44

The Value of Life

John Harris

Suppose that only one place is available on a renal dialysis programme or that only one bed is vacant in a vital transplantation unit or that resuscitation could be given in the time and with the resources available to only one patient. Suppose further that of the two patients requiring any of these resources, one is a 70-year-old widower, friendless and living alone, and the other a 40-year-old mother of three young children with a husband and a career.

Or suppose that following a major disaster medical resources were available to save the lives of only half those for whom medical care was vital for life. Or, less dramatically, suppose that in the next two years, only half of 200 patients waiting for surgery that will alleviate severe discomfort can be accommodated in the only available hospital. Suppose further and finally that all candidates stand an equal chance of maximum benefit from any of the available treatments. Whom should we treat and what justifies our decision?

Many will think that in the first case preference should be given to the young mother rather than the old friendless widower, that this is obviously the right choice. There might be a number of grounds for such a decision. Two of these grounds have to do with age. One indicates a preference for the young on the grounds that they have a greater expectation of life if they are restored to health. The other favours the young simply because their life is likely to be fuller and hence more valuable than that of the older person. Another consideration to which many will want to give some weight is that of the number of people dependent on or even caring about a potential victim. It is sometimes also considered relevant to give weight to the patient's probable usefulness to the community or even their moral character before a final decision is made. And of course these considerations may be taken together in various combinations.

In the case of a major disaster related problems arise. If say a policy of triage[1] has identified the only group of victims to be treated, those for whom medical intervention will make the difference between life and death, but there are still not enough resources to help all such persons, then, again, many will hold that the right thing to do is help the young or those with dependants and so on first.

Those who believe that they ought to select the patient or patients to be saved on any of the above criteria will believe that they must show preference for some types or conditions of person over others. Another available strategy is of course to decline to choose between people in any way that involves preferring one patient, or one sort of person, to another. Perhaps the easiest way of declining to show such a preference is to toss a coin or draw lots to decide who shall be helped. I want to consider what might count as a good reason for preferring to help

Original publication details: John Harris, "The Value of Life," pp. 87–102 from *The Value of Life*, London: Routledge, 1985.

Bioethics: An Anthology, Third Edition. Edited by Helga Kuhse, Udo Schüklenk, and Peter Singer.

some patients rather than others where all cannot be helped and also whether our intuitive preference for saving the younger and more useful members of society can be sustained.

I The Moral Significance of Age

Many, perhaps most, people feel that, in cases like the one with which we began, there is some moral reason to save the 40-year-old mother rather than the 70-year-old widower. A smaller, but perhaps growing, group of people would see this as a sort of 'ageist' prejudice which, in a number of important areas of resource allocation and care, involves giving the old a much worse deal than the younger members of society. This is an exceptionally difficult issue to resolve. A number of the ways of thinking about the issue of the moral relevance of age yield opposed conclusions or seem to tug in opposite directions.

I want first to look at an argument which denies that we should prefer the young mother in our opening example. It is an anti-ageist argument so that is what I will call it, but it is not perhaps the usual sort of argument used to defend the rights of the old.

The anti-ageist argument

All of us who wish to go on living have something that each of us values equally although for each it is different in character, for some a much richer prize than for others, and we none of us know its true extent. This thing is of course 'the rest of our lives'. So long as we do not know the date of our deaths then for each of us the 'rest of our lives' is of indefinite duration. Whether we are 17 or 70, in perfect health or suffering from a terminal disease, we each have the rest of our lives to lead. So long as we each fervently wish to live out the rest of our lives, however long that turns out to be, then if we do not deserve to die, we each suffer the same injustice if our wishes are deliberately frustrated and we are cut off prematurely. Indeed there may well be a double injustice in deciding that those whose life expectation is short should not benefit from rescue or resuscitation. Suppose I am told today that I have terminal cancer with only approximately six months or so to live, but I want to live until I die, or at least until I decide that life is no longer worth living. Suppose I am then involved in an accident and, because my condition is known to my potential rescuers and there are not enough resources to treat all who could immediately be saved, I am marked among those who will not be helped. I am then the victim of a double tragedy and a double injustice. I am stricken first by cancer and the knowledge that I have only a short time to live and I'm then stricken again when I'm told that because of my first tragedy a second and more immediate one is to be visited upon me. Because I have once been unlucky I'm now no longer worth saving.

