

ARTICLE

Age and Illness Severity: A Case of Irrelevant Utilities?

Borgar Jølstad^{1*}  and Niklas Juth² 

¹Akershus University Hospital, The Health Services Research Unit – HØKH, Lørenskog, Norway and Centre for Medical Ethics (CME), Institute of Health and Society, Faculty of Medicine, University of Oslo, Oslo, Norway and ²Karolinska Institutet, Center for Healthcare Ethics, dep. of LIME, Stockholm, Sweden
*Corresponding author: E-mail: borgarjolstad@gmail.com

Abstract

Illness severity is a priority setting criterion in several countries. Age seems to matter when considering severity, but perhaps not small age differences. In the following article we consider Small Differences (SD): small differences in age are not relevant when considering differential illness severity. We show that SD cannot be accommodated within utilitarian, prioritarian or egalitarian theories. Attempting to accommodate SD by postulating a threshold model becomes exceedingly complex and self-defeating. The only way to accommodate SD seems to be to accept some form of relevance view, where some age differences are irrelevant. This view can accommodate SD, but at the expense of consistent priority orderings. Severity thus becomes unsuitable for systematic decision-making. We argue that SD should be dismissed and that we should accept a continuous relationship between severity of illness and age.

What is illness severity and why should we care?

Given the ubiquities of health problems and scarcity of health resources, priority setting is a necessity. One natural desire is for our resources to be used as effectively as possible; given a restricted health care budget, we want to “produce” as “much health” as possible. This strategy amounts to a maximizing ethics with a focus on the cost-benefit ratio of interventions. In several countries, this maximizing ethic is adjusted by a focus on the notions of need or *illness severity*¹ (Barra and others 2020; Franken and others 2015). These notions reflect the view that it does not only matter that we get as much health as possible but also who the recipients of these benefits are. For instance, in the Norwegian priority setting system, illness severity is a criterion alongside cost and benefit (Barra and others 2020). That is to say, Norway is willing to spend more on the same health benefits when they accrue to those who are severely ill. This use of an illness severity criterion leads to the question of what it means to be severely ill. How do we judge that an illness, or state of illness, is more severe than another? Plausible candidates include the level of pain, quality of life, life years lost, risk of

¹Severity of illness and illness severity are used interchangeably throughout the article.

death, and the age of the patient, but there is as yet no consensus on what severity is or should be (Barra and others 2020). There seems to be a close relationship between the desire to prioritize those who are worse off and prioritizing those who are severely ill. Clearly those who have severe illnesses are worse off in some sense, but in precisely what sense is not obvious. Current operationalizations differ as to the role that age plays in illness severity. In the Norwegian system illness severity is operationalized as absolute shortfall: the severity of an illness is a function of how many good life years are lost. Age is thus indirectly a factor when assessing severity in the sense that younger patients typically have more life years to lose than older patients. More generally, it is probable that any severity-measure that operationalizes severity as life years lost will be sensitive to age. This can potentially be considered a form of age discrimination. In the Swedish priority setting system the Human Dignity Principle precludes considering chronological age in priority setting, and thus prevents age from being a factor in severity of illness (Barra and others 2020). The relationship between severity and age is thus not an uncontroversial one. The aim of this article is to explore some intuitions regarding the relationship between illness severity and age that seem to imply that age differences have to be relatively large to merit a differential judgement of illness severity. We will argue that this implication can only be accommodated at the expense of a consistent priority ordering.

But first, why care about illness severity? Why not simply get as much health as possible? There are several reasons for caring about illness severity. First, unlimited aggregation leads to unintuitive consequences. An unlimited aggregation view entails prioritizing minor benefits for large numbers of well-off individuals over large benefits for small numbers of worse off individuals, which can seem morally counterintuitive (Gustavsson and Juth 2019). A second reason is that the worse off are morally more important in some way. Such a concern for the worse off is asserted by both prioritarian, egalitarian, and, arguably also, sufficientarian theories of distribution (Hirose 2014). A third reason, somewhat more speculative, is that a public health care system is viewed by many as a form of safety net. Utilizing resources on relatively well-off individuals that could have been spent on someone who has a desperate need may feel wrong in this context (Gustavsson and Juth 2019). All of these reasons highlight the affinity between a focus on severity and a desire to prioritize the worse off.

In this article, two assumptions are made about severity of illness: that severity of illness is (at least partly) a function of how *bad* the illness is for the person,² and that severe illness constitutes a reason to prioritize a patient. It is thus assumed that severity can be viewed in the following two ways: (Severity 1) as a measure of health-related worse off-ness, and (Severity 2) as a marker for which patients should be given priority. Severity seems to be a thick concept with a descriptive dimension (Severity 1) and a normative dimension (Severity 2). We also assume a deprivation account of the badness of death, where how bad it is to die is a function of how much you lose because of death (Nagel 1970; Solberg and Gamlund 2016). It seems that if such a deprivation account on the badness of death is not assumed the notion of worse off-ness is difficult to make sense of in the context of health priorities. This is also how health loss is assessed in current health priority settings. Leaning on the deprivation account ensures that what we are aiming at is not maximizing the number

²In focusing on the badness for the patient, we are intentionally ignoring other factors that some might consider relevant to illness severity, such as whether the patient has children, is socially important, saves lives for a living, etc.

of good life years in the population (this could entail replacing sick individuals with new individuals) but rather maximizing the good life years of already existing members of the population. A further point is that we are here attributing severity to patients, or cases of illness, rather than to illnesses *per se*.³ A given illness can thus, for many different reasons, be more or less severe *for a person*. Lastly, the subject of this article is severity as a priority setting criterion with a moderating role in a largely maximizing priority setting system. Severity is thus a part of a larger decision-making scheme that also includes the cost and benefit of the intervention.

