

AGE-WEIGHTING*

GREG BOGNAR

Harvard University

Some empirical findings seem to show that people value health benefits differently depending on the age of the beneficiary. Health economists and philosophers have offered justifications for these preferences on grounds of both efficiency and equity. In this paper, I examine the most prominent examples of both sorts of justification: the defence of age-weighting in the WHO's global burden of disease studies and the fair innings argument. I argue that neither sort of justification has been worked out in satisfactory form: age should not be taken into account in the framework of the burden of disease measure, and on the most promising formulations of the fair innings argument, it turns out to be merely an indicator of some other factor. I conclude by describing the role of age in theories of justice of healthcare resource allocation.

1. INTRODUCTION

Suppose you have a life-saving drug, and you have to choose between giving it to either

- (A) a 20-year-old patient, who will live for 10 more years if she gets the drug, and then she dies;
- (B) a 60-year-old patient, who will live for 10 more years if she gets the drug, and then she dies.

The drug is indivisible and the person who does not get it will die shortly. Apart from their age, there is no relevant difference between the two patients, and they would both spend the extra ten years in perfect health. Would you give the life-saving drug to *A* or *B*?

* I would like to thank Dan Brock, John Broome, Sam Kerstein, János Kis, Shlomi Segall and two anonymous referees for useful comments. I would especially like to thank the editor of the journal, Luc Bovens, for numerous suggestions for improving the paper.

Although this is a stylized example, the issues it raises are practical. Should age be a factor in the selection of recipients for organ transplants? How should scarce healthcare resources be allocated between treatments which have unequal benefits for different age groups? How should we contain the escalation of medical expenditures, a significant part of which is spent on healthcare for the elderly?¹

If you prefer giving the drug to the younger person, then apparently you are in the majority. In a study in Australia, respondents considered saving four 20-year-olds as important as saving ten 60-year-olds. A US study found that saving one 20-year-old is roughly equivalent to saving seven 60-year-olds. And in a study in Sweden, saving one 30-year-old was considered to be equivalent to saving thirty-four 70-year-olds. In neither of the latter two studies had the age of the respondents an effect on their preferences, and older people were over-represented in the first. People seem to prefer saving the younger person regardless of their own age.²

Why do people have these preferences? One explanation may be that when people consider the questions in these studies, they compare the overall situation of the possible beneficiaries. They might assume that the older person has had a better life in virtue of having lived longer. Since the younger person is worse off, she should get the benefit. If this was the correct explanation, the preferences of the respondents would correspond to a *prioritarian* view of distributive justice. Like prioritarians, the respondents would be interested in overall well-being and giving priority to the worse off, rather than simply age. Age would be regarded only as an indication of overall well-being.

To be sure, this interpretation of the results entails that the respondents make a mistake when they form their preferences by adding a crucial piece of information which is not part of the question: they implicitly assume that the older person has had a better life. Assuming that they would agree that the quality of life matters as well as its quantity, if we asked the respondents to imagine that the older person's life has been going badly enough to be worse overall than the younger person's life, they would presumably change their preference. If we asked them to imagine that

¹ For discussions of these issues, see, for example, Callahan (1987), Daniels (1988), and Veatch (2000).

² See Nord *et al.* (1996) and Nord (1999: 57–61) for the Australian study, and Cropper *et al.* (1994) and Johannesson and Johansson (1997) for the US and the Swedish ones. All of these studies used the *person trade-off method*, in which the respondent is asked to compare hypothetical programs for treating some given number of patients with some characteristic (in this case, a certain age) with treating some other people with some different characteristic (a different age). The aim is to find the number of patients in the latter program for which the respondent is indifferent between the two programs. Note that studies with different methodologies had similar outcomes, thus the results are unlikely to be an artifact of a particular preference elicitation procedure.

the older person's life has been going just badly enough overall to be as good (or bad) as the younger person's life, they would presumably be indifferent.³

On the standard interpretation of these studies, however, people do not make a mistake when they form their preferences. Rather, the results show that people value health benefits differently depending on the age of the recipient: they *weight* benefits according to the age of the beneficiary. Thus, other things being equal, the same health improvement has more value if it goes to the 20-year-old than if it goes to the 60-year-old.⁴

Some studies, moreover, also found that respondents do not merely look at overall quality and length of life when they form their preferences. Rather, the relative value they assign to a year of life increases rapidly from birth to young adulthood and then steadily decreases throughout middle and old age. This is taken as evidence that people do age-weight health benefits.⁵

Nevertheless, we know little about the *reasons* people have for these preferences – and the little we do know is controversial (see Tsuchiya 1999). Shifting the focus from explaining people's preferences to justifying them, researchers have come up with their own defences of age-weighting. Two main types of rationale have been offered. On the first, age-weighting is justified on grounds of efficiency: if rationing takes into account the age of the beneficiaries, it is more likely to maximize benefits. This has been called "productivity ageism". On the second, age-weighting is justified on grounds of equity: on this view, fairness requires that age is taken into account. This is known as "fair innings ageism".

Some philosophers object to standard measures of health benefit on the grounds that they are "ageist" – that is, they discriminate against the elderly. For instance, if we measure health benefits by quality-adjusted life

³ Another explanation may be that people have the preferences they do because they implicitly assume that younger people would enjoy a given benefit for a longer period, since they expect to live longer. Tsuchiya *et al.* (2003) distinguish the possibility that respondents implicitly assume that benefits last a full lifetime (so that giving them to younger people does more good) and the idea that the same health gain has different values at different ages. They find that people do give unequal value to health gains with respect to age, even though sometimes they disregard the difference between fixed and lifetime benefits.

⁴ This interpretation assumes that people's preferences elicited on the basis of questions about life-saving situations are consistent with their preferences with regard to health improvements in general – a common assumption when using standard preference elicitation procedures. There is some more direct evidence for this consistency, although it has also been found that people's preferences vary to some extent with the type of the intervention (see Johri *et al.* 2005).

⁵ For studies in the UK, the Netherlands, the USA and Australia that show that people's preferences can be represented by a "hump-shaped" age-weighting function, see Lewis and Charny (1989), Busschbach *et al.* (1993), Cropper *et al.* (1994), and Nord (1999: 58–61), respectively.

years (QALYs) – which combine the quality of life improvement with the duration of the improvement – we *might* systematically give priority to the young simply because their prospects for improvement are usually better. If, in addition, we weight a given benefit by the age of the recipient, we almost certainly *will* give priority to younger people. Age-weighting is one kind of ageism.

