapped and various environmental goals. There is no agency responsible for the national coordination of all traffic and transportation issues at a concrete, detailed level.

An alternative to this organization would be to have a single national traffic agency that coordinates the four modes of traffic in all the regions of the country in full detail, presumably using CBAs to optimize the whole traffic system. It can however be doubted if such a system would actually lead to a more efficient system than the current one with its multiple negotiations between agencies and companies that represent different interests. All the three reasons referred to above, namely information

processing problems, cognitive limitations, and deficient democracy, are applicable to this example.

Example 4: Decisions on what risks to accept, and what to pay to avoid a risk, are made by various agencies in different policy areas. Willingness to pay for safety, as measured in the marginal cost of saving a life, differs widely between policy areas (Ramsberg and Sjöberg 1997). Some cost–benefit analysts are dissatisfied with this. They claim that all decisions on risk acceptance should be coordinated so that willingness to pay is equalized across the policy areas. Viscusi (2000: 855) is representative when proposing that we should 'spend up to the same marginal cost-per-life-saved amount for different agencies'.

Such proposals require a high degree of coordination across policy areas. This might have been unproblematic if risk management decisions, in each policy area, were completely separated from other decisions in that area. If all decisions affecting risks were 'pure' risk decisions, and other decisions did not affect risks, then it would be feasible to apply a uniform price to risk avoidance. In practice, however, this is not how these decisions are made. Risk issues are usually parts of various complex issues. For an example, consider an industrial company that invests 50 million euros in new and better machinery in order to both increase efficiency and decrease the risk of workplace accidents. The decision-makers do not know how much of that sum they pay for increased efficiency and how much they pay for decreased risk – and neither do they need to know. Risk decisions are so interwoven with other decisions that it is not feasible to make risk decisions in a fully coordinated and centralized way while retaining an uncoordinated and decentralized decision structure for other decisions. Hence, information problems alone are sufficient to put this application of the synoptic view into question.

It should be emphasized that criticism against the synoptic view of decision making is not equal to a general criticism of CBAs. A CBA does not necessarily have a synoptic, all-inclusive perspective. It is also possible to tailor a CBA to a limited decision perspective. However, the internal logic, or impetus, of cost–benefit analysis seems to lead its practitioners in the synoptic direction. Some analysts tend to promote a centralized or highly aggregated structure of decision-making, without proper analysis of the pros and cons of such a structure. As a consequence of this, issues that will in social practice have to be dealt with as coordination problems can be depicted as optimization problems and therefore not analyzed adequately. The over-synoptic temptation in cost–benefit analysis will have to be resisted, and replaced by strivings to find a reasonable balance between the advantages of optimization and the dangers of centralized decision-making.

3. OPTION CHARACTERIZATION PROBLEMS

The outcome of a decision analysis such as a CBA is sensitive to how the decision alternatives are described and characterized, or in other words, what aspects are selected for evaluation. How far we follow alternatives into the future depends on our predictive capacity (Section 3.1) and on the extent to which we treat our future decisions as under our present control (Section 3.2). Another philosophically interesting aspect of option characterization is the extent to which we include positive consequences of immoral acts in the analysis (Section 3.3). The inclusion/exclusion decisions have to be made in a way that is unbiased between the alternatives, but this is often difficult to achieve (Section 3.4).

3.1 Prediction

In order to assess the alternatives in a decision we need to determine, for each of these alternatives, the future developments that will follow if it is chosen. However, it is well known from the parallel area of technological forecasting that predictions are highly uncertain. A major reason for the many failures of technological forecasting is that the development of technology depends on social factors, such as decisions by individuals and social groups on whether and how to make use of technological options. New social developments and their interactions with technology tend to come as surprises. Therefore, previous forecast-oriented approaches to technology assessment have largely given way to more limited approaches that do not profess to follow the consequences of a particular course of action into the future (Palm and Hansson 2006).

