John Broome

Structured and Unstructured Valuation

Abstract: Economists can value things for cost-benefit analysis using either a structured or an unstructured approach. The first imposes some theoretical structure on the valuation; the second does not. This paper explains the difference between the approaches and examines the relative merits of each. Cost-benefit analysis may be aimed at finding what would be the best action, or alternatively at finding which action should be done in a democracy. The paper explains the difference, and argues that the appropriate aim is the first. Given that, it comes down in favour of the structured approach to valuation. As an example, it discusses different approaches to valuing life in economics.

1. Introduction and an Example

Economists can value things for cost-benefit analysis in two different ways: they can adopt a *structured* or an *unstructured* approach. In this paper I shall first explain the difference between the two approaches. Then I shall examine the relative merits of each. One conclusion I shall draw is that we need to distinguish two different aims that cost-benefit analysis could pursue: it could aim to discover what ought to be done, or alternatively it could aim to discover what would be the best thing to do. These are not necessarily the same. I shall argue in favour of the latter. I shall support the structured approach to valuation, because it turns out to be appropriate for this latter aim.

As an example, I shall concentrate on valuing human life. To value life, some economists adopt a structured approach and others an unstructured approach, so this application illustrates the distinction nicely. A classic cost-benefit problem involving life and death is the question of how much it is worth spending on roads to make them safer. If more money is spent on safety, the benefit will be that some people who would have died on the roads will be saved. But if this money is not spent on safety, it will be available for other uses, which will benefit people in other ways. The problem of road-safety is a matter of balancing one sort of benefit against the other.

A very schematic version of the problem is illustrated in Figure 1. The diagram is divided into two halves. The left half represents the option of spending on safety; the right half the option of not spending on safety. I assume the society at present contains only two people. If money is spent on safety, I assume both of them will survive to old age. But if the money is not spent, one of the people will soon be

Analyse & Kritik 16 (1994), S. 121-132

Figure 1
Road safety

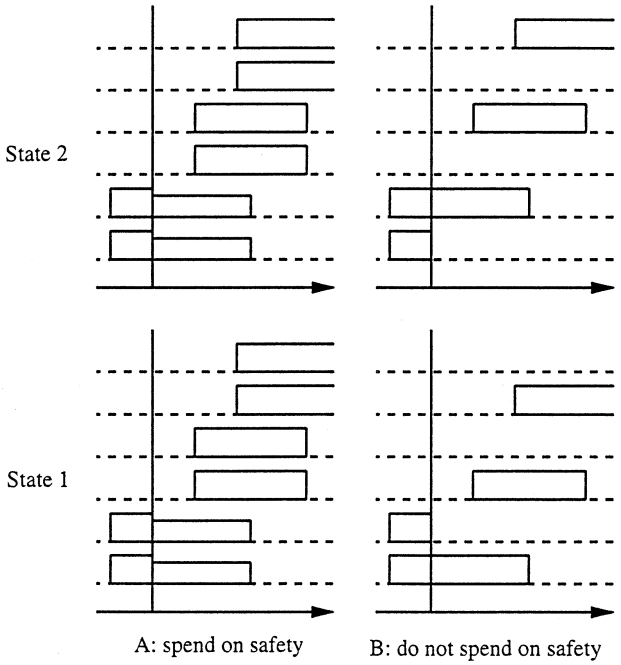

killed on the roads. One person will certainly die, but it is not known which of the two it will be. The uncertainty is represented in the diagram by two 'states of nature'. Nature may be such that, if money is not spent on safety, the first person will die, or it may be such that, if the money is not spent, the second person will die. On the other hand, if the money is spent, the state of nature makes no difference. The diagram shows the results of both options in both states, and it shows that the results of spending on safety are the same in both states.

