



RESPONSE

FAIR INNINGS AND TIME-RELATIVE CLAIMS

BEN DAVIES

Keywords

egalitarianism,
welfare,
healthcare,
rationing,
age

ABSTRACT

Greg Bognar has recently offered a prioritarian justification for ‘fair innings’ distributive principles that would ration access to healthcare on the basis of patients’ age. In this article, I agree that Bognar’s principle is among the strongest arguments for age-based rationing. However, I argue that this position is incomplete because of the possibility of ‘time-relative’ egalitarian principles that could complement the kind of lifetime egalitarianism that Bognar adopts. After outlining Bognar’s position, and explaining the attraction of time-relative egalitarianism, I suggest various ways in which these two kinds of principle could interact. Since various options have very different implications for age-based rationing, proponents of such a rationing scheme must take a position on time-relative egalitarianism to complement a lifetime prioritarian view like Bognar’s.

Greg Bognar has recently offered a possible justification for ‘fair innings’ distributive principles that would ration access to healthcare on the basis of patients’ age.¹ Bognar considers three common justifications for age-based rationing, and argues that they all fail to correspond sufficiently with plausible moral intuitions. He then suggests a prioritarian justification of his own.

In this article, I agree that Bognar’s account is a plausible ingredient of the best justification for age-based rationing of healthcare. However, even if we establish that such an account is the most promising focus for proponents of age-based rationing, there is considerably more justificatory work to do. Section 1 describes Bognar’s account, including some of its putative strengths as a justification of age-based rationing. The prioritarian view he outlines has the strength of individual claims to life-extending interventions depend on how many years they have already lived. His proposal thus sits within a well-established family of egalitarian principles that consider as morally relevant the amount of some good people will have over the course of their lifetime. However, Section 2 notes that another, less established set of egalitarian principles has the strength of

people’s claims depend not on their lifetime shares of a good, but on their shares *at particular times*.

As Section 3 notes, such ‘time-relative’ principles might operate in tandem with the kind of lifetime prioritarian view (or, indeed, some other lifetime egalitarian principle) that Bognar outlines. Moreover, many – though not all – such principles are consistent with some kind of age-based rationing, and different time-relative principles will have different implications for in what circumstances, and to what degree, age-based rationing may be permissibly implemented. As such, proponents of age-based rationing who advocate Bognar’s prioritarian principle as the most defensible justification of that position must also consider the further issue of what, if any, time-relative position to adopt before they can offer a fully worked-out justification of their preferred rationing system.

1. BOGNAR’S AGE-BASED PRIORITARIANISM

In general, prioritarian views take some good *G* that is to be distributed and say that those with more *G* have weaker claims to additional amounts. If two patients will

¹ G. Bognar. Fair Innings *Bioethics* 2015; 29: 251–261.

get an equal benefit from a medical intervention, a prioritarian view will prefer the patient who has benefitted less overall.² Bognar's proposal defines G as additional years of life. Since elderly patients have by definition lived for more years than younger patients, we would weight claims to life-extending treatments according to age.

This view avoids several issues that Bognar raises against competing justifications. Prioritarianism applies age-based rationing to patients of all ages, and so does not single out the very old; Bognar criticises views that apply age-based rationing only to some elderly subgroup of the population because he argues, convincingly, that there is no principled way to identify a threshold at which such rationing should begin.

Second, Bognar's view does not support an absolute priority for younger patients, since prioritarian weightings can themselves be outweighed if the benefit to elderly patients is considerably greater in terms of life-years (either because the young patient will not survive very long even if we treat them, or because there are more elderly patients than young). Again, this seems to accord with common sense reasoning about age rationing.

Finally, Bognar suggests that his view best matches 'empirical data on people's judgements'³ about age-based rationing, which suggest that the public are broadly supportive of age-based rationing, and do not draw any particular threshold at which it should start. As Bognar has acknowledged elsewhere,⁴ the reactions elicited in such surveys of public attitudes are consistent with the possibility that respondents assume that competing patients have had a similar quality of life, such that older patients have had not only more years of life, but *also* more of other goods such as opportunity and welfare. This would be particularly understandable when, in the absence of any information about past quality of life, respondents are told that imagined patients differ 'only in age',⁵ implying that they have had comparable lifetime welfare and opportunity. Still, since Bognar is offering a conditional argument – that *if* we support age-based rationing then we should adopt a prioritarian principle – I will join him in assuming that such surveys really do reflect support for rationing by age as such.