Our inability to predict how different alternatives will develop into the future affects not only cost–benefit analysis but also other policy-guiding practices. It puts limits to our ability to plan rationally for the future, with cost–benefit analysis or with any other method. Arguably, however, this problem is particularly conspicuous for methods like CBA that conform with consequentialist principles. Consequentialism refers to all consequences of an action, including consequences that occur far off in a distant future and/or depend on very complex causal mechanisms. Therefore it is a problem for consequentialism that we cannot know the distal consequences of our actions (Singer 1982; Norcross 1990; Simons 1999; Lenman 2000).[5]

Example 1 above (the road construction project) illustrates how prediction problems complicate the characterization of alternatives in cost–benefit analysis. It is extremely difficult in this case to know what will

[5] In contrast, the causation requirement in legal liability is relatively strict, and liability does not extend into long causal chains (Umari 1999).

https://doi.org/10.1017/S0266267107001356 Published online by Cambridge University Press

happen, if the project does not materialize, to the workers who would otherwise have worked in the project.

CBAs typically involve an attempt to cover consequences as completely as possible. However, as we extend our analysis to include more and more distant indirect consequences, uncertainty tends to increase (Hofstetter *et al.* 2002). The analysis has to stop somewhere before it becomes meaningless.[6] No general method for deciding where to stop is available.

3.2 Control of future decisions

The characterization of future developments that is needed in a CBA is complicated not only by prediction problems but also by the indeterminateness of our control over future decisions. Decision analysis usually has its focus on one-shot decisions, i.e. situations in which a decision-maker has exactly one decision to make in a matter. In practice, we often have a series of future decision points ahead of us that affect the issue at hand. How should such future decisions be dealt with in a CBA intended to guide the decision that comes first?

In decision theory, the treatment of future decisions has mostly been discussed in relation to individual decisions (Rabinowicz 2002). The crucial issue is whether or not one should treat future decisions as under one's control. The obvious alternative to this is to make a probabilistic assessment of how one will act in the future decision situation, just as one would have done if the future decision were to be made by some other person. The following two examples illustrate the problem:

Example 5: A non-smoker considers the possibility of smoking for just one week and then stop, in order to achieve a better understanding of why so many people smoke. When making this decision she may regard herself as being in control of the future decision whether or not to stop after a week. Under this assumption, a trial period of one week does not seem unreasonable. Alternatively, she can make a probabilistic appraisal of what she will do in that situation and will probably find that trying to smoke comes out as a much too dangerous operation.

Example 6: A heavy cigarette smoker considers whether or not to try to quit. Statistics available to her show that among people in her age group and socioeconomic group who have smoked as much as she has, almost everyone who tries to quit fails. When making this decision, she may regard herself as being in control of future decisions whether or not to

[6] In the CBA literature this has been referred to as a 'stopping problem' (Keeney and von Winterfeldt 1986; Wiener 1998; Rascoff and Revesz 2002). On the analogous problem of delimiting legal causation requirements, see Umari (1999: 488).

start smoking again. An analysis based on these assumptions will provide her with good reasons to try to stop smoking. Alternatively, she can make a probabilistic appraisal of her future decisions. From such a viewpoint, quitting may seem to have a too meagre prospect of succeeding to be worth trying.

Probably, most of us would prefer the probabilistic approach to future decisions in Example 5, and the non-probabilistic one in Example 6. In this particular case the difference can be accounted for in terms of cautious decision-making. However, no general rule seems to be available that can be used in all situations to determine whether or not a decision-maker should act as if she is in control of her own future decisions.

The following example shows that the control problem can be highly relevant in cost–benefit analysis:

Example 7: A proposal has been made to allow the construction of a road and a hotel in a previously untouched national park. The negative effects on the park's fauna and flora can be shown to be small. However (we may assume), experience from other parks shows that such decisions tend to lead to a series of small decisions of the same kind that gradually bring about serious damage to the park.

In a traditional CBA of this project, only the direct effects would be included. However, a good case can be made that the indirect effects should also be taken into account. More generally speaking, the treatment of future decision points in a CBA is problematic, and no general recipe for how to deal with them seems to be available.