Severity and age: four cases

A potential and probably commonly held intuition about the relationship between age and severity of illness is that terminal illness is, *ceteris paribus*, less severe when faced by the very old and more severe when faced by the very young. In this regard, consider the following two cases:

Case 1

- Patient A is a 20-year-old who has an expected survival of one year due to an illness.
- Patient B is an 80-year-old who has an expected survival of one year due to an illness.

Is one of these cases more severe, and if so, which one? Next, consider Case 2 below:

Case 2

- Patient C is a 40-year-old who has an expected survival of one year due to an illness.
- Patient D is a 45-year-old who has an expected survival of one year due to an illness.

Is one of these cases more severe, and if so, which one?

First, it might be helpful to standardize the cases further by positing that the patients enjoy the same level of welfare, would prefer to live on, are childless and are equally socially valuable. All of this is meant to ensure that what we are trying to tease out is how bad the illness is from the perspective of the individual patient and not for (for example) a family, or society at large. By assuming that all the patients have the same level of welfare we also mean to focus solely on the length of life, and not on the quality of life. Further, note that all the patients have the same amount of life ahead of them (one year). Hence, the difference between the patients is to what extent they are *worse off* in two ways: (1) they *stand to lose* more or less time compared to a hypothetical “normal” lifespan, and (2) they *have had* more or less time to live.⁴ We thus assume that the 40-year-old would be similar to the 45-year-old in all relevant aspects if she were 10 years older. Our intuitive judgement is that patient A is more severely ill than patient B, but that patients C and D are equally severely ill. Age thus

³We would like to thank an anonymous reviewer for pointing this out to us.

⁴Note that this does not follow from a deprivation account but seems to be an intuitive reason for the badness of early death. Nothing in the article hinges on the reader accepting this as a reason for the badness of dying early.

seems to make a difference for illness severity in case 1 but not in case 2.⁵ Can this difference in intuitive judgement be grounded in moral theory? In a series of interviews on severity that the corresponding author has recently taken part in, several participants expressed views on severity and age of this kind. More specifically, they expressed confidence in assigning differing illness severity when differences in age were large but considered cases equally severe, or were unwilling to make a judgement, when age differences were smaller.⁶ We believe that these intuitions are similar to the intuitions that Francis Kamm claims to have about differences in years saved and whom to aid. She argues for the relative merits of equal chances, a weighted lottery and choosing on the basis of differences in outcome: for a 19-year difference in outcome she leans towards choosing on the basis on this difference, whereas a 2-year difference merits equal chances (Kamm 1993, pp. 287–88). Some differences in years saved are thus enough to merit differential treatment, whereas others are not. An interesting real world example of something similar to this is found in the triage recommendations at the critical care unit at the University of Pittsburgh (University of Pittsburgh 2020).⁷ These recommendations were made while expecting cases of Covid-19 to exceed the resources at the hospital. Patients are here categorized into age groups based on life stages. Whereas being in different life stages can result in differential treatment, a lottery is recommended for other cases. Clearly, some age differences are considered relevant while others are not.⁸ We will therefore assume that this differential judgement is worthy of consideration and investigate what kind of ethical theories can account for the divergence. First, we will consider common ethical frameworks for priority setting in health care (utilitarianism, egalitarianism, prioritarianism, and sufficientarianism). We will then argue that the differential judgement, and its implications, cannot be accommodated within any strictly consequentialist framework but must rather be seen from deontological perspectives. Lastly, we will then discuss problems with these perspectives.

Before we can consider the merits of the different relevant theories, it will be helpful to scrutinize these intuitive judgements. The intuitive judgement that patient A is more severely ill than patient B seems to us to be the strongest. This intuition can be explained in many ways. Illness in old age seems more *natural* (whatever this means), an older person has already had a long life (we are assuming a reasonably happy one) and the person typically does not have many years of life left regardless of illness (at least not compared to what 20-year-olds typically have). Illness, and especially terminal illness, in the young strikes us as tragic. Patient A has not had a chance to live a life and we consider death at this young age as something *unnatural* (whatever this means). Compare this with case 2. Patients C and D are both more like patient A in the mentioned respects. They are both in the prime of their lives, have a lot to look forward to and death at this age strikes us as tragic (but maybe not as tragic as death at 20).

⁵Note that in current priority setting contexts, a 40-year-old would be prioritized over a 45-year-old if the 40-year-old could be helped to a larger extent, for example in an emergency situation, where we have to choose between saving one or the other. In this article, we assume that whatever health gain can be provided is the same for all patients, thus assuming that cost and benefit are equal.

⁶This research, part of the Severity and Priority Setting in Health Care (SEVPRI)-project, is still a work in progress.

⁷We would like to thank an anonymous reviewer for directing our attention to this report.

⁸This resembles a sufficientarian system with multiple thresholds, something that will be discussed later in the text.