2. TIME-RELATIVE VIEWS

Prioritarian – and, indeed, egalitarian – views are predominantly applied across lifetimes. When we are considering which individual to benefit, a lifetime principle will have us consider lifetime shares of some relevant good (age, in Bognar's case). In general, lifetime principles weight claims according to the expected amount of some good G that one will possess over a lifetime.

But egalitarians might also advocate 'time-relative' principles, which weight claims according to how much G one will possess at particular times. For instance, McKerlie⁶ argues that a view which *only* applies some egalitarian principle to lifetimes will ignore morally relevant inequalities between groups or individuals so long as those inequalities are balanced out at another time. If two groups in society live for the same amount of time, it is acceptable from a lifetime position if group A does very well while group B suffers in the first half of their lives, so long as we impose a corresponding inequality that benefits B to the same degree in the second half of their lives, because in *lifetime* terms the two groups will be equal.

Yet, says McKerlie, there seems to be something unjust about *both* of those significant inequalities, even if they are balanced out at another time. He advocates a mixed view, which gives weight to people's claims by considering how good or bad their lives will be overall, but also by considering how well or badly off they will be when they receive benefits (Section 3 considers in greater detail how such a mixed system might work). Similar, though importantly different, views are found in Nord⁷ and Bou-Habib.⁸ At their most plausible – and the description under which I will consider them – time-relative views are concerned with how people are when they *receive* benefits or *incur* losses, using this information to weight their claims to those benefits and against those losses.

The details of such a principle depend on whether our time-relative principle is egalitarian, prioritarian, or sufficientarian. Before discussing these possibilities, it is worth emphasizing that a concern with particular times is not equivalent to a concern with the present. Saying

² Prioritarian views usually consider expected future levels of G as well as past levels; what matters is how well people's lives will go, not how they have gone so far. But when it comes to the life-saving measures with which Bognar is concerned, past levels are nearly equivalent to lifetime levels for whichever patient is not treated, because they will soon die.

³ Bognar, *op. cit.* note 1, p. 260.

⁴ G. Bognar. Age-Weighting. *Econ Philos* 2008; 24: 167–189.

⁵ As they are explicitly in another paper Bognar cites: J.J.V. Busschbach, D.J. Hessing and F.T. de Charro. The Utility of Health at Different Stages in Life: A Quantitative Approach. *Soc Sci Med* 1993; 37: 153–158: 154.

⁶ D. McKerlie. Equality and Time *Ethics* 1989a; 99: 475–491; D. McKerlie. Justice between age-groups: a comment on Norman Daniels *J Appl Philos* 1989b; 6: 227–234; D. McKerlie. Priority and Time *Can J Philos* 1997; 2: 287–309; D. McKerlie. Justice between the young and the old *Philos Public Aff* 2001; 30: 152–177; D. McKerlie. 2013. *Justice Between the Young and the Old* New York: Oxford University Press.

⁷ E. Nord. Concerns for the worst off: fair innings versus severity. *Soc Sci Med* 2005; 60: 257–263.

⁸ P. Bou-Habib. Distributive justice, dignity, and the lifetime view. *Soc Theory Pract* 2011; 37: 285–310.

that we are concerned with what people have at particular times does not mean that we are concerned solely with what people have right now. If a redistributive act has a positive impact (according to some time-relative principle) in the present, but will have a much greater negative impact at various future times, its present impact does not give us overwhelming reason to pursue it. That caveat established, I will now discuss what I take to be the three main time-relative principles: time-relative egalitarianism, prioritarianism and sufficientarianism.

Time-relative egalitarians want people to have equal amounts of *G* at particular times; on the time-relative view I am considering, my claim to a benefit will be proportionately stronger the worse off I am *than others* when the benefit occurs. My claims against a loss will be proportionately weaker the better off I am *than others* when the loss occurs. In considering a distribution at a particular time, then, we would look at individuals' levels of *G* at that time. If there is an inequality, the person who is worse off than others at the particular time has a stronger claim to benefits based on that inequality.