The point is a simple but powerful one. However short or long my life will be, so long as I want to go on living it then I suffer a terrible injustice when that life is prematurely cut short. Imagine a group of people all of an age, say a class of students all in their mid-twenties. If fire trapped all in the lecture theatre and only twenty could be rescued in time, should the rescuers shout 'youngest first!'? Suppose they had time to debate the question or had been debating it 'academically' before the fire? It would surely seem invidious to deny some what all value so dearly merely because of an accident of birth? It might be argued that age here provides no criterion precisely because although the lifespans of such a group might be expected to vary widely, there would be no way of knowing who was most likely to live longest. But suppose a reliable astrologer could make very realistic estimates or, what amounts to the same thing, suppose the age range of the students to be much greater, say 17 to 55. Does not the invidiousness of selecting by birth-date remain? Should a 17-year-old be saved before a 29-year-old or she before the 45-year-old and should the 55-year-old clearly be the last to be saved or the first to be sacrificed?

Our normal intuitions would share this sense of the invidiousness of choosing between our imaginary students by reason of their respective ages, but would start to want to make age relevant at some extremes, say if there were a 2-day-old baby and a 90-year-old grandmother. We will be returning to discuss a possible basis for this intuition in a moment. However, it is important to be clear that the anti-ageist argument denies the relevance of age or life expectancy as a criterion absolutely. It argues that even if I know for

certain that I have only a little space to live, that space, however short, may be very precious to me. Precious, precisely because it is all the time I have left, and just as precious to me on that account as all the time you have left is precious to you, however much those two time spans differ in length. So that where we both want, equally strongly, to go on living, then we each suffer the same injustice[2] when our lives are cut short or are cut further short.[3]

It might seem that someone who would insist on living out the last few months of his life when by 'going quietly' someone else might have the chance to live for a much longer time would be a very selfish person. But this would be true only if the anti-ageist argument is false. It will be true only if it is not plausible to claim that living out the rest of one's life could be equally valuable to the individual whose life it is irrespective of the amount of unelapsed time that is left. And this is of course precisely the usual situation when individuals do not normally have anything but the haziest of ideas as to how long it is that they might have left.

I think the anti-ageist argument has much plausibility. It locates the wrongness of ending an individual's life in the evil of thwarting that person's desire to go on living and argues that it is profoundly unjust to frustrate that desire merely because some of those who have exactly the same desire, held no more strongly, also have a longer life expectancy than the others. However, there are a number of arguments that pull in the opposite direction and these we must now consider.

The fair innings argument

One problem with the anti-ageist argument is our feeling that there is something unfair about a person who has lived a long and happy life hanging on grimly at the end, while someone who has not been so fortunate suffers a related double misfortune, of losing out in a lottery in which his life happened to be in the balance with that of the grim octogenarian. It might be argued that we could accept the part of the anti-ageist argument which focuses on the equal value of unelapsed time, if this could be tempered in some way. How can it be just that someone who has already had more than her fair share of life and its delights should be preferred or even given an equal chance of continued survival with the young person who has not been so favoured? One strategy that seems to take account of our feeling that there is something wrong with taking steps to prolong the lives of the very old at the expense of those much younger is the fair innings argument.

The fair innings argument takes the view that there is some span of years that we consider a reasonable life, a fair innings. Let's say that a fair share of life is the traditional three score and ten, seventy years. Anyone who does not reach 70 suffers, on this view, the injustice of being cut off in their prime. They have missed out on a reasonable share of life; they have been short-changed. Those, however, who do make 70 suffer no such injustice, they have not lost out but rather must consider any additional years a sort of bonus beyond that which could reasonably be hoped for. The fair innings argument requires that everyone be given an equal chance to have a fair innings, to reach the appropriate threshold but, having reached it, they have received their entitlement. The rest of their life is the sort of bonus which may be cancelled when this is necessary to help others reach the threshold.