Despite the possible ethical worries about ageism, it has been argued that the “public values” revealed in people’s preferences should be reflected in healthcare resource allocation decisions and priority-setting. Indeed, sometimes they already are. The World Health Organization (WHO) uses summary measures of population health which include age-weighting. Modifications of cost-effectiveness analysis which incorporate age-weighting for health benefits have been proposed as well.⁶

In this paper, I shall attempt to raise some doubts about age-weighting by presenting some problems with its standard rationales. Section 2 discusses the WHO’s efficiency-based justification. Section 3 looks at some variants of fair innings ageism. Section 4 situates the proposal for age-weighting in theories of justice of healthcare resource allocation. Section 5 concludes.

2. WELFARE INTERDEPENDENCE

The World Health Organization uses disability-adjusted life years (DALYs) to report the burden of disease – that is, the harm resulting from premature mortality and disability from disease and injury for the population in different countries. The DALYs associated with a given condition are the sum of the years of life lost and the years lived with disability due to that condition. The conditions range from facial vitiligo through rheumatoid arthritis to quadriplegia.

The burden of premature mortality from a condition is measured on the basis of *years of life lost* due to that condition. The measure represents the gap between current mortality and an ideal standard of life expectancy. The burden of death at each age is partly determined by the expectation of life at that age. The standard life expectancy is set to 82.5 years at birth for females and 80 years for males; at 50, it falls just short of 34 years for females and 31 years for males. The differences between the sexes occur because of their ostensibly different survival potential.

Standard life expectancies only partly determine the years of life lost from premature death, since years of life lost are obtained by subjecting standard expected years of life to age-weighting and discounting. Thus, the death of a newborn results in 33.12 years of life lost if the child is female

⁶ For instance, weighting QALYs by age has been suggested both by health economists (Rodriguez and Pinto 2000) and philosophers (Kappel and Sandøe 1992).

and 33.01 years of life lost if the child is male; for a 50-year old, the years of life lost are 17.69 and 17.12, depending on the person's sex.⁷

The burden of morbidity from disease and injury is measured in terms of lost quality of life during the time the person is in a worse health state than perfect health. This component of DALYs is *years lived with disability*. Lost quality of life due to different conditions is represented by disability weights between 0 and 1, with 0 for perfect health and 1 for death. Years lived with disability are calculated by weighting the time spent with the disability with its disability weight. But years lived with disability are also subject to age-weighting and discounting. Age-weighting is applied both to the age of the onset of the disability and the duration of the time in diminished health.

DALYs are therefore a composite measure of the harms resulting both from lost quantity of life and lost quality of life. Since they are standardized, the burden of disease in different regions and populations are comparable and health trends can be discerned. For example, according to the projections of the 1996 global burden of disease study, the number of DALYs associated with HIV in both its treated and untreated forms will be just below 2.3 million in 2010 in all established market economies, while it will be well over 20 million in that year in Sub-Saharan Africa (with a slightly larger population). DALYs can also be used as an outcome measure for the purposes of cost-effectiveness analysis of health interventions and policies. The objective associated with the reduction of burden of disease is to *minimize* DALYs, reducing the gap between ideal population health and the actual population health.

Let us consider age-weighting in the context of disability-adjusted life years. Suppose that in my original example both patients are female. If we give the life-saving drug to the 20-year-old, the overall years of life lost, without age-weighting and discounting, will be 78.1 – the sum of the standard life expectancies at ages 30 and 60. If we give the drug to the 60-year-old, the overall years of life lost, again without age-weighting and discounting, will be 79.28 – the sum of the standard life expectancies at ages 20 and 70. The reason for the difference is that the decrease in life expectancy between the ages 20 and 30 is larger than between 60 and 70. For obvious reasons, life expectancies decrease at higher ages; but in the

⁷ Age-weighting and discounting are different operations. When we discount a future benefit, we give it less value the further it is in the future relative to the present. When we age-weight, we give different value to a benefit depending on the age of the beneficiary, regardless of the temporal distance of the present and the time the benefit is realized. Discounting future health benefits and harms is highly controversial. Aware of the controversies, the developers of the DALY measure recommend that the burden of disease should be reported both with and without discounting. Here I put the issues raised by discounting aside. Although data on DALYs without age-weighting are available, the developers clearly don't think similar caution is warranted with respect to it.

life expectancy table used for calculating years of life lost, they decrease at a decreasing rate after childhood.

What about years lived with disability in the example? Since the extra ten years, whomever they go to, will be spent in perfect health, there are no years lived with disability to consider. Thus, the overall DALYs are 78.1 and 79.28.

DALYs would therefore recommend that we give the drug to the younger patient. But this is not because we value years of life lived at different ages differently. We simply minimize the overall burden of premature mortality thereby.

Nevertheless, we might be tempted by the following argument. The difference between 78.1 and 79.28 years of life lost – around one year and two months – is very small, especially given that no matter what we do, there will be almost 80 lost years of life. We might be tempted to argue that such a small difference should not make the difference between life and death. Some philosophers have argued that there are “irrelevant utilities” – harms and benefits which are insufficient to justify a life-or-death choice.⁸ Following their lead, we might come to the conclusion that one year and two months are not sufficient to justify giving the drug to the younger patient. We might instead, for example, prefer to allocate the life-saving drug by tossing a fair coin, giving an equal chance to both patients.

DALY calculations, however, normally do value years lived at different ages differently. The relative value of a year lived at 20 is around one and a half the relative value of a year lived at 60. If we introduce age-weighting in our example, giving the drug to the 20-year-old results in approximately 69.28 (non-discounted) DALYs; giving the drug to the 60-year-old results in slightly less than 77 (non-discounted) DALYs. The difference is over seven years of life lost. Even if we were attracted to giving an equal chance to both patients before, now the burden of premature mortality in the second case is substantially larger than in the first. Arguably, such a large difference should be sufficient to justify giving the drug to the 20-year-old. Thence age-weighting can make the difference between life and death.

Why does the WHO include age-weights in the measurement of the burden of premature mortality, disease and injury? Christopher Murray, one of the developers of DALYs, appeals to two principles that he claims rest on our ordinary notion of fairness. These are as follows:

⁸ See Kamm (1993), although I am not entirely sure what her view would be in this particular case. She says that irrelevant utilities are “concerned with losses or gains of a certain size. These losses and gains are in the appropriate sphere (e.g. life and death), stem from direct need for our resource, and producing them does not involve unfairness or selection as mere means. Here saying ‘a certain type of loss’ is irrelevant really means that a loss of a certain size is irrelevant” (1993: 149). Since there is nothing in the example which is incompatible with this description, the difference in lost years of life may be considered irrelevant utility.

