Priority Setting and Age

Greg Bognar

1 Introduction

The role of age in priority setting is one of the most controversial issues in health policy. It has also been a contentious topic for many years in medical ethics and philosophy, and any discussion of age as a criterion for setting priorities in health care is likely to stir up intense public debate. Age is an easily observable characteristic; hence it is tempting to use it when priorities must be set between different resource uses or patient groups. Indeed, age considerations pervade health systems worldwide. Consequently, there is an urgent need to clarify the role that age can play in health care resource allocation.

Against this background, there have been surprisingly few systematic studies of the justifications for using age as a criterion in priority setting. In this chapter, I provide a broad outline of the range of ideas that have been used to defend the relevance of age. At the end of the chapter, I also reflect on a recent public debate on the role of age in priority setting.

Priority setting is a highly sensitive issue in every country where it comes onto the agenda. Suggestions for using age as a priority-setting criterion are particularly controversial. Part of the reason is that the issue is relatively easy to communicate to the public and to discuss in the media without the use of technical language. As a result, the discussions tend to be rather black and white. The problem of age is usually presented as the narrow question whether age has any legitimate role in priority setting at all. This obscures the fact that few people seem to be opposed to age-based priority setting categorically in every possible setting. Most people are willing to use age as a criterion at least in *some* circumstances and at least in *some* ways. This is shown by empirical studies of public views on priority setting: while

G. Bognar
Stockholm Centre for Healthcare Ethics (CHE) and Department of Philosophy, Stockholm University, Universitetsvägen 10D, Stockholm SE-10 691, Sweden
e-mail: greg.bognar@philosophy.su.se

E. Nagel, M. Lauerer (eds.), *Prioritization in Medicine: An International Dialogue*, DOI 10.1007/978-3-319-21112-1_13

the survey questions and methodologies vary considerably, most studies suggest that a majority accepts that age can have some role in priority setting.[1]

At the same time, a few countries have issued national policy documents with explicit guidance for priority setting in the health sector. Among these countries are Norway, Sweden, and the UK.[2] For the most part, these kinds of guidelines deal with age in an ambiguous manner. For example, while they generally warn against age discrimination and letting age influence priorities, they also recommend criteria that are closely associated with age and discuss circumstances in which age may be a legitimate consideration.[3]

In addition, age is frequently used as an indicator in actual clinical practice.[4] This is partly because chronological age is an objective, easily observable measure, and it is correlated—or at least it is perceived to be correlated—with many factors that are deemed relevant for decision making in the health care sector. For instance, age can be an indicator for the risk of contracting a disease, for the expected severity of the natural course of a disease, for the risk of adverse events from treatment, for the probability of successful treatment, for the duration of benefit if treatment is successful, and so on. Decisions based on such factors can have a profound impact on particular patients or patient groups, yet the role of age in these decisions is seldom explicit. Indeed, even clinicians expressing disagreement with age-based priority setting may themselves use age inadvertently in practice!

Age enters priority setting in numerous other ways. For instance, it influences cost-effectiveness studies as well as disease burden estimates. At the end of the day, considerations of age can pervasively shape policy and practice, having a profound impact on who gets what, when, and how in health care. It is, therefore, a necessary and urgent task to clarify the role of age from an ethical point of view.

At the most fundamental level, age can play two kinds of role in priority setting in health care. On the one hand, it can have a *direct* role in a particular proposal, policy, or set of guidelines for setting priorities. It can be used as an independent criterion. In order to justify its direct role, it must be shown that age in itself is a morally relevant consideration, and hence it is one of the factors that must be taken into account in the deliberation about the use of available resources. Still, it may not

[1] See, for instance, Cropper et al. (1994), Nord et al. (1996), Johannesson and Johansson (1997), Tsuchiya (1999), Eisenberg et al. (2011), Olsen (2013), and Petrou et al. (2013). Among the empirical studies on the views of health personnel, some find support for age-based priority setting and others not (see, e.g., Neuberger et al. (1998), Ryynänen et al. (1999), and Werntoft and Edberg (2009)). At the same time, there is considerable, and ongoing, controversy regarding the question of just how the results of these studies should be interpreted (see Dey and Fraser (2000), Tsuchiya (2000), Tsuchiya et al. (2003), Bognar (2008), and Whitty et al. (2014)). The issue is complex, and the very opposite of black and white!

[2] See Sabik and Lie (2008).

[3] See, for instance, SOU (1995:5) and NOU (1997:18). Moreover, explicit age limits abound in both international and national clinical guidelines. For example, this is the case for the guidelines issued by the European Society of Cardiology (2015) and many guidelines for cancer care and screening (e.g., American Cancer Society (2015)).

