Discrimination and Fairness in QALY-Based Healthcare Allocation

Sean Campbell Sinclair

Submitted in accordance with the requirements for the degree of
Doctor of Philosophy

The University of Leeds

School of Philosophy, Religion and History of Science

Inter-Disciplinary Ethics Applied Centre

March 2018

The candidate confirms that the work submitted is his/her own and that appropriate credit has been given where reference has been made to the work of others.

Acknowledgements

I would like to thank my family for their support and understanding while I was researching and working on this dissertation. Esther, Lauren, and Harry, plus Eric and Irene, it's been great to have you with me on this journey! Special thanks to Esther for chivvying at the right times, for giving me lots of ideas, and for taking on many tasks that I ought to have done. I would also like to thank my supervisors (Chris, Rob and Carl) for really useful comments. If I've developed, it's down to you. Thanks to colleagues at the IDEA CETL for support, including helping with my other work in the final weeks so I could finish this. Thanks to everyone else who gave me feedback on bits of the thesis: it is vastly improved for all those comments and discussions. And thanks to all my other colleagues and friends for all your help and motivational comments! Finally thanks to AHRC for funding my PhD studentship.

Abstract

This thesis addresses the question of what criteria we should apply in healthcare allocation. I accept that benefit maximisation is an important pro tanto aim, but I also argue that other pro tanto aims and constraints are important too.

I argue that we should base our assessment of benefit on deliberative assessments of how much a health outcome constrains the possibilities of living well, usually done by patients in the state unless there are reasons to doubt the reliability of their preferences. However, we should confine our attention to the effects of the disease we are treating or its treatment and ignore symptoms of conditions unrelated to the condition we are treating.

Once this is established, each chapter then defends a different principle that must be weighed against benefit maximisation. First, I argue that we should not allow healthcare allocation policies to be influenced by false or unwarranted judgments regarding groups who lose out by such policies. We should also avoid reinforcing any sense that people with historically disadvantaged traits are excluded. I argue that we should relax our normal criteria in the case of rare diseases, when deciding whether to licence and fund treatments for them on the basis of a principle of inclusiveness. I suggest that perhaps egalitarian considerations give us reason to avoid a sharp cost-effectiveness threshold that discriminates sharply between patients just above and just below the threshold (this could be by instituting a lottery for patients whose cost-effectiveness is near the threshold). I argue that we should give extra weight to QALYs for patients whose quality-adjusted life expectancy, without treatment, is less than society's average. I also argue that we should prioritise patients who have had very little notice of their expected death for life-extending treatment.

Table of Contents

1. Introduction

Governments have finite resources and therefore cannot provide all the healthcare that their citizens might want. Healthcare policy-makers must refuse funding for some treatments, even treatments that are effective. This raises the question, which treatments should we fund and which should we refuse? One popular answer to this has been to fund only the most cost-effective treatments, with the aim of producing the maximum benefit per unit of cash. For example, the UK's National Institute of Clinical Excellence (NICE) assesses treatments by comparing their cost with how much benefit they produce, where benefit is measured in terms of quality-adjusted life years (QALYs).[1] In fact, this was NICE's sole formally-stated criterion from its founding in 1999 until it introduced the end of life premium in 2009.[2] And even after 2009, cost-effectiveness appears to have been NICE's main criterion. Dakin et all found that "Cost-effectiveness alone correctly predicted 82% of decisions" and that few other variables correlated reliably with NICE's decisions.[3]

This thesis addresses the question of what criteria we should apply in healthcare allocation. I do not reject the aim of benefit maximisation. I take it to be one pro tanto consideration. But I will argue that there are other relevant considerations. Essentially, in each chapter, I will defend a different principle that must be weighed against benefit maximisation.

In doing so, I will appeal to a form of Rossian pluralism. W.D. Ross believes that we have a variety of pro tanto[4] duties. For example, suppose I make a promise and I do a wrong to someone. I may therefore have a duty to do what

---

1 I will examine QALYs in more detail in Chapter 3 on how to understand benefit.

2 NICE (2018), Cookson (2013)

3 Dakin et al (2015)

4 Actually Ross calls uses the term "prima facie", but the term "pro tanto" is now generally preferred. Shelly Kagan writes: "A *pro tanto* reason has genuine weight, but nonetheless may be outweighed by other considerations. Thus, calling a reason a *pro tanto* reason is to be distinguished from calling it a *prima facie* reason, which I take to involve an epistemological qualification: a prima facie reason *appears* to be a reason, but may actually not be a reason at all, or may not have weight in all cases it appears to. In contrast, a pro tanto reason is a genuine reason – with actual weight – but it may not be a *decisive* one in various cases." [4] He also says that in this he departs from the "unfortunate terminology" proposed by Ross, which Kagan says has invited "confusion and misunderstanding" (Kagan, 1989, p. 17). I join Kagan in adopting the term "pro tanto" instead of the term "prima facie".

I promised and a duty to repair the wrong. Ross finds that "neither of these reasons is reducible to the other".[5] Of course, these two duties may conflict; an action may be wrong in respect of its being the breaking of a promise but right in respect of its being the repair of a wrong. But this does not lead to moral paralysis. We must weigh the conflicting duties up. We will find that on some occasions, our overall duty may be to fulfil one pro tanto duty, and on other occasions the other. In defence of his pluralist analysis, Ross argues that "If, as almost all moralists except Kant are agreed, and as most plain men think, it is sometimes right to tell a lie or to break a promise, it must be maintained that there is a difference between [pro tanto] duty and actual or absolute duty".[6]

So in what follows, I accept that benefit maximisation is an important pro tanto aim, but I also argue that other pro tanto aims and constraints are important too. I will return to the question of how we should weigh them up in Chapter 9.

1.1 Summary of chapters

In Chapter 2 on discrimination, I discuss a quite general constraint on discrimination[7] in healthcare allocation. If we're going to refuse funding for some treatments, we will have to discriminate between the patients who win and those who lose. We therefore need to avoid unfair discrimination. It may seem that such an innocuous criterion as cost-effectiveness is immune from accusations of unfair discrimination, but in fact it is not straightforwardly invulnerable. For example, suppose a new skin cancer treatment has been launched. For white patients the cost-effectiveness is just below the cost-effectiveness threshold, giving them a prima facie case for funding, but for black patients the cost-effectiveness is just above the threshold. Should we treat the white patients and turn the black patients away? There might be reasons to say "No" (implying that we should either fund both, or refuse both, producing less benefit either way). In chapter 2, I argue that the way an act is motivated plays a critical role in determining whether or not it constitutes unfair discrimination. For example, in the case of the skin cancer treatment,

---

[5] Ross (2000), p24

[6] Ross (2000), p28

[7] In my usage, it is not analytic that discrimination is impermissible. For example, it may be permissible for an employer to discriminate between job candidates by intelligence where that is required for the job.

discrimination might be unfair if healthcare allocators are motivated by negative attitudes to the group they discriminate against, using the slight difference of cost-effectiveness as a pretext for discriminating between the two groups of patients (perhaps simultaneously funding other patient groups whose cost-effectiveness is at a similar raised level). On the other hand, if healthcare allocators are not motivated by negative attitudes and they discriminate here because the cash saved could make a big difference elsewhere, then there might be circumstances in which such discrimination could be defended.

Then in Chapter 3 on how to understand benefit, I consider a question which many benefit maximisers take for granted. What kind of benefit should we aim to produce? The UK's NICE and many other health systems assume that benefit consists of preference satisfaction. But Daniel Hausman argues that this misses the point - health states are valuable or not according to how much they constrain the possibilities of living well. I endorse Hausman's account in outline, but I propose three modifications. First, Hausman does not offer an account of wellbeing, but I argue that if we don't know what it means to say someone lived well, we can't say whether or not a given health state constrains the possibilities of living well. I therefore offer an account of wellbeing. Secondly, I argue that Hausman needs more descriptive parameters; we cannot adequately characterise all health outcomes with the two descriptive parameters he proposes. And thirdly, I argue that Hausman over-simplifies the way our overall evaluations of outcomes derive from their non-evaluative features. Hausman cannot capture the value of all health outcomes by evaluating just a few outcomes and then deriving an algorithm which implies the value of other outcomes. Nevertheless, I also acknowledge that in the real world of policy-making, generalisations will be necessary, so there is a balance to be struck between getting preferences regarding all possible health outcomes and being guided by general principles.

Then in Chapter 4 on rare diseases, I argue that certain egalitarian considerations highlighted in an argument of John Taurek's may sometimes outweigh benefit maximisation. As an intuitive example, suppose the entire population is threatened by a fatal flu virus and we have to choose between two vaccines. One is more than twice as effective as the other, and costs twice as much, so it is more cost-effective. By using this very effective vaccine we'll give ten years extra life expectancy to half the population. By contrast, with the other vaccine, the cheaper one, we'll only give 4½ years extra to each recipient - but we'll be able to afford to give the vaccine to everyone. Many people will say we should choose the vaccine that gives a bit of benefit to everyone rather than the one that gives a lot of benefit to only some people, even though the result will be that we produce 10% less benefit.

