



The Economic Value of Life

John Broome

Economica, New Series, Vol. 52, No. 207 (Aug., 1985), 281-294.

Stable URL:

<http://links.jstor.org/sici?sici=0013-0427%28198508%292%3A52%3A207%3C281%3ATEVOL%3E2.0.CO%3B2-7>

Economica is currently published by The London School of Economics and Political Science.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://uk.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://uk.jstor.org/journals/lonschool.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

The Economic Value of Life

By JOHN BROOME

University of Bristol

INTRODUCTION

Here is a stylized version of the problem of valuing life. A government has to decide between two alternative states of affairs it could bring about. In one, lives are saved at some cost in resources. In the other, the resources are used for some different purpose. Both states are to be thought of as complete histories for the world. They contain no uncertainty; the histories are fully determined.

The morally significant differences between the states can be classified under three headings. First, there is *length of life*: at least one person lives longer in one state than in the other. Second, there is *wellbeing*: since life-saving uses up resources, on the whole people will be worse off in the state where lives are saved. But there will usually be exceptions. Saving a life usually benefits several people, including the person who is saved and his dependants. The person may also contribute to the economy's production during his extra years of life, and so benefit the rest of us, but this has to be set against the demands he makes on the world's resources for his own consumption.

The third heading is *population*: the two states may have different populations. In one sense they certainly must have: there is a *time* when their populations must be different. Immediately after a person's life is saved the world's population is one more than it would have been. But in this paper when I speak of population I mean *timeless* population: all the people who live at some time or other. In this sense our alternative states do not necessarily have different populations, but often they will have. Saving a life often affects timeless population. The person who is saved may later have children, who may start a whole line of descendants. Or if somebody is not saved and dies his or her spouse may remarry and have more children. Or if a baby dies its parents may have another child instead. And so on. Sometimes the effect may be to change not the *numbers* of people who live but only their *identities*; the states may contain different people but the same number altogether. I call this too a difference in population.

Because it involves no uncertainty, the stylized choice I have described is simpler than the choices that governments face in practice. If a government opens a kidney unit, say, it does not normally know whose lives will be saved as a result, or even how many. In making its decisions it will therefore need some way of handling the uncertainty. But I am leaving uncertainty out of this paper because I have discussed it elsewhere (Broome, 1982) and the problems are hard enough without it. The stylized choice is only a first step towards practical decision-making.

It is true that some authors (e.g. Mishan, 1971; Jones-Lee, 1976) have claimed that one particular type of uncertainty—where it is uncertain *whose* life is at stake—can actually make decision-making easier. It allows us to replace the problem of valuing lives with what they believe is the easier problem

of valuing risks to lives. But I hope I have shown (Broome, 1978, 1982) that this must be a mistake, because we cannot value a risk to someone's life unless we can also value his actual life. The reason can be put very briefly: if it is worth the government's spending $\pounds x$ to save a person from a chance p of dying, ordinary decision theory tells us it is worth its spending $\pounds(x/p)$ to save him from certain death.¹

We have, then, two alternatives. How are we to decide which is better? What makes this question difficult is the likelihood that the alternatives will have different (timeless) populations. But all the work I know on valuing life, except for the article by W. B. Arthur (1981), discussed in Section III below, has simply ignored differences of population. One consequence is this. Suppose it should happen that one of the alternatives contains all the same people as the other and some extra people as well.² Then the good contained in the extra people's lives will not be taken account of in comparing the alternatives. For instance—this is the most important example—suppose a person is saved and later has children. Then the good this person enjoys in the remainder of his life will be counted by most economists as a benefit of saving him, but not the good enjoyed by his children and their descendants. I shall call this "the usual practice" in valuing life.

One motivation for this practice is perhaps the present fashion for valuing life by means of people's *willingness to pay* for safety. People are willing to pay for an increased chance of living longer, but no one can be found who is willing to pay for a chance of being born. But of course, this is not by itself a reason not to count the good of unborn descendants among the benefits of life-saving. Unborn people should not be disenfranchised just because they have so far acquired no purchasing power.

We need to decide, on good grounds, whether the usual practice is right or wrong. Until we have done so we cannot even start to fix a value on life. This is true even for cases where saving a life happens to make no difference to timeless population. The usual practice counts, as a benefit of saving a life, the good of the person who is saved but not the good of his descendants. If this is wrong it might be for either of two reasons. It might be wrong not to count the good of descendants, in which case young people who might yet have children ought to be valued much more highly than they usually are. Or it might be wrong to count fully the good of the person who is saved (this will be made more precise in Section II), and that will make a difference to the value of anybody's life whether or not saving him will alter the population. Section II shows that this is actually quite a plausible conclusion to come to. In this way, considering changes in population can tell us something about the value we should attach to life-saving even when it does not change the population.

In this paper I shall try to assess the usual practice, and I shall argue in the end that it cannot be justified. I shall also enquire whether there is any justifiable alternative we might put in its place, and argue that at the moment we do not know of one. I believe that at the moment we have no sound basis for valuing life.

Various principles have been proposed for comparing two states of affairs that may have different populations.³ I shall proceed by trying out some of the most important ones, to see what they have to say about valuing life.