[4] See Kapp (1998).

be—and it usually isn't—the only relevant factor, and it may not provide the decisive consideration. But it cannot be ignored.

This implies that when a set of guidelines for priority setting recommend age as an independent criterion, then if two patients are equal on every other criterion, a difference in their age will result in priority being assigned to one over the other.

On the other hand, age can have an *indirect* role in priority setting. Rather than serving as an independent priority-setting criterion, it can be used as an indicator or proxy for some other, morally relevant factor. In this case, it enters the deliberation on setting priorities indirectly, representing some other factor that is not measurable or tractable in any better way. For instance, age may be a rough indicator of expected health benefit: other things being equal, a younger patient may benefit more from some intervention. Thus, it might be argued that she should be given priority for some scarce resource. But this isn't because she is younger; it is because this is the way to maximize expected benefit. Even though age is used here as a criterion of *decision making*, it is not used as a criterion for the *moral justification* of the decision. The moral justification of the decision is provided by considerations of benefit maximization.

Introducing these distinctions has the advantage of getting to the crux of the moral issue right away. Any proposal to use age as a criterion for priority setting must explain whether it considers age a morally relevant consideration in itself, or it regards it only as an appropriate indicator for some other morally relevant factor. Priority setting is the application of moral theories and moral principles for resource allocation in health care. So it must be based on defensible ethical arguments. Since age-based considerations already permeate decision making in the health care sector, we cannot avoid addressing this issue.

2 Indirect Views

Suppose you are the only surgeon in the emergency room when two patients are brought in. They both need immediate life-saving surgery, but you can only operate on one of them. There is no relevant difference between the two patients (or at least you are unaware of any), except that one is 30 years old and the other is 70 years old. The surgery can restore both of them to full health and would not decrease their life expectancy. In the rest of their lives, their quality of life would be equally high. Which of these two patients should you save?

Most people agree that the younger patient should be saved.[5] There are different ways to justify this choice. For instance, you might argue simply that saving the younger person is more likely to maximize benefits: since the 30 year old person can expect to survive longer into the future than the 70 year old, saving her life does more good. This is a consequentialist justification: it appeals to the value of the outcomes. The best-known consequentialist moral view is *utilitarianism*. Assuming that the 30

[5] See, for instance, Cropper et al. (1994), Nord et al. (1996), and Johannesson and Johansson (1997), for studies in the USA, Australia, and Sweden, respectively.

years old would survive for many more years, and the 70 years old would survive for only a few, utilitarians would agree that the younger patient should be saved.

This sort of consequentialist justification has been called *utilitarian ageism*.[6] But utilitarianism is not really concerned with age. In utilitarian ageism, both the "utilitarianism" and the "ageism" bits are slightly misleading. For one thing, what matters in this example for utilitarianism is not age, but life expectancy. The younger patient should get priority because she can expect to live longer. In utilitarian ageism, age is primarily an indicator for period life expectancy. (Period life expectancy is life expectancy at different ages.) Since life expectancy typically decreases as people grow older, age can be a rough indicator of the magnitude of benefit from treatment.

But ultimately what utilitarianism is concerned with isn't life expectancy either. Rather, it is well-being. The fact that the younger person has a greater life expectancy is relevant only insofar as life expectancy itself can be taken as an indicator—in this case, a rough indicator of expected well-being. Therefore, utilitarian ageism gives a double role to age: it is an indicator of life expectancy, which is itself an indicator of well-being. But in both, the role of age is indirect. It provides no independent consideration. In itself, it is morally irrelevant that one patient is 30 and the other is 70.

I will refer as *indirect views* to justifications that seek to establish an indirect role for age in priority setting. Utilitarian ageism is one example.

Although utilitarian ageism can justify the choice of saving the younger person in the example, many philosophers find it problematic. They point out that utilitarian ageism inherits the problems of utilitarianism in general. One of these problems is that utilitarianism is insensitive to the distribution of benefits. If you can provide the same improvement in well-being to a person who is badly off or to another person who is well off, then utilitarianism does not imply that you should benefit the worse off person. Since the size of the benefit is the same, benefiting the well off person is just as good as benefiting the badly off person. But many people would agree that you should benefit the worse off person. You ought to be concerned with the inequality between the two people.

The problem for utilitarian ageism can be illustrated by slightly modifying our example. Suppose that the two patients who are brought into the emergency room would be able to survive for 10 years only. That is, the 30 year old patient will die at 40 if you save her life, and the 70 year old patient will die at 80 if you save her life. Suppose also that in the remaining 10 years their lives would be equally good. Thus, the size of the benefit for these two patients is the same.