3.3 Moral exclusion

It is a common criticism against consequentialism that it counts positive consequences of immoral acts in the same way as positive consequences of morally acceptable acts. In hedonistic utilitarianism, the pleasure a murderer derives from his deed is included on the positive side of the hedonic calculation. Similarly, preference utilitarianism will include the satisfaction of his preferences as a positive factor. This is commonly taken to be a problematic feature in utilitarian theory. Various adjustments have been made to avoid it, including the introduction of 'laundered' preferences and second-order preferences (Frankfurt 197; Broome 1991).

Cost–benefit analysis has the same problem as consequentialism with positive consequences of immoral acts. For a somewhat extreme example (from an area in which CBAs are seldom performed), suppose that a CBA is made of a programme against sexual assaults. The sufferings of the victims should clearly be included in the analysis. On the other hand, few would wish to include the perverse pleasures experienced by the perpetrators as

a positive factor in the analysis. Generally speaking, we expect a credible CBA to exclude positive effects that depend on immoral behaviour or preferences. However, it is difficult to delimit the scope of this exclusion rule. The following example shows that the same type of problem can emerge also in dealing with subject-matter of types that are commonly subjected to cost–benefit analysis:

Example 8: In the CBAs of Swedish road construction projects, the time gain of motorists is one of the factors included in the analysis. This applies not only to legal time gains but also to the time gains that motorists obtain by exceeding speed limits.

Since speeding is a criminal offence that kills a large number of innocent persons each year, the inclusion of time gains from speeding is morally controversial. Presumably, cost–benefit analysts would not consider including the usefulness of a new road for illegal street-racing in their analyses. The reason why time gains from 'ordinary' speeding is included in the analysis is quite obvious: although speeding is a serious crime in view of its consequences, it is socially much more accepted than other crimes that kill innocent victims.

As this example shows, it is far from evident where to draw the line of moral exclusion. Furthermore, the cost–benefit analyst will have to decide whether to base her decisions concerning moral exclusion on her own moral values, on those of the customer, on moral consensus in society, or perhaps on an independent moral analysis.

3.4 Bias

It is clearly a desideratum that characterization problems such as those introduced in Sections 3.1–3.3 should be dealt with uniformly, so that all decision alternatives are treated in the same way. This is necessary to avoid bias. However, further specifications are needed of what it means to be non-biased in this respect.

The most obvious delimitation criterion for cost–benefit analysis is maximal inclusiveness:

(1) Include all consequences that can be identified.

This is indeed the standard answer to accusations that cost–benefit analysis leaves out important factors from the analysis.[7] In practice it is seldom

[7] This proposal seems to be implied by Amartya Sen (2000). The other major response to this accusation is that a CBA is not the last word, but has to be followed by a discussion that includes aspects not covered in the CBA. This response is particularly often used

https://doi.org/10.1017/S0266267107001356 Published online by Cambridge University Press

applied, since – as we have already mentioned – some consequences are extremely difficult to express in a form that is accessible to quantified analysis. The following variant describes better what ambitious cost–benefit analysts try to accomplish:

(2) Include all consequences that can be identified and quantified.

However, both (1) and (2) may give rise to biased selections of consequences. This will happen if there are systematic differences between positive and negative effects so that one of the two categories is easier to identify (or in the case of 2: easier to quantify) than the other.

In environmental economics, several mechanisms have been pointed out that can give rise to such biases. Hence, some authors claim that in the analysis of a proposed regulation, much more focus is usually placed on possible secondary effects that are negative than on those on the positive side (Rascoff and Revesz 2002; Hofstetter *et al.* 2002). So-called attentiveness effects of a regulation are instances of possibly neglected positive effects. For example, the promulgation of a health and safety regulation might make people more sensitive to needs of safety that go beyond the specific regulation (Rascoff and Revesz 2002: 1811). Another important class of potential positive effects of regulation that are difficult to include in a CBA is innovation effects. In a CBA the estimates of compliance costs are normally based on presently available technology. It would not seem responsible to assume that yet unknown technologies will reduce these costs. Yet, this is what often happens. Good-faith estimates by companies of their compliance costs may tend to be too high, because firms do not fully anticipate cost–saving measures that they will discover once they devote resources to compliance.[8] According to the controversial Porter hypothesis, environmental regulations typically