The attraction of the fair innings argument is that it preserves and incorporates many of the features that made the anti-ageist argument plausible, but allows us to preserve our feeling that the old who have had a good run for their money should not be endlessly propped up at the expense of those who have not had the same chance. We can preserve the conclusion of the anti-ageist argument, that so long as life is equally valued by the person whose life it is, it should be given an equal chance of preservation, and we can go on taking this view until the people in question have reached a fair innings.

There is, however, an important difficulty with the fair innings argument. It is that the very arguments which support the setting of the threshold at an age which might plausibly be considered to be a reasonable lifespan equally support the setting of the threshold at any age at all, so long as an argument from fairness can be used to support so doing. Suppose that there is only one place available on the dialysis programme and two patients are in competition for it. One is 30 and the other 40 years of age. The fair innings argument requires that neither be preferred on the

grounds of age since both are below the threshold and are entitled to an equal chance of reaching it. If there is no other reason to choose between them we should do something like toss a coin. However, the 30-year-old can argue that the considerations which support the fair innings argument require that she be given the place. After all, what's fair about the fair innings argument is precisely that each individual should have an equal chance of enjoying the benefits of a reasonable lifespan. The younger patient can argue that, from where she's standing, the age of 40 looks much more reasonable a span than that of 30, and that she should be given the chance to benefit from those ten extra years.

This argument generalized becomes a reason for always preferring to save younger rather than older people, whatever the age difference, and makes the original anti-ageist argument begin to look again the more attractive line to take. For the younger person can always argue that the older has had a fairer innings, and should now give way. It is difficult to stop whatever span is taken to be a fair innings collapsing towards zero under pressure from those younger candidates who see their innings as less fair than that of those with a larger share.

But perhaps this objection to the fair innings argument is mistaken? If seventy years is a fair innings it does not follow that the nearer a span of life approaches seventy years, the fairer an innings it is. This may be revealed by considering a different sort of threshold. Suppose that most people can run a mile in seven minutes, and that two people are given the opportunity to show that they can run a mile in that time. They both expect to be given seven minutes. However, if one is in fact given only three minutes and the other only four, it's not true that the latter is given a fairer running time: for people with average abilities four minutes is no more realistic a time in which to run a mile than is three. Four minutes is neither a fair threshold in itself, nor a fairer one than three minutes would be.

Nor does the argument that establishes seven minutes as an appropriate threshold lend itself to variation downwards. For that argument just is that seven is the number of minutes that an average adult takes to run a mile. Why then is it different for lifespans? If three score and ten is the number of years available to most people for getting what life has to offer, and is also the number of years people can reasonably expect to have, then it is a misfortune to be allowed anything less however much less one is allowed, if nothing less than the full span normally suffices for getting what can be got out of life. It's true that the 40-year-old gets more time than the 30-year-old, but the frame of reference is not time only, but time normally required for a full life.[4]

This objection has some force, but its failure to be a good analogy reveals that two sorts of considerations go to make an innings fair. For while living a full or complete life, just in the sense of experiencing all the ages of man,[5] is one mark of a fair innings, there is also value in living through as many ages as possible. Just as completing the mile is one value, it is not the only one. Runners in the race of life also value ground covered, and generally judge success in terms of distance run.

What the fair innings argument needs to do is to capture and express in a workable form the truth that while it is always a *misfortune* to die when one wants to go on living, it is not a *tragedy* to die in old age; but it is, on the other hand, both a tragedy and a misfortune to be cut off prematurely. Of course ideas like 'old age' and 'premature death' are inescapably vague, and may vary from society to society, and over time as techniques for postponing death improve. We must also remember that, while it may be invidious to choose between a 30- and a 40-year-old on the grounds that one has had a fairer innings than the other, it may not be invidious to choose between the 30- and the 65-year-old on those grounds.

If we remember, too, that it will remain wrong to end the life of someone who wants to live or to fail to save them, and that the fair innings argument will only operate as a principle of selection where we are forced to choose between lives, then something workable might well be salvaged.