All this leads us to ask whether there is something about the old that leads us to consider an illness less severe. Maybe there is a cutoff – say, at 80 – where death is “natural” or at least not as bad? This would be in line with the “fair innings” argument where “everyone is entitled to some ‘normal’ span of health ... The implication is that anyone failing to achieve this has in some sense been cheated, whilst anyone getting more than this is ‘living on borrowed time’” (Williams 1997: 119). The “fair innings” is a threshold after which additional years of life have less moral importance. The fair innings argument can account for our intuition by positing that patient B has had a full life (or at least a chance at one), whereas none of the other patients have. Patients A, C and D are worse off than patient B (severity 1) in a sense that makes them more worthy of moral concern by being prioritized (severity 2). On a strict “fair innings”-threshold interpretation patient B thus stands out as less severely ill than the other three patients who are equally severely ill because neither of them has reached the threshold.⁹

But now compare cases 1 and 2 with the following:

Case 3

- Patient E is a 20-year-old who has an expected survival of one year due to an illness.
- Patient F is a 50-year-old who has an expected survival of one year due to an illness.

Is one of these cases more severe, and if so which one?

Case 4

- Patient G is a 77-year-old who has an expected survival of one year due to an illness.
- Patient H is an 80-year-old who has an expected survival of one year due to an illness.

Is one of these cases more severe, and if so which one?

Our intuitive judgement here is that patient E is more severely ill than patient F, but that patients G and H are equally severely ill. This cannot be accounted for by the fair innings idea strictly conceived. If one assumes that we are all entitled to 80 years of life, then patient H is the only one whose illness is less severe. Our intuitive judgement is that the difference between patient G and H is not *large enough* to warrant a differential judgement of illness severity and that the difference between patient E and F is clearly significant. The intuitive judgements seem to lead to the following principle:

Small differences (SD): small differences in age are not relevant when considering differential illness severity.

SD implies that there is no continuous relationship between age and severity. The severity of an instance of illness, rather than reflecting an underlying ranking of states of affairs (or a consistent ranking of the severity of illness based strictly on how the illness in question affects age of dying), seems to be relative to other instances of illness.

⁹If this seems strange, it is perhaps because if the 45-year-old has been cheated in comparison to the fair innings, then it seems that the 20-year-old has been cheated to a larger extent. John Harris, the originator of the fair innings argument, discusses this in *The Value of Life* (Harris 1985).

Terminal illness does not seem more severe for a 20-year-old than for a 25-year-old, the illness of a 25-year-old does not seem to be more severe than the illness of a 30-year-old etc., but the illness of a 20-year-old seems to be more severe than the illness of, for instance, a 50-year-old. We thus have a case where $20 \not\asymp 25$, $25 \not\asymp 30$, $30 \not\asymp 35$, $35 \not\asymp 40$, $40 \not\asymp 45$ and $45 \not\asymp 50$, but $20 > 50$. If these intuitive judgements are taken at face value, it seems that the “not more severe than”-relation is not a transitive one. The same will follow for the equally severe as-relation.

One problem when trying to make sense of SD is that it flies in the face of the worse off-aspect of severity. Assuming you are not a “hard epicurean” it is difficult to argue that death at 40 does not leave you worse off than death at 45: you have lost an additional 5 years (assuming that you would otherwise live for the same amount of time), and you have had 5 fewer years. If we are to keep both SD and admit that death at 40 is worse than death at 45, then severity cannot be a function of worse off-ness alone. This difficulty will be explored later on, but for now the problem seems to be that severity has a priority aspect and unless one is a consequentialist, worse off-ness (severity 1) and moral significance (severity 2) do not completely overlap. This claim will be explored further in what follows.

Utilitarianism and prioritarianism

From the perspective of utilitarianism, the morally right thing to do is to maximize the net sum of goodness over badness. Utilitarianism is both *consequentialist* in that what matters is the resultant states of affairs and *aggregative* in the sense that all utility is counted, no matter the size of the utility. We will here ignore the differences between different forms of utilitarianism and simply assume that there is some form of health-related good that can be maximized, most naturally good life years. Whether the good accrues to someone who is worse off or not does not matter in itself according to utilitarianism. Axiologically, utilitarianism is committed to a neutral weighting between individuals, and typically assumes a continuous ranking of states of affairs. All of this ensures that from a utilitarian point of view the state of affairs of dying as a 20-year-old (rather than as a 21-year-old) is, *ceteris paribus*, neither better nor worse than dying as an 80-year-old (rather than as an 81-year-old), but the same can be said for the difference between 40 and 45 (or even 40 and 41). If we were able to avoid one of these states of affairs we should, all else being equal, be indifferent between the patients. Where all the patients have one year to live, there is no reason to prioritize the worst-off, since they cannot be helped to a higher degree. A year of life is, all else equal, equally worthy of consideration regardless of whom it accrues to. On the most natural reading of utilitarianism, given the neutrality to worse off-ness, the notion of severity thus seems superfluous. Or, at least, it has no independent moral weight – at most it is a way to talk about levels of value. What matters is the utility that can be produced. Any amount of badness and good would simply be factored into the calculation, and whether we call it severe or not is irrelevant.

If one is attracted to the idea of maximizing *and* prioritizing the worse off, then *prioritarianism* seems to be a natural next step. From a prioritarian perspective we should give more, but not absolute, priority to helping those who are worse off (Hirose 2014; Parfit 1991). Prioritarianism builds on the idea that the claims of the worse off are morally weightier to some degree. Well-being is claimed to have marginally decreasing moral importance. This is typically assumed to take the form of a concave value function of moral goodness, such as this (see Figure 1):



Figure 1. Example of a prioritarian value function.