towards the accusation that cost–benefit analysis leaves out issues of distributive justice. Hence David Copp (1987: 71) claims that distributive justice is 'outside the scope of CB [cost–benefit] analysis'. According to David Schmidtz (2001: 153), when a proposal fails the test of CBA, then no further discussion is warranted. However, if the proposal passes the test of CBA, then it may nevertheless fail a subsequent test of fair distribution. This is a strange asymmetry, since it prohibits measures that have a very positive distributive effect but slightly higher total costs than total (non-distributive) benefits.

[8] In a review of six OSHA rulings, the OTA found that regulatory costs were often overestimated, in part because OSHA deliberately reported estimates on the high side in order to minimize conflicts with industry on these estimates (Office of Technology Assessment 1995). The price of tradable SO_2 emission permits introduced under the American Clean Air in 1990 act were several times below *ex ante* estimates (Ellerman and Montero 1998). A brief review of about a dozen cases of environmental regulation in the US indicated that costs for emission reduction were consistently overestimated, generally by more than 100%. However, costs for environmental cleanup were underestimated. (Goodstein 1997) See also Hammitt (2000).

stimulate innovations to such extent that gains are created that offset the costs of regulation (Porter 1990; Xepapadeas and de Zeeuw 1999).

Another type of systematic bias is of particular relevance for (2): some types of consequences are difficult or impossible to quantify. This includes risks of cultural impoverishment, social isolation, and increased tensions between social strata. These difficulties can lead to a bias in a CBA in which an alternative with mostly quantifiable negative consequences is compared to one whose major drawbacks are non-quantifiable (Hansson 1989).

For our present purposes it is not necessary to determine the size of these or other types of potential biases. It is sufficient to observe that biases can arise when alternatives differ in how accessible to analysis their (negative and positive) consequences are. Based on this, (1) and (2) should be replaced by a criterion that excludes evaluations creating bias. At least provisionally, the sensible modification would seem to consist in going as far as we can in the inclusive direction without succumbing to bias:

(3) Include those, and only those, consequences that can be identified and meaningfully assessed in a reasonably uniform manner for all options under consideration.

(3) is an attempt to strike a balance between the two requirements to include as many aspects as possible and to treat all alternatives in the same way. It does this by including exactly those aspects that can be included without treating alternatives differently.

4. VALUATION PROBLEMS

By far the most discussed philosophical issue in cost–benefit analysis is the treatment of incommensurable values (Section 4.1). In addition, two value-related assumptions that are usually taken implicitly for granted are in need of explicit discussion, namely the transferability of values across contexts (Section 4.2) and the justifiability of interpersonal aggregation (Section 4.3).

4.1 Incommensurability

Most of the philosophical discussion about cost–benefit analysis has been concerned with the difficulties involved in assigning an economic value to that which we conceive as invaluable, such as a human life or an animal species. It is more or less in the nature of cost–benefit analysis to be open to criticism for valuing the invaluable. When we are only concerned with commodities that have an uncontroversial monetary value (and we do not wish to take externalities into account), we can usually rely on actual

market values. Recourse is usually taken to cost–benefit analysis only when goods are involved that do not have a market value.

Human lives do not have a monetary price in the common sense of the word. A cost–benefit analyst who assigns a monetary value to the loss of a human life does not thereby imply that someone can buy another person, or has the right to kill her, for that price. Essentially, lives and money are incommensurable, and the values of lives included in a CBA are for calculation purposes only. However, incommensurability between life and money is only one of the many incommensurabilities that have to be dealt with in cost–benefit analysis. Death, disease and environmental damage are not easily commensurated. There is no definite answer to the question of how many cases of juvenile diabetes correspond to one death, or of what amount of human suffering or death corresponds to the extinction of an antelope species.