Utilitarian ageism implies that saving the life of the 30 year old patient is just as good as saving the life of the 70 year old patient. But for many, this is not the right implication. They would insist that even in this case the 30 year old patient should be given priority. It would be *unfair* not to choose the 30 year old patient.

Obviously, those who take this view need to explain why not saving the younger patient straightaway is unfair. They can choose from competing theories of fairness. One explanation may be that the unfairness is due to the *inequality* between the lives

[6] See Nord et al. (1996).

of the two people: it is unfair to the 30 year old patient, who is worse off in terms of lifetime well-being, if the 70 year old, who is already better off, is saved. It is unfair to increase the inequality between the two patients.

The sort of moral justification that appeals to inequality has been influential in health economics and health policy in the form of the "fair innings" argument.[7] It is a well-known fact that life expectancy at birth varies with social position: the better off you are, the greater your life expectancy. This remains true when the quality of life is also taken into account. That is, the better off have greater *quality-adjusted life expectancies* (QALEs), both at birth and later, than the worse off, poorer members of society. According to the fair innings argument, these inequalities must be addressed.

One way to address these inequalities is to try to increase the QALEs of the worse off. For instance, the mean quality-adjusted life expectancy at birth may be considered as a sort of threshold—commonly called the *fair innings threshold*. Health policy can then be designed around this threshold in order to reduce the discrepancy in QALEs between the better off and the worse off socioeconomic classes or groups. The intuitive idea is that everyone, regardless of their initial position in society, should have an equal chance to live a long and healthy life. It is unfair if people's QALEs differ merely because of the circumstances of their birth.

Since the better off have greater QALEs, health disparities can be reduced by selecting policies which equalize QALEs by sacrificing some of the overall health (or longevity) of the population. Quality-adjusted life expectancy can be used as a measure for the overall health of the population and for the inequality in health within the population. Additional years of life can be given different weights according to how well off people are: additional years to the better off have smaller weights than additional years to the worse off. Thus, this view can help quantify the equity-efficiency trade-offs between population health and equality in health—or between benefit maximization and fairness.

It should be clear that the role of age in the fair innings argument is indirect. It is merely an indicator that can help design policies to reduce unfair inequalities. In itself, age is not a morally relevant consideration.[8]

The fair innings argument focuses on inequality between full lives. It takes a whole-life perspective. Many philosophers accept that this is the right perspective to take when it comes to fairness. A person may be badly off now, but she may become better off later on. Perhaps she is badly off now only because she has sacrificed some of her current well-being for greater well-being later in her life. This is the *argument from compensation*: a person who is badly off at some time can be compensated by advantages at some other time in her life. In such circumstances, the inequality that obtains between her and others at some particular time is not necessarily unfair. Because of this, you need to take a whole-life perspective for assessing the inequality. Inequality is a concern between full lives.[9]

[7] See Williams (1997).

[8] Views that are similar in important respects are proposed by Ottersen (2013) and NOU (2014:12).

[9] For the argument, see Nagel (1979).

This argument, however, implies that you need not be concerned about inequalities that obtain between people at particular times, at least as far as their lives are equally good overall. In this case, it is not unfair that some of them right now live in poverty and poor health while others are healthy and affluent. But this implication is troubling. For instance, it suggests that it should not be a matter of concern if the elderly now live in poor conditions with inadequate health care as long as they used to be sufficiently well off earlier in their lives such that there is no overall inequality between them and others. Or it should not be a matter of concern if the children alive today get a bad start in life as long as their lives get better later on to make up for their current deprivation in terms of equality between full lives. The whole-life perspective ignores inequalities between people at different ages, or stages, of their lives. This seems wrong.

There are different strategies that try to avoid this implication. One is to give up the whole-life perspective or at least amend it with some other principle for the allocation of resources that applies to particular times or time periods. So you might accept, say, one moral principle that aims for equality between full lives and another moral principle that applies to particular times. One proposed candidate for the latter is the *time-specific priority view*.[10] The priority view, or *prioritarianism*, holds that the right course of action or policy is that which maximizes weighted well-being, where the weights are given by a function that increases with higher levels of well-being at a decreasing rate. In practice, this means that a given benefit will have greater value if it goes to a person who is worse off, and the worse off a person is, the greater the value of the same benefit.[11] Time-specific prioritarianism differs from the "standard" version in that it considers the well-being of a person at a particular time or stage of life, whereas the latter considers overall lifetime well-being. Thus, time-specific prioritarianism takes a *sub-lifetime*, rather than a whole-life, perspective. Hence on this proposal, sub-lifetime shortfalls in well-being are not ignored: the time-specific priority view directs you to make people who are worse off at a particular time better off.

On the time-specific priority view, it becomes unfair if the elderly live in poor conditions, even if they are just as well off as others in terms of their lifetime well-being. And it is unfair if children get a bad start in life even if their lives get better later on to make up for their deprivation.