Since conferring benefits at particular times can have effects at later times, we will also need to consider how relatively well or badly off people will be at other times when the effects of our redistribution are felt; if rectifying an inequality at one time would lead to greater time-relative inequalities elsewhere, all things considered we should leave things as they are. However, this does not mean that every time we assess a time-relative claim, we must benefit the person who will face the greatest sum of lifetime inequality. If we can make *P* better off now, reducing the current inequality between her and *Q*, and can do *nothing either way* to reduce the inequalities between *P* and *Q* in the future, time-relative egalitarianism says that we should benefit *P* now even if *Q* will be worse off, to a greater degree, and more often. This is because the job of time-relative distributions is not to compensate people for future or past inequalities; it is to remedy inequalities at particular times. Since in the scenario under consideration we reduce a time-relative inequality and do not make any others worse, we should benefit *P*. As I will argue in Section 3, however, we might *additionally* apply a lifetime principle that tells us to benefit *Q* because of her greater lifetime inequality.

The difference between time-relative egalitarianism and prioritarianism is that while the former is inherently comparative, the latter is concerned with people's absolute states. Time-relative prioritarians say that people have stronger claims the worse off they are at particular times. Of course, these claims will compete with one another and so be apt for comparison; but our concern is fundamentally with how badly or well off a person is, rather than how much worse or better off she is than others. On this view, then, we weight claims to *G* at particular times according to how much *G* people have at those times. The

difference is that the *degree* of weighting will depend not on how much worse off individuals are than one another, but on how badly off they are on some absolute scale. This means, *inter alia*, that prioritarian time-relative weighting will have a smaller effect in well off societies than in badly off societies; time-relative egalitarians would place considerable weight on the claim of the worse off individual in a distributive decision between two well off people.

Like time-relative egalitarianism, time-relative prioritarianism as I understand it says that the strength of one's claims depends on one's situation when a benefit or loss will be felt. So one's claim to a benefit is proportionately stronger the worse off one is absolutely when the benefit takes place; and one's claim against incurring a loss is proportionately weaker the better off one is absolutely when the loss will be incurred. And as with time-relative egalitarianism, we may also want to consider unintended ramifications of our actions at future times. *P* has a strong prioritarian claim to a benefit at time t_1 because he will be badly off at t_1 . However, we can only give *P* that benefit by forcing losses on *Q* at t_2 , when she is even worse off than *P* is at t_1 . Assume further that *P* will suffer no ramifications at t_2 if we do not benefit him at t_1 . In this case, we may say that, all else being equal, the initial transfer is overruled on time-relative prioritarian grounds by the strength of *Q*'s claim against her future losses.

Finally, time-relative sufficientarianism says that we should be particularly concerned by whether people have enough at particular times. There is some disagreement about whether sufficientarians can also have concerns with the distribution of goods among people who do have enough. Casal⁹ says that it is inherent to sufficientarianism that it is indifferent about such distributions; Shields¹⁰ disagrees. If sufficientarians were concerned with distributions beyond sufficiency, they might apply a rule like prioritarianism to govern such decisions; what would then be distinctive about time-relative sufficientarianism is that it gives *special* weight to those who do not have enough, not that it has no interest in those who do. While time-relative prioritarianism would suggest a continuous correlation between the strength of our claims, and our position when we incurred benefits and losses, time-relative sufficientarianism would say that there was a 'discontinuous' relationship between the two, with the break in continuity occurring at the point where one did not have enough. If Casal is right, on the other hand, time-relative sufficientarians would say that people have no time-relative claims – although they might have other claims – unless they would not have enough when a benefit or loss was incurred.

⁹ P. Casal. Why sufficiency is not enough. *Ethics* 2007; 117: 296–326.

¹⁰ L. Shields. The prospects for sufficientarianism. *Utilitas* 2012; 24: 101–117.

There are also various questions about how we should decide between individuals who are below the sufficientarian threshold, wherever that is: whether we maximise the number of people who have enough; apply one distributive principle above the threshold and another below; or simply apply prioritarian reasoning throughout, but with a discontinuity in the strength of people's claims when they are below the threshold. Whatever our decision on these issues, time-relative sufficientarians will say that we should give special weight to a person's claim to a good if it will prevent them from dropping below a sufficiency threshold, or bring them above it, at a particular time or times.