I. THE TOTAL PRINCIPLE (TP)

For each person who exists in alternative *A*, work out how good his life is in *A* (aggregating the good and bad his life contains in whatever way is appropriate). Add up the goodness of everyone's life in *A*. Do the same for alternative *B*. The better alternative is the one with the greater total.⁴

Suppose in alternative *A* a person's life is saved so that he lives longer than he does in *B*. Among the benefits of saving him the TP will include all the extra good his life contains as a result. It will also include, equally, all the good enjoyed by any children he may have after his life is saved, and by all their descendants. The TP, then, disagrees with the usual practice in valuing life. Either we shall have to give up the usual practice or else we shall have to find good reasons for rejecting the TP.

And actually, there are good reasons for rejecting it. Chief among them is Derek Parfit's "repugnant conclusion" (1984, pp. 381-390). Given any state of affairs, even one where everyone's life is extremely good, it is always possible to imagine another state where people's lives are hardly good at all but where the population is so large that the total of good is actually greater than it is in the first state. The TP says that the second state is better—a conclusion that seems repugnant.

Another reason for rejecting the TP is closer to our concerns. It is worth looking at because it helps point the way to an alternative principle that offers an intuitively attractive foundation for the usual practice. The TP rules out a very plausible belief many people hold about the modern world: the belief that it is good to save life but also good to reduce the number of babies born. If it is good to reduce the number of babies born, then according to the TP that can be only because the good contained in a person's life is generally outweighed by the harm done to others by the demands he makes on the world's resources. But if this is true then, according to the TP, it is generally going to be better to let a person die than to save him. It will depend a bit on how productive the person is and what stage of life he has reached, but generally the good contained in the rest of his life will be outweighed by the harm done by his claim on resources. Saving him could generally be made worthwhile only by particular "side-effects" on other people that his death would have: the grief of relatives, for instance. And it does not seem plausible that side-effects like these can be all that makes the difference between saving life and making babies, making it right to encourage one and discourage the other. It seems that the crucial consideration should be the good of the person who is saved.

Take a particular example. According to the TP, contraception is pretty much on a par morally with infanticide. Both will have about the same effect on total good. The only moral difference is that killing a baby is likely to distress its parents and others.⁵ But it is very hard to believe that these side-effects are really the only moral difference between contraception and infanticide. The baby's own good must surely be crucial.

I think it seems fairly obvious where the TP goes wrong here. It attaches the same value to a person's good whether he already exists (or will exist in either alternative), or whether he is someone who might or might not be brought into existence. Suppose there are two ways of adding some amount

of good to the world: the life of someone already alive might be prolonged so as to add a certain amount of good, or some new person might be created whose life would contain this amount of good. The TP is indifferent between these alternatives. But this seems wrong. It seems to be treating people as mere receptacles of good, so that it is worth bringing them into the world simply for the sake of the good they will enjoy. This seems to get things the wrong way round. We count happiness, for instance, as a good because we think it good for people to be happy. We do not think happiness worth creating for its own sake, so that it is worth bringing people into the world simply to be bearers of happiness (compare Narveson, 1973, p. 73). In a sense, the TP does not value people at all. The only reason for having people around, it says, is to be the possessors of good. For this, one person will serve as well as another. There would be nothing gained by saving a person's life if instead that person could be replaced by someone else whose life would be just as good. The TP does not value at all, for its own sake, the continuing existence of a person.

What makes the TP unattractive could be explained by saying that saving a person's life benefits that person, whereas creating a person does not seem to benefit him at all.⁶ A state where a person continues to live is a better state for him than one where he dies, but a state where a person lives is neither better nor worse for him than one where he never lives at all. What seems wrong with the TP is that it favours changes—creating people—that benefit nobody. This intuition suggests that when we compare alternatives we should restrict ourselves to counting only good that benefits *somebody*, and this means good that belongs to people who exist in both the alternatives being compared. Section II examines some principles that embody this restriction.

II. RESTRICTED PRINCIPLES

The Person-restricted Principle (PRP)

Compare the alternatives *A* and *B* person by person. For each person check whether his life is better in *A* or in *B* or in neither. Then take all the people whose lives are better in *A* and add up the amounts by which they are better. Take all the people whose lives are better in *B* and add up the amounts by which they are better. The better alternative is the one with the greater total.⁷

The PRP only takes account of the good of people who exist in both alternatives. Consider somebody who exists in only one, say *A*. This person's life is neither better in *A* than in *B*, nor in *B* than in *A*. For his life to be better in *A*, for instance, it would have to be true that

$$(\exists X_a)(\exists X_b)(\text{His life in } A \text{ is good to degree } X_a \text{ \& his life in } B \text{ is good to degree } X_b \text{ \& } X_a \text{ is greater than } X_b.)$$

But this cannot be true because it cannot be true for any X_b that this person's life in *B* is good to degree X_b . Since his life does not exist in *B* it cannot have any property in *B*.⁸

So the PRP does the job we were wanting done at the end of Section I. This means it supports the usual practice for valuing life. If a person's life is saved the PRP tells us to count as a benefit the extra good he enjoys in his extended life but not the good in his descendants' lives; he himself exists (at

some time) in both the state where he is saved and the state where he is not, but his descendants exist only in the former.