As is evident from the figure there is more value associated with increasing the well-being of someone lower on the well-being scale. It is also clear that the relationship between well-being and moral value is continuous. Due to the continuity of worse off-ness and of the function of worse off-ness and moral importance we end up with a system where those who are worse off are given priority proportional to how badly off they are. But it is obvious that this system cannot accommodate SD. If patient C is worse off than patient D (and we are assuming this), then it is obvious that we should prioritize patient C. Patient C is thus more severely ill than patient D on both the worse off- and priority-aspects of severity (severity 1 and 2). Prioritarianism, like utilitarianism, assumes a continuous relationship between worse off-ness and moral importance, and therefore cannot accommodate SD.

Egalitarianism

The four most influential egalitarian theories of distribution are *telic egalitarianism*, *luck egalitarianism*, *sufficientarianism*, and *Rawls maximin or leximin principles* (Hirose 2014). For the purposes of this article, we will disregard luck egalitarianism, because illness would in most cases be considered a form of brute luck.¹⁰ We will first consider telic egalitarianism and maximin/leximin.

Telic egalitarianism is a family of theories that claim that inequality contributes to the badness of an outcome over and beyond the effects that this inequality has on specific individuals. According to Parfit (1991) telic egalitarians believe that equality has intrinsic value or, rather, that the more inequality, the worse the outcome (all else being equal). In our examples, this means that in addition to the badness contributed by the illness of the patients there is a separate badness contributed by the fact that the illness is unequally distributed in the population. The illness of patient A is thus a worse outcome than the illness of patient B both because patient A is worse off and because the illness of patient A contributes more to the separate badness of inequality. The same can be said for the illness of patient C and D. Patient C is both worse off than patient D and contributes more to the badness of inequality. The difference in moral importance caused by differential worse off-ness is thus even greater according to telic

¹⁰Some instances of illness, such as illnesses caused by smoking or obesity, can arguably be considered instances of option luck. We are disregarding this complicating factor for the purposes of this article.

egalitarianism, and there is no reason to disregard the difference in worse off-ness between patient C and D. Telic egalitarianism is thus also unable to accommodate SD.

Rawls' theory on maximin or leximin dictates giving either absolute (maximin) or lexical (leximin) priority to the worse off (Rawls 1999).¹¹ On a maximin reading of case 1, patient A is not just more morally important than patient B, she is the only morally important person in this case. On a leximin reading, patient A is morally more important as long as something can be done for her, and only when this is not possible does patient B acquire moral importance. But the same can be said for the remaining three cases. There is, following the logic of Rawls theory, no reason not to give weight to the differences in worse off-ness, and every reason to do so. Rawls' theory thus seems, to an even larger extent than previously discussed theories, to be unable to accommodate SD.

The reason that these theories cannot accommodate SD seems to be that these theories are consequentialist;¹² moral importance (priority) is a function of consequences (either worse off-ness or worse off-ness and the badness of inequality). The two views on severity (worse off-ness and being marked for priority) are either collapsed into one, or on the utilitarian reading rendered superfluous. The ones who should be prioritized (severity 2) are the ones who are worse off to the largest degree (severity 1) (on the utilitarian reading worse off-ness means a greater opportunity to do good) and there is no room for SD.

Sufficientarianism

Sufficientarianism is a family of theories that claim that what is morally important is that everyone has enough, or more accurately that we should prioritize “those whose well-being is below a certain threshold” (Hirose 2014: 112). The already discussed fair innings argument can thus be viewed as an essentially sufficientarian argument, albeit one with a high threshold. This theory might be able to make sense of the intuitive judgements of the cases if we interpret them as solely pointing to the lower severity of the very old and higher severity of the young. To be able to account for all of the cases we would have to assume a sufficientarian theory with multiple thresholds. The result might be something akin to a curve that looks like this (see Figure 2):




Figure 2. Example of a sufficientarian two-threshold value function.

¹¹We are here disregarding the complications around the focus on the worst-off groups and worst-off individuals.

¹²We are here using the term consequentialist in the Rawlsian (Rawls 1999) sense to include all theories that view the right as following from the good, rather than in Kagan's (Kagan 1992) sense, where the difference between consequentialist and non-consequentialist theories consists in whether they limit moral

where age is on the X-axis and degree of severity is on the Y-axis. This would imply considering the illnesses of the young as very severe,¹³ the illnesses of the middle-aged as medium severe and the illnesses of the old as less severe.¹⁴ This could account for some of our intuitions but would lead to a new set of problems. Say that we define young as under-30-year-olds. Then what about 31-year-olds? Is the difference between a 30-year-old and a 31-year-old large enough to merit a difference in kind? And, similarly, is the difference between a 79-year-old and an 80-year-old relevant when considering the severity of an illness? The thresholds seem arbitrary and the value difference between people on different sides of the threshold seems too sharp. This problem would also appear if we attempted to avoid the discontinuity-difficulties by claiming that the relationship between severity and age takes the form of categories (for instance 10–20, 21–30, 31–40 years old, etc.). We could perhaps avoid the problem of sharpness by making the thresholds sloped. The curve would then look something like this (see Figure 3):




Figure 3. Two-threshold value function with sloped thresholds.