Again, I assume that future times influence our decision insofar as they are affected by it; for time-relative sufficientarians, the fact that someone will suffer insufficiency in the future is not a special reason to benefit them now if that benefit will not help avoid that insufficiency. On the other hand, if we must decide between someone who currently lacks enough, but soon will be better off, and someone who currently has enough but for whom a benefit would see off a future insufficiency, we might prefer to benefit the latter person even though they are currently better off.

The broad distinction between time-relative and lifetime views can thus be restated as follows. Imagine – simplifying things considerably – that we have a series of benefits and losses to confer on people. Lifetime views will say that the strength of people's claims depends entirely on how their life goes overall; in particular, it makes no difference when they receive the benefit or incur the loss. The person with the worse life has a stronger claim even if we can only benefit him at a time when he is very well off, and the corresponding loss to the better off person will be incurred when she is very badly off. A time-relative view says the reverse; our claims to benefits and against losses vary according to our situation when those benefits and losses will be felt.

If neither of these views looks plausible in isolation, we might look instead to a broader position that tries to combine the weightings of our lifetime and time-relative positions into a single overall weighting. This, however, turns out to be more complex than that may sound, because there are a number of ways of combining such principles. As such, I argue that even if Bognar's principle is included in the best justification for age-based rationing, proponents of that view will need to take a stance on time-relative distribution for a complete justification of their view.

3. LIFETIME AND TIME-RELATIVE PRINCIPLES

Bognar points out that because a time-relative principle 'only takes into account how badly off people are at particular times', and 'poverty and illness are more prevalent

among the elderly',¹¹ time-relative principles cannot provide support for age-based rationing. As such, he says that proponents of age-based rationing cannot adopt a time-relative egalitarian principle in isolation. As Bognar acknowledges,¹² this is not McKerlie's view; McKerlie does not suggest that we should only give weight to time-relative considerations – indeed, as far as I am aware nobody holds this view – but insists that time-relative considerations must play some role in our allocations, giving weight to people's claims along with lifetime considerations.

Recall that Bognar's prioritarian principle is supposed to be the best defence of age-based rationing of healthcare. This is because it takes age itself as a relevant good to be (indirectly) distributed, and weights our claims to interventions that will deliver more of that good according to how much we have already had. As Bognar notes, time-relative views will not tell us to apply lesser weight to the same benefit in a continuous manner as patients get older. But this does not mean that supporters of age-based rationing can reject time-relative views even if age-based rationing is justified, because a number of time-relative views are compatible with age-based rationing so long as they are held *along with* a lifetime principle. Assuming that proponents of age-based rationing should indeed adopt something like Bognar's lifetime prioritarianism, what is then crucial in deciding the precise nature of a scheme of age-based rationing is which time-relative principle is adopted, if any, and how it interacts with lifetime prioritarianism.

A time-relative principle can interact in various ways with a lifetime principle: it can act as a *constraint*; it can operate on an *equal* basis in terms of weighting people's claims; it can itself be *constrained* by lifetime principles; it can act as a *trump* in a lifetime distribution; or it can act in different ways depending on the *context*. This will depend both on what commitments we conceive of a time-relative principle as reflecting, and which distributive principle – equality, priority or sufficiency – we choose to implement at times.

While it is true that a time-relative principle itself will give us no reason to prefer younger patients in healthcare allocations, and that a distribution with a time-relative component will give greater weight to elderly people than one without, several of the possibilities above would still allow for age-based rationing if combined with lifetime prioritarianism. I will not consider all possible iterations of how the three kinds of time-relative principles – egalitarian, prioritarian and sufficientarian – might operate in each of these contexts, but I will now draw

¹¹ Ibid.

¹² Bognar *op. cit.* note 1, p. 261, fn 29.

out some key examples to give some sense of the breadth of possibilities available here.