However, this Taurek-derived consideration is quite general, applying beyond rare diseases. I go on to offer a more specific defence of orphan drugs policies. These are policies to facilitate access to medicines for patients with rare diseases. Drugs for rare disease patients tend to be expensive, because of problems with economies of scale in R&D, so they would not get funded without such policies. If healthcare allocation policy is driven purely by benefit-maximising principles, most rare disease patients will be permanently excluded from access to new treatments (in fact, these diseases would be "orphaned" by drug developers - as we saw before orphan drugs policies were introduced). In contrast, I argue that a consideration of egalitarianism or inclusiveness gives us reason to introduce orphan drugs policies, such as relaxing the cost-effectiveness thresholds for rare disease patients. I offer an analogy with decisions about where to locate GP clinics and fire stations and schools. In such decisions, we do not only aim to produce the greatest benefit. If we did, the facilities would mostly be built in locations accessible to large populations, viz., towns and cities. However, a proportion of facilities are built in rural areas, at some cost in terms of benefit maximisation. We are similarly concerned to achieve close to universal coverage for postal services, phone networks, power networks, and so on. So, although we do not completely ignore questions of benefit maximisation, we also respect a consideration to do with inclusiveness, which gives us a pro tanto reason to share out the important benefits of society amongst all of society's members. Rare disease patients can argue that a similar pro tanto consideration applies in their case. Once again, we see that benefit maximisation should not be our only aim in healthcare allocation.

Then in Chapter 5 on disability discrimination, I consider another, well-known challenge against focusing exclusively on benefit maximisation, this being that we could end up discriminating against disabled patients. The problem starts because most people, especially non-disabled people, consider that people with disabilities have a lower health-related quality of life than someone without disabilities. This valuation implies that a treatment that extends life produces less benefit in someone disabled than someone who is not disabled, and this in turn means that (assuming the valuations are correct) the treatment will be less cost-effective in disabled people. If we focus exclusively on maximising benefit, this could result in a treatment being funded for non-disabled people and not for disabled people. Most people would find this highly objectionable. We might prefer not to be disabled but that doesn't mean we want to see disabled people stand less chance of life-saving treatments. I argue that this is another case where benefit maximisation must give way to other considerations. But what principled case can we make for a qualification on benefit maximisation in this case?

I start with a general argument that certain benefits we could produce are simply not relevant in determining who should get treatment, even as pro tanto considerations. As an extreme example, take non-health benefits. We should ignore them. We shouldn't save the wealthy patient instead of the poor patient just because the wealthy person is living *la dolce vita*; we should focus on people's *health-related* quality of life only. Also, we should not take account of the good people can do for others. Frances Kamm points out that we should not favour the great philosopher over the non-philosopher for a life-saving treatment just because we can expect valuable discoveries from the philosopher.[8] So in a number of cases, we seem to respect a principle involving what Kamm calls "separate spheres".[9] Healthcare policy-makers may only take account of health-related benefits in allocating health. This may mean producing less benefit than we could.

In a similar vein, it can be argued that some health benefits and health impairments are also irrelevant. For example, Kamm argues that if we have a chance of saving someone's life, but we have to choose between them and saving someone else's life *and* curing a third person's sore throat, the opportunity to cure the sore throat should not sway us: "I believe it would be wrong to deprive A of his 50% chance to be saved simply in order to get the extra utility of curing C's sore throat associated with saving B."[10] The chance to cure a sore throat seems too trivial to enter into our consideration of whose life to save.

Finally, in relation to disabilities, I offer an argument that certain disabilities should be ignored. To argue this, I offer a new defence of the separate spheres principle and argue that we must ignore some of the benefits that we could produce.

---

[8] Kamm (1993), p. 258

[9] Kamm (1993), p. 259

[10] Kamm (1993), p. 101

In Chapter 6 on age discrimination, I address another possible consideration that could count against benefits: the age of patients. Age discrimination is illegal in most countries, but Norman Daniels argues that some age-based allocation policies are not really discriminatory, despite appearances. He argues that it is possible for a stable system which discriminates by age to avoid treating people differently, as long everyone of each age always gets the same entitlements. The reason is that everyone can hope to pass through each age. So when you consider each person's life as a whole, they end up getting the same entitlements. Since such a system does not treat people differently, from a whole life perspective, there is no discrimination: "Differential treatment by age, over time, is not unequal treatment of persons, even if it is unequal treatment of age groups on each occasion, at each moment.[11] Finally, with a lot of caveats, Daniels argues that it might be in people's interests to have a healthcare system that spends more per unit of benefit on younger people than elderly people, because younger people have not yet had a chance to complete their most important projects. I endorse this as a pro tanto consideration, even though such a policy could result in healthcare producing less benefit overall. However, in response to Daniels, I argue that citizens' self-interested concerns cannot explain everything a healthcare system ought to do. In a case where the interests of a minority diverge from the interests of the majority such that the majority has no interest in relaxing the cost-effectiveness threshold for younger patients, it difficult for Daniels to explain why the majority should fund healthcare for the minority, because this would involve interpersonal transfers from the majority to the minority. Daniels' explicitly excludes such interpersonal transfers from the scope of his account. This limitation means Daniels' account has counterintuitive consequences for certain cases. I conclude that although self-interest might be one consideration determining the healthcare priorities we ought to adopt, it is not the only consideration.

Then in Chapter 7 on fair innings, I offer another age-related factor to be considered when we formulate our healthcare priorities. Fair innings theories traditionally imply favouring the young over the old, such as saving the 20 year-old rather than the 70 year-old. This exposes them to accusations of unfair discrimination. Daniels adopts one response to such accusations; I take a different approach. I argue that at base, fair innings intuitions do not

---

[11] Daniels (1988), p. 42. This idealisation ignores' start-up' problems that arise when we begin such policies, as Daniels clarifies in Daniels (1989).

motivate favouring younger patients over older patients, at least not directly. Rather they motivate favouring patients who are expected to die young over patients who are expected to die old. This is on the basis that patients expected to die old are better off. Klemens Kappel & Peter Sandoe argue that, "To give the liver to the older person rather than the younger is like giving money to the rich instead of the poor."[12] However, if an older patient is expected to die at a younger age than a younger patient, then the younger patient would be "richer" in the way they describe, and so (contrary to their assumption), fair innings considerations might justify treating the older patient before the younger. So a fair innings theory need not entail direct age discrimination. But again, if the patient expected to die young stands to gain less benefit than the patient expected to die old, by adopting a fair innings policy we may produce less benefit than we could.

And finally in Chapter 8, I consider whether a "rule of rescue" obligation is ever applicable in healthcare allocation. By this I mean (roughly) an obligation to fund life-extending treatment for someone who faces imminent death, without applying our normal criteria in terms of cost-effectiveness or benefit maximisation. I start by considering some cases from outside healthcare allocation where we fail to maximise welfare. These are the cases such as the Chilean government's 2010 rescue of the miners stuck down the Copiapó mine, or the Australian government's 1997 rescue of the lone yachtsman Tony Bullimore, lost in the southern ocean after his boat had capsized. In such cases, governments spend well above their normal cost-effectiveness threshold to rescue the prospective victims, and thereby end up producing less benefit from that budget than they could have done otherwise. But can such "rule of rescue" behaviour be defended? I argue that extant defences fail, but I go on to offer a new defence which vindicates this behaviour, with implications for healthcare allocation that include non-benefit maximisation.

So although I argue that we have a pro tanto reason to maximise benefit, I also argue that a number of other considerations are important too. Sometimes these other reasons will pull against benefit maximisation, or against each other. As noted, this means there will be times when we have to decide how these conflicting considerations balance out. I address this question in chapter 9 on trade-offs. I consider whether we should rely on general trade-off algorithms, or instead deliberate over the best policy case by case. Alan Williams and Richard Cookson, both health economists, defend an approach

---

[12] Kappel & Sandoe (1992)

based on general trade-off algorithms, i.e. general principles which state how conflicting considerations are to be traded off across a great variety of cases. In contrast, the philosopher Stephen Toulmin argues that we should judge how conflicting considerations trade off in individual cases without trying to state a general trade-off principle.[13] I will argue that we should pursue Williams & Cookson's strategy of stating general trade-off principles, but without assuming it will always be possible. In line with this, I consider a process called Multi-Criteria Decision Analysis (MCDA), which involves getting policy-makers and stakeholders to deliberate over criteria and their relative weights with a view to producing an algorithmic policy-making model. I argue that on the evidence available about MCDA, the process of agreeing an algorithm can produce better quality deliberation, so even the opponents of algorithms have some reason to support them.

I then conclude the thesis with an outline of policy proposals.

1.2 Method

It will be noted that in many of the arguments I summarise above, I rely on intuitions, such as the intuition that the Chilean government ought to rescue the miners. This raises a question of method. How do I go about determining what the right thing for policy-makers to do is? Here it is helpful to set my approach in the context of the alternative ways I *could* have derived my answers. Arras (2016) maps three broad options for deriving conclusions in bioethics: high theory, anti-theoretical casuistry, and mid-level theorising. I endorse Arras's view that the right approach involves mid-level theorising. I endorse his objections to the other two alternatives. To take them in turn, Arras first objects to high level theorising. As examples of this approach, he cites:

> Joseph Fletcher's (1974) and Peter Singer's (1999) utilitarian approaches to the whole spectrum of bioethical issues, [and] Alan Donagan's explicitly Kantian argument for informed consent in medical practice and research (1977).[14]

Arras offers several objections to this approach as a basis for reaching conclusions in bioethics. One is that there is no prospect in sight of theorists

---

[13] Actually Toulmin does not address healthcare allocation specifically, but his argument is a general argument which applies across a variety of areas of policy-making.