Say the line from T1 to T2 represents ages 30–40 and the line from T3 to T4 represents ages 70–80. This would remove the problem of sharpness but would reintroduce a continuous relationship between age and severity on parts of the line (the sloped lines). One could then claim that there is, all else being equal, no difference in severity between a 40- and 60-year-old. But one would be forced to accept that there is a difference in severity between a 30-year-old and a 30-year-and-a-day-old. What could be the reason for accepting the relevance of such small differences on certain parts of the line while ignoring large differences on other parts of the line? Reintroducing continuity on parts of the line seems more arbitrary than accepting it on all parts of the line.

In addition to the problems associated with arbitrariness and sharpness it also seems like a multiple threshold system will become very complex if we are to accept SD. If we

aggregation. Note that within Rawls' theory it is the maximin/leximin decision procedure that is clearly consequentialist in this sense. Rawls' system as a whole, especially considering his principle on basic liberties, is plausibly a hybrid view.

¹³We are ignoring the difficult questions regarding the disvalue of death for the very young, such as in Jeff McMahan's (McMahan 2019) Time-relative Interest Account. We are simply stipulating that the graphs start at 20 years.

¹⁴This would essentially be a simplified version of a life stages view. We believe that the arguments discussed here would apply to a more complicated life stages view as well.

make a series of cases such as 30 vs. 60, 40 vs. 70, etc. we would have to include thresholds between 30 and 60, 40 and 70, etc. It seems like no number of thresholds would be enough to accommodate our intuition. Introducing enough thresholds would approximate a continuous line, and thus fail to accommodate our intuition.

These two contrasting difficulties, between on the one hand having to find more or less arbitrary cut-offs and accepting that a difference of one year (or even a day) is relevant for considerations of illness-severity on the other, seem to be impossible to avoid within a consequentialist system with an underlying complete ranking of states of affairs. This leaves us with the options of either dismissing SD and our intuitive judgements or exploring (at least partly) deontological views. We first turn to deontological views.¹⁵

Severity as relevant worse off-ness

In the current Norwegian priority setting system funds are allocated on the premise of optimizing the numbers of Quality Adjusted Life Years (QALYs) per monetary unit, with extra weight on QALYs accrued to those with the most severe conditions.¹⁶ Even though they are weighed more heavily, the QALYs accrued by curing or treating, for instance, lethal cancer are weighed against the QALYs gained by curing or treating less severe conditions. If one were to choose between an intervention that saved the life of one patient from lethal cancer and an intervention that cured the once-a-month moderate headache of a thousand individuals the answer is not straightforward and would involve calculating the QALYs gained (weighed by severity) for each intervention. We believe that many would object to this calculation on the grounds that once-a-month moderate headaches are irrelevant when compared to the cancer patient's prospect of shortening her life, no matter how many people suffer from these headaches.¹⁷ In *Morality, Mortality* (1993) Frances Kamm introduces the notion of *irrelevant utilities*. An irrelevant utility is a utility that should not be considered in the process of deciding between competing claims for a resource. The notion of irrelevant utilities is dynamic. Whether a utility is irrelevant is dependent on what it is measured up against: in the context of who should live, a broken arm might be irrelevant. A broken arm, however, is not an irrelevant utility when measured against a broken leg.

The notion of irrelevant utilities seems to be able to buttress the intuition expressed in SD. Consider the option (or impossible choice) of either saving the life of one person, or the life of another plus, say, curing the intermittent headache of a third person. Is it right, in the context of choosing who lives, to prioritize based on the added utility of curing the intermittent headache? This issue goes to the heart of the current priority-setting framework and might also aid us in our quest to make sense of SD. If we assume that severity of illness has normative implications, then the question of who is more severely ill may not reduce to a question about who is worse off (severity 1), but also a question of who should be given higher priority when distributing scarce resources (severity 2). When claiming that the terminal illness of a 20-year-old is more severe than the terminal illness of an 80-year-old we are not *only* saying that the 20-year-old is worse off, but also that this worse off-ness is significant when deciding whom to aid. When we are claiming that the terminal illness of a 40-year-old is equally severe as for a

¹⁵Deontological in the sense of limits or constraints on moral aggregation. See Kagan (1992).

¹⁶A QALY is a life year adjusted for the health-related quality of life in that year. A QALY number is thus a function of the number of years and the health-related quality of life in those years.

¹⁷Varieties of this view can be found in Kamm (1993), Scanlon (2000) and Temkin (2014).

45-year-old, we might be claiming that even though the 40-year-old is worse off than the 45-year-old (severity 1) this difference in worse off-ness is not significant when deciding whom to aid (severity 2). The extra years are in this context something akin to an irrelevant utility, or rather a form of *irrelevant worse off-ness*. We can thus conceptualize *severity as relevant worse off-ness (SRW)* as illness severe enough to justify different priority setting. An advantage to conceptualizing the relationship between severity and age in this way is that we can assume a full ranking worse off-ness while keeping our intuitions that seem to display intransitivity. What we mean by a full ranking of worse off-ness is that it should be possible to rank all people from worst to best off in such a way as to make the worse off-relation a transitive relation (probably on a linear scale). On this scale the soon-to-be-dying 40-year-old is worse off than the soon-to-be-dying 45-year-old and would also be worse off than a soon-to-be-dying 41-year-old. If we consider some additional years as irrelevant utilities the question of whether the 40- or 45-year-old is more severely ill is not a question reduced to who is worse off, but rather whether one is worse off enough to make unequal treatment legitimate. The question of who is more severely ill will then be a function of whether there is a *morally relevant* difference in worse off-ness *in a priority setting decision*.