If time-relative sufficientarianism acted as a constraint on Bognar's lifetime prioritarianism, this might give us some *prima facie* reason to prefer younger patients in many distributive decisions, but to at least weaken that weighting if the older person would fall, or was already below, some sufficiency threshold. Bou-Habib suggests that our distributive decisions should be broadly lifetime egalitarian, but that we must give special consideration for those whose capacity for autonomy is under threat.¹³ If we abandon Bou-Habib's focus on autonomy, we can also include states that are plausibly below a sufficientarian threshold, but which do not necessarily threaten one's capacity for autonomy, such as severe pain. Similarly, we might substitute a concern for lifetime equality with lifetime priority.

Another option is to give equal weighting to lifetime and time-relative priority, as McKerlie favours.¹⁴ Such a view would have us weight people's claims according to how bad their lives will be, *and* how bad things will be at particular times.

Admittedly, this gives us no clear guidance in some cases. For instance, P and Q are nearing the ends of their respective lives, of equal length. P has had a bad life overall, but things look as though they will go fairly well until her death (though not enough to fully compensate her lifetime position); Q is in the reverse situation, with a good life overall but bad times to come. Without further specification, McKerlie's view gives us no guidance about how we should weight these considerations. Still, in cases with which Bognar is centrally concerned, where one patient will die, each person's time-relative position is plausibly equal. On the equal weighting view, then, we would simply decide the tie by appealing to lifetime priority, giving a result identical to a view which only considered the lifetime position. The difference is that in allocations of goods that do not involve life and death, time-relative claims are relevant.

This view also highlights the fact that the 'currencies' of our lifetime and time-relative principles need not be the same. I have said that it makes little sense to apply time-relative reasoning of any kind to age as such. But one could apply *lifetime* prioritarian weighting to numerical age, and also have additional prioritarian weighting for those who were badly off in some other good, such as welfare, at particular times.

Another alternative is to start with time-relative considerations, and then adjust them according to lifetime considerations. We might insist, for instance, that those who are below some sufficiency threshold at particular

times should be given a strong preference in our allocations, but that when we are deciding between people in this category – say, people who would die without intervention – lifetime priority should weight claims. Again, lifetime priority here could consider purely numerical age so that we prefer to save a younger person's life over an older person's life, or it could consider other lifetime goods.

Indeed, such a view seems to fit better with our widespread views on healthcare allocation than would an appeal to lifetime priority in isolation. An appeal to time-relative considerations allows us to restrict those who are competing for a benefit in the first place. For instance, the financial resources spent on extending an elderly patient's life could not only be used to extend a younger person's life, but might also just *improve* the life of a younger person who does not have any particular need in the foreseeable future. If our elderly patient has had a very good life, a pure lifetime view would prefer to benefit the younger person, and not see the patient's current state as relevant *except* insofar as it made her life worse overall. Perhaps this is right. But insofar as people seem to have some concern with the fact that someone is (or will be) *in pain*, or *dying*, it looks as though time-relative considerations may have a place in our broader distributive theory.

Many advocates of age-based rationing do not apply it, or at least not as starkly, to pain relief or palliative care.¹⁵ And indeed, Bognar's argument applies lifetime prioritarian reasoning only to additional life-years. But unless we, perhaps implicitly, incorporate some kind of time-relative view, it is not clear why we shouldn't also apply age-based rationing to other kinds of healthcare, such as pain relief. So even if we should not advocate a time-relative view with regard to life-saving interventions, the same thinking does not warrant abandoning it altogether. Perhaps distributive justice is sufficiently complex that time-relative considerations are irrelevant to some decisions, but not to others.

One might claim that the reason for this disparity is that it is obvious that the value of pain relief does not diminish just because one has had a longer life, whereas the value of additional life-years does so diminish. But this presents a dilemma for a view that rejects time-relativity altogether, depending on exactly how we flesh it out. On the one hand, such a claim might mean that the actual utility to an agent of pain relief does not diminish with age; pain relief is fundamentally valuable to agents because it relieves their pain at particular times,

¹⁵ See e.g. M.P. Battin. Age Rationing and the Just Distribution of Health Care: Is There a Duty to Die? *Ethics* 1987; 97: 317–340; D. Callahan. 1995. *Setting Limits: Medical Goals in an Aging Society* Washington; Georgetown University Press, p. 72; Lazenby. *op. cit.* note 16; J. Bidadanure. In Defense of the PLA. *Am J Bioeth* 2013; 13: 25–27.