The results are represented by graphs. Each graph has time on its horizontal axis, and a vertical bar marks the present. Each dotted line represents a person, and standing on each is a little graphlet that shows the person's level of wellbeing at all times in her life. A person's graphlet starts when she is born and continues till she dies. Each graph, then, shows a distribution of wellbeing across people and across time. To keep the diagram simple, I have assumed that people's wellbeing generally stays constant, so the graphs are generally horizontal straight lines. The presently living people are the bottom two in each graph; they were born before the present. The graphs show that, if money is spent on safety, both the existing people will continue to live, but they will suffer a reduction in their wellbeing. This is because the money will not be available for other beneficial purposes. If the money is not spent on safety, one person's life is shown ending immediately,

but the survivor will have a better life than she would have had if the money had been spent on safety.

Besides the two people alive now, the diagram shows their children and grandchildren. I intend it to represent something that is typical of young people: if a young person's life is saved, typically she will later have children and indeed a whole line of descendants, who would never have existed had she died. If money is spent on safety, the diagram shows two new people in each succeeding generation. But if the money is not spent, one of the presently living people dies before having children, so the diagram shows only one new person in each generation. (Evidently, these people reproduce by parthenogenesis.)

That is a schematic illustration of the problem facing a cost-benefit analyst. Practical problems are much more complicated, of course, and often there will be more than just two options available. The analysis proceeds by setting a value on each of the options, and it favours the one that turns out to be the most valuable. As I say, we can distinguish structured and unstructured approaches to making the valuation. I shall start by explaining the unstructured one.

2. The Unstructured Approach

The unstructured approach sets out from the preferences of the various people who are alive at the time of decision making. Each of these people will have preferences about the options that are available, and these are the basic data for the unstructured approach. They are the people's preferences about the options taken as a whole: for instance, their preferences between spending on safety or not spending on safety. I call this the unstructured approach because the economist making the evaluation does not herself take any account of the internal structure of the options. For instance, she takes no account of the way each option distributes wellbeing across time or across states of nature.

At least, she takes no direct account of these things. She does take account of people's preferences between the options, and in forming their preferences the people will presumably take account of the options' internal structure. In this way the structure enters the economist's evaluation indirectly. If an option has uncertain results, then each person's preferences will presumably take account of the uncertainty. In life-and-death cases, each option presents each person with a life of some quality and length, so in these cases the people's preferences will presumably reflect the values they attach to different sorts of lives. Because of uncertainty, in practice each option will present each person with not just one life but a portfolio of possible lives: a different one in each state of nature. The people's preferences should reflect their assessments of these portfolios. For example, in Figure 1 each person faces two alternatives: on the one hand living a moderately good life for sure, and on the other a chance of dying immediately and a chance of living a very good life. Each person's preferences between the two options should reflect her assessment of the quality of the lives she might lead, and of the uncertainty that faces her.

Because the unstructured approach ignores the structure of the options, it assumes no particular theory about the form of people's preferences. It does not assume they will conform to expected utility theory, for instance. Expected utility theory requires a person's preference to conform to some particular formal conditions when the person is faced with uncertainty, but the unstructured approach does not assume those conditions. Nor does it impose any theory about how a person values her future life: it assumes nothing about how she takes account of its prospective length and quality.

People's preferences amongst the options are the basic data of the unstructured approach, but this approach needs more than the basic data to arrive at an overall valuation for the purposes of cost-benefit analysis. For one thing, it needs a way of aggregating people's preferences together, or weighing them against each other. It needs to do some 'interpersonal weighing', I shall say. The subject of this paper is the basic data, not the interpersonal weighing, but I need to say something about the interpersonal weighing all the same. Particular views about it are often implicit in the way the basic data are collected in practice. This can obscure the distinction between unstructured and structured approaches, so I need to explain it.

The basic data are the people's preferences. Each person's preferences can be represented by a utility function. A utility function is a function that assigns a number called a 'utility' to each of the options in such a way that one option gets a higher number than another if and only if it is preferred. It represents the order of preferences by means of utilities. The way cost-benefit analysts aggregate different people's preferences together is to add up utilities across people. (This is not generally how they themselves describe what they do.) Each person's utility function assigns a utility to each option. Adding these utilities across people gives a total utility to each option, and cost-benefit analysis favours the option that has the highest total.