However, it is worth noting that prioritarianism is concerned with shortfalls in well-being in absolute terms; strictly speaking, it is not concerned with how people fare compared to one another. Therefore, inequalities are unfair on prioritarianism because the badly off are badly off in absolute terms, not merely because they are worse off than others. Prioritarianism is an egalitarian view only in this broader sense. It differs from "standard" egalitarianism in that it is not comparative.

There are many questions that proposals that separate the whole-life and sub-lifetime perspectives have to answer. What is the relation between the principles for the allocation of resources between full lives and the principles that apply to specific

[10] It is proposed by McKerlie (2013).

[11] See Parfit (1995).

times? What happens if their recommendations are in conflict? Can theories that combine the whole-life and the sub-lifetime perspectives remain coherent?

Here I have to set these questions aside, but I do want to make a couple of points. First, note that there is no entailment between the whole-life and sub-lifetime perspectives on the one hand and particular principles of resource allocation on the other. No principle seems to fit better one or the other perspective, and vice versa. The proposal I have examined is to take an egalitarian view when it comes to full lives and a prioritarian view when it comes to particular times. But other combinations are equally possible: you can accept egalitarianism, prioritarianism, or indeed utilitarianism, on both levels, or any combination of two of these (or other) views on the whole-life and sub-lifetime levels. To be sure, not all of these combinations will be equally plausible. But the distinction between the two perspectives and the differences between the moral principles are independent of one another.

The other remark I want to make is to highlight that none of the views discussed so far takes age into account directly. Principles that are proposed for the sub-lifetime perspective, including the time-specific priority view, apply to inequalities or disadvantages at particular times. They are not concerned with age unless age can be taken as an indicator of disadvantage or shortfall in well-being. The sub-lifetime perspective does not make age in itself relevant, even if inequalities at particular times often take the form of inequalities between people from different age groups. Plainly, inequalities at particular times occur between people from the same age groups just as well. The issues between the whole-life and the sub-lifetime perspectives are orthogonal to the problem of the role of age as a criterion for priority setting.

Naturally, this will leave those who believe that age in itself is morally relevant unsatisfied. They argue that the moral issue between saving the 30 year old and the 70 year old does not turn merely on who has had more or less well-being throughout their life or at particular times. They hold that even if the 30 year old has already had an overall better life, it might still be unfair not to save her. According to this view, age is not merely an indicator for some other factor. It is morally relevant in itself.

3 Direct Views

I will refer as *direct views* to justifications that seek to establish a direct role for age in priority setting. When age has a direct role, it provides an independent moral consideration, to be taken into account with others, in health care resource allocation. On these views, an additional unit of time can have different values depending on the age of the person who receives it.

I have argued that the whole-life perspective has the implication that inequalities between people or shortfalls in well-being at particular times are ignored. If inequalities at particular times are ignored, then inequalities between people in different age groups will be ignored. Thus, for instance, if there is persistent inequality between the old and the young or children and the middle-aged, then they will be revealed as inequalities at particular times. The whole-life perspective will ignore them just as it ignores inequalities at particular times.

One strategy to avoid these implications was to introduce principles of resource allocation for the sub-lifetime perspective in addition to principles that apply to full lives. But, as I pointed out in the last section, there is a worry that such views just lead to inconsistencies. Another strategy is to think of the difference between the whole-life and the sub-lifetime perspectives in a different way. The proposal is that the problem of resource allocation between different people at different stages of their lives can be analyzed in terms of resource allocation within the life of a single person. You can consider how a person would prudently allocate resources for her full life and derive principles of fairness that apply to parts of life. This is the strategy followed by the *prudential lifespan account*.[12]

The strategy is inspired by the following thought. In real life, we all make trade-offs between different times in our lives—we all make *intrapersonal* trade-offs. (The most common example is saving: sacrificing some amount of present well-being in order to promote well-being in the future.) These trade-offs should be prudent: rational and without bias toward the near future. So we can use the idea of prudential *intrapersonal* trade-offs to guide our views about permissible *interpersonal* trade-offs. By prudently allocating resources over your life, you maximize your well-being over your lifetime. By allocating resources over different life stages of different people in a similar manner, you maximize the well-being of all the people involved. That is, if interpersonal resource allocation is designed analogously to intrapersonal resource allocation, it will make everyone as well off as possible over their full lives. And when interpersonal trade-offs are designed analogously to intrapersonal trade-offs, everyone is treated equally over their full lives.