In a CBA, multi-dimensional decision problems are reduced to uni-dimensional ones. The common way to do this, technically, is to assign monetary values to all types of consequences, even those that are incommensurable with money. Therefore the problem of incommensurability appears as a problem of assigning sums of money to units in the analysis that do not have a monetary price. However, if we removed money from the analysis we would still have to deal with comparisons between deaths, diseases, and environmental damage. The basic, underlying, problem is not limited to valuation in monetary terms. The essential problem – or perhaps even dilemma – is that we need to comparatively evaluate entities that we conceive as incomparable. Such 'impossible' comparisons are unavoidable components of many of the decisions that we have to make in different social sectors. The problem does not come with cost–benefit analysis, it is only more clearly exhibited when a CBA is performed in order to guide the decision.

This being said, it should be conceded that money has connotations not shared by non-monetary units that can sometimes be used for the same or similar purposes, such as QALYs (quality adjusted life years). The use of money instead of some other unit may therefore send a message that can be conceived as desecrating the value of life. Nevertheless, the crucial issue is the very act of comparison, not the currency in which it is expressed.

One of the most common methods used to derive calculation values for non-market goods is contingent valuation. This means that the values are based on people's answers to questions of the type: 'How much would you be prepared to pay for saving the giant panda from extinction?' The presumption is that the sum of everyone's answer to that question determines the value that the non-extinction of the giant panda should be assigned in a CBA. It turns out, however, that our answers to such questions do not give good indications of our priorities. Hence, Beattie *et al.* (1998) found that many respondents tend to report an amount that would not

seriously disturb their normal expenditure and savings patterns, typically a sum in the range £50–200 per annum. Respondents were also insensitive to the magnitude of the risk reduction (cf. Vadnjal and O'Connor 1994; Hart and Latacz-Lohmann 2001). No other, more reliable method seems to be available for eliciting calculation values from questionnaire respondents.

Some methods used in cost–benefit analysis, including contingent valuation, tend to give more influence to affluent people since they can pay more than others to have it their way (Copp 1987). Although this is a common feature in CBAs, it is methodologically easy to avoid for instance by relating (actual or hypothetical) payments by an individual to that individual's income.

4.2 Transferability across contexts

In CBAs, cost estimates are regularly transferred across contexts. This applies, in particular, to estimates of life values.

Example 9: In a CBA of mammography performed in 1992, the American FDA used values of life to determine, in monetary terms, the economic benefit from saving a life with mammography. The life values they used were derived from estimates of how much more male workers are paid when working in occupations with a high risk of fatal accidents (Heinzerling 2000: 205–6).

Example 10: In 2000, the American EPA performed a CBA of a new standard for arsenic in drinking water. Here as well, values of life were used that were derived from studies of how much male workers receive in compensation for risks of fatal accidents (Heinzerling 2002: 2312).

In both these cases, it would have been possible to use life values derived from the very context of the CBA in question. Women could have been asked how much they are prepared to pay for mammography, given realistic assumptions about the risk reduction it gives rise to. Their willingness to pay for reduced risks could then be used in a CBA for mammography. Although the use of such values would not have been unproblematic, it would at least have been much closer to the relevant context than the life value that was actually used. Similarly, people could have been asked how much they were prepared to pay for reduced levels of arsenic in drinking water, given realistic assumptions about the health effects of such a reduction. As was noted by Lisa Heinzerling, from whose work both of these cases are taken:

> Because a thriving market for bottled water exists and because a desire for safe and healthy drinking water drives much of this market, the EPA's choice to use wages for risky jobs as a proxy for the value of clean drinking water

> is problematic. There is a commercial market in the very item the EPA was valuing, but the EPA ignored it (Heinzerling 2002: 2324).

The assumed transferability across contexts that is illustrated in these examples is in fact an essential condition without which cost–benefit analyses as we know them today cannot be justified. If all values used in a CBA had to be derived from the precise context of the particular CBA, then the practice of performing CBAs would come close to that of performing opinion polls on the topic to be analyzed. We might for instance inform people of the risks associated with arsenic in drinking water and then ask them how much they are prepared to pay for arsenic-free drinking water. We might ask women how they value mammography. Why is this not better than performing a CBA with economic values derived from other contexts?[9] Once transferability across contexts is given up, we seem to enter a slippery slope in which the characteristic features of CBAs as we know them today would be lost.