Because a utility function only has to represent the order of a person's preferences, many different utility functions will represent the same preferences. This means that, before they aggregate preferences by adding utilities, cost-benefit analysts have a choice of which utility function to pick for each person. Naturally, they pick the one they think best serves the purpose of aggregation or interpersonal weighing. I shall give two examples.

One example is the use of 'healthy-year equivalents', which has recently been gaining ground in health economics (see, for instance, Mehrez/Gafni 1989). For each of the options we are interested in, we can define its healthy-years equivalent (or 'hye') as the number of years of healthy life that the person finds indifferent to the option. The option, remember, may involve some uncertainty: it may be a portfolio of various possible patterns of life. We can assume that the person prefers a longer healthy life to a shorter one. So one option will have a greater hye than another if and only if it is preferred. It follows that hyes are utilities. More exactly, the function that assigns to each option its hye is a utility function; it is one of the many utility functions that represent the preferences. Hyes are used by health economists to compare alternative treatments; a treatment that produces

more hyes in total is supposed to be better than one that produces fewer. So hyes are added across people. Evidently, these economists believe hyes are appropriate utilities for the purpose of interpersonal weighing. The effect is to count a hye to one person as equally as valuable as a hye to another. I am not proposing to assess the merits of this idea; I am only offering it as an example of the unstructured approach. It is unstructured because hyes represent each person's preferences about the options as a whole, incorporating uncertainty and all their other features.

A second example is the use of willingness to pay. Pick one of the options and call it 'the status quo'. Then for each option we can define a person's 'willingness to pay' for the option relative to the status quo as the amount of money the person would be willing to pay to have the option instead of the status quo. Let me put this more precisely. If the option came about, then in each state of nature the person would have some particular amount of money. Imagine this option came about, but as well as that imagine the person had a fixed sum of money taken away from her in every state of nature, deducted from what she would otherwise have had. The amount of deducted money that would make her indifferent to the status quo is her willingness to pay for the option. Her willingness to pay for the status quo is zero, and her willingness to pay for some options may be negative.

Normally a person would be willing to pay more to have one option than another if and only if she prefers the former to the latter. So the function that assigns to each option her willingness to pay for it is a utility function for her. It gives a scale of utility in terms of money. It is only one such scale; there is a whole range of 'money-metric utilities', as they are called. First, there is a whole range of willingness-to-pay scales of utility, because each choice of a particular option to serve as the status quo will yield a different scale. And there are other money-metric utilities besides those. One, known as the 'equivalent variation' of the option, is minus the amount of money the person would be willing to pay to have the status quo rather than the option.

Many cost-benefit analysts evaluate an option by adding together different people's willingnesses to pay for the option. They must think that willingness to pay is an appropriate basis for interpersonal weighing. In effect, they must think a dollar to one person is as valuable as a dollar to anyone else. Willingness to pay is a second example of an unstructured approach to valuation.

3. The Structured Approach

That is the unstructured approach, then. By the structured approach, I mean an approach where the economist takes account of the structure of the options facing the people, and imposes some theory about how the value of an option depends on its structure. This is a vague description so far, and a matter of degree: there are highly structured approaches and less structured approaches.

In practice, every cost-benefit analysis is structured to some degree. Imagine a completely unstructured cost-benefit analysis of, say, a new road. Let it use, say,

willingness to pay as the basis for interpersonal weighing, and let us take the option of not building the road as the status quo. To do an unstructured analysis, we would find out from each person what she is willing to pay to have the road, leaving it to her to put together all the different costs and benefits to her of having it. She would have to make her own assessment of the damage the road will do to the environment, of how often she is likely to use the road, of how much time it will save her when she does use it, of how likely she is to get killed while using it, and so on. All of these things will determine her preferences, and how much she is willing to pay to have the road. Few people will make a good job of that sort of assessment, and no cost-benefit analyst in practice uses such raw preferences about the options. Instead she first produces a breakdown of the various benefits and harms that will result from the road. She calculates the saving in time, the change in people's risk of accident, the effects on the environment and so on. Then she elicits values for the various different components. In all of this, the cost-benefit analyst makes theoretical assumptions. For instance, she makes assumptions about what is valuable about a new road; she assumes its value is to save time and improve safety, say, rather than to let people enjoy the sound their tyres make on the road surface. These are no doubt innocuous assumptions, but still they impose some sort of structure.