Principle 1. The burden calculated for like health outcomes should be the same.

Principle 2. The non-health characteristics of the individual affected by a health outcome that should be considered in calculating the associated burden of disease should be restricted to age and sex.⁹

These principles, taken together, establish information constraints on the calculation of DALYs. How bad a health outcome is evidently depends on a lot of variables. Some of these variables, like the etiology of the disability, are biomedical. These are relevant to the description of a health state. But DALYs do not merely describe health states; they measure their *burden*. According to the second principle, variables like income, level of education or the availability of social support are irrelevant for this exercise.

Murray (1996: 7) uses the following example to illustrate the motivation for the principles. Suppose two patients who are identical in every respect except that one is rich and the other is poor arrive simultaneously at the emergency room, and there are insufficient medical supplies to treat both of them. He claims that it would be unfair to choose to save either of them on the basis of their wealth. Hence the principles should reflect this unfairness.

The problem with this example is that claims about how people's health should be *evaluated* cannot be inferred from claims about how they should be *treated*. We might agree with Murray that fairness requires that we should not discriminate between the patients on the basis of their wealth, thus they should have an equal chance to be saved. Nonetheless, we might also claim that more benefit would be produced by saving the rich patient (on the assumption that she has more well-being), or that the benefit produced by saving the poor patient should have more weight (on the assumption that benefits to the worse off have more value). Our view about how the patients ought to be treated should not distort the measurement of the value of the outcomes.

Some philosophers suggest that fairness is about mediating people's conflicting claims for some good – in this case, saving their lives.¹⁰ If this

⁹ The principles are quoted from Murray and Acharya (1997: 709); see also Murray (1996: 6).

¹⁰ See, for example, Broome (1999). For the argument of the previous paragraph, see also Broome (2002). Broome also argues that health outcomes cannot be evaluated independently of overall well-being, since health and other components of well-being are not separable (2002: 93–100). The question of how bad the consequences of disease or injury are for a person cannot be answered unless we take into account the consequences of other factors for the person's well-being: paraplegia is less bad if public buildings are wheelchair accessible, mental illness is less bad if there is social support. Thus, when we evaluate health, we have to evaluate overall well-being. Broome, therefore, rejects

view is correct, Murray's principles cannot rest on the notion of fairness. I want to suggest, however, that the principles rest instead on the notion of *impartiality*. Fairness concerns how people should be treated; impartiality concerns how their interests should be taken into account. Thus, one way a choice can be unfair is if it is not based on an impartial appraisal of the interests of the people affected. On this interpretation, the principles give shape to some ideas about how we should evaluate health benefits by setting constraints on which consequences of disease or injury are relevant. If two people are in the same health state, it would not be impartial to calculate the burden of their disability differently – especially on the basis of a non-health characteristic like income.

Separating the roles of fairness and impartiality in motivating the principles is also important for the following reason. Its developers recommend DALYs both as a measure of population health and as a tool for cost-effectiveness analysis on the basis of which resource allocation decisions can be made. Measuring the burdens of ill-health in an impartial way is important for both purposes. But to suggest that the underlying principles already incorporate our concern for fairness is to disregard fairness claims that may arise independently. In order to make just resource allocation decisions, an impartial appraisal of the burdens of benefits is certainly necessary. However, it may in itself be insufficient, for other considerations may need to be taken into account. Therefore, whether or not DALYs are appropriate as a measure of burdens, they cannot be assumed without further argument to be an appropriate tool for decision making.

How can the constraints articulated in Murray's principles be defended? In particular, how can we justify Principle 2 that allows age to be taken into account? (I will ignore sex.)

One argument starts from the fact that people's productivity varies at different ages. Young adults contribute more to economic output than children and the elderly do. On this argument, DALYs are interpreted as a measure of the *social value* of health in terms of its contribution to human capital and economic prosperity. From this perspective, however, many other variables besides age are relevant for the evaluation of health: income and level of education are obvious examples. Thus, this argument cannot defend Principle 2.

Murray also rejects this argument, but he does not reject the basic idea that age-weighting makes it possible to incorporate the social value of health into the measurement of the burdens of ill-health. His own argument begins by noting that people promote the well-being of others

Principle 2, and he must also think that Principle 1 is at best merely vacuously satisfied – since given that we have to measure overall well-being, there are hardly ever any health outcomes that are alike. For the sake of the argument, however, I shall assume that Principle 1 is not merely vacuously satisfied.

in many ways – by providing care, household services and so on. These contributions cannot be captured if we focus on economic output and human capital only. Nonetheless, *welfare interdependence*, as he calls it, has a crucial role in society. If we broaden our focus to take all sorts of welfare contributions into account, it follows that, on the one hand, it is not necessarily the case that people with higher income or level of education contribute more to the welfare of others. Income or education are therefore not relevant variables. On the other hand, young adults tend to contribute more to the well-being of others by providing for their children and taking care of the elderly. Some of this contribution is economic, but most of it is non-monetary. Consequently, age is the most suitable variable to capture welfare interdependence.

One may worry that Murray's conclusion is too strong. For one thing, even if we can exclude income and education, other variables may be relevant to capture welfare interdependence. Why should we not take into account, for instance, whether one is a parent? On the other hand, welfare interdependencies are highly variable across cultures. In some of them, older generations provide substantive family support, including helping parents to raise their children. Emphasizing the dependence of the elderly on their adult children may be indicative of a form of cultural bias.

Murray also argues that using age as a variable is not discriminatory in the way using income or education level would be. Most people live through different ages: with the passage of time, they proceed from childhood to adulthood to old age. Thus, age-weighting does not discriminate between individual lives, only between periods within lives. Even though there certainly are variations in the extent to which particular individuals provide for the well-being of others, by and large it is true that people at different ages contribute differently. Age-weighting remains impartial with respect to the characteristics of particular persons.

Thus, there are two claims embedded in the argument from welfare interdependence: that age is the appropriate variable to take welfare interdependence into account, and that using this variable is not discriminatory.