SRW thus has appeals but is not without problems of its own. Perhaps most notably, SRW works in the context of pairwise comparison, but not when confronted by a context where multiple decisions are to be made over time. To see why this is so imagine that you are making decisions on which treatments to fund in a publicly funded health care system. In addition to questions of cost and effectiveness, severity of illness is a priority setting criterion. Treatment 1 can give 30-year-old patients with a terminal illness an additional year of life. The cost is high but given the severity of the illness (surely terminal illness in 30-year-old patients is severe) you decide to fund the treatment. Now imagine that treatment 2 is like treatment 1 but aimed at 35-year-old patients. Illness at 35 is equally severe as illness at 30 ($30 = 35$), so (all else being equal) you should fund the treatment. Then treatment 3 shows up, aimed at 40-year-old patients. Our intuition is that terminal illness at 40 is equally severe as illness at 35 ($35 = 40$). Then treatment 4, and all the way up to treatment n aimed at 90-year-old patients, shows up. We thus end up assigning the same severity to 30- and 90-year-olds, even though we wish to claim that the illness of the 30-year-old is more severe than that of the 90-year-old; we want to say $30 > 90$ but end up saying not $30 \succ 90$ or alternatively $30 = 90$. Note that this problem is not dependent on a specific degree of overlap. If one thinks that 5 years is more than enough to consider one case of illness more severe than another, the same problem can be shown to result from any difference (one year for instance) given enough treatment options (treatment for 31-year-old patient, 32, 33 ... 80, etc.). SRW is thus an intransitive relation.¹⁸

Interestingly, the “collapse” of the severity relation would not come up if we only considered, for example, a 20-, 50- and 80-year-old patient.¹⁹ Let us stipulate here that the age difference has to be at least 20 years for a differential judgement of illness severity to be merited. We would then have no problem claiming that the 20-year-old is more severely ill than the 50-year-old, who is more severely ill than the 80-year-old. But once we include more alternatives, for instance a 35-year-old and a 65-year-old, the

¹⁸Strictly speaking the “more severe than”-relation and the “equally severe as”-relation within SRW will be intransitive. “More severe than” will be a transitive relation.

¹⁹How large the differences have to be depends on the intuitive judgements of the reader.

severity ordering collapses²⁰ and we are no longer able to prioritize the 20-year-old over the 80-year-old. This is clearly a violation of the independence of irrelevant utilities (my preference of apples over oranges should not be dependent on whether there are pears present). A few clarifications are in order. First, the problems just mentioned are dependent on the severity-relation being a total ordering. This means that, given any two patients A and B with any possible level of illness related worse off-ness, the illness of patient A is either more severe than the illness of patient B ($A > B$), less severe than the illness of patient B ($A < B$) or (either precisely or imprecisely) equally severe as the illness of patient B ($A = B$). This essentially means that all patients can be compared in terms of severity of illness.²¹ This might not be a completely uncontroversial assumption, but it is difficult to see how a severity-relation that fails to compare all relevant alternatives can function as a criterion for systematic priority setting. Second, intransitivity and sensitivity to irrelevant alternatives are only issues for the normative aspect of severity (severity 2). We are thus not arguing that worse off-ness is an intransitive relation.

The need for a complete ranking for priority setting purposes also rules out the concept of parity. According to Ruth Chang (Chang 2016) two items can be comparable in the sense of being *on a par*; they are comparable without being either more, less or equally X. Two instances of illness could thus be related such that neither is more severe than the other, without this implying that they are equally severe, yet still be comparable. The important aspect of parity for our purposes is that parity is an intransitive relation. A being on a par with B, which is on a par with C, does not imply that A is on a par with C. This intransitivity will give rise to the same problems regarding multiple decisions as SRW.²²

Some might be tempted to save SD by assuming that severity is what Temkin has called an essentially comparative concept. The essentially comparative view “implies that an outcome may have a certain value relative to one alternative-set, and another value relative to another alternative-set” (Carlson 2013). In our context this would entail that when we are comparing a 20- and 80-year-old this is one alternative set, with its own relevant factors. When comparing a 20-, 30- ... 80-year-old, this is another alternative set, with its own factors. However, in the context of priority setting we need a consistent frame of reference, a “privileged alternative set” (Carlson 2013). Systematic use of the more severe-than relation seems difficult, and probably unjust, if there is no common reference frame for priority setting.

What might a priority setting system, sensitive to severity and age, look like if one were to accept some form of relevance view while still maintaining transitivity and insensitivity to irrelevant alternatives for the purposes of multiple decisions? Here is a suggestion. Let A, B, C, D and E be five patients who are 20, 35, 50, 65 and 80

²⁰In the sense that all the illnesses are equally severe, and neither is more severe than the others.

²¹Depending on your definition of commensurability, full commensurability might also require precise cardinal comparability as well. This is probably also necessary if severity is to function as a systematic priority setting criterion on a large scale. Say that we claim that the illness of a 20-year-old is, *ceteris paribus*, more severe than the illness of a 40-year-old. We can do this without specifying how much more severe, but if severity is to function as a systematic criterion the difference has to be cardinally measurable and commensurable with other criteria of our health care setting. These difficulties will not be addressed in the current article.