But the structured approaches I am thinking of impose more structure than that. What sort of structure do I mean, and what sort of theory should it be based on? There are many different possibilities, depending on precisely what we take to be the objective of our valuation, and I shall come to that later. Clearly, too, whatever theory we impose will need justification. But here are some examples of structures we might impose. First, we could impose some theory of rationality on people's preferences. Doing so would imply that we wished to base our valuation on what people's preferences would be if they were rational, rather than on people's actual preferences. For instance, we could impose expected utility theory, which many people take to be an account of rationality for preferences. Here is how this would work. We would evaluate an uncertain option, from the point of view of each person, by first evaluating the various outcomes that would result from it in the various states of nature, and then calculating their expectation, perhaps adjusted for the degree of riskiness. This structured evaluation would start from the value of outcomes. By contrast, an unstructured evaluation would be based on the people's preferences about the options themselves, taken as a whole.

Another example is that we could impose some sort of structure across time. For instance, we could evaluate each person's life as an aggregate of how good her life is at each time. Perhaps we might assume it is the simple total of the good contained in her life, or perhaps a total discounted for time, or perhaps some other sort of aggregate. Another example: we might impose expected utility theory at the level of our own evaluations. We might insist that the value we attach to an uncertain prospect should be the expectation of the value we assign to the outcomes that will result from it.

Like the unstructured approach, the structured approach will usually be based ultimately on people's preferences. But they will be preferences about different

things. Instead of preferences about options treated as a whole, they will be preferences about parts or aspects of options. For instance, a structured cost-benefit analysis of a road will rely on people's preferences about such things as saving time and the loss of a natural landmark, instead of their preferences about the existence of the road together with everything that goes with it. Another example: if we adopt expected utility theory, we might rely on people's preferences about what happens in each of the possible outcomes of an uncertain prospect, rather than their preferences about the prospect as a whole.

A good example of a structured approach in practice in economics is the use of quality-adjusted life years (qalys) to evaluate treatments in medicine. The value of a life is calculated as its length in years, adjusted for the quality of those years: a year in a wheelchair, say, or blind, is counted for less than a year in good health. Qalys impose a structure across time.

Qalys are commonly used by health economists, whereas transport economists generally adopt a less structured approach. To bring out the difference, let us compare how a typical transport economist and a typical health economist would approach the example of road safety in Figure 1. Both would ignore the people in future generations who are not yet alive; I shall mention them later, but for the time being I shall ignore them too. The transport economist would start by finding out which of the two options the people prefer as individuals, without worrying about why they prefer them. We cannot tell their preferences from the diagram, so we do not know what conclusion the transport economist will come to in this case. The health economist would look at the qalys that will result from either option for the two people, proceeding as follows. She will look at each state of nature separately. Since the states are symmetrical, her results will be the same for both. In each state, she will look to see which option delivers more qalys in total to the two people. As I happen to have drawn the diagram, the safety option wins easily in both states. Next the health economist will aggregate across states of nature. In this case that is easy, because safety is better in both states. So the health economist will be in favour of spending money on safety.

People's preferences are implicitly embedded in the health economist's calculations, but they are overlaid by a fair amount of theory. Preferences enter through the scale of wellbeing or quality of life that the health economist will use. I have been taking this scale for granted, but now is the moment to say that is generally derived from people's preferences about different qualities of life. Evidently, the transport economist's method relies on preferences in a more raw form.