Thus, for example, if it is prudent to give more weight to flourishing in your middle years as opposed to your old age, then it is justified to give priority to benefiting people in their middle years rather than in their old age. The old cannot complain, since they had priority when they were in their middle years. There is no unfairness. If people would rationally prefer to have access to life-saving resources when they are 30 years old rather than when they are 70, then it is not unfair to use life-saving resources to save a 30 year old rather than a 70 year old. In sum, if it is *prudent* to allocate resources in a particular way within one life, then it becomes, on this view, *fair* to allocate resources in the corresponding way between different people.

This is a different way of thinking about fairness. In the prudential lifespan account, fairness is not a matter of inequality or disadvantage. It is a matter, instead, of prudential (or rational) justifiability to each person.

To be sure, the prudential lifespan account needs to be formulated on an abstract level to yield useful conclusions. A thought experiment can help here. You can imagine that a rational person tries to determine how she should allocate a fixed amount of resources over her full life. To do this, she must ignore her present age, and she must assume that she will live through all life stages. So the person should be placed behind a "veil of ignorance."[13] Otherwise, the solution to the allocation problem could not be generalized.

[12] This account is introduced by Daniels (1988, 2008).

[13] The thought experiment is borrowed from Rawls (1971).

The trade-offs that rational people would accept behind the veil of ignorance can be expected to maximize their lifetime well-being. They can be only *expected to* do so, since the lives of different people will in fact go differently. Some people, for instance, will die prematurely. So once the veil is lifted, people will end up in different positions. Nevertheless, the idea is that it is not unfair to set priorities this way, because no person, if placed in an impartial situation behind the veil of ignorance, could object to the principles that determine the trade-offs between different age groups. The principles can be justified to each person.

What sort of trade-offs would people agree to behind the veil of ignorance? They might agree, for instance, that it would be better if fewer resources are spent on the very old when those resources can be spent on benefiting the young. Thus, they might agree that different age groups should be entitled to different amounts of resources. In this way, age becomes directly relevant. The prudential lifespan account provides one kind of justification to use age as an independent criterion in priority setting.

Should we accept the prudential lifespan account? Some considerations suggest that it is less useful for providing guidance in priority setting than it might initially seem. One criticism of the account is that it would leave too little for the elderly and especially for the very old. This might be in conflict with our moral intuitions. People behind the veil might give less priority to benefits at extreme old age, since they have to distribute a fixed amount of resources and it makes sense to make sure you have enough at earlier life stages.[14]

Perhaps that is so—but it is hard to say. In my view, the main problem is that it is difficult to come to definite conclusions from the thought experiment involved in the prudential lifespan account. How would you distribute resources over your life? How much would you leave for extreme old age? Well, the only answer, it seems to me, is that *it depends*—it depends, for instance, on how much you are supposed to be able to distribute. If the resources are sufficient, you might want to allocate the same amount to every life stage. If there are fewer resources, perhaps you would consider good health to be more important at particular life stages: in young adulthood, for example, when most people are responsible for young children, or maybe in early childhood, in order to have a good start in life. So perhaps you would be willing to make trade-offs between life stages.

The answers crucially depend on the assumptions that the prudential lifespan account makes. Behind the veil of ignorance, you must assume that you will live through every life stage. You must assume there is no premature mortality. But if you knew that there was some probability of dying at each life stage, you would likely be willing to make different trade-offs. You may be willing to accept more risk at some life stages in exchange for higher well-being at others. The assumption that there is no premature mortality drastically limits the usefulness of the prudential lifespan account. But if the assumption is dropped, it's impossible to draw any specific conclusions from the thought experiment.

[14] This objection is made by McKerlie (2013), among others.

Another, related assumption is that you must allocate a fixed amount of resources behind the veil. Your share of resources is determined independently of the allocation problem. But it's hard to see what justifies this assumption. If you know that there is premature mortality, you will recognize that some people will die before they have used up their full share of resources. Why should those resources not be redistributed and added to the shares of those who survive? (After all, this would be better than wasting those resources.) Once again, the assumption drastically limits the usefulness of the account, but it seems impossible to draw any specific conclusions from the thought experiment in its absence.

At this stage, it is tempting to return to a less complex account of the role of age in priority setting. Recall the example of the 30 year old and the 70 year old patients in the emergency room. Many people agree that it is unfair if you do not choose to save the life of the 30 year old. We have been looking at different attempts to explain the unfairness. One idea was that the unfairness is due to inequality: it is unfair if the younger patient ends up with a much shorter life than the older person. This conception of fairness is comparative. Another conception tied fairness to justifiability to each person. Putting limits on resources is not unfair, on this view, just in case rational individuals taking an impartial perspective would agree to do so. This conception of fairness is not comparative. However, it was difficult to derive any specific conclusions from it.