While 'old age' is irredeemably vague, we can tell the old from the young, and even the old from the middle-aged, so that, without attempting precise formulation, a reasonable form of the fair innings argument might hold; and might hold that people who had achieved old age or who were closely approaching it would not have their lives further prolonged when this could only be achieved at the

cost of the lives of those who were not nearing old age. These categories could be left vague, the idea being that it would be morally defensible to prefer to save the lives of those who 'still had their lives before them' rather than those who had 'already lived full lives'. The criterion to be employed in each case would simply be what reasonable people would say about whether someone had had a fair innings. Where reasonable people would be in no doubt that a particular individual was nearing old age *and* that that person's life could only be further prolonged at the expense of the life of someone that no reasonable person would classify as nearing old age, then the fair innings argument would apply, and it would be justifiable to save the younger candidate.

In cases where reasonable people differed or it seemed likely that they would differ as to whether people fell into one category or the other, then the anti-ageist argument would apply and the inescapable choice would have to be made arbitrarily.

But again it must be emphasized that the fair innings argument would only operate as a counsel of despair, when it was clearly impossible to postpone the deaths of all those who wanted to go on living. In all other circumstances the anti-ageist argument would apply.

So far so good. There are, however, further problems in the path of the anti-ageist argument and some of them are also problems for the fair innings argument.

Numbers of lives and numbers of years

One immediate problem is that, although living as long as possible, however long that turns out to be, will normally be very important to each individual, it seems a bad basis for planning health care or justifying the distribution of resources.

Suppose a particular disease, cancer, kills 120,000 people a year. Suppose further that a drug is developed which would prolong the lives of all cancer victims by one month but no more. Would it be worth putting such a drug into production? What if, for the same cost, a different drug would give ten years' complete remission, but would only operate on a form of cancer that affects 1,000 people? If we cannot afford both, which should we invest in? Or what if there is only one place on a renal dialysis programme and two patients who could benefit, but one will die immediately without dialysis but in six months in any event. The other will also die immediately without dialysis but with such help will survive for ten years. Although each wants the extra span of life equally badly, many would think that we ought to save the one with the longer life expectancy, that she is the 'better bet'.

All of these cases are an embarrassment for the anti-ageist argument, for our reaction to them implies that we do value extra years more. But how much more?

Extra life-time versus extra lives

> If we choose to save one person for a predicted span of sixty years, rather than saving five people each for a predicted span of ten years, we have gained ten extra life years at the cost of overriding the desires of four extra people.[6]

So far we have looked at the issue of whether we should count length of life or desire to live as the most important factor when deciding which of two people should be saved. If all things are equal, there can be no reason to prefer one to the other and so we should choose in a way that does not display preference, by lot, for example. The question that seems so difficult is what, if any, difference should length of life make to such choices?

The anti-ageist argument says that it should make no difference, but the cases we have just been examining seem to pull the other way. And if we are persuaded by such cases this seems to imply that we do think length of life or life expectancy gives additional value to lives and so constitutes a factor which must be given some weight. One consequence of this is that we should think it more important to save one 10-year-old rather than five 60-year-olds (if we take 70 as an arbitrary maximum).[7] Equally, it would be better to save one 20-year-old rather than two 50-year-old people, for we would again save ten life years by so doing. Or one 15-year-old rather than two 45-year-olds (a saving of five life years) and so on.

It is just at this point that the anti-ageist argument seems to require resuscitation, for there is surely something invidious about sacrificing two 45-year-olds to one 15-year-old. To take the 'life years' view

seems to discount entirely the desires and hopes and life plans of people in middle age, whenever an importunate youngster can place herself in the balance against them. But we do not normally think it better to save a 15-year-old rather than a 45-year-old when we cannot save both, so why should we think it better to save a 15-year-old rather than *two* 45-year olds?

For those who do favour saving one 15-year-old rather than two 45-year-olds, there is another difficulty. The life-time view seems to commit us to favouring total life-time saved rather than total number of people saved, with bizarre consequences. Suppose I could prolong the lives of 121,000 people for one month? This would yield a saving of 121,000 life months. Alternatively I could develop a drug which would give ten more years of life to 1,000 people. This would yield a saving of 120,000 life months. Thus, on the time-span view, we should choose to extend the lives of 121,000 people by one month rather than 1,000 people by ten years each. So, what started out by looking as though it constituted an objection to the anti-ageist argument actually supports it in some circumstances. For, while we should favour length of life, where numbers of lives balanced against one another are equal, we should favour numbers of lives where, summed together, they yield a greater contribution to the total amount of life-time saved.