This example illustrates a conspicuous division between the way human life is typically valued in two areas of economics: transport and health. For some reason the traditions are different in the two areas. The main difference is that health economists generally use a structured approach and transport economists an unstructured one. (But remember I have mentioned already a move towards an unstructured approach within health economics, using healthy-years equivalents.) There are other differences too, of course, but I think they are relatively minor. One difference often mentioned is that qalys are not intended for cost-benefit analysis of treatments, only for 'cost-effectiveness analysis'. Qalys are intended to

serve as a measure of the benefit that results from a treatment, but not as a measure that is comparable with the cost. When we evaluate various alternative treatments, we can work out the benefit of each in terms of qalys. We can also work out the cost of each in terms of money. We can then calculate the qalys each produces per unit cost. This will tell us the best way of distributing our resources across different treatments: we should concentrate on those that produce the most qalys per unit cost. That is the purpose of cost-effectiveness analysis. It will not tell us whether any particular treatment is worth its cost. That would require comparing the benefit in terms of qalys with the cost in terms of money. It would require us to decide the money-value of a qaly, which health economists generally decline to do.

Health economists typically see their job as finding the right way to distribute whatever resources are given over to health, rather than judging what total resources should be given over to health. On the other hand, transport economists are more ambitious. Not only do they tell us how the resources that are to be devoted to transport should be divided between different uses, they also tell us whether each particular use is worthwhile on its own. That is the idea of cost-benefit analysis. I do not know why the two departments of economics have this different culture. In any case, the difference is small. To start doing cost-benefit analysis, health economists have only to fix a money value on a qaly. They could do this by willingness to pay or some other means. Indeed, theoretically the difference is nonexistent. Theoretically, a cost is nothing more than a diminution in some people's quality of life – the quality of taxpayers' lives or the lives of whoever pays. My example of road safety shows that: the cost of safety shows up as a diminution of the people's wellbeing. In theory, then, money costs should show up in a qaly calculation, just like changes in people's state of health. They do not show up in practice because health economists deal only with quality of life in terms of health. But there is no reason in theory why they should accept that limitation.

4. Arguments for the Two Approaches

Which is the better approach to valuation: the structured or the unstructured? I shall start by mentioning some disadvantages of the unstructured approach, and mention an advantage of it later.

One disadvantage is the well-known incoherence of people's preferences in the face of uncertainty and, indeed, even not in the face of uncertainty. People's preferences suffer from framing effects, embedding effects, preference reversal and all the other freaky effects that psychologists have discovered, to an extent that they cannot possibly be considered rational. All this is so well known by now that I do not need to spell it out (see, for instance, Tversky/Kahneman 1986). I will mention one example I came across recently. In a recent survey, M. W. Jones-Lee, G. Loomes and P. R. Philips (1994) asked a sample of people how much they would be willing to pay to reduce their risk of suffering an accident on the roads

that would lead to a minor permanent disability. The subjects were asked how much they would pay to reduce the risk from 24/100,000 to 20/100,000, and also how much they would pay to reduce it from 24/100,000 to 12/100,000. 37% of them said they were willing to pay only the same for the larger reduction in risk as for the smaller one, and 9.4% said they were willing to pay less for the larger reduction. No plausible account of rationality can be reconciled with these preferences. Discoveries like these led these three authors to adopt a structured approach to evaluating risks of injuries. They did not base their valuation directly on people's willingness to pay for changes in risk, but on a procedure that depended on expected utility theory. Evidently they did not wish to base their valuations on patently irrational preferences. This seems a reasonable attitude for them to take; I shall give some grounds for it later.