[14] Arras (2016), citing Fletcher (1974), Singer (1999), and Donagan (1977)

agreeing *which* high level theory is right. Another objection is that high level theories do not obviously entail the kind of concrete conclusions regarding specific cases that we need in bioethics. He gives the example of Norman Daniels. He says:

> [Daniels'] work on the theory of just health care has constituted an attempt, sustained over several decades, to develop an explicitly Rawlsian account of just access to health care and the social determinants of health. Although Daniels at first held out the hope that his theory, based upon a robust account of equal opportunity, could provide the requisite guidance for social policies bearing on access to health care and rationing, he now explicitly acknowledges that philosophical theory is not sufficiently fine-grained for such concrete policy making and must be supplemented by justly structured political deliberation.[15]

A third objection is that high level theories have counterintuitive implications. Arras says of Peter Singer's allegedly weak response to an objection based on one such implication:

> Utilitarians often try to finesse such objections by means of various rule-based strategies - e.g., all of us will ultimately be better off if we are allowed to favor close relatives in certain circumstances - but such attempts to "save the phenomena" of ordinary morality often lack plausibility.[16]

This brings us to the other end of the spectrum. Arras summarises the casuistic method as follows:

> moral certitude (or our best approximation thereof) is to be found in so-called paradigm cases, where our intuitions are most strongly reinforced. Moral analysis of a given situation begins, then, with a scrupulous inventory of the particular facts of the case ... This nexus of particulars is then compared with the details operative in one or more paradigm cases - i.e., clear-cut examples of right or wrong conduct. ... Eventually, as we discover in the analogous common law tradition, we end up working our way through many related cases over time, and along the way generate a sophisticated typology of cases and governing paradigms that provide us with a rich repository of values for social criticism.[17]

---

[15] Arras (2016), citing Daniels (1996) and Daniels (2007)

[16] Arras (2016)

[17] Arras (2016)

Some theorists take this method to an extreme, which Arras calls "strongly particularist casuistry". As an example, Arras considers Toulmin:

> For strongly anti-theoretical casuists like Stephen Toulmin, the suspicion of theory extends even to mid-level bioethical principles, which, he argues, serve no justificatory function. In contrast to mainstream bioethical thinkers like Beauchamp and Childress - and even in contrast to more mainstream casuists, like his co-author, Albert Jonsen (1995) - for whom justification involves, inter alia, bringing actions or policies under various specified ethical principles or maxims, Toulmin contends that moral principles serve only an heuristic function; that is, they serve primarily to remind us of salient features of past decisions. Principles are, as it were, the ribbon we wrap around decisions we have already come to on the basis of particularistic casuistical reasoning.[18]

And here, Arras says,

> hard-core bioethical casuistry converges with the epistemology of moral particularism as forcefully elaborated in the work of Jonathan Dancy ... moral particularists like Dancy [deny] that the moral valence of any particular element must remain constant from one case to another. In other words, they would contend that in some situations lying might be positively good, not merely a bad to be outweighed by some other element of the situation.[19]

Arras summarises the position as follows:

> For strong particularists and hard-core bioethical casuists, then, justification in ethics will not depend upon bringing a set of facts under a suitably interpreted general principle; rather, justification will be a matter of all the discrete elements of a particular decision fitting together or "adding up" holistically in the right way. In some cases, lying will have a positive moral valence, while in others, truth-telling may have a negative valence; everything depends upon the particular constellation of circumstances presented by the case. At least with regard to the business of moral justification, then, there's no room for generalizations of any sort within this kind of strong particularist epistemology.[20]

Arras's objection is that

---

18 Arras (2016), citing Jonsen (1995)

19 Arras (2016)

20 Arras (2016)

> analogical reasoning is not self-directed. It requires principles or maxims, a sense of what's ethically relevant ... Generalizations or principles also provide us with the crucially important understanding of what's morally relevant and why, which drives analogical reasoning forward. ... It is unclear that Dancy's reliance upon nuanced moral perception and narrative epistemology can really provide us with a plausible, let alone serviceable, notion of moral justification. If someone asks us for a moral justification of our stand on a particular issue, they are most likely going to be unsatisfied with such responses as: "All the facts just seemed to me to add up in a way that yields this conclusion," or "My heightened skills of moral perception indicate that this is the correct judgment to make about this particular constellation of facts." Instead, we will most likely want to hold out for some sort of inference or argument that moves from some sort of moral generalization (e.g., "lying is wrong") to an all-things-considered judgment about this particular instance of lying (which might gain justificatory support from other elements of the situation).[21]

I find this objection convincing. Moral discourse is essentially reason-giving, and it is not clear that a hard-core particularist approach can explain why our reason giving goes as it does. For Dancy and Toulmin, an observation that a situation has feature x is not enough to justify the thought that we have reason to do something or not do it, by itself; but in fact, we often take such a single feature to clearly count for or against an action, with no further analysis required. And even granting Dancy that, say, lying can change moral valence in some cases, this is not to concede that *all* morally relevant features are changeable in this way. To the extent that we can precisely characterise features that carry their positive or negative weight universally, we should do so. It will represent a philosophical insight and will help policy-makers and other moral agents do the right thing more often. On the other hand, we should also grant Arras that moral agents should assess "the particularities of moral situations in all their individuality and complexity" and not carelessly assume that because a situation superficially satisfies some general principle, the right thing to do is obvious.

So, in the following, I aim to state mid-level principles involving the features I have mentioned above - rarity, disability, and so on. My method is somewhat casuistic, but I do not treat intuitions regarding small sets of cases

---

[21] Arras (2016)

as entirely decisive. Rather, I pursue reflective equilibrium.[22] I consider intuitions regarding specific cases, develop a mid-level theory encompassing those cases, and then modify either the theory or my judgments regarding individual cases until the implications of the theory match my judgments regarding specific cases. Sometimes this method will require me to give up intuitions regarding a case for sake of a principle that explains a lot of other cases and is intuitively compelling in its own right.

I should also note that I do not rely on my intuitions alone. In many places I also appeal to existing policies and legislation, as well as data from public opinion surveys - for example, in the chapters on discrimination, rarity, disability, fair innings, and rule of rescue. I hope my conclusions are made more robust by this appeal to the intuitions of a wide, mixed audience, many of whom have had a chance to think carefully about the issues in question.

---

[22] Rawls (1971)

2. Is discriminating by prognosis wrong in the same way as racist discrimination?

When we allocate healthcare by cost-effectiveness, we sometimes refuse funding for a treatment because it does not produce as much benefit as other treatments we could fund with the same money, such as when it produces fewer life years or a smaller quality of life benefit.[23] In fact, I will consider principles entailing such policies in the next two chapters. But before that, I must consider an objection to any such principle from John Harris. He argues that funding a treatment for one patient and not another because we expect the first patient to benefit more is wrong in the same way as racism,[24] sexism etc. In the following, I start by giving Harris's argument. I find that Harris does not do enough to show that such discrimination[25] is wrong in the same way as racism/sexism. What's needed to decide the issue is a clearer account of what makes acts of racist/sexist discrimination wrong. I work towards such an account by offering prima facie reasons for thinking that the way an act is motivated plays a critical role in determining its moral status. I then set out to defend an account from Larry Alexander according to which the reason wrongful discrimination is wrong is that it is based on false or unwarranted judgments regarding the group who lose out by the discrimination. I then notice that the account does not imply an attractive result for certain intuitively wrongful acts of discrimination. I consider one option in response, which is to adopt a more inclusive account offered by Richard Arneson. But this has counterintuitive consequences of its own. Instead, I find that Alexander can fend off my problematic cases by highlighting a variety of wrong-making features of the relevant acts of discrimination which he can recognise consistently with his account. These features do not make the

---

[23] Life years and quality of life will often be combined into a single measure of benefit, namely quality-adjusted life years or QALYs. I will cover these in more detail in a later chapter.

[24] I primarily use the terms "racist" and "racism" in this paper to apply to acts, not people. I use them to refer to discrimination that is morally wrong because it is based on race. I leave open for now whether it's possible for discriminating by race to be permissible in some circumstances. The same comments apply, mutatis mutandis, to the term "sexism".

[25] In my usage, it is not analytic that discrimination is impermissible; for example, it may be permissible for an employer to discriminate by intelligence in recruiting for a job that needs it.

discrimination intrinsically wrong in the same way as racism/sexism, but they make it wrong in other ways. I therefore adopt Alexander's account. I conclude that, contrary to Harris, discriminating by life expectancy and quality of life is not wrong in the same way as racism/sexism, as long as it is not motivated by unwarranted judgments regarding the patients who lose out by the discrimination. However, as a caveat, I find that the discussion might have highlighted one respect in which discrimination by life expectancy can be problematic.

2.1 Harris's Argument

In a series of papers, John Harris has challenged discrimination in healthcare allocation based on age, disability, quality of life and life expectancy. For example, he says:

> Surely the principle governing any distribution of public resources must be equality: surely each is entitled to the same concern, respect, and protection as is accorded to any. People are equal and equally worth treating or saving and equality is not health status dependent. When we say all are equal we exclude discrimination on the basis of all the usual suspects: race, gender, religion, and so on. The moral principle outlawing discrimination protects (or should protect) all persons equally. People's lives and fundamental interests should be given equal weight regardless of race, creed, colour gender, age, life expectancy, or quality of life so long as that quality of life is worth having for the person whose life it is. ... The principle of equality has the advantage of very wide appeal and acceptance, and versions of it are enshrined in many national constitutions throughout the world - for example, those of the United States of America and France and in various declarations of human rights.[26]

This argument starts by citing some uncontentious principles, opposing racism and sexism, and concludes that these entail some other principles such as anti-discrimination by age or life expectancy. However, our anti-racist principles do not directly entail that discrimination by life expectancy is wrong. So how does Harris infer his more contentious principle from the more familiar ones? He does so by appealing an underlying "moral principle

---

[26] Harris (2005)

outlawing discrimination". He claims this moral principle outlaws discrimination not only by the "usual suspects" such as race and gender, but also by quality of life and life expectancy. So there are three key stages in his argument: (1) He cites some widely agreed principles regarding "the usual suspects" such as race and gender. (2) He appeals to an underlying general principle to explain them, "the moral principle outlawing discrimination". Harris read this as outlawing *all* discrimination. (3) He shows that this underlying general principle entails the further specific principles he wants to defend, debarring discrimination by life expectancy, quality of life, etc.