Consider the second claim first. It seems to me that age-weighting *can* be discriminatory, especially if, for instance, the relative weight of childhood is smaller than the relative weight of young adulthood, as it is in DALYs. If a child suffers from the same disease for the same period as a young adult, the burden associated with her condition will typically be smaller, and treating the adult may lead to a larger overall reduction of DALYs. One might believe that an older person cannot reasonably complain if her interests are given a smaller weight than the interests of a younger person, because the older person has already benefited from having her interests count for more when she was younger. But this reasoning is not available in the case of the child: she has not been

compensated by having her interests count for more at some other time (and she might not ever be compensated if she dies prematurely). Thus, the child could reasonably complain that her interests are given lower weight than the interests of her older peers. Unless some further argument is given, impartiality is violated in her case. Treating her differently on this basis would be unfair.

Leaving aside the worries about age-discrimination, I will concentrate on the first claim in the rest of this section. If welfare interdependence is taken into account, DALYs provide more than a summary measure of the badness of premature mortality and the consequences of disease and injury. Because of age-weighting, they also represent the social value of health. The social value of the health of a 20-year-old is greater than that of a 60-year-old, since young adults contribute more to the well-being of others. Other things being equal, depression when you are 20 is not worse for you than depression when you are 60; rather, it tends to be worse *for others*. It tends to be more burdensome for society.

Whether this rationale is adequate turns on how we identify the consequences of disease and injury. And here we run into problems. Here's why.

The WHO distinguishes between impairment, disability and handicap. An *impairment* is a loss or abnormality of psychological, physiological or anatomical function, described in biomedical terms. It is the direct consequence of disease or injury. A *disability* is a loss or restriction of ability to carry out an activity which is considered normal for humans as a consequence of an impairment. A *handicap*, in contrast, is the disadvantage resulting from an impairment or disability which limits or prevents the individual to fulfil her role in her economic, social and cultural environment.¹¹ For instance, suppose a chemical imbalance in the brain causes depression. This impairment may lead to a disability characterized by loss of motivation and inability to concentrate, which, in turn, may lead to unemployment and social isolation – a handicap.

It is not entirely clear how handicap and well-being differ from one another. The notion of handicap captures the disadvantage that results from a disability, which varies with the social role and circumstances of the person. Perhaps one could say that a handicap reflects a particular *way* a person's life goes worse for that person. Thus, a handicap represents the consequences of a given disability only, such that different disabilities result in different handicaps. But the disadvantage that results from different disabilities cannot be partitioned this way. The same handicap

¹¹ I use these definitions as they appear in Goerdts *et al.* (1996: 100). More recently, the WHO revised its classification. The revised classification emphasizes functioning, including the impact of impairments as well as social and environmental factors. For more details, see <http://www.who.int/icidh/>.

may be the consequence of several disabilities. Co-existent disabilities can amplify each other's impact, and the removal of one disability may not lead to a reduction of disadvantage. I doubt that a distinction between handicap and well-being can be maintained.

Disability-adjusted life years attempt to represent disability. Suppose, however, that we tried to interpret DALYs some other way. For instance, we might suppose that DALYs measure handicap. In this case, they would capture the full consequences of health for well-being. These consequences would vary greatly with age, income, level of education, availability of support services, and many other factors. One relevant factor would be the support of others. That is, the way young adults contribute to the well-being of others would already be taken into account in DALYs.

This is because, on this interpretation, different disability weights would have to be assigned not only to treated and untreated forms of diseases and injuries, but also to the same disease or injury depending on whether the person is rich or poor, young or old, educated or not. The effects of welfare interdependence would already be incorporated into these disability weights. Weighting them by age would be counting them twice over.

On the other hand, suppose that DALYs measure impairment. But different impairments have different consequences for a person's well-being. Thus, if DALYs were interpreted this way, focusing on the medical condition rather than the badness of the condition, they could not incorporate the consequences of health for well-being at all. They would not measure the value of health or the burdens of ill-health.

As a matter of fact, Murray *et al.* (2002: 735–7) argue that disability weights are a measure of health states and not a measure of the value or utility of those health states. But few commentators agree that it is possible to order and compare health states on any basis other than their value – at least beyond trivial comparisons. On the one hand, a person in full health is definitely healthier than a person with a disability. On the other hand, it is difficult to see how you can compare a pair of complex health states involving different disabilities unless you are prepared to appeal to their value. One way of doing that is to compare the health states with respect to their impact upon the person's well-being – which depends, among other things, on the economic, social and cultural environment. In any case, if DALYs are intended to serve as a measure of health rather than a measure of the value of health, then welfare interdependence should have no place in them when they are presented as a summary measure of population health.

DALYs attempt to strike a middle ground between impairment and handicap by focusing on disability. Welfare interdependence, however, should be relevant only to handicaps. Because of age-weighting, it turns out that depression is less bad if you are 60 than if you are 20. But this is not

because, other things being equal, it is less of a *disability* to be depressed when you are 60 than to be depressed when you are 20; it is because, other things being equal again, it is a less bad *handicap* to be depressed at 60 than at 20.

Consequently, the middle ground of disability is unstable ground on which to justify age-weighting. DALYs attempt to include one sort of factor in determining the burdens of ill-health and exclude others. This is arbitrary. Either DALYs represent handicap, in which case many other variables besides age should be taken into account, or they represent disability, in which case age should be excluded along other variables.

Murray is aware of this problem. At one point, he suggests that disability weights represent average handicap (1996: 38–9). I am not sure how that “average” is to be understood. When disability weights were assigned for different conditions, assessors from different regions of the world were told to think of the average individual in their own society in their average social environment. But disability weights are uniform across all regions. If they do represent global average handicap, then it would seem that they conceal huge disparities between different parts of the world. To refer back to my earlier example: the overall consequences for well-being of living with HIV in Sub-Saharan Africa, even if one receives treatment, are very different than living with HIV in established market economies. If DALYs represent average global handicap, they underestimate the burden of HIV in the former.

DALYs attempt to measure the burden of disease by the impact of ill-health on people’s well-being, taken in a limited sense as disability. The variations of the impact with age either has no place in such a measure or they should be taken into account along with other sources of disadvantage associated with ill-health – but in this case, they lead to the double-counting, which certainly violates impartiality. Justifying age-weighting as a way to account for welfare interdependence does not succeed.

3. FAIR INNINGS

Some may prefer giving the life-saving drug to the 20-year-old person in my example because they think this is what fairness requires; it would be *unfair* to give it to the older person or selecting the recipient randomly. Several philosophers have argued for this view, using the *fair innings argument*.

The argument is intended to establish the conclusion that if different persons compete for life-extending resources and they are identical in every relevant respect except for their age, then it is justified to give priority to the younger person even if the net benefit accruing to that person is not larger than the net benefit that would be received by the older person.