The less complex account that I am about to introduce is based on yet another conception of fairness. In the example of the 30 years old and the 70 years old, it would be unfair not to save the 30 years old, not because she would end up living less than the 70 years old, but because 30 years is not enough to have a *complete life*. It would be equally unfair not to save the 30 year old patient even if there was no 70 year old patient that she has to compete with for a scarce resource. This sort of unfairness is not comparative.

This view has also been called the *fair innings argument*.[15] (It is not the same view, however, as the one I discussed in the previous section. They are different views that, confusingly, go by the same name. It is also confusing that they are both called *arguments*, rather than specific views on how resources should be distributed, even though that's what they are.) It is based on the idea that there is a length of life that should be considered a full or complete or reasonable lifespan. A complete life lasts long enough to contain all the most important experiences of life: growing up, finishing your education, falling in love, building a career, starting a family, and see your children grow up and start families on their own. Suppose that for such a complete life, 70 years are necessary. It is a tragedy to die younger than this (and the younger one dies, the more tragic the death is), but it is not a tragedy to die once you have reached this threshold. Therefore, when scarce resources must be allocated, you should make sure that people reach this age: it is their *fair innings threshold*. Over this age, their claims on society's resources diminish.

[15] It is introduced, although not unequivocally accepted, by Harris (1985). For a view that is similar in many respects, see Callahan (1987).

This view is able to explain the common moral judgment in the examples about the 30 years old and the 70 years old. It is unfair to save the older patient because she has already had her "fair innings." She has reached the threshold. If you cannot save both of the patients, you should save the patient who otherwise would not have a complete life. This is true even in the variant of the example where the two patients can only survive for another 10 years. Since the older patient has reached the threshold, the younger patient should be saved, even if she cannot herself reach the threshold.

To be sure, it would be unfair, according to this view, to give priority to one patient over another when neither has reached the fair innings threshold; and it would be unfair to prefer one patient over another when both patients have reached the fair innings threshold. The fair innings argument applies to conditions of resource scarcity between the young and the old who have reached the threshold.

In practice, the view would imply, for instance, that people over the fair innings threshold become ineligible for costly life-extending treatment. In their case, medical care should focus on palliative care and the maintenance of quality of life. Costly curative and life-extending treatments should be provided only to those who have not reached the threshold. (Perhaps they could be provided to those over the threshold when all other claims have been satisfied, but this is usually unlikely to be the case, given the facts about resource scarcity.)

Of all the views discussed so far, the fair innings argument gives the most central role to age. On this view, age is not merely an indicator for a complete life; rather, the notion of a complete life is defined in terms of age. (This is so even if, as defenders of the view might point out, the fair innings threshold need not be a particular age—the threshold could be left somewhat vague or defined as a range to allow for some individual variation.) Having a complete life takes a certain amount of time. No 30 year old can have all the experiences that normally make up a complete life. The fair innings argument is not about having enough lifetime well-being, but about having enough time for a complete life.

Nevertheless, the fair innings argument is not without problems. One immediate question is why there should be a fair innings threshold at all. Why not hold, instead, that those who have lived longer should have relatively less priority across all ages?[16] If age can make a difference to what is fair when resources must be allocated between the young and the old, why shouldn't it be relevant each time when people from different age groups compete for resources? If you have to choose between saving the life of a 30 years old and a 40 years old, why should the younger patient not have priority?

The fair innings threshold has been defended by an analogy.[17] Suppose two people are given the chance to run a mile, which most people can do in 7 minutes. One of these people is given only 3 minutes and the other is given only 4. In this case, it is not true that the second person is given a fairer running time than the first person:

[16] This sort of view is proposed, for instance, by Lockwood (1988).

[17] See Harris (1985: 92–93).

it is just as impossible to run a mile in 4 minutes as it is in 3 minutes. The unfairness is the same to both of these people.[18]

It is not hard to see how the analogy is meant to work. According to the fair innings argument, if your choice is between saving a 30 years old and a 40 years old, it is unfair to give priority to the younger patient. Just like in the case of the runners, it is equally unfair if they cannot reach their fair innings. Therefore, there is no justification for saving the 30 years old straightaway (as opposed to, for instance, giving them equal chances by tossing a coin to decide whom to save).

The problem with this defense is that it provides no independent argument for the fair innings threshold. Even though neither of the runners can hope to finish the whole mile in their time, they might value the ground they can cover in their allotted time. They might prefer to have as long as possible. So the first person does have, it seems, a stronger complaint in comparison to the other. Similarly, if what is valuable is to have the most important experiences that a complete life can offer, then the 30 years old has undeniably had less of a chance for a complete life. It is not implausible to argue that she should have priority.