The irrationality of our preferences is a practical disadvantage of the unstructured approach. There is a more theoretical disadvantage too. I can illustrate it with the road-safety example once again, though to use that example I shall have to make a supposition that may seem far-fetched. So far, I have said nothing about the probabilities of the two states of nature. Let us suppose the two people assign them different probabilities. Suppose each is confident in her own driving ability; each knows someone will be killed if there is no spending on safety, but each believes it is more likely to be the other person. Let us suppose each forms her preferences on the basis of her own expected wellbeing, calculated according to her own probabilities. Provided each is optimistic enough about her own chances of survival, each will prefer the unsafe option. An unstructured approach, then, based on people's preferences, will have to come down in favour of this option. But this is clearly the less good one. The two existing people get less wellbeing in total, and the wellbeing is unequally distributed between them. The future people certainly do not count against the unsafe option. There are fewer of them in that option, and we have perhaps grown accustomed to thinking it is a good idea to have fewer people. But that is because we fear that having too many people will reduce the general standard of life. In this example that is not so. Therefore, we certainly have no reason to be against the existence of those people, even if we are not in favour of it either. There are really no grounds for concluding that the unsafe option is the better. Yet this is the one recommended by the unstructured approach. I have already said an evaluation that used qalys would reach the opposite conclusion. This is because it would override the differing probability judgements of the two people. It is these differing judgements that cause the problem, and most structured valuations would set them aside. Here is another reason for preferring a structured approach.

But the same example reveals simultaneously an advantage of the unstructured approach. Although it is clear that the safe option is better, there is also one thing to be said for the risky option: both people prefer it. The present population is unanimously in favour of it. So a democrat must surely be in favour of the risky option. If the structured approach is in favour of the safe option, it is undemocratic. Since the unstructured approach works straight from people's preferences, it has democracy on its side.

5. What Is Best and What Ought to Come About

The conflicting arguments for and against the unstructured approach show we need to make a distinction that is often missed by welfare economists. When we are faced with a choice of options, there is the question of which is best and there is the question of which ought to come about, and the answers to these questions need not be the same. Any democrat must recognize that. When there is a choice, you may believe that one of the options on offer is the best. But on the other hand, you may believe a different option is supported by most people. If so, and if you are a democrat, you may therefore believe this second option is the one that should come about. So the option you believe should come about is not the one you believe is best. That may also be the situation of a democrat faced by the example of safety on the roads. Because both people prefer the risky option, the democrat may believe the risky option should come about. Nevertheless, for the good reasons I gave, she ought to believe the safe option is the best.

Since there is sometimes a difference between what should come about and what is best, which should cost-benefit analysis aim at? I think it should aim at what is best. Democracy is a matter of having a satisfactory democratic process. If the process works as it should, it will ensure that what comes about is democratic. It is the process that should bring about what should be brought about. One essential input into the democratic process, without which it will not work properly, is discussion and judgement about what is best. People who participate in the democratic process at all levels make their decisions and place their votes partly on the basis of what they believe is for the best. If they are to be well-informed, they need to be informed by difficult and careful analysis of the relative goodness of the choices. This is just what good cost-benefit analysis should provide.

Is there any case on the other side, for aiming cost-benefit analysis at what should come about? The case would have to be that cost-benefit analysis is itself part of the democratic process. Democracy is a mechanism whereby the actions of individual people together determine what happens. When economists do their cost-benefit calculations, they would have to see themselves as part of this mechanism. This seems quite implausible to me; it is the economists and not the public who produce the result of the analysis. At best, economists might work out what conclusion a truly democratic process would arrive at if it was allowed to operate; they cannot see their work as a truly democratic process in its own right. So at best they could short-circuit democracy. But if they try to do this, they will subvert democracy. Suppose a decision is made on the basis of an analysis, which concludes that democracy would have led to this decision had it been allowed to operate. This is not a democratic decision. If welfare economists want to support democracy, they should concentrate on providing information about the relative goodness of the alternatives available, and leave the democratic institutions to maintain democracy.

The appeal of unstructured valuations can only be to democracy. They are based on raw preferences unfiltered by theory imposed from outside. People's raw preferences amongst options are unlikely to be good indicators of what is best.

They will certainly not be if they are irrational. That is why I said in Section 4 that the well-established irrationality of people's preferences gives good grounds for adopting the structured approach. But even if a person's preferences were rational, they might not be good indicators of what is best for several reasons. One is disagreement about probabilities. I showed why in Section 4. In the example of Figure 1, I showed that both people might prefer the risky option because of their differing judgements of probability, and yet the safe option is the better. There may be no irrationality of lack of information embedded in the preferences; the people might have reached their different judgements by setting out from different prior probabilities. So if cost-benefit analysis aims at what is best, as I think it should, the structured approach is the right one.