Different formulations of the argument specify different relative weights to ages.¹²

A crucial premise is that justice or fairness can be applied to quantities of life. As a person grows older, she approaches or reaches her “fair share” of life – her fair innings. Hence, at least relative to others who have had less of her share of life, her claim on resources is diminished or relinquished entirely. The rationale that arguments from fair innings provide for age-weighting is that standard evaluations of the benefits of healthcare resource allocation fail to take account of requirements of fairness. One way to take these requirements into account is to introduce age-weights.

The fair innings argument was introduced by John Harris (1985: 87–102; 1988), although he gives it a very limited role. Harris argues that since everyone values going on living equally, everyone suffers the same injustice if their lives are cut short. In deciding whom to save, what matters is not the magnitude of the benefits, but the magnitude of the injustice. This implies that age should not be taken into account. This is the *anti-ageist argument*, which gives support to the following principle:

Anti-Ageist Principle. Each person should have an equal claim for a life-saving resource, even if some person’s prospect with respect to the quality or length of her life is worse than the prospects of others.

There is one exception, though. Some people have had a reasonable life – for instance, they have lived three scores and ten years. When they reach 70, they had their fair innings, and the remaining years are “bonus time”. They cannot reasonably complain that an injustice is done to them if resources are diverted to save the life of those who are under the threshold. This is Harris’s fair innings argument, which gives support to another principle:

Harris’s Fair Innings Principle. If it is impossible to save the lives of both a person who has had her fair innings and another person who has not, then the life of the person who has not had her fair innings should be saved.

Conjoined, the two principles determine the following values for evaluating alternatives in life-saving situations. If the people affected are

¹² I put aside the complication that benefits to children are often thought to have smaller relative value than benefits to young adults. Thus, I leave open the question of how fair innings arguments could be modified in order to be compatible with this view. Note also that fair innings arguments are usually put forward in the context of the distribution of resources for saving or extending lives. It is often unclear whether they are intended to apply only to this narrow context, or (assuming that a distinction can be maintained between the two) more generally to the distribution of all healthcare resources.

	Young gets the drug		Old gets the drug	
	lives	dies	lives	dies
Young	1	0	0.9	0.1
Old	0.1	0.9	1	0

FIGURE 1.

all under the fair innings threshold, saving any of them has uniform value, since, by the Anti-Ageist Principle, we are prohibited to differentiate by prospective quality or length of life. But if a person has had her fair innings, saving her has strictly smaller value (perhaps none at all). This difference becomes important when we have to choose between people under and over the fair innings threshold.

There are many difficulties with Harris's view.¹³ I want to concentrate on the claim that the scope of the Fair Innings Principle can be limited the way Harris intends. Note that he assumes that there is a clear-cut distinction between life-saving and other contexts. He claims that the principle applies if and only if it is "clearly impossible to postpone the deaths of all those who wanted to go on living" (1985: 94).

Suppose you have to choose between giving the only drug you have to an older person who has had her fair innings and a younger person who has not. If you give the drug to the younger person, she will certainly survive, while the older person will die with a 0.9 probability. If you give the drug to the older person, she will certainly survive, and the younger person will also survive with a 0.9 probability. Giving the drug to the older person makes it very likely that the deaths of both persons can be postponed, although there is a slim chance that the young person dies. The older person is saved for certain, although she already had her fair innings.

The situation is illustrated in Figure 1. The numbers in the cells indicate the probability of the different outcomes, depending on who gets the drug.

I am not sure whether we should regard this situation as one in which it is "clearly impossible to postpone the deaths of all". Suppose it is such a situation: it is clearly impossible to make sure that no-one dies, regardless

¹³ See, for instance, Broome (1988) or Kappel and Sandøe (1992). Their objections mainly target the Anti-Ageist Principle, which for the sake of the argument I take for granted here. For a view that is similar in many respects to Harris's, see Callahan (1987).

of what you do. If you give the drug to the older person, it is highly likely that both persons will survive – but it is not certain. So Harris's Fair Innings Principle should apply.

But if it does, then it prohibits giving the drug to the older person. You should give the drug to the younger person. But if you do that, the older person is highly likely to die.

On the other hand, suppose that this is not a situation in which Harris's Fair Innings Principle applies. It is not clearly impossible to postpone the deaths of all, since if you give the drug to the older person, both of them, with a bit of luck, can survive. Hence your choice should be determined by the Anti-Ageist Principle only. It tells you to give an equal chance of getting the drug to both persons. But if you do that, there is a significant probability that the older person will die.

You might believe that you should aim to minimize the risk of anyone dying. You can do that easily by giving the drug to the older person, or at least giving her a higher chance of getting the drug.¹⁴

These counter-intuitive results are not merely artifacts of the Anti-Ageist Principle. They are due to the assumptions that a fair innings is a fixed threshold, beyond which a person had a full, reasonable life, and that the years remaining after one reached her fair innings are a kind of bonus, giving rise to no obligations of justice on others.

Of course, as Harris notes, setting a fair innings at three scores and ten years is a simplification. Besides quantity, the quality of life also matters. Reasonable people might disagree whether someone had her fair innings. Hence Harris further limits the scope of his Fair Innings Principle by insisting that it only applies if reasonable people agree whether someone had her fair share of life. But whenever they do, there will be a cut-off point somewhere, and we run into these problems.

There is, however, another way of interpreting the notion of fair innings. On the view I have been discussing, a fair innings is a more or less fixed threshold of life years beyond which the person has no claims for having her life prolonged, given that she competes for the available resources with others who have not reached this threshold. On an alternative view, a person's claim on resources steadily diminishes the longer she lives: she has had *more of* a fair innings. Justice requires that

¹⁴ Since their prospects are different, you do not actually give an equal chance of *survival* to the younger and the older person if you give them an equal chance of *getting the drug*. Equalizing their chances of survival would require giving the drug to the older person with probability 0.9. (In this case they would both have a 91% chance of living.) So it may be objected that this is what the Anti-Ageist Principle requires. But the principle also says that your choice should be independent of the prospects of the persons. If you give a higher chance of getting the drug to the older person, the younger person may reasonably complain that her chances of survival are diminished merely due to the older person's worse prospects – who might, as in this case, have already had her fair innings!

when we allocate scarce resources we take into account the share of life that a person has had relative to others. We could formulate a fair innings principle for this view as follows:

Life-Share Fair Innings Principle. It is not unjust to allocate life-extending resources on the basis that the younger a person is, the weightier, other things being equal, that person's claim is on available resources.¹⁵

This view is difficult to reconcile with the rationale for age-weighting from the notion of a full, reasonable life. On the natural understanding of that notion, a person's life has either been full and reasonable or not. Thus, some other rationale must be given to motivate the view that the weight of a person's claim diminishes as she ages.