However, the PRP may need some modification. And it will turn out that this may bring it even closer to what economists do in practice. Consider this example. In alternative *A* a baby dies. His parents soon replace him with another child who lives happily till he is 80. Resources that could have been used to save the baby are used instead to save a 20-year-old woman, who lives happily until she is 80. In alternative *B* the baby is saved and lives happily until he is 80, but the woman dies. The side-effects on other people are about equal in the two alternatives. The PRP favours *B*. *B* gives 80 years of good life to the baby, but *A* gives only 60 to the woman. The fact that in *A* the baby is replaced counts for nothing.

Many people will, I think, doubt that the PRP has got the right answer here. We can see what it is about this principle that brings it to its conclusion. It attaches great importance to a person's continuing identity, to the fact that he remains the same person throughout his life. In the example let us look forward 70 years. Seventy years on, a person will be alive and enjoying life. Depending on our choice between *A* and *B*, he will be either the baby or his replacement. To the PRP it matters crucially which he is. If he is the replacement, the good he is enjoying does not count in the moral calculation. But if he is the baby it counts, and so does all the rest of the good he enjoys throughout his life. Once a person has got a toe-hold on existence, by getting born or conceived or whatever else makes a person, then the whole of his future life gets weighed in the scales. The TP, we found, made light of the difference between contraception and the death of a baby. The PRP, on the other hand, values the difference at the whole of the good in a person's life. It values it, for instance, more than it values the life of a 20-year-old.

It is natural to think that continuing identity cannot be as important as all that. If a person 70 years on is enjoying life it does not seem it should make such a vast moral difference who that person is, the baby or his replacement. It seems wrong to treat the baby's good and the replacement's so differently, counting one fully and the other not at all. This thought could take us in either of two directions: we might decide to count the replacement's good to some extent, or the baby's less than fully. The former is a move back towards the TP, which counts equally all good whoever it belongs to. We shall presumably not want to go back as far as that, and count the replacement's good fully and equally with the baby's, because we have already found reasons for rejecting the TP. One reason, indeed, was precisely that it reckons life-saving and replacement equally valuable. Instead, we shall want to find some halfway house, where the replacement's good counts, but not as much as the baby's. Even this, though, would certainly mean that, in assessing the value of saving a person's life, we should, contrary to the usual practice, have to give some weight to the good of his descendants. And it conflicts with the intuition that drew us to the PRP in the first place: that creating a new person is no benefit to the person, so that the good he will enjoy should not be a factor in our moral calculations. If this intuition is right, we ought not to count the replacement's good at all, and I shall go no further along this path.

The other way to go is to combine this intuition about creation with the intuition that identity cannot be as important as the PRP makes it. In working

out how this can be done we shall get some guidance if we first look for some sound basis for the latter intuition. There is perhaps a basis to be found in a reductionist theory of personal identity propounded by Derek Parfit (1984, Part III). I am the same person as the person I shall become at some future date. According to Parfit, what makes this true is the fact that between him then and me now there are a number of particular relations. Among them are psychological relations such as memory and intention: this person will remember doing some of the things I do now, and will carry out some of the intentions I now have. And Parfit argues that what is really significant is not strictly the fact of identity itself, but the existence of these psychological relations that partly constitute identity. Identity is indirectly significant only because of these relations. But the further apart in time are I now and the person I shall become, the weaker will be the psychological relations between us. Consequently, the further apart we are in time, the less is the moral significance that should be attached to our identity.

I am not going to discuss the truth of this reductionism. But I am going to consider what happens if we combine it with the intuition about creation that underlies the PRP. In comparing two alternatives, the PRP counts only the good of people who exist in both. People who exist in only one are not counted because they are neither better nor worse off in either. In our example the PRP does not count the replacement's good. It does, however, count the baby's good, even the good he enjoys at times—70 years in the future, for instance—when he exists in only one state. But *at these times* he, like the replacement, is neither better nor worse off in either state. Where he differs from the replacement is that at these times he is still the same person as a person (the baby) who exists at *another* time (actually now) in both states. That is why the PRP treats his good at those times differently from the replacement's. But we are now supposing that this fact of intertemporal identity may have limited moral significance. If, to take an extreme, it had no moral significance at all,⁹ then no difference in treatment could be justified; given that we are not going to count the replacement's good, we ought not to count the baby's good either, at times when he exists in only one state. According to the view I have described, on the other hand, intertemporal identity has a moral significance that declines with increasing separation in time. If so, partial discounting of good may be appropriate. Good that comes to the baby at a time when he exists in only one state should be discounted, and the more remote is the time from a time when he exists in both states, the bigger the discount should be. This brings us to a new general principle.

The Restricted Principle with Discounting (RPD)

For each person who exists in alternative *A*, compute as follows a "discounted total" of his good in *A*. First add up, undiscounted, all the good he enjoys at times when he also exists in *B*. To this add the good he enjoys at other times, discounted to an appropriate extent determined by the distance between the time when he enjoys it and the nearest time when he exists in *B*. If he never exists in *B* his discounted total (in *A*) is nought. Now add up all the discounted totals for everyone who exists in *A*. Then calculate in a symmetrical way discounted totals for everyone who exists in *B*, and add them up. The better state is the one with the greater total.