In order to defend transferability across contexts, one would have to claim that it is *better* to use values from a certain context (such as wages that compensate for workplace risks) in a CBA concerning another context (such as mammography), than to use values derived in the context of the CBA in question. I can think of two major types of defence of such a claim.

First, it could be claimed that values should be transferred against contexts for practical reasons, in particular in order to reduce costs of investigation. This would be a weak defence since it defends transferred values only as approximations of contextual values. Transferred values cannot be defended in this way in cases where contextual values would yield a widely different result. (We do not know which these contexts are, but arsenic in drinking water may be a case in point.)

Secondly, it could be claimed that transferred values must be used in order to obtain consistency between different CBAs. If different values are used in different CBAs, then decision-makers basing their decisions on these different CBAs would arguably make inconsistent decisions. This argument is based on the value assumption that our evaluations of a consequence should be the same irrespective of the context in which that consequence appears. We can call this the *sameness thesis*. According to this widely held view, a life lost in a workplace accident and a (hypothetical) life lost due to arsenic in drinking water should be assigned the same value.

The sameness thesis is very difficult to defend from a normative point of view. When evaluating consequences we do not normally separate them from their contexts, and it is far from obvious why we should do so. To

[9] The transferability issue has a parallel in general utilitarianism. If there is no general value that can be transferred, but all values are holistic and context-dependent, then the additive structure of utilitarianism is not meaningful.

take just one example, we may choose to pay more per life saved in a law enforcement programme that reduces the frequency of manslaughter than we would pay for most other life-saving activities. Policy analysis should be adjustable to our choices in these matters, not dictate them according to the sameness thesis or any other fixed principle.

It should also be noted that in order to motivate transfer of values across contexts, it is not sufficient to argue for the context-independence of these values. Arguments must also be given as to why certain values and not others should be chosen to be transferred across contexts. I have not seen any such arguments and in particular no arguments for the common practice of using risk-related wage differentials to determine values of life. Most of us have very limited information about the dangers in different employments. Therefore, our choices of employment do not seem to be a good guide as to how we value risks economically.

4.3 Interpersonal aggregation

In a CBA, all costs and all benefits are combined in one and the same balance. This means that a disadvantage affecting one person can be fully compensated for by an advantage of the same size that affects some other person. In other words, *interpersonal compensability* of advantages and disadvantages is assumed (Hansson 2004).[10] This is an assumption that cost–benefit analysis shares with utilitarianism. We can express it as a weighing principle, as follows:

The collectivist weighing principle

An option is acceptable to the extent that the sum of all individual costs that it gives rise to is outweighed by the sum of all individual benefits that it gives rise to.

This is not the only way in which costs can be weighed against benefits. Another possibility is to perform the weighing individually for each affected person, and require a positive balance for each person:

The individualist weighing principle

An option is acceptable to the extent that the costs affecting each individual are outweighed by benefits for that same individual.

[10] Interpersonal compensability should not be conflated with the related but distinct issue of interpersonal comparability. Even if a benefit is greater than a harm, it need not cancel out the harm. Interpersonal comparability does not imply interpersonal compensability, but they are nevertheless closely related since the former is a necessary prerequisite for making the latter operative.

https://doi.org/10.1017/S0266267107001356 Published online by Cambridge University Press

Individualist weighing has a strong tradition in social practices that have their origin in the physician–patient relationship. This is perhaps best exemplified by the ethical principles employed in clinical trials.

Example 11: A physician selects patients for a clinical trial with a new, experimental treatment. Such a treatment involves risks and benefits that have to be weighed against each other. If the physician based this decision on a conventional CBA, then she would include a patient in the study if the risk to this patient is outweighed by the total social benefit. The total social benefit includes the expected gains from the study for future patients. With such a criterion, a patient can be included in the trial even if the risks by far exceed the expected gains to her personally. This, of course, is not how such decisions are made. Instead, they are made in accordance with the individualist weighing principle. A patient is not offered the opportunity to participate in a clinical trial unless it is believed that the risks to which she will be exposed are outweighed by the expected advantages for her of the experimental treatment (Hansson 2006).