One rationale might be that claims on distributable resources are generated by *need*.¹⁶ Thus, in the context of the allocation of life-extending resources, one factor that determines the strength of a person's need is age. Other things being equal, the older a person is, the less she needs the benefits of life-extending treatment such that the weight of her claim upon available resources steadily diminishes. On this view, age-weighting ensures that we remain sensitive to differences with respect to need.

It should be clear that "need" is not used in this argument in its ordinary sense. If a person is about to be killed, we do not think he needs to be saved less if she is older. If there are two patients with a headache, we are unlikely to believe that one needs a pain reliever more just because she is younger.

Frances Kamm (1993: 234–5; 2002: 692) suggests that when we allocate life-extending resources, we should distinguish between need and urgency. *Need* is determined by how badly a person's life *will have gone* if she is not saved. In contrast, the *urgency* of helping a person is determined by how badly a person's life *will go* if she is not saved. Note that "urgency" is not used in its ordinary sense here either, in which it refers to how soon the employment of the resource is necessary. On Kamm's definitions, need is based on an evaluation of the overall benefit a person enjoys throughout her life, depending on whether she gets the resource; urgency refers to the prospective benefit that the person enjoys from the time of allocating the life-extending resource, depending on whether she gets it. When we consider need, our evaluation concerns a person's whole life, but when we consider urgency, our evaluation looks only at the future. Need and urgency are two different perspectives with respect to which we can evaluate competing claims.

¹⁵ This formulation follows Lockwood (1988: 50).

¹⁶ See Broome (1999: 119) for the suggestion that needs may generate fairness-claims.

Kamm argues that sometimes fairness requires that we decide on the basis of need. This is true when the candidates for a resource do not differ with respect to urgency, but it might be true even if thereby we help the less urgent. Consider the case in which there is no difference in urgency. For illustration, my original example will once again do. It is equally urgent for *A* and *B* to get the life-saving drug, since if they do not, they will die. We might also consider need, in which case we should

give to those who, if not helped, will have had less of the good (e.g. life) that our resource can provide (at least if they are equal on other health dimensions) before giving to those who will have had more of it even if they are not helped. (Kamm 2002: 695)

If *A* does not get the drug, she will have lived only 20 years, while if *B* does not get the drug, she will still have lived 60. Hence *A*'s need seems to provide a stronger claim on the life-saving drug. But this is not necessarily so. Need is based on an evaluation of how well overall a person's life goes relative to the lives of others. In order to decide whether *A* or *B* is needier, we need to know more. If the older person's life will have been going badly enough to be worse overall than the younger person's life, then she is needier than the younger person. If her life will have been going just well enough overall to be as good as the younger person's life, they are equally needy.

That is to say, it is not *life as a good*, but the *goods in life* that matter: if *A* is needier, that is not because she will have had less life, but because she will have had fewer goods in her life. The *ceteris paribus* clause in the Life-Share Fair Innings Principle cloaks this distinction, since it can be read in different ways. On one reading, other things are equal only if the candidates for the resource will have been – that is, would overall be – equally well off. Then, according to the principle, the claim of the younger person gets more weight. But if the candidates for the resource would overall be equally well off, they are equally needy. Their needs generate claims with the same weight. This is in conflict with the principle. Therefore, the concept of need cannot give a rationale for the Life-Share Fair Innings Principle. The principle must be motivated by other considerations.

On another reading, other things might be equal even if the candidates would not be overall equally well off. It might be that the younger person would be worse off – that is, if not given the resource, she will have had less good in her life (perhaps merely due to having lived less). Therefore, she is needier. Or it might be that the older person would be worse off overall, despite having lived longer. In that case, she is needier. Of course, normally younger people tend to have had fewer goods in their life – especially compared with substantially older people. So often we can take age as an indication of need. But then the benefit that we have to allocate is not valued differently according to age; it is valued differently according

to need. The two are different, since we care about age only insofar as it can stand as an indicator for something else. What we ultimately care about, in this case, is need.

Kamm might not disagree. She says that “age can be given some weight in decision making in virtue of its role in identifying the neediest candidates” (1993: 248). Thus, if people of different ages compete for resources, we should favour the young, since their need is likely to be greater. But if this is so, then the Life-Share Fair Innings Principle is redundant. Appealing to other considerations, perhaps together with using age as an indicator, are sufficient to do the work.

Similar problems arise for other proposals which may be employed to defend the Life-Share Fair Innings Principle. For instance, we might try to defend the principle by appealing to equality. Alan Williams (1997) argues that the fair innings argument can be quantified in order to determine the equity-efficiency trade-offs involved in any attempt to reduce health inequalities.

Williams assumes that a fair innings is determined by the mean quality-adjusted life expectancy (QALE) at birth. Here we consider actual life expectancy at birth in a given country. This contrasts with standard life expectancies used in DALYs, which reflect some ideal, biologically determined survival potential. Actual life expectancies vary with societies and populations. They are adjusted for quality, since a long life with poor quality deserves consideration just as much as a shorter life with higher quality. QALEs represent the “normal” span of life people should enjoy in their society.

Since the better off tend to have greater QALEs, health disparities can be reduced by selecting policies which equalize QALEs by sacrificing some of the overall health of the population. This objective applies to the distribution of healthcare resources in all settings, not only in the context of saving or extending lives. The degree to which people are willing to sacrifice overall health reflect their aversion to inequality. On this view, the objective is to equalize health outcomes throughout people’s lives.

For example, suppose you have to allocate funds between pediatric and geriatric care services. You know that pediatric care services would be utilized mainly by poorer people, and they would contribute to longer and better quality lives for the worst off. But you also know that more better-off people will need geriatric care services, and the overall improvement in the health of the population would be larger if they were given priority. Based on fair innings considerations, you may nevertheless choose the former program even at the expense of lower population health in order to reduce disparities in quality-adjusted life expectancies.

This view focuses on the disparities among QALEs of social classes. Different weights are attached to additional life years gained by people according to the social class they belong to. People from the worst-off

class who exceed their fair innings tend not to be penalized the way their better off contemporaries are. Hence, Williams's version of the fair innings argument is not a proposal to weight health benefits to reflect their differing value due to the age of the recipient. Rather, it takes age – more precisely, quality-adjusted life expectancies – as a proxy equalizandum and age-weighting as an instrument of equalization.