Harris's conclusion is that we should institute an allocation system which avoids discriminating between patients according to the cost-effectiveness of the treatments they need, such as a lottery or waiting list.

2.2 Comment On Harris

But things are not as simple as Harris thinks. Not all discrimination is wrong. To see this, bear in mind that Harris claims he's referring to universal principles (not principles that only apply in healthcare allocation). But if it was universally unacceptable to discriminate on any factor, employers would be wrong to select one candidate over another on the basis of their ability to do the job. To treat all with the kind of equal respect Harris thinks they deserve, we'd have to have a lottery between job candidates regardless of their ability.

Harris might say recruitment for jobs is a different matter, since that is a case of two parties engaging in a contract to swap benefits with each other, whereas he's referring to a case where government is distributing benefits to the needy, and in this case governments must treat all their citizens equally.

But to go back a step, Harris seems to base his argument on universal principles. He appeals to our anti-racist, anti-sexist intuitions and to universal declarations of human rights which apply to all spheres of human activity. These do not just apply to governments distributing benefits. They apply to everyone, including employers.

Nevertheless, perhaps Harris would like to restrict his claim to cases where government distributes benefits without expectation of a return. For example, he says "Surely the principle governing any distribution of public resources must be equality".[27] But even in such cases, we can find discrimination. When government-owned universities admit students, we allow them to discriminate between applicants to determine who gets places and grants and who doesn't. Candidates for undergraduate and postgraduate places are no doubt evaluated on their ability, and some are rejected.

Nevertheless, in defence of Harris, it has to be admitted that some recruitment processes fall short in terms of the universal principles that Harris invokes. Some forms of discrimination are acceptable and some aren't. I take it that Harris would argue that discrimination by prognosis in healthcare allocation is like the bad kinds of discrimination in the relevant respect, and not like the good kinds. How could he make this out? Perhaps one way is in terms of Kant's injunction that we should "So act that you treat humanity, whether in your own person or in the person of any other, always at the same time as an end, never merely as a means."[28] Harris can make use of this idea that humans have value in themselves. Our actions should express respect for the value of the people we deal with, thereby treating them as ends. Good discrimination in recruitment expresses the view that the candidate who loses out is valuable, despite not being suitable for the post (or at least, the discrimination is consistent with that view). In contrast, when a recruiter rejects a candidate for belonging to a certain ethnic group, the recruiter does not merely express the view that the candidate is not suitable for the job, but also the view that the candidate was an inferior being in virtue of their ethnicity. This does not treat the candidate as having value in themselves, so it breaches the Kantian injunction. In a somewhat similar vein, Harris might want to argue, when a patient with poor prognosis is rejected for a treatment that could help them, healthcare policy-makers express the view that the patient's life is not as valuable as the life of another patient who gets funded.

But in response, healthcare policy-makers will say that, despite discriminating by prognosis, they express no disrespect for patients as persons. The patient is not identical with their prognosis, any more than the job candidate's Kantian value as a person is identical with their potential value to the prospective employer in the employer's workplace. If discrimination by the latter can be consistent with the Kantian injunction, then so can discrimination by the former. But perhaps Harris would respond that a patient's personhood involves their life as much as it involves their ethnicity

---

[27] Harris (2005)

[28] Kant (1996) with a modified translation from Kerstein (2009)

or gender, and their life incorporates their life expectancy and quality of life. So in refusing funding for one patient in favour of another on grounds of life expectancy or quality of life, we express the view that the first patient's life is less valuable than another's, thereby expressing the view that that patient is less valuable.

Harris would here rely on an intuition that a patient's life expectancy and quality of life are aspects of them as a person. I do not share this intuition, but I admit my intuitions are not clear, and anyway it would be better to find a resolution of this debate which avoids appeals to clashing intuitions. I will aim to resolve the impasse by going back to the original question of what wrongful discrimination involves. We need to articulate a better explanation of what makes acts of racist/sexist discrimination wrong. Such an explanation will help us determine whether or not it is permissible to discriminate by the traits we focus on in healthcare allocation, such as life expectancy and quality of life.

In looking for this better explanation, my strategy is to start with the traits we may not discriminate by, and look for a common factor they share. In the UK, people with the following characteristics are legally protected from various kinds of discrimination:

age;

disability;

gender reassignment;

marriage and civil partnership;

pregnancy and maternity;

race;

religion or belief;

sex;

sexual orientation[29]

In many jurisdictions apart from the UK, these characteristics, or large subsets of them, are commonly dealt with simultaneously in legislation. Groups of these traits are also often discussed collectively in the academic literature. Philosophers such as Harris, Edmonds, Segall, Garcia, Arneson, and Alexander seek a necessary condition for racism, sexism and the like because

---

[29] Citizens Advice (2018)

these philosophers think the moral reason we have for avoiding racism is closely related to the moral reason we have for avoiding sexism.

If we can identify a characteristic which is shared by these traits, this would be a strong candidate for helping to explain what is wrong with discrimination by these traits.

2.3 Wrongful Discrimination Is Badly Motivated

I propose to explore the claim that when an act of discrimination is wrong in the same way as racism/sexism, what makes it wrong in that way is that it is motivated badly. I will start with some cases that constitute prima facie grounds for that view. I will then try to characterise the problematic attitudes more precisely, considering accounts from Arneson and Alexander.

As an initial stab, consider an employer who discriminates against people with a certain accent for no reason other than having a negative stereotype about people with that accent. This seems morally wrong in the same way as racism/sexism. Compare this with the case of a retail employer who has found that counter assistants with the accent sell 5% less (for example, suppose this is the Manchester United FC shop and the accent in question is from Liverpool, home of Manchester United's greatest rivals; or suppose this is a Welsh rugby kit shop, and the accent in question is from the USA, not known for its rugby prowess). In these circumstances, if the employer has to choose between two candidates who are otherwise equal, but one of whom has the accent, it might not be racist for the employer to favour the one without the accent, if this is done for purely business reasons and not because of any negative attitudes towards the accent. Thus the moral status of an employer's discriminating by accent can depend on the nature of the attitudes motivating their discrimination.

We can also imagine a case within healthcare allocation where the moral status of a decision is determined by the attitudes of the decision-makers. Suppose there is a gene predominant amongst the English which makes the treatment for disease X less cost-effective than for those without the gene. Suppose the Scottish healthcare allocation committee, the SMC, refuse funding for people with that gene. Contrast two cases. In one, there is no evidence of anti-English bias on the part of committee members. In the other, there is such evidence. For example, suppose the Scottish board is populated with anti-

English extremists who've been found guilty of burning the English flag etc. Also, suppose it is known that Thai people also have a gene which reduces the cost-effectiveness of the treatment to the same extent, but the cost-effectiveness threshold is relaxed in their case on the grounds that to discriminate against Thais would be racist. The English are held not to be a race and so no such provision is made for them. An Englishman living in Scotland would have reasonable grounds for complaint in the second case but maybe not in the first. What makes the difference is not the decision but the motivation.

In fact, it will often be impossible to judge the moral status of acts of indirect discrimination without consideration of the attitudes that motivate them. Janet Radcliffe Richards quotes a real life case which provides materials for this point, the case of London Transport bus garages in 1969, when male drivers protested about women being allowed to drive the buses. Their grounds were that women were not strong enough. Radcliffe Richards says "they eventually conceded that the women could drive the little, single decker buses, but not the big ones, which were obviously men's work".[30] This clearly seems like a case of direct discrimination. Now let's change it to a case of indirect discrimination. Suppose these drivers, in an effort to find firmer grounds for their proposed restrictions, concede that gender was irrelevant. Instead they insist that the relevant requirement was strength, and switch to insisting that strength tests be included in the recruitment process. How we would make out an accusation that this proposal was sexist? The difficulty is that the men might be correct that strength is important for safety; stronger drivers might be better able to control the buses in emergencies. In that case, it won't be immediately obvious that the insistence on strength is sexist. On the other hand, strength will not be a decisive consideration, because it will only be one requirement amongst many. Just in terms of safety, the requirements might include strength; reaction time; situational awareness; night vision; ability to cope with distractions. And safety will not be the only consideration. Other factors might include passenger comfort (smooth stop/starting); punctuality/reliability; and people skills. For each criterion, each candidate will fall on a spectrum between strong and weak, strong on some criteria and weak on others. So how do we apply such criteria? Presumably candidates will first be assessed against the essential criteria, such as a current driving licence. Then the best candidate will be selected from those who remain. But

---

30 Radcliffe Richards (2000)

how do we say that one candidate is best, if no candidate is strongest on every criterion? The answer must involve at least two steps: assessing each candidate on each criterion, and weighting the criteria. Of course neither assessment nor weighting will necessarily be simple.[31] But whatever the details, weightings will be required. And my point is that the weightings question will be open, in that different weightings will be defensible. One source of openness will be epistemic limitations: it will not be clear whether strength or reaction times contribute more to safety. Another source of openness will be the debatability of priorities within safety. For example, should we prioritise passengers or pedestrians? Should we accept many injuries for the sake of avoiding one death? Another source of openness will be the debatability of priority for safety vs priority for other concerns, given that a driver who is very cautious about safety might incur costs in terms of service punctuality and passenger comfort.