Example 12: Although fish is generally speaking healthy food, contaminants in fish caught in certain waters give reason to recommend limits in fish consumption. Such recommendations are based on the positive and negative health effects on the individual (and in the case of pregnant or breast-feeding women, on corresponding effects on the child) (Knuth *et al.* 2003). It would be regarded as inappropriate to base such recommendations on a full CBA that included other factors, such as the effects of diminished fish consumption on employment in the fishing industry or on regional economics.

Why do we perform CBAs, with collectivist risk-weighing, when deciding on road projects and safety engineering, but use other types of calculations, based on individual risk-weighing, when deciding on clinical trials and dietary advice? These differences do not have their origin in principled normative argumentation but in the diverging social traditions in different policy areas. To put it bluntly, in some policy areas we have a tradition of sacrificing individual interests for the sake of collective goals, whereas in other areas individual interests have a much stronger protection.

It is a problem for cost–benefit analysis to motivate why we should employ total (collective) aggregation instead of alternative methods that protect individuals against sacrifice of their interests for collective goals.

5. CONCLUSION

In summary, we have identified ten classes of philosophical problems that affect the practical performance of cost–benefit analysis: Bias in

https://doi.org/10.1017/S0266267107001356 Published online by Cambridge University Press

topic selection, dependence on the decision perspective, dangers of super synopticism and undue centralization, prediction problems, the indeterminateness of our control over future decisions, the need to exclude certain consequences for moral reasons, bias in the delimitation of consequences, incommensurability of consequences, difficulties in defending the essential requirement of transferability across contexts, and the normatively questionable but equally essential assumption of interpersonal compensability. I do not wish to claim that all CBAs are affected by all these problems. However, most CBAs seem to be affected by some of them. This is sufficient reason to pay much more attention than has previously been done to the fundamental problems of cost–benefit analysis. It is also sufficient reason to avoid exaggerated claims on behalf of cost–benefit analysis, or of course of any other methodology that shares these problems.

We started out by noting that cost–benefit analysis is a form of applied consequentialism, and that it can therefore serve as a testing-ground for consequentialism. The preliminary tests reported here can be summarized as follows: 'The device malfunctions in several ways, and the fundamental construction may be flawed. However, we have no better alternative to recommend. It can therefore be used, but only with great caution.'

REFERENCES

Anderson, E. 1988. Values, risks and market norms. *Philosophy and Public Affairs* 17:54–65.

Beattie, J., J. Covey, P. Dolan, L. Hopkins, M. Jones-Lee, G. Loomes, N. Pidgeon, A. Robinson, and A. Spencer. 1998. On the contingent valuation of safety and the safety of contingent valuation: Part 1–Caveat Investigator. *Journal of Risk and Uncertainty* 17:5–25.

Broome, J. 1991. *Weighing Goods*. Blackwell.

Copp, D. 1987. The justice and rationale of cost–benefit analysis. *Theory and Decision* 23:65–87.

Ellerman, A.D. and J.-P. Montero. 1998. The declining trend in sulphur dioxide emissions: implications for allowance prices. *Journal of Environmental Economics and Management* 36:26–45.

Frankfurt, H.G. 1971. Freedom of the will and the concept of a person. *Journal of Philosophy* 68:5–20.

Gochfeld, M. 2004. Risk–risk balancing for hazardous waste workers: alternative work, traffic fatalities and unemployment. *Risk Analysis* 24:347–8.

Goodstein, E. 1997. Polluted data. *The American Prospect Online* 8(35). www.prospect.org/print/V8/35/goodstein-e.html.

Hammitt, J.K. 2000. Are the costs of proposed environmental regulations overestimated? Evidence from the CFC Phaseout. *Environmental and Resource Economics* 16:281–301.

Hansson, S.O. 1989. Dimensions of risk. *Risk Analysis* 9:107–112.

Hansson, S.O. 2004. Weighing risks and benefits. *Topoi* 23:145–52.