¹³ Bou-Habib *op. cit.* note 8, pp. 298–309.

¹⁴ e.g. McKerlie 2013, *op. cit.* note 6, pp. 88–117.

not because it gives them a life that is freer of pain. That is probably correct, but seems to provide precisely the basis *for* having a separate time-relative principle, even if that is only targeted at pain relief. Moreover, the correlative principle applied to life-years would either have to demonstrate that additional life-years *of any quality* are less valuable for those who have had longer lives *of any quality*; or it would have to rely on general claims about the diminishing value of additional years.

On the other hand, the claim about a disparity between life-years and pain relief might be about moral worth: while the *moral* importance of pain relief does not diminish with additional life-years, the moral importance of additional life-years does. But this requires argument. One might say that it is obvious that the moral value of additional life-years should depend on how many life-years one has already had, because the similarity of currency makes it obvious that older patients have had more of a relevant good, whereas a long life is no guarantee of a life free of pain. This still leaves open the question of why we should not in principle prioritize pain relief according to lifetime experience *of pain*; if we are going to prioritize on a purely lifetime basis in all cases, why not prioritize consistently? And if we think there is some difference between these two cases, isn't it plausible that one difference is that pain relief is fundamentally concerned with one's situation at a particular time, whereas life extension is at least partly concerned with one's lifetime share?¹⁶

Moreover, even if the empirical evidence Bognar cites supports rationing on the basis of age alone, the responses elicited all concern life-extending interventions but do not at all address the prioritization of other kinds of healthcare. Perhaps the public take precisely the view suggested above, seeing many health interventions, particularly pain relief, from a time-relative perspective, but viewing life-extension from a lifetime perspective, and so apply lifetime prioritarian reasoning to life-extension more readily than to other forms of healthcare.

Some concern with time-relative sufficiency might also explain an apparent preference, found in one paper Bognar cites¹⁷ (and not addressed in others), to apply

age-based rationing to group-level decisions, and a reluctance to apply them in cases of individual competition. Although there are various explanations for this, one possibility is that a concern with time-relative sufficiency – expressed as the protection of a particular core of welfare, opportunity or interests otherwise conceived – conflicts with a lifetime prioritarian view, such that where no individual can be identified as losing out on that protected core, people are happier to defer to lifetime priority than when an individual can be so identified. Perhaps people see life-extension predominantly under the guise of adding life-years when considered at a more abstract group level, but find its time-relative component more compelling when they consider competition among individuals.

This may look like an incoherent set of preferences, but they can be made coherent. For instance, one might worry about the idea that the age of a patient should automatically decide between two individuals, even if we think that youth should have some weight. One option would be to use patients' age to determine their *weighted chances* in cases of irresolvable conflict by holding a weighted lottery to decide which patient should be treated. The reasoning is as follows: if two patients are of the same age, they should each have an equal chance of being selected for treatment. If patients' ages differ, then allowing additional weight to youth would not automatically have us choose the younger patient – effectively moving straight from 50:50 chances to zero chance for the older patient – but to shift the proportional chance according to age difference: e.g. 80:20.¹⁸ The rationale would be that the latter system continues to recognise that each patient has a strong time-relative claim that is not fully commensurable with their lifetime claim.

The obvious problem with this proposal is that it is practically infeasible to hold a weighted lottery every time there is a conflict between two individuals. Weighting chances by group membership is one way – albeit an imperfect way – of approximating this procedure. A group-weighted principle would allocate greater weight to funding for conditions and targeted health programmes that respectively harm and benefit the young, but it would not exclude or even give less of a chance to individual patients on the ground of their old age.

Here is a simplified example. Condition 1 affects only young patients, and condition 2 affects only elderly patients. We have one hundred patients in each group, but only enough resources to treat one hundred in total. If the two groups were both young, we might split

¹⁶ This is not to rule out the possibility that time-relative considerations could affect claims to life-extending healthcare. After all, such interventions do not only extend our lives, they also *save* our lives at particular times.