This formulation of the RPD does not specify what is an "appropriate extent" to discount. That is determined by the importance of identity. If identity is actually all-important, then no discounting is appropriate, and the RPD reduces to the PRP. According to the reductionist view, on the other hand, a person's good should be discounted by a greater amount the further it comes from a time when the person exists in both alternatives. Exponential discounting might sometimes be a reasonable approximation.

In our baby example reductionism suggests more radical discounting. The psychological connections between a baby and later stages of his life are rather weak, much weaker than those between a 20-year-old and later stages of her life. So the baby's future good ought to be discounted much more than the 20-year-old's. Consequently the RPD is likely to favour saving the 20-year-old. This is quite a plausible response to the example.¹⁰ Several authors (e.g. Glover, 1977, Chapter 12; Singer, 1979, pp. 133-139) have argued that the death of a baby is a lesser evil than the death of an older person. Their reasons have to do with a baby's lack of autonomy and full personhood. The reason I have just given—that a baby is not closely integrated psychologically with later stages of his life—is a related one.

The RPD discounts a person's future good according to its distance from the nearest time when he is alive in both alternative states. This is the time when, if life-saving is in question, he will die if he is not saved.¹¹ The discounting is not related to the *present* (the moment when a decision is to be made): the present plays no part in the RPD. This is as it should be. In comparing the goodness of states of affairs, it cannot make any difference what the date happens to be when we compare them.¹² But it does mean there may be a conflict between the RPD and people's present preferences, even if the preferences are rational, because people may rationally form their present preferences by discounting to the present. This needs to be looked at carefully.

In forming preferences between states of affairs, people sometimes discount their future good. According to reductionism this may be rational.¹³ Good that comes to a person has a different sort of significance, to the person, from good that comes to other people. If the person does not care about other people, it is the total of his own good that determines his preferences. But when a person looks ahead to the good that will come to him many years in the future, it may seem to him that the difference between this future good and other people's good is less important to him now than the difference between his present good and other people's. His attitude to his future self may to some degree resemble his attitude to other people. And reductionism suggests that this may be rational; the fact that his distant future good is *his* may genuinely have diminishing significance the further in the future the good comes, because of the weakening of his psychological relations with his future self. Consequently, when a person separates out, from all the good in the world, the good to which he will now attach the special significance he attaches to his own good, he may reasonably not include his future good fully but instead may discount it. To him now his future good has only partly the status of his own good, and partly the status of other people's. His present preferences, then, will be determined by the total of his good discounted back to the present.

But this total discounted to the present is not the "discounted total" of the RPD. In the RPD the good a person enjoys at a time when he is alive in

both alternatives is not discounted at all, and the rest is discounted back to the time when (if life-saving is in question) his life is saved. This will not normally be the present. The preferences the RPD has on a person's behalf, as it were, may therefore differ from the person's own present preferences. But this again is as it should be. A person's future good does not fully register in his present preferences because, if the account I have given is right, his future good has for him partly the status of other people's good, which he does not care about. But when we assess the actual goodness of a state of affairs, other people's good counts as much as his. Just because his future good is like other people's good, that is no reason not to count it fully. To the extent that people discount their future good in forming their present preferences, moral judgments must go beyond present preferences.

What does all this have to say about the usual practice in valuing life? Well, the most usual practice—though I have not mentioned this before—incorporates some discounting. And the way it does so may in practice conform very roughly to what the RPD recommends. At least, this may be true if people form their preferences in the way I described above, by discounting in accordance with the diminishing importance of the identity between their present and future selves. To see why, think about some proposed project that will save a person's life ten years in the future. The resulting benefit will actually be spread out over the period from ten years ahead onwards, during which the person enjoys the good that death would otherwise have deprived him of. But in practice a cost-benefit analysis will locate the benefit at a single time ten years ahead, and assign it a value at that time. The value will be derived from separate studies of what people are willing to pay to escape a risk of dying. In all the studies I know, this means a risk of dying very soon.¹⁴

If people are rational according to the above account, what they are willing to pay should be determined by the total of their future good if they survive, discounted to the time they are making the decision. And in the studies this is also the time when, if they do not pay, they are in danger of dying. So the value of saving a person's life that emerges from these studies should in principle (abstracting from all problems apart from discounting) be the total of the person's future good discounted to the time when his life is saved. In our proposed project this time is ten years ahead. The value, located at that time, of saving a person's life at that time will be the total of his future good discounted to that time. Now, a cost-benefit analyst valuing the project ten years earlier may or may not discount this value again back to the time of valuation. If he does not, then the result, so far as discounting is concerned, will fit the RPD:¹⁵ if a person's life is saved his later good will be discounted back to the time when it is saved, but it will not be discounted further.

The upshot of this section, then, is that Restricted Principles, and especially the RPD, seem to offer some support for the usual practice in valuing life.