Given this openness, it will be arguable that the strength criterion should not be given more weight than currently, but it will also be arguable that it *should* be given more weight than currently. So the mere fact that drivers are arguing that the strength criterion should be given more weight is not sufficient to show that the drivers' proposal is sexist. How can it be sufficient, when they are arguing for a position which is in fact defensible? This puts us in a quandary. If, despite the defensibility of the proposal, we want to say that the driver's proposal is sexist, how will we make out the accusation? An analysis in terms of the drivers' attitudes provides a way of diagnosing the flaw in their position: the problem is that it is motivated by prejudice against women. Perhaps it will be helpful to put this in terms of an agent's reasonableness. There is a difference between a position's being defensible and its proponent's being reasonable. When a child refuses to let a sibling play with her toy, her position might be defensible (it's her toy) but it might be that she is not being reasonable. When Southern whites instituted literacy tests for prospective voters, in order to make it difficult for blacks to register, the tests might have been defensible in some minimal way,[32] but proponents of the tests were not

---

[31] For example, candidates with reaction times below a certain threshold might be rejected whatever their performance on other criteria, or we might demand higher scores from them on strength and situational awareness.

[32] For example, see Klarman (2006). The tests discriminated against black people, since blacks tended to be less well-educated (also the rules about who had to take the tests were discriminatory, e.g. in terms of "grandfather rights"). Nevertheless, certain views about the value of democracy can be enlisted in support of literacy tests. For example, some people think the primary value of democracy is that it is "more reliable in helping participants discover the

being reasonable. Reasonableness involves taking a balanced view of relevant considerations and ignoring irrelevant ones. As I will argue in a moment, this involves not having false or unjustifiable negative judgments about people. So if London bus drivers demanded that a lot of weight be placed on the strength criterion, it might be that their demand was defensible but they were not being reasonable; they were not assessing the job requirements in a balanced way. Specifically, if the main reason they wanted so much weight put on strength tests was a desire to keep men and women in their traditional gender-based roles, then we could say they were being unreasonable about the safety issue.[33] In light of evidence that the male drivers were driven in this way by sexist stereotypes, women who lost out could argue that the drivers were only arguing for more weight to be given to the strength criterion because they wanted to keep women out, and not because they genuinely thought safety was as important as they were saying. If the men were then somehow forced to admit that this was why they put the proposal, and also made to see that their sexist views were indefensible, they would then be obliged to admit the proposal was not justified, since on their view of the matter, safety considerations alone were not enough to justify it (one *could* reasonably argue that safety concerns were that important, but the drivers don't think so).

To generalise, there are many cases like this where we must consider people's attitudes before we can tell whether they are being racist or sexist, because people can often rationalise their sexist or racist discrimination in terms of legitimate considerations. In fact, more often than not, those who discriminate racistly or sexistly will themselves be convinced that they are motivated by legitimate considerations, as the evidence of widespread implicit

---

right decisions." (Christiano, 2015). If this view is right, then perhaps literacy tests can be defended on the grounds that you can't understand what you're voting for unless you can read such items as newspapers and party manifestos. Along somewhat similar lines, J.S. Mill advocated giving extra votes to "all graduates of universities, all persons who have passed creditably through the higher schools, all members of the liberal professions" (Mill, 1861). I don't share this view; I merely aim to show that it is defensible in the minimal sense that an argument can be marshalled in support of it.

[33] I take it that such an account would be applicable if various counterfactuals were true. For example, in the nearest possible world in which the male drivers do not desire to keep women out, do they still want the strength criterion to be given so much weight? Unfortunately, in the real world, it's difficult to tell whether such a counterfactuals are true, which might prompt the question: What real world evidence could there be that such desires were operative? One possible response is that this is not a relevant question; I'm investigating what wrongful discrimination consists of, not how we know it has occurred. But my slightly more constructive response would be that we do ordinarily attribute desires and attitudes to each other in order to evaluate each other's actions, and I see no greater difficulty here than elsewhere.

bias suggests. For example, consider a study in which 238 male and female academic psychologists were asked to evaluate job applications. The applications were identical except for the gender of the applicant:

> All the curriculae vitae actually came from a real life scientist at two different stages in her career, but the names were changed to traditional male and female names. Both men and women were more likely to vote to hire a male job applicant than a female job applicant with an identical record.[34]

Thus female academic psychologists exhibited implicit bias against female job applicants. If asked their reasons for rejecting a female job applicant, I take it that these female academics would not have pointed to the gender of the applicant as one of their reasons. They would have been able to construct a defence of their rejection based on items on the CV.[35]

This provides another illustration of the point that in many cases it will be useless to consider the observable circumstances in which discrimination occurs and determine whether legitimate reasons for the discrimination can be found there. We must understand the attitudes of those who discriminated.

So, in view of this evidence that attitudes can be relevant, I will now consider accounts of racism/sexism in terms of the agent's propositional attitudes.

2.4 Alexander's Solution: False Or Unjustified Negative Judgments

To outline what follows, I first introduce the account I wish to defend, an account from Larry Alexander according to which the reason wrongful discrimination is wrong is that it is based on unwarranted judgments. However, I then introduce cases of wrongful discrimination which Alexander's account seemingly fails to capture, cases which involve negative attitudes which are not judgments. In response, I consider an alternative "attitudinal" account offered by Richard Arneson. This account is more

---

34 Steinpreis (1999)

[35] Of course they couldn't have justified the discrimination versus the otherwise identical male CV if challenged on it directly. But if they'd been asked to consider the female CV in isolation, I suggest they could have picked out weaknesses in it which could have been taken as reasons not to hire the candidate.

encompassing than Alexander's; as well as ordinary judgments, Arneson's account encompasses a certain category of attitudes called deep-seated aversions, and as a result Arneson's account has intuitively satisfying implications for the cases that were problematic for Alexander. However, the difficulty for Arneson is that his account looks ad hoc, in that it is difficult to find independent reasons for deprecating unwarranted deep-seated aversions. Also Arneson's account has counterintuitive implications of its own. I then return to the problematic cases and find that although Alexander's account cannot explain our intuitions, Alexander can appeal to other considerations to explain our intuitions. I therefore conclude that Alexander's account is the right account.

So, to start on Alexander, on his account there are two ways that racism and sexism can be wrong:

1. Some acts of racist and sexist discrimination are "intrinsically wrong" in virtue of the unjustified judgments which motivate them. Such acts are wrong regardless of the context.[36]

2. Other acts of racist and sexist discrimination are only "contingently wrong" according to culture and context, e.g. according to their consequences.

For his account of intrinsic wrong, Alexander says:

> Discrimination may be intrinsically wrong because it is based upon biases, the incorrect judgments of lesser moral worth, or upon the shallow aversions or inaccurate negative stereotypes that are produced by such judgments. Discrimination may be intrinsically wrong because it is based on an unjustifiable ideology of moral role.[37]

Now, before I go any further, I need to clarify an aspect of Alexander's account. In the above passage, Alexander disparages "incorrect judgments of lesser moral worth" but he also disparages "an unjustifiable ideology of moral role." I take it that with both these descriptions, Alexander intends to capture a discriminator's judgment that the person discriminated against is inferior in some way, such that the person does not have the same rights as the discriminator or others that the discriminator regards as peers. The difficulty

---

[36] "Intrinsically wrong" is the phrase used by both Arneson and Alexander to characterise the acts they focus on, and I will use the same terminology.

[37] Alexander (1992), p. 218. When Alexander says something is intrinsically wrong, I interpret him as meaning that there is what Nozick would call a "side-constraint" prohibiting it. The prohibition is not contingent on consequences, for example.

is that different judgments will get disparaged according to which of Alexander's two descriptions we go by. There are various subtleties,[38] but the most important difference between the two descriptions is that false but justified judgments would be disparaged by an account based on the first description but not an account based on the second. The accounts would also reach different conclusions regarding true but unjustified judgments. Thus to avoid ambiguity, we need to be clear whether we are considering an account based on one or other description, or some disjunction or conjunction of them. I will address the rest of this chapter to an account which disparages unjustified judgments as well as false judgments (whether justified or unjustified). If there are judgments which are not truth-apt but which are apt for justification,[39] then the account I address disparages unjustified judgments of that kind (but not justified judgments of that kind, since they are neither unjustified nor false). This position seems to most closely reflect the intention of the passage above and is also, I contend, defensible.

So, according to Alexander's "judgmental" account, intrinsically wrong discrimination is always to be explained by a false or unjustifiable judgment of lesser moral worth. For example, Alexander says of the biases of the Nazis against the Jews that they were

> intrinsically morally wrong because Jews are clearly not of lesser moral worth than Aryans. When a person is judged incorrectly to be of lesser moral worth and is treated accordingly, that treatment is morally wrong regardless of the gravity of its effects. It represents a failure to show the moral respect due the recipient, a failure which is by itself sufficient to be judged immoral.[40]

However, Alexander takes a somewhat different line regarding what he calls "deep-seated aversions". Deep-seated aversions are dispositions to avoid people with a certain trait or to feel uncomfortable dealing with them, e.g. out of disgust or irritation. These aversions and attractions are not biases because they are not based on judgments of differential moral worth, though these aversions and attractions may be rooted in ancient biases that now operate

---

[38] For example, I set aside any special connotations of the word "ideology".

[39] For example, I suggest it is not inconsistent to hold that moral judgments are non-cognitive attitudes, and therefore not truth apt, but yet also to hold that there are ways of justifying them, such as the ways seen in ordinary moral discourse.