²²Chang (2016) suggests that there are pragmatic pressures to accept commensurability where there might be no real commensurability. Our desire to make health care priority setting systematic might function as such a pragmatic pressure.

years old respectively, all with an expected survival rate of one year due to illness. Let us once again stipulate that 20 years is enough to warrant a differential judgement of illness severity. If we had to choose between giving an extra year of life to either A or B on the grounds of severity, we would not be permitted to choose A because she is younger, but we would be permitted, or perhaps even required, to save A over C, D or E. So as far as illness severity goes $A = B$ & $A > C, D$ and E. Similarly for B: $B = A$ and C & $B > D$ and E. For C: $C < A, C = B$ and D & $C > E$. So, C should be given priority over E, be treated equally to B and D, and be prioritized less than A. This would let us make single decisions, but how would this play out as a priority setting criteria with multiple decisions? Let cohort stand for each group of patients who are the same age. Each cohort X would stand in one of three relations to all other cohorts: $X = Y, X > Y$ or $X < Y$. Let us call the $X = Y$ relations equality-relations, the $X > Y$ relations more severe than-relations and the $X < Y$ less severe than-relations. For any cohort X, the number of these relations will be a function of the cohort's age. A young cohort will have few (if any) less severe than-relations, comparatively few equality-relations, and many more severe than-relations. A middle-aged cohort will have more less severe than-relations, the highest number of equality relations and fewer more severe than-relations. An old cohort will have many less severe than-relations, comparatively few equality-relations, and few (if any) more severe than-relations.



Figure 4. Ratio of more severe than- to less severe than-relations as a function of age.

Figure 4 shows how the ratio of more severe than-relations to less severe than-relations changes as a function of age. In other words, it maps how severe terminal illness is in a cohort relative to all other cohorts. The graph would look different if we included equality relations because the middle-aged have comparatively more equality-relations than either the young or the old. Prioritizing based on these relations would mean giving higher priority to 40-year-olds than to 50-year-olds because the 40-year-olds should be given priority relative to a larger part of the population than the 50-year-olds. This is so even though if this were a singular decision, we would not choose between them based on age. This would paradoxically reintroduce the continuous function of age and severity on a priority setting level. We would in effect be mindful of the fact that priority setting is a continuous process and always bear in mind not just singular decisions, but the larger framework. It seems that this does

not preserve the spirit of the relevance view. Or, to put it somewhat differently, it “saves” the relevance view by making it practically irrelevant.

It thus seems that the relevance view and SRW can accommodate SD, but that SRW is not suitable for priority setting on a systemic level. The possibility of devising a cohort-based system remains but seems to make both our intuitive judgements and SD irrelevant. The last possibility to be discussed is whether we should dismiss our intuitive judgements and SD.

Should we dismiss SD?

So far, we have explored two ways of making sense of SD, neither of which seems satisfactory. A multiple threshold model will either be plagued by arbitrariness or approximate a continuous line. SWR is troubled by intransitivity and sensitivity to irrelevant alternatives and is as such not suited for systematic priority setting. A third alternative is to dismiss SD and accept a linear relationship between age and severity.

What reasons do we have for dismissing SD? First, there is the possibility that the intuitive judgements regarding our four cases are not about relative age differences per se, but rather stem from other attributes that we have failed to adequately control for. We might simply not be comfortable making relative judgements of worse off-ness unless the differences are obviously large. If faced with a 40- and 45-year-old we might not be sure whether the 45-year-old has really either had more or stands to lose more in terms of good life years (for instance, the 45-year-old may have had more life-years with really poor quality of life compared to the 40-year-old). When faced with a 20- and 80-year-old, the difference is, or at least seems, obvious. Thus, even though we have assumed all else equal to a maximum extent in the cases, and thus made it clear who is worse off, our intuitions might not be sensitive to this kind of “fine tuning”. If this is true, then our intuitions might be a result of uncertainty regarding worse off-ness rather than a result of beliefs regarding relevance. A similar argument can be made regarding life stages. Even though we have stipulated that the patients are similar in all relevant aspects other than age we might not be able to exclude intuitions based on ideas of the life stages of the different age groups. We might for example be unable to make a differential judgement of severity between the 40- and 45-year-old because they are, at least in our minds, in the same stage of life. Excluding this by stipulating that the patients are alike in all relevant aspects makes theoretical sense but might be difficult in practice. These are essentially debunking arguments; our intuitive judgements of X might be better explained without postulating Y (Tersman 2008). Debunking is a contested field, but the specific intuitions discussed in this article seem especially suitable for a debunking explanation. This is so, again, because there are obvious confounders, such as uncertainty regarding value difference between “close enough” alternatives.

A second reason for dismissing SD is that there seems to be something wrong with the very notion of irrelevance when applied to life years. Life years (and preferably good life years) are, in a very basic sense, all we have. A year can mean watching your child learn to walk, mastering a skill, saying a proper goodbye to your loved ones, reconciliation and so much more. Even a day can be significant in the scheme of an entire life. From this perspective, claiming that any significant span of time is insignificant in the scheme of justice may seem perverse.

A third reason, perhaps the weightiest, is that SD forces us to accept either some form of exceedingly complex threshold model or a conceptualization of severity that is

intransitive and sensitive to irrelevant alternatives. In light of these consequences, accepting a continuous relationship between severity and age might not seem so unappealing.