¹⁷ E. Nord et al. The significance of age and duration of effect in social evaluation of health care. *Health Care Anal* 1996; 4: 103–111. Although subjects had some preference for life-saving *projects* directed at the young (p. 106) only a minority (17.6%) said that age as such should play a role in individual allocations, while roughly equal numbers preferred no interpersonal age priority, or preferred only to ration healthcare for the 'very old' (p. 105).

¹⁸ I take the idea of weighted lotteries from discussion of whether group size should affect our life-saving decisions in e.g. J. Timmerman. The Individualist Lottery: How People Count, But Not Their Numbers. *Analysis* 2004; 64: 106–112; B. Saunders. A defence of weighted lotteries in life saving cases. *Ethical Theory Moral Pract* 2009; 12: 279–290.

resources to treat fifty in each group. The group weighting principle would tilt things in favour of condition 1, offering funds for, say, eighty young patients and only twenty elderly. Each young patient thus has a greater chance of being treated than each elderly patient, though some young patients will die where some elderly patients live.

This would rule out directly deciding between two patients on the basis of age, but would still give a clear role to time-relative claims as well as lifetime claims. Younger patients would have a greater chance of benefiting because treatments for them would have greater weighted funding; but older patients would still have considerable chance of benefitting from life-extending treatments. This might mean that sometimes an individual older patient was chosen over a younger patient in an allocation at the very local level of patient care; but the principle of age weighting would nonetheless be present in our system as a whole.

An alternative to weighted lotteries is to count age as one consideration among many in deciding priority of some treatments. This will not be practical for all life-saving medical interventions – in particular, it would be infeasible in the case of emergency treatment – but might work for interventions that are allocated by waiting lists, such as access to organ transplants.

Time-relative principles might be made inconsistent with age-based rationing, but only, as far as I can see, if a time-relative principle acts as a trump. One can imagine, for instance, a distributive scheme that in general prefers to allocate benefits on the basis of lifetime priority – again, this could be on the basis of numerical age – but which does not allow that principle to determine or even count in allocations below a particular threshold. If death brought one below the threshold, then this would rule out age-based rationing in life-extending decisions. Perhaps something like this idea is contained in the ‘rule of rescue’, which tells us that ‘Our moral response to the imminence of death demands that we rescue the doomed’;¹⁹ but it should be clear that it is not the only way that one can apply time-relative principles.

So while it is true that time-relative principles cannot justify age-based rationing, there are various possibilities for including a time-relative principle when allocating health-related (and other) goods, many of which are con-

sistent with some form of age-based rationing, or lifetime weighting on the basis of some other good. Insofar as one is attracted to age-based rationing, then, Bognar is right that one should not *only* apply a time-relative principle; indeed, I think this position is implausible whether or not one supports age-based rationing. I have also suggested that proponents cannot apply a time-relative principle as a trump against lifetime claims. But one might apply time-relative principles on an even footing as lifetime principles, give one or the other some degree of non-absolute priority, or do all three regarding different goods.

4. CONCLUSION

Bognar argues that if we are to advocate healthcare rationing on the basis of numerical age alone, we should adopt his lifetime prioritarian view, which weights individuals’ claims to additional life-years according to how many they have already had i.e. how old they are. Although I have expressed some reservations – which Bognar appears to share to some degree – about the justification for age-based rationing, I agree that this is currently its most plausible defence. However, even if we embrace lifetime priority of age, we must decide whether to adopt time-relative distributive principles. I have outlined several principles that have our claims depend in part on how we will be when benefits and losses occur, and suggested various ways in which such time-relative principles might interact with lifetime principles. Since many of these could affect how we apply age-based rationing without ruling it out, proponents of rationing need to commit to more than lifetime prioritarianism. Bognar has offered proponents an important element of a full theory, but there is more to be decided before age-based rationing is fully vindicated.

Acknowledgements

My thanks to Greg Bognar, Matteo Mameli and an anonymous reviewer for helpful feedback on this article.

Ben Davies has recently submitted his PhD in Philosophy at King’s College London. It addresses and refines central objections to anti-aging enhancement. His research interests include distributive justice, issues around death and dying, and subjective theories of value.

¹⁹ A. Jonsen. Bentham in a Box: Technology Assessment and Health Care Allocation *Law Med Health Care* 1986; 14: 172–74: 172.