Unfortunately, however, this support seems itself to be unsupported. Restricted Principles seem to be quite untenable. Derek Parfit (1984, pp. 351–379) presents objections to them that are apparently quite conclusive. I do not have the space to present these reasons here, but I shall briefly mention one. Suppose we have to choose between two alternative states that both contain the same number of people. Suppose that the people who exist in both

states are equally well off in either. But suppose there are some people who exist only in state *A*, and others who exist only in *B*, and that the people who exist only in *A* are better off in *A* than the people who exist only in *B* are in *B*. (Imagine that resources could be used to save one or other of two people. Both will have equally good lives if saved and both will have children. But the children of one will have better lives than the children of the other.) Restricted Principles compare states only by how good they are for people who exist in both. So they would be indifferent between *A* and *B*. But actually, *A* seems obviously better.

Apart from Restricted Principles, I believe no other principle has been proposed for comparing alternatives with different populations that is consistent with the usual practice in valuing life. If Restricted Principles have to go, then so must the usual practice. What can be put in its place? Only one other principle has played any part in the literature on valuing life, and that is the one I turn to next.

III. THE AVERAGE PRINCIPLE (AP)

Calculate the average goodness of lives in alternative *A*, averaging over everyone who exists in *A*. Do the same for *B*, averaging over everyone who exists in *B*. The better alternative is the one with the bigger average.¹⁶

The AP certainly does not support the usual practice in valuing life. It says that adding a person to the population is a good thing if and only if his life will be better than average. If a person's life is saved and he then has children, the children count as a benefit of saving his life if their lives are better than average, and a cost if worse. On the other hand, apart from the question of children, prolonging a person's life is unambiguously a benefit according to the AP, provided only that his life is not a bad one. Prolonging a life increases the amount of good it contains as a whole, and so helps to raise the average goodness of lives.

General criticisms of the AP are to be found in McMahan (1981), Parfit (1984, pp. 420-472), Sumner (1978) and in many other places. I shall not try to add to them. One of the difficulties about discussing the AP is that it has many different versions besides the one above, and I cannot deal with all of them. But although the AP is common elsewhere in economics, in the literature on valuing life I believe it has appeared only once, in an important paper by W. B. Arthur (1981).¹⁷ Since Arthur's paper is also the first to take account of the effect of life-saving on population, I shall devote this section to a discussion of it. As I shall explain, Arthur seems to have in mind a different version of the AP from the one above, but that is a convenient one to start from all the same.

Arthur works with a simplified model. He supposes that the population and the economy are growing proportionally, and that average living standards do not change over time. He sets out to value a life-saving programme that is repeated for each generation. He is thinking of something like a cure for heart disease, which will prolong a constant fraction of the lives in each generation. He compares a state of proportional growth that has the cure with one that does not. If the cure saves some people who later have children, the former

state will have a faster-growing population and hence a faster-growing economy. Lives in the former will on average be longer. But on the other hand, each year of life will be less good because some resources will be used up on the life-saving programme.¹⁸ If these effects exactly cancel, so that an average life contains just as much good in either state, then Arthur would take the cure to be just on the edge of being worthwhile.

This application of the AP may seem very plausible. But that is only because of Arthur's assumption (which he intended as a mere simplification) that in proportional growth all generations are equally well off. If instead lives got progressively better or worse as time passed, then the average would always be swamped by the indefinitely large number of people who will live in the distant future. Suppose, for instance, that there is a steady increase in productivity (independent of population) so that future lives are better than present ones. Suppose some life-saving programme is available but only at a heavy cost. Suppose the cost is so heavy that adopting the programme will make generation worse off on average than they would have been without it, even taking account of their increased life expectancy. Suppose that each person who will exist whether or not the programme is adopted will individually be made worse off by it, even those whose lives it prolongs. Despite all this, the AP favours this programme. It will undoubtedly make the average life better¹⁹ (assuming that some of the people saved in each generation later have children so that population grows faster) because it will increase indefinitely the number of people living in the indefinite future when life is better anyhow.

This implication of the AP is hard to accept. The programme makes nobody better off and some people worse off. The AP is willing to sacrifice the good of existing people and people who will exist anyway for the sake of people who, but for the programme, would never have existed at all. This is precisely what we found intuitively unattractive about the Total Principle.²⁰ Indeed, in this respect the AP is even less attractive than the TP; what it values in the new people brought into existence is not really even the good they enjoy, but the mere fact that their good is above the average. The AP would *also* be willing to sacrifice the good of existing people and people who will exist anyway for the sake of *preventing* the existence of people whose lives, though good, would be less good than the average.

Arthur himself would not, I think, favour the life-saving programme I have described. He would not apply the AP in that way. In comparing the average goodness of lives in the two alternatives, he would not take the average over everyone who ever lives. Instead, he would make the comparison generation by generation: for each generation he would compare the average life of that generation in one alternative with its average life in the other. There is a new version of the AP implicit in this. It might be hard to formulate in general,²¹ but in our example it is plain what this version would say: the life-saving programme makes each generation worse off on average, so the programme ought not to be adopted. This conclusion is a plausible one.

But actually, the generation-by-generation version of the AP is certainly unacceptable because it takes no account of the size of the generations. Another example will show this. Suppose life is going to be good for a century or so and then for some reason bad for another century. Suppose we have a choice of two policies. Both bring into existence the same number of people altogether.