The more accurately quality-adjusted life expectancy at birth tracks how well off a person is going to be throughout her life, the more accurate instrument age-weighting might be to approximate equality of outcomes or to indicate differential need. Whether using it is appropriate is partly an empirical question. But, as we have seen in Section 2, the measurement of health is difficult and controversial. Thus, age-weighting as an instrument of equalization is bound to leave us uncertain as to whether we target the important inequalities and how successful we are in achieving our aim. It is also more difficult to influence health inequalities than, for instance, inequalities of income and wealth – hence age-weighting may be a grossly inefficient policy instrument. Furthermore, it is a controversial question whether information on inequalities in terms of QALEs, QALYs, or similar measures allows us to identify injustices in society which are morally important to rectify. This is a large and complicated issue beyond the scope of this paper.¹⁷

One consideration against using age-weighting as an instrument in the service of equality or other objectives might be this. One of the ways in which this paper's opening example is unrepresentative of decisions for the allocation of life-extending resources is that whoever gets the life-saving drug will derive the same net benefit: both the 20-year-old and the 60-year-old would live for another ten years with the same quality of life. Often, however, younger people will derive more benefit from the same resource since their remaining life is longer or their quality of life is higher or both. Thus, if we take into account overall expected net benefits, the young will tend to get priority simply in virtue of having better prospects.

The magnitude of expected benefits will be one factor in the allocation of life-extending resources on any sensible view. This already, so to speak, stacks the cards against the elderly. So far, the only rationale we have found for age-weighting health benefits is that by age-weighting we can promote some other value with which age is arguably connected. If we are concerned with these other values, we should try to take them into account directly.

¹⁷ For a more detailed discussion of these and further arguments, see Hausman (2007). See also Daniels *et al.* (2004).

4. AGE-WEIGHTING AND JUSTICE

The familiar device of the veil of ignorance is often used in discussions on justice in health (see Daniels 1985, McKie *et al.* 1998). How would rational individuals choose to allocate healthcare resources from an impartial position? What role would they give to age?

In order to answer these questions, Norman Daniels argues that the *interpersonal* problem of just distribution across different age-groups can be reduced to the *intrapersonal* problem of the “prudential lifetime allocation” of healthcare resources (1985, 1988). Problems of distribution of resources between age-groups can be addressed by regarding them as problems of distribution between life-stages.

In Daniels’s view, we appeal to the prudential reasoning of hypothetical deliberators placed behind a veil of ignorance. The veil constrains the available information. Thus, the deliberators do not know their age, socioeconomic status and income, family situation, genetic history and health prospect, and their particular conception of the good. They do know, however, the age and disease profiles and the technological stage of their society. Daniels suggests that in this situation people would select institutions for the distribution of health resources which protect people’s range of opportunities at each stage of life. The relevant notion of opportunity is *age-relative*: the range of opportunities and the particular opportunities which are important for people change as they move through different stages of life.

Behind the veil, people have no reason to assume that their lifespan differs from what is normal in their society. Thus, they have to balance the amount of resources which increases their chances of reaching the normal lifespan with the amount they save for interventions which extend their life once they have exceeded the normal lifespan. If the circumstances require, they may prudentially limit the resources available to marginal extensions of their lives. Thus, in some cases they may opt for rationing healthcare resources by age. In doing so, however, they do not value benefits differently merely because they come at different times in their lives. The prudential allocation of resources over a lifetime does not entail that later life years are discounted.¹⁸

Perhaps one could argue for a stronger claim. Prudential deliberators behind the veil of ignorance will realize that the resources necessary to protect their age-relative opportunity range differ at various stages of life. This is partly because the opportunities typical for various life-stages are different, partly because people face different risks at different life-stages,

¹⁸ See Daniels (1985: 104–7). As he notes, later life years may be discounted on a particular, and controversial, view of personal identity. I cannot address this issue here. See Parfit (1984), Daniels (1996) and McMahan (2002).

and partly because maintaining and restoring a person's health at a level which is appropriate for her age-relative range of opportunities require different resource expenditures. No doubt the institutional arrangement prudential deliberators would select will reflect these facts. Were the decision makers behind the veil, however, to assign decreasing value to later years, they would distort the trade-offs they have to make between life-stages. Rather than confronting the problem of resource allocation across time head-on, they would prejudice their judgments in favour of earlier life-stages. Once the veil is lifted and they reach their old age, they would have to confront the fact that the difficulty of protecting their range of opportunity is exacerbated because they selected institutions which include age-weighting in addition to other forms of rationing.

Daniels would not disagree: on his view, neutrality regarding time is a requirement of prudence. With respect to the choice behind the veil of ignorance, it requires that we do not give lesser weight to benefits which we receive at later stages of our life. The deliberators behind the veil are concerned with their well-being over their whole life (1988: 56–61). But even if prudence had no such requirement, the deliberators, it seems to me, would want to avoid distorting the trade-offs between stages of their lives by discounting later periods of life. Although a veil of ignorance argument does not exclude the possibility of an agreement on age-based rationing, it does not establish the case for weighting benefits by age.

5. CONCLUSION

I began this paper with empirical findings which ostensibly show that people value health benefits differently depending on the age of the beneficiary. Interpreting the reasons for people's age-weighting preferences is a difficult task, and there is no guarantee that we have understood them properly. I have examined two sorts of justification and argued that neither has been worked out in a satisfactory form. One justification was a defence of productivity ageism in the framework of DALYs. Although the idea that people's welfare contributions or productivity in its broadest sense should be taken into account when we evaluate the burdens associated with premature mortality and disability has initial plausibility, I argued that on closer inspection it becomes apparent that age has no place in such a measure. This is because it either fails to capture important aspects of productivity, or it introduces too much into measures which attempt to reflect only one dimension of well-being – the impact of health.

Another sort of justification appeals to the idea of fair innings. I argued that a fair innings principle is difficult to formulate coherently; moreover, on its more promising formulations, age turns out to be merely an indicator

of some other factor. Although these factors may be relevant to fairness, it remains unclear whether age is an appropriate indicator for them.

Nevertheless, there are several ways in which age may be relevant to the distribution of healthcare resources. Since young people tend to have better health prospects, policies which take into account the magnitude of benefits are likely to favour them. Treatments and interventions (including preventive services) which primarily target the young may also be more cost-effective, and a healthcare system which gives priority to them may be more efficient. Moreover, in virtue of having lived longer, older people – at least from an overall, lifetime perspective – may be better off than younger people, and if our aim is to give priority to the worse off, we will tend to give priority to the young. All of these factors may give an indirect, but pervasive role to age in healthcare resource allocation.