The fair innings argument cannot account for this judgment. In order to do that, we can reintroduce prioritarianism, albeit in a formulation that differs from that which I gave above. Recall that prioritarianism is the view that the right course of action or policy is that which maximizes weighted well-being, where the weights are determined by a function that increases with higher levels of well-being at a decreasing rate. As I explained, this means that a given benefit has greater value if it goes to a person who is worse off, and the worse off a person is, the greater the value of the same benefit. Now instead of well-being, prioritarianism can be applied to life-years. In this application, an additional year has greater value if it goes to a person who is worse off in terms of years of life—that is, younger—and the younger the person is, the greater the value of the additional year of life.[19]

This view can justify the judgments that many people have in the cases that we have discussed. It implies that when you have to choose between saving a 30 years old or a 70 years old, you should save the 30 years old, even when each of them can survive for only 10 more years. And when you have to choose between saving a 30 years old or a 40 years old, the view also implies that you should save the 30 years old. By attributing different values to additional years of life, prioritarianism applied to life-years gives a central role to age. This view might provide the best basis for such a role.

4 Priority Setting and Age in Practice

The recent experience of Norway provides an illustration of the controversy over the use of age in health policy[20]. In November 2014, the third Official Committee on Priority Setting in the Health Sector presented its report, laying out a new,

[18] In fact, some people can run a mile within four minutes. I will follow Harris in ignoring this complication here.

[19] This view is defended in Bognar (2015). See also Bognar and Hirose (2014).

[20] See also chapter "Recent Developments on the Issue of Health-Care Priority Setting in Norway".

comprehensive framework for setting priorities in Norway.[21] As part of this framework, three new criteria were proposed: first, a health-benefit criterion, according to which the priority of an intervention increases when the expected health benefits (and other relevant welfare benefits from the intervention) are greater; second, a resource criterion, according to which the fewer resources an intervention requires, the greater its priority; and, third, a health-loss criterion, according to which the priority of an intervention increases when the expected lifetime health loss of the beneficiary is greater. The committee emphasized that these criteria must be considered together and recommended that they are applied throughout the health sector.

In its mandate, the committee was specifically asked to consider whether age should have "intrinsic value" in priority setting: that is, whether it is morally relevant in itself and could be used as an explicit, independent criterion. The committee concluded that age should not serve as an independent criterion. To support its conclusion, the committee argued that the relevant concerns indicated by age are already taken into account by the three proposed criteria. In other words, the committee denied that age in itself is morally relevant. At the same time, the committee agreed that age can legitimately influence priorities through its proposed criteria. It agreed that it may influence priorities through the health-benefit criterion, for example, due to the correlation between age and the risk of contracting disease, the risk of increased severity of disease, the risk of adverse events from treatment, and so on. The committee also agreed that age may correlate with decreased health loss, although it emphasized that the correlation may often be weak, since many other factors can determine the magnitude of health loss.

The report attracted considerable attention in the national media and generated a lively debate. The question of age was central from the outset. The day the report was released, Norway's largest newspaper featured a 23 year old patient with multiple sclerosis on the front page and declared that the committee recommended that "young people should be prioritized over the elderly in the health queue." In the same issue, a 72 year old man was reported to find the recommendations "unfair" and "discriminatory."[22] Age has continued to be in the limelight in the debate over the report. It has frequently been claimed that age should never influence priorities, often accompanied by the erroneous claim that age has not played any role in priority setting before. More nuanced positions have also been put forward. For example, it has been argued that age can be relevant to priority setting at the macro level and for preventive measures, while it should be irrelevant at the clinical level and for curative services. Similarly, it has been argued that children should have priority over adults, while age should not influence priority among adults, or at least not when the difference in age is small. It has also been argued that the three proposed criteria, put together, allow age to have too much influence.

The debate in Norway has demonstrated once again how controversial the issue of age can be—even in a country with a long tradition of systematic priority setting

[21] NOU (2014:12).

[22] Dommerud and Olsen (2014).

and robust public debate. People sharply disagree on how age should influence priority setting. Their sharp disagreements are reflected in the philosophical debate on the justification of the use of age in resource allocation. As I have tried to show in this chapter, there is a wide range of ideas and proposals in this area. The issues raised by age are likely to remain for a long time on both the philosophical and the public agenda.

Acknowledgments Financial support from the Swedish Research Council and from the Swedish Research Council for Health, Working Life and Welfare is gratefully acknowledged (2014–4024). I would also like to thank an audience at the KTH Royal Institute of Technology (Stockholm) for valuable comments.