[40] Alexander (1992), p. 158-9. In line with my selected interpretation of Alexander, when I says "When a person is judged incorrectly to be of lesser moral worth...", I read him as saying "When a person is judged unjustifiably to be of lesser moral worth ..."

subconsciously. Nor are these aversions and attractions based on moral ideals or on stereotypes - those who have them do not believe that the unwanted (or wanted) associations are morally forbidden (or required), or that the target group likely has some other trait that is straightforwardly relevant to the association.[41]

Alexander gives the following examples:

> Aversions and attractions based on race, ethnicity, or gender are only part of this category of preferences for and against certain types of people. More common perhaps are aversions and attractions based on aesthetics - how others look, sound, or even smell. Physical appearance is a frequent basis for preferences, either categorically ("I want to be around handsome people in all contexts") or in particular contexts ("I hire only good looking dancers for my troupe").[42]

Alexander does not deem acts motivated by deep-seated aversions to be intrinsically wrong:

> Discrimination based on deep-seated aversions ... may be wrong, but it is not intrinsically so. Rather, particular types of such discrimination will be wrong in particular cultures, historical eras, and contexts, and not wrong in others.[43]

I read him as being led to this position because deep-seated aversions are not judgments, therefore a fortiori they are not judgments of lesser moral worth, and therefore do not get captured by Alexander's account of intrinsic wrong. Thus they fall into the second category of contingently wrong acts of discrimination; they will only be wrong for contingent reasons associated with particular times and places. In certain circumstances, it can be permissible. Alexander does not characterise these circumstances in detail, but he offers the following as an example of the contingencies which point in one direction or the other:

> [some deep-seated aversions] are most likely the products of biases, ideals, and stereotypes that have become buried in the subconscious: they will tend, if widespread, to reinforce conscious biases, ideals, and stereotypes, and to be experienced by their victims in the same way the victims experience conscious biases, ideals, and stereotypes. If they are

---

[41] Alexander (1992), p. 165

[42] Ibid.

[43] Alexander (1992), p. 218

> uniform, so that the same groups tend to be preferred and dis-preferred, and widespread, and if they relegate the dis-preferred to less desirable positions in the socioeconomic hierarchy, their social effects may be quite devastating. If they are idiosyncratic and variable, uncommon, or context-specific - "I'm uncomfortable around Italians in my private club but not at work" - rather than categorical - "I prefer to avoid Jews in all contexts" - and do not dis-prefer the already relatively disadvantaged, their adverse social effects may be relatively minimal.[44]

Alexander does not explicitly say whether such actions with minimal adverse social effects are wrong or not. However, Alexander does say that deep-seated aversions are only contingently wrong and that if they are "idiosyncratic and variable, uncommon, or context-specific", and they do not have wider "social effects", and they do not do not dis-prefer the already relatively disadvantaged, then discrimination motivated by those deep-seated aversions is permissible.

Why should we accept Alexander's distinction between the attitudes which may be disparaged as grounds for discrimination (judgments) and the attitudes which may not be so disparaged (deep-seated aversions). Why should we deem all discriminatory acts based on false or unjustified judgments of lesser moral worth to be intrinsically wrong in virtue of this defect in their motivation, but not deem any acts motivated by deep-seated aversions to be intrinsically wrong, despite the latter's apparent equal lack of epistemic credentials? To answer this, it will help to start by examining judgments, and see what's good about them when they are good and what's bad about them when they are bad. Then we can see why deep-seated aversions can't be bad in the same way. So, we have seen that Alexander evaluates judgments in terms of their semantic status, and also in terms of their epistemic status: judgments can be bad in virtue of being false or bad in virtue of being unjustifiable. Starting with their semantic status, it is uncontroversial that true judgments are better than false ones, at least in respect of their semantic status. Because of this, if I discriminate against someone on the basis of a falsehood, we can say that my practical reasoning has gone wrong in a certain respect. If the disfavoured party loses out, that flaw in the decision-making procedure provides them with grounds for a reasonable complaint, in proportion to the scale of the harm, and perhaps in proportion to the scale of any epistemic carelessness on my part (such as lazily

---

[44] Alexander (1992), p. 166

ignoring relevant considerations). Something similar can be said about the epistemic status of my judgments. If I discriminate against someone on the basis of an unjustified judgment, my practical reasoning has gone wrong in that respect, and again, the victim of my discrimination has a reasonable complaint.

So, why is lack of justification a flaw in the case of judgments and not in the case of deep-seated aversions? I believe this can be explained with a principle along the following lines:

> An attitude's lack of justification only counts against the attitude if there are other possible attitudes which compete with it and which would be better justified (where one attitude competes with another if a rational agent could not hold both simultaneously).

My argument for this principle would be that if there is no better attitude I could have that is incompatible with my current attitude, what reason do I have to drop my attitude in favour of an alternative?

This principle can be also defended in terms of its implications for specific cases. If the Nazis could offer no good reasons for their contempt, and there were good reasons to be uncontemptuous, then being contemptuous seems clearly inferior to being uncontemptuous in respect of justifiability. Suppose, on the other hand, I love red cars and hate blue cars. Suppose also that although there is no good reason for me to love red cars, nor are there any good reasons for me to love blue cars. In that case, my love of red cars is not rationally flawed.[45] Neither colour-loving attitude is inferior to the other in terms of justifiability. In view of such cases, I conclude that an attitude's lack

---

[45] Perhaps because my attitudes to car colours are not judgment-sensitive attitudes in Scanlon's sense, whereas the Nazis' contempt is a judgment-sensitive attitude. Scanlon defines judgment-sensitive attitudes as attitudes that an ideally rational person would come to have whenever that person judged there to be sufficient reason for them and that would, in an ideally rational person, "extinguish" when that person judged them not to be supported by reasons of the appropriate kind. Hunger is obviously not a judgment-sensitive attitude; but belief is, and so are fear, anger, admiration, respect, and other evaluative attitudes such as the view that fame is worth seeking. ... judgment-sensitive attitudes constitute the class of thing for which reasons in the standard normative sense can sensibly be asked for or offered. ..... Because of this dependence on judgment, these are things we can properly be "held responsible" for in several central sense of that phrase: they can properly be attributed to us, and we can properly be asked to defend them". (Scanlon, 1998, pp. 20, 21 & 22). It is worth noting that on this account, judgment-sensitive attitudes are judgment-sensitive in the ideally rational person and not necessarily in actuality. So, even if my belief that I will win the lottery is resistant to all counter-evidence, it may still be a "judgment-sensitive attitude" if the ideally rational person would abandon it in my circumstances. A judgment-sensitive attitude counts as judgment-sensitive whoever holds it.

of justifiability only constitutes a flaw in cases where there are reasons telling in favour of some competing attitude. I therefore take Alexander to be suggesting that discrimination is only wrong in the same way as racism/sexism if it is motivated by a false or unjustifiable judgment, where there is some competing judgment which could be justified.

Finally, it should be noted that Alexander only impugns false or unjustifiable negative judgments regarding people. This helps avoid a counterexample he would face otherwise. Consider a case in which I buy a Skoda from Jack, rather than a Porsche from John, because I believe (wrongly) that the Skoda is a better/faster car. Then I would have an false negative judgment, which leads me to discriminate (in some sense) between Jack and John.[46] Intuitively, in such a case, I do not seem to have discriminated wrongfully. But if Alexander's account was not restricted to judgments regarding people, it would imply that I have. Alexander's account avoids this problem. Nevertheless that leaves another problem: Is this restriction not ad hoc? Why are false or unjustified judgments regarding people relevantly flawed, but false or unjustified judgments regarding objects are not equally flawed? One response is that the attitudes explaining my judgment are flawed in virtue of my failure to recognise the moral status of the person I discriminate against, my failure to recognise their possession of the same human rights as everyone else. This kind of mistake is not possible regarding cars, since they do not possess rights. In addition, to express a negative judgment about someone by discriminating against them is to express a kind of insult about them. As a general rule, it is wrong to express a false or unjustified insult about someone: it hurts them and affects their status in society. (And of course, when the insult is expressed by means of an act of discrimination, some kind of injury is added to the insult, in the form of a harm or loss of benefit). So I am only entitled to insult someone if my insult is true and justified (and possibly not even then).

2.5 Objection: Counterintuitive implication

However, it looks like Alexander's account fails to capture some paradigm racist/sexist acts which, intuitively, would be wrong regardless of the wider

---

[46] Thanks to Rob Lawlor for this case.

context in which they are performed. This is because the account does not count acts motivated by deep-seated aversions as intrinsically wrong discrimination. To reprise, Alexander distinguishes two categories of aversions, shallow and deep-seated. Shallow aversions are produced by incorrect judgments of lesser moral worth. As per the quotation above, shallow aversions are "intrinsically wrong". However, for Alexander, discrimination based on deep-seated aversions is not intrinsically wrong. And in "particular cultures, historical eras, and contexts" it will be not wrong in any way at all. Let's suppose for the sake of argument that discrimination based on deep-seated aversions would not be wrong on this account in circumstances where people with the discriminated trait have not been disadvantaged by discrimination historically and where there will be no negative consequences for others with the trait (such as, in a modern western community, discrimination based on an idiosyncratic aversion to people with blue eyes). But in the same circumstances, any discrimination based on unjustifiable judgments of lesser moral worth would still be wrong.