But one puts many more of them into the lean years, the other many more into the fat ones. Suppose, though, that the former policy is slightly cheaper to run, so that it makes life in each century slightly better than the other one does. The generation-by-generation version of the AP will prefer the former policy because it makes each generation on average better off. But plainly this is really the worse policy because it puts more people into the bad years; it makes the worse-off generations larger.

My slight modification of Arthur's example, then, has left the AP in a bind. The version of it I stated at the start of this section delivers an unacceptable conclusion, and the version whose conclusion is more acceptable is itself unacceptable on other grounds. Arthur is unique in having brought the AP to bear on valuing life, but it seems that his work gets its plausibility only from his very special assumption that all generations are equally well off. Quite apart from all this, there are very good reasons given by the authors mentioned on p. 289 for rejecting the AP anyway. I do not think it can help us with valuing life.

IV. CONCLUSIONS

Saving life very often changes the timeless population of the world. So before we can value the saving of life we need to be able to value changes in population. Indeed, Section II above shows that until we can do this we cannot even be secure in valuing acts of life-saving that happen not to alter population.

I have examined several principles for valuing changes in population. Some—Restricted Principles—seemed to support what is at present the usual practice in valuing life: to count as a benefit of life-saving the good added to the life of the person who is saved, but not the good of people who are added to the population as a result. But it turns out in the end that none of the principles I considered seems acceptable. I gave some grounds for this conclusion, but for a full account I referred elsewhere, particularly to Parfit's (1981) *Reasons and Persons*. Other principles have been proposed besides the ones I mentioned, but, as Parfit also shows, they all seem unacceptable too. Population theory is at present at an *impasse*.

I think we have to conclude, first, that the usual practice has no sound basis and, second, that we have no soundly based alternative to put in its place. I think at the moment we are not in a position to set an economic value on life.

No doubt many people will be impatient with the arguments I have been using. Decisions have to be made, they will say, so we simply have to have some value for life. And they will argue that it is manifestly inefficient not to have a uniform value established for all of the government's decision-making. If two branches of government use different values, then the resources used on life-saving would save more lives if they were redistributed between the branches.

It is one thing to say there is a manifest inefficiency; it is another to know how to do better. One point at least is plain: establishing a uniform value for life would *not* be the best thing to do. Not all lives can be equally valuable; age at least must make a difference. In distributing life-saving resources, maximizing the number of lives saved is not the right objective. But the message

of this paper is that we do not know what the right objective is. Should it be to maximize the good that is added to the world, or the good that is added to the lives of existing people, or what? How do we weigh babies against adults, or older people against young people who are likely to have children? If there are inefficiencies at the moment, these questions will need answering before we can get rid of them.

If we fix no definite economic value on life, the decisions will still get made as they always have. Like many other hard decisions, they have to be made without the guidance of clear criteria. If they are to be made well, what we most need to improve is the process by which they are made. We need sensitive and humanitarian decision-makers, who will face up to the full difficulty of life-and-death decisions. But putting a money-value on life helps to make the decisions seem mechanical and easy. We do not want our rulers to be sheltered by their experts from a full appreciation of their responsibilities.

ACKNOWLEDGMENTS

Some years ago I received comments from Richard Lecomber on some writings of mine that have since developed into this paper, and he spent some time discussing them with me. As I expected he would, he showed me my ideas in a new light. I particularly want to record this debt because, very sadly, Richard has recently died.

In writing this paper I have also greatly benefited from discussions with Derek Parfit on various points. I first came to understand the importance of the connection between population theory and valuing life as a result of reading two papers, one by Bruce Chapman (1980) and one by J.A. Mirrlees (1982). I have received very helpful comments from Brian Arthur, Tony Brewer, Geoff Cupit and this journal's referee. The paper was written while I was a Visiting Fellow at All Souls College, Oxford, supported by a grant from the Social Science Research Council. I thank the College for its generous hospitality, and the SSRC for its support.

NOTES

¹ x is given by $U(L, B - x) = pU(D, B) + (1 - p)U(L, B)$, where U is the government's decision-theoretic utility function, B is the budget it has for life-saving and other projects, and L and D stand for whether the person is alive or dead. To save him from certain death it is worth spending y where $U(L, B - y) = U(D, B)$. Provided x and y are small enough for this spending to have no significant effect on $\partial U / \partial B$ (so the government has no significant risk aversion over these amounts of money, which is certainly a good assumption), U can be approximated by a first-order Taylor expansion. A little algebra shows that $y = x/p$.

² In practice, very few actions will simply add people to the population. Suppose someone's life is saved and he then has children. The children will probably marry people who would otherwise have married someone else, so in the next generation some children will be added to the population but also some children will not be born who otherwise would have been. And so on. Since the literature on valuing life ignores all these effects, I am not sure what attitude most economists would take to them. So I shall leave them aside, and in examples concentrate on people who are simply added.

³ The most thorough and authoritative treatment of this subject is in Derek Parfit's *Reasons and Persons* (1984, Part IV), which also gives references to the extensive literature.