Conclusion

This exploration of various ways of accounting for our intuitive judgements seems to suggest that we are faced with three options, none of them completely satisfactory. The first option is to accept that the judgements of severity are intransitive and sensitive to irrelevant alternatives. This would involve either a great deal of difficulty when applying the concept in priority setting work or using the concept in a way that makes the relevance view irrelevant for priority setting (as in the cohort system described earlier). Given that transitivity and insensitivity to irrelevant alternatives are typically considered prerequisites for rational decision-making, it also seems like a theoretically problematic option. The second option is to develop a complex threshold model. If the core of the intuition is SD, then the number of thresholds would be high, and theoretically limitless. The third option is to disregard SD and assume a continuous relationship between severity and age. This would force us to accept that the illness of a 40-year-old is, all else equal, more severe than the illness of a 41-year-old. Considering that time spent alive matters to us, this seems like a better option than the alternatives.²³

Declaration of competing interests. We have no competing interests to declare.

References

- Barra, Mathias, Mari Broqvist, Erik Gustavsson, Martin Henriksson, Niklas Juth, and others. 2020. Severity as a Priority Setting Criterion: Setting a Challenging Research Agenda. *Health Care Analysis*, 28.1: 25–44 <<https://doi.org/10.1007/s10728-019-00371-z>>.
- Carlson, Erik. 2013. Intransitivity. In *International Encyclopedia of Ethics* (American Cancer Society) <<https://doi.org/10.1002/9781444367072.wbiee410>>.
- Chang, Ruth. 2016. Parity, Imprecise Comparability and the Repugnant Conclusion. *Theoria*, 82.2: 182–214 <<https://doi.org/10.1111/theo.12096>>.
- Franken, Margreet, Elly Stolk, Tessa Scharringhausen, Anthonius de Boer, and Marc Koopmanschap. 2015. A Comparative Study of the Role of Disease Severity in Drug Reimbursement Decision Making in Four European Countries. *Health Policy*, 119.2: 195–202 <<https://doi.org/10.1016/j.healthpol.2014.10.007>>.
- Gustavsson, Erik, and Niklas Juth. 2019. Principles of Need and the Aggregation Thesis. *Health Care Analysis*, 27.2: 77–92 <<https://doi.org/10.1007/s10728-017-0346-6>>.
- Harris, John. 1985. *The Value of Life: An Introduction to Medical Ethics* (London: Routledge) <<https://doi.org/10.4324/9780203005606>>.
- Hirose, Iwao. 2014. *Egalitarianism* (London: Routledge) <<https://doi.org/10.4324/9781315772004>>.
- Kagan, Shelly. 1992. The Structure of Normative Ethics. *Philosophical Perspectives*, 6: 223–42 <<https://doi.org/10.2307/2214246>>.
- Kamm, Frances M. 1993. *Morality, Mortality. Volume I: Death and Whom to Save From It*, BRL:KIE/42748 <<https://repository.library.georgetown.edu/handle/10822/743570>> [accessed 7 June 2021].

²³We would like to thank two anonymous reviewers for helpful comments. The thesis of this article was presented at the Philosophy of Priority Setting seminar on 1 June 2021. We thank the participants for valuable discussion and comments. We would also like to thank Mathias Barra and Carl Tollef Solberg for helpful comments on a previous version of this article. The work of the authors is funded by two research grants: Norwegian Research Council 2020-303724: Severity and priority setting in health care, and Swedish Research Council 2021-01266: Distributive justice in health-care – why should severity of ill health matter? (JUST SEVERITY).

- McMahan, Jeff.** 2019. In Defense of the Time-Relative Interest Account: A Response to Campbell. In *Saving People from the Harm of Death* (Oxford University Press), pp. 267–78 <<https://doi.org/10.1093/oso/9780190921415.003.0020>>.
- Nagel, Thomas.** 1970. Death. *Noûs*, 4.1: 73–80 <<https://doi.org/10.2307/2214297>>.
- Parfit, Derek.** 1991. *Equality or Priority* (University of Kansas, Department of Philosophy) <<https://kuscholarworks.ku.edu/handle/1808/12405>> [accessed 7 June 2021].
- Rawls, John.** 1999. *A Theory of Justice: Revised Edition* (Cambridge, MA: The Belknap Press of Harvard University).
- Scanlon, Thomas.** 2000. *What We Owe to Each Other* (Cambridge, MA: Harvard University Press).
- Solberg, Carl Tollef, and Espen Gamlund.** 2016. The Badness of Death and Priorities in Health. *BMC Medical Ethics*, 17.1: 1–9 <<https://doi.org/10.1186/s12910-016-0104-6>>.
- Temkin, Larry S.** 2014. *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford University Press).
- Tersman, Folke.** 2008. The Reliability of Moral Intuitions: A Challenge from Neuroscience. *Australasian Journal of Philosophy*, 86.3: 389–405 <<https://doi.org/10.1080/00048400802002010>>.
- University of Pittsburgh, Department of Critical Care Medicine, School of Medicine.** 2020. *Allocation of Scarce Critical Care Resources During a Public Health Emergency* <Allocation of Scarce Critical Care Resources During a Public Health Emergency> [accessed 23 December 2021].
- Williams, Alan.** 1997. Intergenerational Equity: An Exploration of the “Fair Innings” Argument. *Health Economics*, 6.2: 117–32 <[https://doi.org/10.1002/\(SICI\)1099-1050\(199703\)6:2<117::AID-HEC256>3.0.CO;2-B](https://doi.org/10.1002/(SICI)1099-1050(199703)6:2<117::AID-HEC256>3.0.CO;2-B)>.