The problem for Alexander is that this looks like a counterintuitive result. Suppose I have a deep-seated aversion to a certain accent. I find it grating. People with this accent have not suffered discrimination in the past. I interview a candidate for a job and she is by far the best candidate, but I reject her in favour of someone else because she has this accent. I do not harbour any judgments of lesser moral worth about her. I hide my reasons for rejecting the candidate so there are no further negative consequences for her or for others with the same accent. Such discrimination may seem intrinsically wrong. But Alexander looks forced to say it is not, because my aversion to the accent does not look epistemically flawed in the way Alexander needs. It is a mere deep-seated aversion, not apt for the kind of epistemic credentials which Alexander could use to disparage it. So Alexander must say that it is not intrinsically wrong. In fact, in certain circumstances, it will not be wrong at all.

One defence for Alexander is to argue that generally, a boss can legitimately demand a candidate they will be comfortable working with. If the boss has a deep-seated aversion against the candidate and this will be a barrier to smooth and happy working relations, that could be a legitimate reason to reject the candidate. But then an objector can ask Alexander to consider the case where the recruiter will not be working alongside the candidate. For example, the candidate will be working at the bottom of the ocean, so live communication will be impossible and all contact will be written. But then Alexander can ask why a deep-seated aversion would bring a recruiter to discriminate against such a candidate, when there will be no contact between them? The best explanation of discrimination in such a case would be that the recruiter incorrectly judges the candidate to be of lesser moral worth, in which case Alexander can account for the discrimination as wrong in the same way as racism/sexism. But in response, perhaps the objector can suggest that Alexander is making convenient assumptions about human nature here. So let's consider for the sake of argument a situation in which all communication between the recruiter and the job-holder will be written, and the recruiter has an extreme, deep-seated aversion against the candidate's accent, and the recruiter post-rationalises her aversion in terms of some minor flaw in the CV and therefore rejects the candidate. This is not an easy case to imagine, but Alexander must deal with it to make his case. Is this not a case of discrimination which is wrong in the same way as racism/sexism?

2.6 Arneson's Solution: Unwarranted Negative Attitudes

It might seem that the solution to this difficulty is to adopt a more inclusive account which captures the problematic cases. Richard Arneson offers just such an account. However, I will show that there are reasons to reject it. Arneson's account is like Alexander's in that it says the moral status of an act is partly determined by the epistemic status of the attitudes which motivate it.[47] However Arneson not only disparages discrimination motivated by unwarranted judgments as intrinsically wrong, but also discrimination motivated by deep-seated aversions. Arneson would class my aversion to the accent as an "unwarranted animus" and therefore count my discrimination as intrinsically wrong.

In detail, Arneson's all-encompassing "attitudinal" account says:

> Discrimination that is intrinsically morally wrong occurs when an agent treats a person identified as being of a certain type differently than she otherwise would have done because of unwarranted animus or prejudice

---

[47] To be clear, Alexander spends most of his paper focusing on preferences: specifically on the question of when discriminatory preference are intrinsically or extrinsically wrong. However, this is in order to explain the wrongfulness of discrimination. For example, he says: "I am going to approach the question of what makes discrimination wrongful by examining discrimination as an expression of various types of preferences" and concludes: "What makes wrongful discrimination wrong? ... Discrimination may be intrinsically wrong because it is based upon biases, the incorrect judgments of lesser moral worth, or upon the shallow aversions or inaccurate negative stereotypes that are produced by such judgments."

> against persons of that type. ... one is led to defective conduct toward the other by unjustified hostile attitudes toward people perceived to be of a certain kind or faulty beliefs about the characteristics of people of that type.[48]

Clarifying one of his terms, Arneson says "Animus is hostility or, more broadly, a negative attitude, an aversion."[49] Since Arneson's account encompasses acts motivated by such attitudes, it is more inclusive than Alexander's. Arneson and Alexander diverge specifically with respect to acts motivated by deep-seated aversions. I therefore now consider Arneson's position on deep-seated aversions. Arneson would count a deep-seated aversion as an unwarranted animus, and would therefore deem discrimination driven by such an aversion to be "intrinsically morally wrong". For example, Arneson says:

> the treatment of ugly persons may often qualify as wrongful discrimination because the ugly person is being treated worse than he would be if he were not a member of a group that excites revulsion, and the revulsion may qualify as an unwarranted hostile attitude.[50]

For this reason, more acts satisfy Arneson's account than Alexander's. Specifically, Arneson's all-encompassing "attitudinal" account would entail that an act motivated by a deep-seated aversion is intrinsically morally wrong, while Alexander's more restricted "judgmental" account would not. As per the above quote, on Arneson's account an act of discrimination against someone ugly which is motivated by revulsion would be intrinsically wrong. In contrast, as we saw above, Alexander says that discrimination based on "aversions and attractions based on aesthetics - how others look", motivated by a deep-seated aversion, is not intrinsically wrong.

---

[48] Arneson (2006). Arneson offers another way that discrimination can be wrongful based on beliefs formed in a epistemically faulty way. For example: "I simply am lazy in forming beliefs. I harbor no animus against Korean-Americans, but I discriminate against them on the basis of negative beliefs about the characteristics of Korean-Americans that I lazily absorb from the prevailing culture." I do not address this condition in Arneson's account in detail, but none of my arguments hang on that omission.

[49] Arneson (2006). Arneson offers another way that discrimination can be wrongful based on beliefs formed in a epistemically faulty way. For example: "I simply am lazy in forming beliefs. I harbor no animus against Korean-Americans, but I discriminate against them on the basis of negative beliefs about the characteristics of Korean-Americans that I lazily absorb from the prevailing culture." I do not address this condition in Arneson's account in detail, but none of my arguments hang on that omission.

[50] Arneson (2006), p. 803

So Arneson can tally with intuitions in the case of the accented candidate that posed the problem for Alexander. However, Arneson faces two problems. First, he faces counterexamples. My unwarranted attitudes can result in a variety of discriminatory but permissible acts. For example, I get my hair cut at barber A not barber B because I am averse to barber A's voice. I get my kebabs at A's takeaway not B's takeaway because I find chef A's freckles disturbing. In each case these acts of discrimination will leave the losers less well off than they would have been; perhaps it will put them out of business (for example, if I am a frequent user buying on behalf of my large family, and their business is only marginally profitable). Yet my lack of warrant for my aversions provides no reason for me to re-consider such habits. My acts are not irrational or unjustified or impermissible just because the attitudes motivating them are unwarranted.

A similar difficulty will beset any attempt to explain what is bad about an unwarranted negative attitude in terms of its consequences, or in terms of any other feature of an act other than the lack of warrant itself. The other considerations with which such accounts are bolstered will only be associated with the lack of warrant contingently, so such accounts will not avoid the possibility that an act motivated by an unwarranted attitude will not have this other feature, or the possibility that an act *not* motivated by an unwarranted attitude will nevertheless *have* this other feature. So to disparage unwarranted attitudes, we must criticise the attitudes themselves for their lack of warrant.

The second problem for Arneson is not that his account fails as a set of necessary and sufficient conditions, but that it lacks explanatory power. Specifically, even in cases of wrongful discrimination motivated by unwarranted aversions, the unwarrantedness of the aversions does not help explain the wrongfulness of the discrimination, because we have been given no independent reason to think that lack of warrant is a flaw in an aversion. For example, suppose I have a deep-seated aversion to an accent and discriminate against someone on that basis. The discrimination is motivated by "unwarranted animus or prejudice", so Arneson will say it is intrinsically wrong. But as argued above, I contend that the acceptability of a deep-seated aversion is not compromised by its lack of warrant, assuming that there are no competing attitudes that could be justified. My being indifferent to the accent or loving the accent would be no better warranted than my being averse to it.[51] All possible attitudes towards the accent are in the same boat with respect to warrant. So why is the acceptability of my aversion in any way compromised by its lack of warrant?

To give an analogy, my love of chips lacks warrant.[52] But intuitively, that is not a problem. My attitude would be no better warranted if I had an attitude of hatred for chips or of indifference to chips. The fact that none of these attitudes to chips would be warranted is not a defect in them, since they are not at any disadvantage to each other as a result of their lack of warrant. They are not apt for warrant. Similar points could be made regarding myriad other preferences. For example, my not liking the flavour of beetroot; my liking one individual's sense of humour and not another's; my preferring bright red cars to dark blue cars; my liking a certain regional accent and not another. It might be impossible to justify these attitudes in the sense of giving reasons for them, and impossible to find any other kind of warrant for them. Yet this does not seem to be a defect in them. My lack of warrant for my love of red cars and aversion to blue cars is no reason to abandon those attitudes.

(I should clarify that Arneson's account does not imply that an unwarranted aversion to chips would be flawed in virtue of its lack of warrant; his account is focused specifically on aversions to people, not to such things as chips or blue cars. My point is rather that if Arneson was right that a lack of warrant for an aversion was a rational flaw in the aversion, a similar reproach would be applicable in other domains in which aversions are seen. If no such reproach seems applicable in other domains, the suspicion has to be that what seems like a rational flaw in aversions to people is actually not a flaw, but the result of a kind of halo effect from the intuitions we originally set out to explain. The attitudes *seem* flawed to us because the discrimination they cause seems wrong.)

So, it seems that Arneson is trying to apply epistemic norms where they are not relevant. Arneson would disparage my attitude towards the accent for its lack of warrant, but he has given no more grounds for this than he has for disparaging my dislike of red cars for *its* lack of warrant.

Having said that, Arneson's account gains some plausibility because it has the right implications for cases such as my recruitment case given above. But

---

[51] Putting pragmatic reasons aside; presumably strong feelings either way would carry certain pragmatic costs.