Hansson, S.O. 2006. Uncertainty and the ethics of clinical trials. *Theoretical Medicine and Bioethics* 27:149–67.

Hart, R. and U. Latacz-Lohmann. 2001. The indifference curve, motivation and morality in contingent valuation. *Environmental Values* 10:225–42.

Heinzerling, L. 2000. The rights of statistical people. *Harvard Environmental Law Review* 24:189–207.

Heinzerling, L. 2002. Markets for arsenic. *Georgetown Law Journal* 90:2311–39.
Hofstetter, P., J.C. Bare, J.K. Hammitt, P.A. Murphy, and G.E. Rice. 2002. Tools for comparative analysis of alternatives: competing or complementary perspectives? *Risk Analysis* 22:833–51.
Hornstein, D.T. 1993. Lessons from Federal Pesticide Regulation on the paradigms and politics of environmental law reform. *Yale Journal on Regulation* 10:369–446.
Keeney, R.L. 1990. Mortality risks induced by economic expenditures. *Risk Analysis* 10:147–59.
Keeney, R.L. 1994. Mortality risks induced by the costs of regulations. *Journal of Risk and Uncertainty* 8:95–110.
Keeney, R.L. and D. von Winterfeldt. 1986. Why indirect health risks of regulation should be examined. *Interfaces* 16:13–27.
Knuth, B.A., N.A. Connelly, J. Sheeshka, and J. Patterson. 2003. Weighing health benefit and health risk information when consuming sport-caught fish. *Risk Analysis* 23:1185–97.
Lenman, J. 2000. Consequentialism and cluelessness. *Philosophy and Public Affairs* 29:342–70.
Norcross, A. 1990. Consequentialism and the unforeseen future. *Analysis* 50:253–56.
Office of Technology Assessment. 1995. *Gauging Control Technology and Regulatory Impacts in Occupational Health and Safety: An Appraisal of OSHA's Analytic Approach*, US Congress, OTA-ENV-635, Washington DC.
Palm, E. and S.O. Hansson. 2006. The case for ethical technology assessment (eTA). *Technological Forecasting and Social Change* 73:543–58.
Porter, M. 1990. *The Competitive Advantage of Nations*. Free Press.
Rabinowicz, W. 2002. Does practical deliberation crowd out self-prediction. *Erkenntnis* 57:91–122.
Ramsberg, J. and L. Sjöberg. 1997. The cost-effectiveness of life saving interventions in Sweden. *Risk Analysis*, 17:467–78.
Rascoff, S.J. and R.L Revesz. 2002. The biases of risk tradeoff analysis: towards parity in environmental and health-and-safety regulation. *University of Chicago Law Review* 69:1763–836.
Sagoff, M. 1988. Some problems with environmental economics. *Environmental Ethics* 10:55–74.
Schmidtz, D. 2001. A place for cost–benefit analysis. *Philosophical Issues* 11:148–171.
Sen, A. 2000. The discipline of cost–benefit analysis. *Journal of Legal Studies* 29:931–52.
Simons, K.W. 1999. Negligence. In *Responsibility*, ed. E.F. Paul, F. D. Miller Jr, and J. Paul, 52–93. Cambridge University Press.
Singer, M. G. 1982. Incoherence, inconsistency, and moral theory: more on actual consequence utilitarianism. *Southern Journal of Philosophy* 20:375–91.
Sunstein, C.R. 2005. Cost–benefit analysis and the environment. *Ethics* 115:351–85.
Umari, B.A. 1999. Is tort law indifferent to moral luck? *Texas Law Review* 78:467–92.
Vadnjal, D. and M. O'Connor. 1994. What is the value of Rangitoo Island?, *Environmental Values* 3:369–80.
Viscusi, W.K. 2000. Risk equity. *Journal of Legal Studies* 29:843–71.
Wiener, J.B. 1998. Managing the iatrogenic risks of risk management. *Risk* 9:39–82.
Xepapadeas, A. and A. de Zeeuw. 1999. Environmental policy and the competitiveness of capital. *Journal of Environmental Economics and Management* 37:165–82.