Unfortunately the force of the comparison between extending the lives of 120,000 people for one month or 1,000 for ten years was to encourage us to think that life-time saved was more important than numbers of lives saved. Its support for this conclusion now seems less decisive. What it seems to indicate is a very complicated calculus in which allocation of resources would be dependent on the amount of life-time such allocation could save. It would also lead to some bizarre orderings of priority, and not necessarily to those envisaged by enthusiasts of such a scheme.

One such enthusiast, Dr Donald Gould, produced the following scenario:

> Calculations are based on the assumption that all who survive their first perilous year ought then to live on to the age of 70…In Denmark for example, there are 50,000 deaths a year, but only 20,000 among citizens in the 1–70 bracket. These are the ones that count. The annual number of life years lost in this group totals 264,000. Of these 80,000 are lost because of accidents and suicides, 40,000 because of coronary heart disease, and 20,000 are due to lung disease. On the basis of these figures, a large proportion of the 'health' budget ought to be spent on preventing accidents and suicides and a lesser…amount on attempting to prevent and cure heart and lung disease. Much less would be spent on cancer which is predominantly a disease of the latter half of life, and which therefore contributes relatively little to the total sum of life years lost…No money at all would be available for trying to prolong the life of a sick old man of 82.[8]

The first thing to note about Gould's scenario is, that while deaths before the age of 70 may be the only life years *considered to have been lost*,[9] it does not follow that there is no reason to attempt to *gain* life years by prolonging the lives of the over-seventies if that seems feasible. For example, if a reasonable prognosis is that the life of the 70-year-old could be prolonged for five years by some intervention, then that is still a gain of five life years. This can have important consequences, for it means that it would be quite wrong to write off all care for the over-seventies. Suppose a simple procedure would add one year to the lives of all septuagenarians. This would yield a huge gain in life years spread over a whole population. Suppose, as is perhaps likely, the number of septuagenarians in Denmark was over 260,000, then the number of life years saved by adding a further year to their lives would exceed the total to be gained by all the measures to prevent accidents, suicides, heart disease and so on. This would then become the chief priority for health care spending.

Gould starts his calculations after 'the first perilous year' but this cut-off point would require justification. We might well conclude, persuaded by his general line of argument, that neonatal and postnatal care would have the first priority for resources.

The life-time position then can support a wide variety of practices and may lead to a policy of achieving small gains in lifespan for large numbers of people rather than to the sorts of substantial gains for those individuals with most to lose that its supporters seem to have principally in mind.

Threshold of discrimination It is tempting to think that we might be able to get over some of the problems

of the life-time position by arguing that we can discount small gains in time as below the level of discrimination, in the sense that the benefit to the individual which accrued from living for a comparatively short period of extra time was nugatory. This might solve a few of the problems for the life-time position which arise from the necessity it imposes of favouring one group of people over another, wherever and whenever they are sufficiently numerous that the total life-time saved by rescuing them, even for a negligibly short period, exceeds that which might be saved by rescuing another smaller group who would live longer individually, but shorter collectively. However, the problem will remain wherever the amount of life-time to be saved is just enough to be worth having (or is thought so to be by those whose time it is) but seems a poor return on the investment required to procure it or in terms of other savings, including savings of longer individual life-time, that might be made instead.

People versus policies We are strongly inclined to believe that where, for example, we can prolong the lives of 120,000 people by one month or 1,000 people by ten years that we should do the latter and that it is better to use a scarce resource to save the life of someone who is likely to live on for at least ten years rather than that of someone who will die in six months in any event. This inclination makes it look as though what we must in fact value is length of life-time rather than simply saving lives. But valuing life-time can be as dangerous to our moral intuitions as is the anti-ageist argument. Again, it might be tempting to believe that a policy of devoting resources to saving individual lives for as long as possible was better than simply maximizing life-time saved. There might be a number of different grounds for such a belief. One such ground would be the expectation that procedures which could prolong individual lives by a substantial period would lead to a greater saving of life-time in the long term than would procedures which merely postponed death for a month or so. But in the absence of any strong evidence for such a conclusion this expectation would be at best an act of faith and at worse a pious hope. Is there any way out?