⁴ Henry Sidgwick (1907, pp. 414-416) gave the TP his support, having been perhaps the first to distinguish it from other versions of utilitarianism. As I have stated it, it takes no account of the distribution of good between people. If equality is valuable the TP may be adjusted accordingly (see Dasgupta, 1983). But since equality is not central to our problem I shall ignore it throughout this paper.

⁵ Compare Jonathan Glover (1977, pp. 162-164). Glover seems willing to accept that side-effects may be all that makes the moral difference between contraception and infanticide. I shall return to the question of the value of a baby's life in Section II, particularly on p. 287.

⁶ I say "does not seem to" because there has been disagreement on this point. Jefferson McMahan (1981) and Derek Parfit (1984, pp. 487-490) suggest that creating a person might

defensibly be said to benefit him. I am not going to argue about this, because the argument would merely be about the meaning of "benefit"; McMahan suggests it might be given a "noncomparative" meaning. In stating the principles in Section II I have deliberately avoided the word. The substantive question is whether the fact that a person will have a good life if he lives is a reason for bringing it about that he does live. It is at least plausible to think not, and Section II follows up the implications of that thought. If, on the other hand, we think it *is* a reason, then it is not important to decide whether that is because creating a person counts as benefiting him or because of something else. Either way, we shall want to stick to the TP or some similar principle.

⁷ This principle, or something like it, was first proposed by Jan Narveson (1967).

⁸ There has been some debate around this point, but I think this argument is sound. Notice that it has nothing to do with the impossibility of referring to or identifying a person in a world where he does not exist. It has sometimes been supposed that this is the problem, so that R. M. Hare (1975, p. 220), for one, seems to believe that the argument can be answered by supplying a way of making the reference. But that is not the problem at all. There is usually no difficulty about referring to a person in a possible world where he does not exist; we establish the reference in the actual world where he does exist. When we imagine a world in which Mrs Thatcher never existed, there is no difficulty about knowing *who* never existed because we know who Thatcher is from our world. But in that world her life would not have been good or bad to any degree.

The reasons for my conclusion is the same as the reason why Queen Victoria's reign was neither shorter nor longer than it would have been if Queen Victoria had never come to the throne. A non-existent reign is not like a very short reign. It is true that sometimes a non-existent life does seem like a very short life. Indeed, it may be hard to draw a line between them. If the zygote that actually became Thatcher had never been implanted, would Thatcher have had a life or not? This question is debatable, and indeed it may have no definite answer. But the right conclusion to draw is not that I was wrong to say that *if* Thatcher had not existed her life would not have been good to any degree: the right conclusion is that the truth of the *if* clause may sometimes be difficult to establish and even indefinite. It might also be a reasonable inference that the PRP is dubious because it makes so much depend on whether or not a person exists in both the alternatives being compared, a question that might sometimes have a doubtful or indefinite answer. Perhaps it is wrong to attach great moral significance to something so unclear (cf. Roupas, 1978, p. 167). I shall be showing in this section, however, that it is possible to reduce the weight put on this question without abandoning the spirit of restricted principles.

⁹ Some authors have claimed that the reductionist theory of identity actually implies that identity can have no moral significance. The sort of relations—memory and so on—that reductionism takes to constitute identity are not, these authors think, the sort of relations that can give identity moral significance. Several opponents of reductionism have used this claim against it (e.g. Butler, 1975, p. 102; Swinburne, 1973–1974, p. 246), and some friends of reductionism have given it a qualified endorsement (Parfit, 1984, pp. 307–312; Perry, 1976). If identity has no moral significance, then the RPD (below) will compare alternatives only on the basis of good enjoyed by people at times when they are alive in both alternatives; it will discount totally all other good. Saving a life it will take to confer no benefit apart from side-effects. The Epicureans would have embraced this conclusion. Some economists have embraced it too. R. F. F. Dawson (1967, p. 8), for instance, says: "Losses due to a death are measured in terms of losses to those remaining alive: his relatives and the community in general. In other words the costs are considered *ex post*, that is based on the costs to the population after the accident." (To be fair, I must mention that Dawson later changed his mind about this (see Dawson, 1971).)

¹⁰ And in this way the RPD can overcome the objection to the PRP mentioned in note 8.

¹¹ A person's life may end at different times in different states. It is also possible for it to begin at different times. But the facts of human reproduction confine this latter difference to such a small range that I shall ignore it.

¹² The date cannot make any difference to which state of affairs is better, but it might, even so, make a difference to which is the right one for the government to bring about. We do not all necessarily have a duty always to bring about the best state of affairs. Perhaps, for instance, parents have a special duty to promote the good of their children, even if it brings about a state of affairs that is on balance not the best. Similarly, a government might have a special duty to promote the good (or even to satisfy the present desires) of presently living people. I find this implausible; I suspect a government's duty is to be impartial between living people and people yet to be born. But I shall circumvent this political question by talking only about the goodness and badness of states of affairs, and not about the government's duty.

¹³ See Parfit (1984, p. 317). There are, of course, many well-known opinions to the contrary (e.g. Sidgwick, 1907, p. 381; Pigou, 1932, p. 25).

¹⁴ For a survey see Blomquist (1982).