References

American Cancer Society (2015) Screening recommendations by age. http://www.cancer.org/healthy/toolsandcalculators/reminders/screening-recommendations-by-age. Accessed 9 Apr 2015

Bognar G (2008) Age-weighting. Econ Philos 24:167–189

Bognar G (2015) Fair innings. Bioethics 29:251–261

Bognar G, Hirose I (2014) The ethics of health care rationing: an introduction. Routledge, New York

Callahan D (1987) Setting limits: medical goals in an aging society. Simon & Schuster, New York

Cropper ML, Aydede SK, Portney PR (1994) Preferences for life saving programs: how the public discounts time and age. J Risk Uncertain 8:243–265

Daniels N (1988) Am I my parents' keeper? An essay on justice between the young and the old. Oxford University Press, New York

Daniels N (2008) Justice between adjacent generations: further thoughts. J Polit Philos 16:475–494

Dey I, Fraser N (2000) Age-based rationing in the allocation of health care. J Aging Health 12:511–537

Dommerud T, Olsen T (2014) Prioriteringsutvalget: 'De Som Taper Flest Leveår, Bør Få Mest'. Aftenposten, 12 Nov 2014, pp 4–5

Eisenberg D, Freed GL, Davis MM, Singer D, Prosser LA (2011) Valuing health at different ages: evidence from a nationally representative survey in the US. Appl Health Econ Health Policy 9:149–156

European Society of Cardiology (2015) ESC clinical practice guidelines list. http://www.escardio.org/Guidelines-&-Education/Clinical-Practice-Guidelines/ESC-Clinical-Practice-Guidelines-list. Accessed 9 Apr 2015

Harris J (1985) The value of life. Routledge & Kegan Paul, London

Johannesson M, Johansson P-O (1997) Is the valuation of a QALY gained independent of age? Some empirical evidence. J Health Econ 16:589–599

Kapp MB (1998) De facto healthcare rationing by age. J Legal Med 19:323–349

Lockwood M (1988) Quality of life and resource allocation. In: Bell JM, Mendus S (eds) Philosophy and medical welfare. Cambridge University Press, Cambridge, pp 33–55

McKerlie D (2013) Justice between the young and old. Oxford University Press, New York

Nagel T (1979) Equality. In: Mortal questions. Cambridge University Press, Cambridge, pp 106–127

Neuberger J, Adams D, MacMaster P, Maidment A, Speed M (1998) Assessing priorities for allocation of donor liver grafts: survey of public and clinicians. Br Med J 317:172–175

Nord E, Street A, Richardson J, Kuhse H, Singer P (1996) The significance of age and duration of effect in social evaluation of health care. Health Care Anal 4:103–111
NOU (1997) Prioritering på Ny: Gjennomgang Av Retningslinjer for Prioriteringer Innen Norsk Helsetjeneste. Statens forvaltningstjeneste, Oslo
NOU (2014) Åpent og rettferdig – Prioriteringer i Helsetjenesten. Departementenes sikkerhets- og serviceorganisasjon, Oslo
Olsen JA (2013) Priority preferences: 'end of life' does not matter, but total life does. Value Health 16:1063–1066
Ottersen T (2013) Lifetime QALY prioritarianism in priority setting. J Med Ethics 39:175–180
Parfit D (1995) Equality or priority? The Lindley Lecture, Department of Philosophy, University of Kansas, Lawrence
Petrou S, Kandala NB, Robinson A, Baker R (2013) A person trade-off study to estimate age-related weights for health gains in economic evaluation. Pharmacoeconomics 31:893–907
Rawls J (1971) A theory of justice. Harvard University Press, Cambridge, MA
Ryynänen O-P, Myllykangas M, Kinnunen J, Takala J (1999) Attitudes to health care prioritisation methods and criteria among nurses, doctors, politicians and the general public. Soc Sci Med 49:1529–1539
Sabik LM, Lie RK (2008) Priority setting in health care: lessons from the experiences of eight countries. Int J Equity Health 7:4
SOU (1995) Vårdens Svåra Val. Socialdepartementet, Stockholm
Tsuchiya A (1999) Age-related preferences and age weighting health benefits. Soc Sci Med 48:267–276
Tsuchiya A (2000) QALYs and ageism: philosophical theories and age weighting. Health Econ 9:57–68
Tsuchiya A, Dolan P, Shaw R (2003) Measuring people's preferences regarding ageism in health: some methodological issues and some fresh evidence. Soc Sci Med 57:687–696
Werntoft E, Edberg AK (2009) The views of physicians and politicians concerning age-related prioritisation in healthcare. J Health Organ Manag 23:38–52
Whitty JA, Lancsar E, Rixon K, Golenko X, Ratcliffe J (2014) A systematic review of stated preference studies reporting public preferences for healthcare priority setting. Patient 7:365–386
Williams A (1997) Intergenerational equity: an exploration of the 'fair innings' argument. Health Econ 6:117–132