The fallacy of life-time views

Suppose various medical research teams to be in competition for all research funds available and that one team could demonstrate that it was capable of producing an elixir of life that would make anyone taking it immortal. Suppose further that the entire world medical research budget, if applied to this end, would produce just enough elixir for one dose, and that nothing less than a full dose would have any effect at all. The life-time view suggests that all the money should go to making one person immortal rather than, say, to an alternative project by which another team could make everyone on earth live to a flourishing 80![10]

But there is an obvious fallacy in this argument which reveals a defect in the whole life-time approach. Making one person immortal will produce a saving of no more life years than would the alternative of making everyone on earth live to a flourishing 80. So long as the world itself and its population lasts as long as the immortal (and how – and where – could he last longer?) there would be no net increase in life years lived. Indeed, so long as there is either a stable or an increasing world population, from the life years point of view, it matters not at all who lives and who dies, nor does it matter how many years anyone survives. For, so long as those who die are replaced on a one-for-one or better than a one-for-one basis, there will be no loss of life years. Nor will there be any gain in life years when particular individuals live for longer. For if the overall world population is stable then prolonging the life of particular individuals does not increase the total number of life years the world contains. And if the world population is increasing then it is highly unlikely that prolonging the lives of particular people will fuel that increase. Indeed the reverse is more likely to be the case with the survival of people beyond child-bearing age having a retarding effect on the rate of increase.

In the context of a stable or of an increasing world population, any idea that any policy which did not have the effect of increasing the population in fact made any contribution to the amount of life-time saved would be an illusion.

We do not then have always to calculate the probable net saving in life-time of any particular policy or

therapy, before knowing what to do, and can revert to the more customary consideration of the numbers of lives that might be saved or lost. This, however, highlights once again the problems of whether lives that can only be saved for relatively short periods of time (that can only be prolonged by a few months say) are as worth saving as those for whom the prognosis in terms of life expectancy is much longer. A manoeuvre that seems to capture our intuitions here involves modifying the life-time view into a worthwhile life-time view.

Worthwhile life-time

While to many just staying alive may be the most important consideration, and while they may even wish to continue to live even at appalling cost in terms of pain, disability and so on even, as we have seen,[11] where their lives are hardly worth living, they of course prefer to live worthwhile lives. So that, while any life might be better than no life, people generally expect medical care to concern itself not simply with preventing death but with restoring worthwhile existence.

Many sorts of thing will go to diminish the worth of life just as many and various considerations go to make life valuable and these will differ from individual to individual. For the moment we are just concerned with the question of how life expectancy operates as one of these.

If someone were sentenced to death and told that the execution would take place at dawn the next day, they would not, I imagine, be excessively overjoyed if they were then informed that the execution had been postponed for one month. Similarly if the prognosis for a particular disease were very accurate indeed, to be told that one had only seven months to live would not be dramatically less terrible than to be told one had six months to live. There are two related reasons for this. The first is simply that the prospect of imminent death colours, or rather discolours, existence and leaves it joyless. The second is that an almost necessary condition for valuing life is its open-endedness. The fact that we do not normally know how long we have to live liberates the present and leaves us apparently free to plan the future without having to be constantly aware of the futility of so doing.[12] If life had a short and finite (rather than indefinite) future, most things would not seem to be worth doing and the whole sense of the worth of life as an enterprise would evaporate.[13]

In the light of these considerations many people would not much value such short periods of remission, and support for policies which could at best produce such small gains might well be slight. However, some might well value highly the chance of even a small share of extra time. So far from emptying their life of meaning, it might enable them to 'round it off' or complete some important task or settle or better arrange their affairs. It might, so far from being of no value, be just what they needed to sort their life out and make some sort of final sense of it.

We have frequently noted the extreme difficulty involved in discounting the value of someone's life where we and they disagree about whether or not it is worth living, and we have also noted the injustice of preferring our assessment to theirs when so much is at stake for them. In view of all this it would be hard to prefer our judgement to theirs here.