What makes it right is that it can be aimed at what is best. This gives us some guidance about the sort of structure a cost-benefit analyst should impose: it should be structure that aims at what is best. The analyst needs a theory of good, therefore. She needs, for instance, a theory about what is the best action to take in the face of uncertainty. She needs a theory about what makes a life good: a theory about how the wellbeing that comes in a life should be aggregated together to determine the overall goodness of a life. She needs a theory about how the length of a life contributes to determining the goodness of that life, and so on.

6. Valuing Existence

One final note. In the example of Figure 1, there are people who exist in one of the options being compared but not in the other. How are these people to be taken into account in a cost-benefit analysis? So far I have simply ignored them, but I want to mention them now.

It is within the spirit of the unstructured approach to count only the preferences of people who exist at present. These people are the only ones who have a vote in a democracy; people who do not yet exist do not vote. Stephen Marglin expresses the idea: "I want the government's social welfare function to reflect only the preferences of present individuals. Whatever else democratic theory may or may not imply, I consider it axiomatic that a democratic government reflects only the preferences of the individuals who are presently members of the body politic." (Marglin 1963, 97) Marglin gives no direct weight to the wellbeing of people who do not yet exist. Their interests may be represented indirectly in the preferences of existing people, if the existing people happen to care about them. Since the special appeal of the unstructured approach is its effort to satisfy democracy, it is natural for it not to count future people.

There is a second reason why those who adopt the unstructured approach generally do not count future people (see Drèze 1992, and Broome 1992). Commonly, these people use willingness to pay as their measure of people's preferences. People not yet born have no purchasing power and consequently no willingness to pay for anything, so naturally they do not get counted. But this is a weak reason for not counting them. Willingness to pay may be a useful money

measure of utility for weighing up benefits to one person against another. But the arguments for willingness to pay – whatever they may be – do not say anything about which people should be counted in the first place. Should we count people not yet alive? That must be a separate question. The fact that we have settled on a particular measure for comparisons cannot preempt that question.

I have argued in favour of the structured approach because it can be directed at what is best. If we are concerned with what is best, people not yet born should certainly be counted. Compare two options that are equally good for everyone now alive, and for everyone who will ever live, bar one. Suppose one option is better than the other for one person who is not yet born but who certainly will be born, whichever option is chosen. Then the one that is better for this person is plainly better. Plainly, this person who is not yet born should be counted.

If a person is definitely going to exist, there is no doubt we should count her wellbeing in evaluating the options. But in Figure 1 the very existence of particular people is in question; the different options have different populations. What we have to evaluate is, not making a future person better or worse off, but bringing a future person into existence. This is a tremendously difficult problem that needs to be faced before cost-benefit analysis of life-and-death problems can be done properly. But I am not going to try and deal with it in this paper.

Bibliography

Broome, J. (1992), Reply to Blackorby and Donaldson and Drèze, in: *Recherches Economiques de Louvain 58*, 167–71

Drèze, J. (1992), From the 'Value of Life' to the Economics and Ethics of Population: The Path Is Purely Methodological, in: *Recherches Economiques de Louvain 58*, 147–66

Jones-Lee, M. W./G. Loomes/P. R. Philips (1994), *Valuing the Prevention of Non-Fatal Road Injuries: Contingent Valuation vs Standard Gambles*, typescript

Marglin, S. A. (1963), The Social Rate of Discount and the Optimal Rate of Investment, in: *Quarterly Journal of Economics 77*, 95–111

Mehrez, A./A. Gafni (1989), Quality-Adjusted Life Years, Utility Theory and Health-Years Equivalents, in: *Medical Decision Making 9*, 142-9

Tversky, A./D. Kahneman (1986), Rational Choice and the Framing of Decisions, in: *Journal of Business 59*, 250-78, reprinted in K. Cook/M. Levi (eds.) (1990), *The Limits of Rationality*, Chicago, 60–89