¹⁵ So the RPD opposes this second discounting, discounting the value of future life-saving. In this I agree with the RPD. I do not think a life saved at one date can differ in value from a similar life saved at another date. Some of the arguments in this section help to explain why. But the subject has ramifications I cannot pursue now. For a contrary opinion see Arrow (1983).

¹⁶ The origins of the AP are obscure, but it seems to have been invented by economists. See the brief survey in Sumner (1978, p. 107).

¹⁷ Arthur actually defines his criterion for evaluation as "the representative person's expected lifetime welfare" (1981, pp. 56-57; my emphasis). But the representative person turns out to be the average person. As I shall explain, if different generations were to differ in their wellbeing, Arthur might want to pick many representative (average) people, one from each generation.

¹⁸ The cost in resources is actually quite complicated. Besides the direct cost there is the cost of supporting the old people whose lives have been prolonged beyond retirement, and the increased investment required to make the capital stock keep up with the population. But these complications need not concern us.

¹⁹ If the increase in productivity goes on for ever, the average will actually be infinite. So technically we shall have to calculate as follows. Take the average goodness of lives for the next T years if the programme is adopted and compare it with the average for those T years if it is not. Then make T progressively larger. We can be sure that, provided T is large enough, the former average will be bigger.

²⁰ And the example is just as much an objection to the TP as to the AP.

²¹ The theory of optimal growth sometimes takes, as the objective to be maximized, the integral over time of people's average utility at a time. This is a rough approximation to this version of the AP.

REFERENCES

- ARROW, K. J. (1983). *The Tanner Lectures on Human Values*. Oxford (to be published).
- ARTHUR, W. B. (1981). The economics of risks to life. *American Economic Review*, 71, 54-64.
- BLOMQUIST, GLENN (1982). Estimating the value of life and safety: Recent developments. In *The Value of Life and Safety* (M. W. Jones-Lee, ed.) Amsterdam: North-Holland, pp. 27-40.
- BROOME, JOHN (1978). Trying to value a life. *Journal of Public Economics*, 9, 91-100.
- (1982). Uncertainty in welfare economics and the value of life. In *The Value of Life and Safety* (M. W. Jones-Lee, ed.) Amsterdam: North-Holland, pp. 201-216.
- BUTLER, JOSEPH (1975). Of personal identity. In his *The Analogy of Religion* (1736); reprinted in *Personal Identity* (John Perry, ed.). Los Angeles: UCLA Press, pp. 99-106.
- CHAPMAN, BRUCE (1980). Coase, cost and causation. Law and Economics Workshop Series, no. WSII-15. Faculty of Law, University of Toronto.
- DASGUPTA, PARTHA (1983). Lives and welfare. *Population Growth and Economic Development* (D. Gale Johnson and Ronald Lee, eds).
- DAWSON, R. F. F. (1967). *Cost of Road Accidents in Great Britain*. Road Research Laboratory Report LR79. London: Ministry of Transport.
- (1971). *Current Costs of Road Accidents in Great Britain*. Road Research Laboratory Report LR396. London: Department of the Environment.
- GLOVER, JONATHAN (1977). *Causing Death and Saving Lives*. Harmondsworth: Penguin.
- HARE, R. M. (1975). Abortion and the golden rule. *Philosophy and Public Affairs*, 4, 201-222.
- JONES-LEE, M. W. (1976). *The Value of Life: An Economic Analysis*. Oxford: Martin Robertson.
- MCMAHAN, JEFFERSON (1981). Problems of population theory. *Ethics*, 92, 96-127.
- MIRRLEES, J. A. (1982). The economic uses of utilitarianism. In *Utilitarianism and Beyond* (Amartya Sen and Bernard Williams, eds). Cambridge University Press.
- MISHAN, E. J. (1971). Evaluation of life and limb: A theoretical approach. *Journal of Political Economy*, 79, 687-705.
- NARVESON, JAN (1967). Utilitarianism and new generations. *Mind*, 76, 62-72.
- (1973). Moral problems of population. *The Monist*, 57, reprinted in *Ethics and Population* (M. Bayles, ed.) Shenkman, 1976, pp. 59-80 (page references are to the reprinted version).
- PARFIT, DEREK (1984). *Reasons and Persons*. Oxford University Press.
- PERRY, JOHN (1976). The importance of being identical. In *The Identities of Persons* (Amelie Oksenberg Rorty, ed.) Berkeley, Calif.: University of California Press.
- PIGOU, A. C. (1932). *The Economics of Welfare*. London: Macmillan.
- ROUPAS, T. G. (1978). The value of life. *Philosophy and Public Affairs*, 7, 154-183.
- SIDGWICK, HENRY (1907). *The Methods of Ethics*. London: Macmillan.
- SINGER, PETER (1979). *Practical Ethics*. Cambridge University Press.
- SUMNER, L. W. (1978). Classical Utilitarianism and the population optimum. In *Obligations to Future Generations* (R. I. Sikora and Brian Barry, eds.) Philadelphia: Temple University Press.
- SWINBURNE, R. G. (1973-1974). Personal identity. *Proceedings of the Aristotelian Society*, 74, 231-247.