Perhaps the problem would in reality be a small one. These dilemmas only arise where we cannot both help some people to live for relatively short periods *and* at the same time help others to live for much longer ones. Where there is no such conflict there is no question that we should go on helping people to stay alive for just so long as they want us to. However, the fact that hard cases are rare does not mean that we can turn our backs on them.

Fair innings or no ageism?

We have then two principles which can in hard cases pull in opposite directions. What should we do in the sorts of hard cases we have been considering? First, we should be clear that while the very old and those with terminal conditions are alike, in that they both have a short life expectancy, they may well differ with respect to whether or not they have had a fair innings. I do not believe that this issue is at all clear-cut but I am inclined to believe that where two individuals both equally wish to go on living for as long as possible our duty to respect this wish is paramount. It is, as I have suggested, the most important part of what is involved in valuing the lives of others. Each person's desire to

stay alive should be regarded as of the same importance and as deserving the same respect as that of anyone else, irrespective of the quality of their life or its expected duration.

This would hold good in all cases in which we have to choose between lives, except one. And that is where one individual has had a fair innings and the other not. In this case, while both equally wish to have their lives further prolonged one, but not the other, has had a fair innings. In this case, although there is nothing to choose between the two candidates from the point of view of their respective will to live and both would suffer the injustice of having their life cut short when it might continue, only one would suffer the further injustice of being deprived of a fair innings – a benefit that the other has received.

It is sometimes said that it is a misfortune to grow old, but it is not nearly so great a misfortune as not to grow old. Growing old when you don't want to is not half the misfortune that is not growing old when you do want to. It is this truth that the fair innings argument captures. So that while it remains true, as the anti-ageist argument asserts, that the value of the unelapsed possible lifespan of each person who wants to go on living is equally valuable however long that span may be, the question of which person's premature death involves the greater injustice can be important. The fair innings argument points to the fact that the injustice done to someone who has not had a fair innings when they lose out to someone who has is significantly greater than in the reverse circumstances. It is for this reason that in the hopefully rare cases where we have to choose between candidates who differ only in this respect that we should choose to give as many people as possible the chance of a fair innings.

Notes

1 Triage is a policy for coping with disasters where resources are insufficient to provide the normal standard of care for all. It involves dividing survivors into three groups: those who will die in any event, those who will live in any event, and those for whom care will make the difference between life and death. Care is then given only to this last group. The argument is that this is the most economical use of resources where resources are insufficient to help all.

2 This may be a rash assumption because of the voluntary nature of many risks.

3 Of course if I don't value it because it is so short as to be scarcely worth having then the point does not apply in such a case.

4 I owe this objection to Tom Sorrel and am greatly in his debt here and elsewhere in this chapter for his generous and penetrating criticisms and comments.

5 No non-sexist form is available here, nor is one desirable since a different formulation would lose the resonance of the phrase.

6 Jonathan Glover, *Causing Death and Saving Lives* (Harmondsworth: Penguin, 1977), p. 221.

7 I'm assuming 70 as the full measure of life expectancy of healthy people and that all candidates are healthy in the sense that there is no reason to regard their life expectancy as less than average.

8 Quoted by Jonathan Glover (see note 6), p. 221. I am indebted here and elsewhere to Glover's stimulating discussion of these matters.

9 Because any figure of life expectancy will be arbitrary and one has to be taken.

10 The elixir of life example which prompted this argument about the fallacy of life-time views in stable or increasing populations I owe to Tom Sorrel, whose formulation of it I largely use.

11 See chapter 2 in *The Value of Life* (1985), from which this extract is taken.

12 Many people have argued of course that it is always futile to plan for the future because the inevitability of our world's ultimate destruction makes everything futile.

13 For the record we should note that small gains in life-time will only seem to be worthless to those who gain them if it is known that they will be short. If the potential beneficiaries are kept in ignorance of the fact that they can be granted only a short remission then the extra time will not be clouded by the futility deriving from its short duration and the gain, though small, will be as worthwhile as any other segment of their lives of comparable duration. Of course the deception may not be justified.