Thus, if different age-groups compete for scarce resources, there may be circumstances in which age is a relevant factor for resource allocation decisions. But we have to be careful in specifying what these circumstances are. By giving a default answer, age-weighting hinders rather than helps this task.

REFERENCES

- Broome, J. 1988. Good, Fairness and QALYs. In *Philosophy and medical welfare*, ed. J. M. Bell and S. Mendus, 57–73. Cambridge: Cambridge University Press.
- Broome, J. 1999. Fairness. In *Ethics out of economics*, 111–22. Cambridge: Cambridge University Press.
- Broome, J. 2002. Measuring the burden of disease by aggregating well-being. In *Summary measures of population health: Concepts, ethics, measurement, and applications*, ed. C. J. L. Murray, J. A. Salomon, C. D. Mathers and A. D. Lopez, 91–113. Geneva: World Health Organization.
- Busschbach, J. J. V., D. J. Hessing and F. Th. de Charro. 1993. The utility of health at different stages in life: A quantitative approach. *Social Science and Medicine* 37, 153–8.
- Callahan, D. 1987. *Setting limits: Medical goals in an aging society*. New York: Simon & Schuster.
- Cropper, M. L., S. K. Aydede and P. R. Portney. 1994. Preferences for life saving programs: How the public discounts time and age. *Journal of Risk and Uncertainty* 8, 243–65.
- Daniels, N. 1985. *Just health care*. Cambridge: Cambridge University Press.
- Daniels, N. 1988. *Am I my parents' keeper? An essay on justice between the young and the old*. New York: Oxford University Press.
- Daniels, N. 1996. Problems with prudence. In *Justice and justification: Reflective equilibrium in theory and practice*, 284–301. Cambridge: Cambridge University Press.
- Daniels, N., B. Kennedy and I. Kawachi. 2004. Health and inequality, or, why justice is good for our health. In *Public Health, Ethics, and Equity*, ed. S. Anand, F. Peter and A. Sen, 63–91. Oxford: Oxford University Press.
- Goerdt, A., J. P. Koplan, J.-M. Robine, M. C. Thuriaux and J. K. van Ginneken. 1996. Non-fatal health outcomes: Concepts, instruments and indicators. In *The global burden of disease: A comprehensive assessment of mortality and disability from diseases, injuries, and risk factors in 1990 and projected to 2020*, ed. C. J. L. Murray and A. D. Lopez, 99–116. Cambridge, MA: Harvard School of Public Health on behalf of the WHO and the World Bank.
- Harris, J. 1985. *The value of life*. London: Routledge & Kegan Paul.

- Harris, J. 1988. More and better justice. In *philosophy and medical welfare*, ed. J. M. Bell and S. Mendus, 75–96. Cambridge: Cambridge University Press.
- Hausman, D. M. 2007. What's wrong with health inequalities? *Journal of Political Philosophy* 15, 46–66.
- Johannesson, M. and P.-O. Johansson. 1997. Is the valuation of a QALY gained independent of age? Some empirical evidence. *Journal of Health Economics* 16, 589–99.
- Johri, M., L. J. Damschroder, B. J. Zikmund-Fisher and P. A. Ubel. 2005. The importance of age in allocating health care resources: Does intervention-type matter? *Health Economics* 14, 669–78.
- Kamm, F. M. 1993. *Morality, mortality. Volume I: Death and whom to save from it*. New York: Oxford University Press.
- Kamm, F. M. 2002. Health and equity. In *Summary measures of population health: Concepts, ethics, measurement, and applications*, ed. C. J. L. Murray, J. A. Salomon, C. D. Mathers and A. D. Lopez, 685–706. Geneva: World Health Organization.
- Kappel, K. and P. Sandøe. 1992. QALYs, age and fairness. *Bioethics* 6, 297–316.
- Lewis, P. A. and M. Charny. 1989. Which of two individuals do you treat when only their ages are different and you can't treat both? *Journal of Medical Ethics* 15, 28–32.
- Lockwood, M. 1988. Quality of life and resource allocation. In *Philosophy and Medical Welfare*, ed. J. M. Bell and S. Mendus, 33–55. Cambridge: Cambridge University Press.
- McKie, J., J. Richardson, P. Singer and H. Kuhse. 1998. *The allocation of health care resources: An ethical evaluation of the 'QALY' approach*. Aldershot: Ashgate.
- McMahan, J. 2002. *The ethics of killing: Problems at the margins of life*. Oxford: Oxford University Press.
- Murray, C. J. L. 1996. Rethinking DALYs. In *The global burden of disease: A comprehensive assessment of mortality and disability from diseases, injuries, and risk factors in 1990 and Projected to 2020*, ed. C. J. L. Murray, and A. D. Lopez, 1–98. Cambridge, MA: Harvard School of Public Health on behalf of the WHO and the World Bank.
- Murray, C. J. L. and A. K. Acharya. 1997. Understanding DALYs. *Journal of Health Economics* 16, 703–30.
- Murray, C. J. L., J. A. Salomon, C. D. Mathers and A. D. Lopez. 2002. Summary measures of population health: Conclusions and recommendations. In *Summary measures of population health: concepts, ethics, measurement, and applications*, 731–56. Geneva: World Health Organization.
- Nord, E. 1999. *Cost-value analysis in health care: Making sense out of QALYs*. Cambridge: Cambridge University Press.
- Nord, E., A. Street, J. Richardson, H. Kuhse and P. Singer. 1996. The significance of age and duration of effect in social evaluation of health care. *Health Care Analysis* 4, 103–11.
- Parfit, D. 1984. *Reasons and persons*. Oxford: Clarendon Press.
- Rodríguez, E., and J. L. Pinto. 2000. The social value of health programs: is age a relevant factor? *Health Economics* 9, 611–21.
- Tsuchiya, A. 1999. Age-related preferences and age weighting health benefits. *Social Science and Medicine* 48, 267–76.
- Tsuchiya, A., P. Dolan and R. Shaw. 2003. Measuring people's preferences regarding ageism in health: some methodological issues and some fresh evidence. *Social Science and Medicine* 57, 687–96.
- Veatch, R. M. 2000. *Transplantation ethics*. Washington, DC: Georgetown University Press.
- Williams, A. 1997. Intergenerational equity: an exploration of the 'Fair Innings' argument. *Health Economics* 6, 117–32.