[52] Thick-cut chips from fresh potatoes. Not French fries.

defending a theory in terms of the intuitions it purports to explain gets things the wrong way round. It makes the account ad hoc. Arneson has identified a feature which happens to be associated with some intuitively wrongful acts, but to explain those intuitions (and the wrongfulness of the acts), he needs to find *independent* reasons to think that his feature is morally relevant. This could be by showing that a principle involving the feature is widely-accepted, or intuitively compelling without argument, or otherwise defensible without "ad hoc" appeal to the intuitive wrongness of the cases picked out by the necessary and sufficient conditions. If, in contrast, our theoretical adequacy conditions merely demand that a theory offer necessary and sufficient conditions, without any demand for such independent explanatory power, we could end up with a theory that states necessary and sufficient conditions for wrongfully racist/sexist discrimination in terms of some trivial feature which is morally irrelevant. As an extreme example, suppose we aim to account for wrongfully racist discrimination, and it happens to be the case that all the individuals who are disadvantaged by such discrimination happen to share a certain physical characteristic, such as the ratio of their ear length to nose length falling within a certain range. This characteristic is not shared by anyone else. In that case, we could state necessary and sufficient conditions for racist discrimination in terms of this characteristic. But such an explanation would not pick out the features of wrongfully racist discrimination which matter, morally. Whatever is wrong with wrongfully racist discrimination, it has nothing to do with this characteristic.

The theoretical adequacy condition I am suggesting here is widely recognised. Most theorists of discrimination have tried to respect it, at least implicitly. They characterise racism and sexism in terms which are independently compelling. In other words, the features mentioned in each account are plausibly reason-giving features in their own right, independently of the anti-discrimination principles to be explained. The norms to which theorists have appealed, implicitly or explicitly, have included moral norms, as in Segall's account (a prima facie plausible moral requirement not to undermine equality),[53] and prudential norms (for example, accounts which

---

[53] Segall (2012)

disparage irrelevant/arbitrary criteria, such as that from Conaghan),[54] and epistemic norms (as Alexander proposes in the account I have considered).[55]

So the problem is that Arneson cannot say what is wrong with all the attitudes he disparages (whereas Alexander can). What Arneson missed is this. The fact that, in certain contexts, the presence of a feature such as warrant is good-making and its absence is bad-making, does not entail that the absence of this feature is *always* bad-making. The question of whether the absence of warrant is bad-making depends on whether the attitude which lacks it needs it. Analogously, we can disparage a knife for being blunt, because it needs to be sharp, but we can't disparage a hammer for being blunt, because it doesn't need to be sharp. Similarly, we can disparage a factual belief for not being warranted, but we can't disparage my disliking red cars for not being warranted.

So Arneson's account has two flaws. First, some discriminatory acts motivated by deep-seated aversions are not wrong in the same way as racism/sexism. Second, lack of warrant is not a bad feature of deep-seated aversions, so pointing out that an act of discrimination is motivated by an unwarranted aversion does not help explain the wrongfulness of the discrimination.

2.7 A Defence for Alexander: Other Considerations Explain the Problematic Cases

I will now set aside the option of exploring alternatives to Alexander, since I believe another strategy is viable. Alexander can fend off my problematic case by pointing out that my discrimination is likely to have a variety of wrong-making features which he can recognise consistently with his account. These features do not make my discrimination wrong in the same way as racism/sexism, but they make it wrong in other ways. Alexander can therefore vindicate our intuitions regarding many variants of the problematic case. Admittedly, that may still leave a small range of cases in which Alexander must say the discrimination is permissible. But in those cases, Alexander's

---

[54] Conaghan (1999) quoted in Halldenius (2005): "Discrimination is narrowly conceived as 'irrational' decision making, unrelated to differences in the productivity enhancing characteristics of particular workers."

[55] Alexander (1992)

result is much less counterintuitive. I will now run through four such extraneous incriminating factors.

The first extraneous incriminating factor involves my role-related responsibilities to others. For example, I will have responsibilities to my boss and other stakeholders to whom I owe contractual duties. So if I have a boss and/or colleagues, they can complain that I have failed in my obligation to do my best to ensure the health of the organisation I work for by employing the best candidate for the job. Alexander can criticise my discrimination on the grounds that I have let down people to whom I owed a duty of care.

The second extraneous incriminating factor relates to candidates who have the hated accent. They have incurred costs on a reasonable expectation of getting a chance based on their job-related abilities. It turns out those reasonable expectations have not been met. My aversion means that a candidates with the hated accent have wasted their time and money submitting applications and attending interviews without realising I had this aversion. The aversion is unusual so the candidate could not have reasonably expected it. Again, these would be grounds for a complaint from the candidate and criticism from Alexander of effectively misleading candidates into wasting their time and money.

The third incriminating factor to be eliminated is not strictly extraneous. It is a case of discrimination which is intrinsically wrong and which Alexander can account for as such. Suppose I rejected this candidate because of their accent, but my aversion takes the form of contempt. Specifically, as soon as I hear this accent, I jump to the conclusion that the person speaking is stupid. This case is home territory for Alexander. My attitude towards this candidate consists of a thick evaluative judgment - a factual judgment that he lacks intelligence, combined with an evaluative judgment that this lack of intelligence makes him contemptible.[56] Both facets of this judgment can be supported or defeated by reasons. The factual judgment should be sensitive to evidence as to whether the candidate really lacks intelligence, and the evaluative judgment should be sensitive to reasons why people lacking intelligence should not be judged as being of lower moral status than everyone else. In addition, in an epistemically ideal world it could probably be shown that my judgment is driven by my underlying dislike of the accent, or by various insecurities, and

---

[56] There are other accounts of the nature of thick judgments, such as Vayrynen (2013), but I believe Alexander could account for my reprehensible attitudes consistently with those alternative accounts.

perhaps on learning this I would see for myself that my judgment had poor foundations. Any time an agent rationalises disgust or irritation or hatred in terms of some belief, such as a belief that a weak handshake indicates unreliability or a particular accent indicates stupidity, that is epistemically reprehensible on Alexander's account, and Alexander can comfortably tally with our intuitions.

The fourth set of extraneous incriminating factors involves traits associated with historical social disadvantage, or "HSD traits,"[57] which have been discriminated against in the past, such that people with them are disadvantaged now. Suppose that people with the accent in question come from a disadvantaged background. Then Alexander can say that my discrimination against a candidate with that accent is wrong. Alexander specifically mentions the case where deep-seated aversions lead to discrimination which dis-prefers the already relatively disadvantaged. Alexander would say that, although discrimination motivated by such a deep-seated aversion is not intrinsically wrong, nevertheless it might be contingently wrong in the circumstances (perhaps based on an egalitarian or prioritarian principle implying that the worst-off ought to be protected from further disadvantage). Although I would argue there are limits on individuals' responsibility for correcting social ills, nevertheless there might at least be supererogatory reasons to try and control my aversion and contribute to improving this individual's prospects.[58] My aversion to the accent might seem trivial in comparison to this opportunity. Surely I can get used to the accent, or find working methods to avoid hearing it?

And the fifth and final extraneous incriminating factor comes down to this: the accented candidate is better than the accepted candidate, and we want to see good people given a chance to show it, whenever feasible. To deny someone such a chance is probably to leave the world a worse place, in some respects, than it could have been. The best candidate won't get to develop their talents still further in a job they would have done well and my business won't thrive as much as it would have done. Again, perhaps I am not morally obliged to do all I can to improve the world, but nevertheless there is a supererogatory consideration in this.

---

[57] I borrow this term from Deborah Hellman. See Hellman (2008)

[58] Or perhaps an imperfect duty. Perhaps I have a duty to act to correct social ills sometimes, but no specific duty to act to correct social ills in any given instance. See Kant (1785), 4:421

Thus in many cases of this kind, Alexander will be able to satisfy our intuitions that my act of discrimination is wrong. Nevertheless, this leaves a small residue of cases in which Alexander must say my discrimination would be permissible. To see this, we must re-conceive the problematic case so as to eliminate the extraneous incriminating factors. For example, we need to suppose that I do not breach any obligation towards my boss or colleagues to maintain the health of the organisation I work for. Perhaps we can suppose that I am a sole trader with no colleagues. We also need to suppose that I do not let down the candidate. Let's suppose that I announce my aversion to the accent in the job ad, making clear that no disparaging judgments are involved, leaving accented candidates to apply at their own risk.[59] We also need to suppose that I do not feel any contempt whatsoever for this candidate, nor do I indulge in any other unjustified negative judgments. My attitude consists of irritation pure and simple. Let's suppose that I am genuinely positive about this candidate's abilities; in fact, I recommend the candidate to another employer. Finally we need to suppose that no HSD factors are in play. Let's suppose that my aversion to the trait is highly idiosyncratic so there are no wider consequences that follow from my discrimination. Also, the accent is from a relatively privileged community, and I know that this candidate will be able to find another job quite easily.

It is difficult to eliminate the last of the extraneous incriminating factors, to do with not letting the best candidate exercise their talents, since this would involve dropping one of the hypotheses of the original case. If there is some residual sense that something is wrong with my discrimination, I suggest it is to do with this.

But otherwise, once we conceive the case clearly in these terms, eliminating extraneous incriminating factors, I suggest that the appropriate response to the re-framed case is: Why shouldn't the owner of the business choose who works there, as long as he does not rely on any unjustified judgments about them, and does not mislead anyone? I'm entitled to stick with my idiosyncratic aversions, and act on them, unless I'm given a specific reason to disparage them. The only reason left in this case is that I'm not letting the best candidate exercise their talents. But it is implausible that everyone who does *that* is committing a wrong of the same kind as racism/sexism (if I employ the

---

[59] Like Arneson's doughnut shop owner, who announces that the successful applicant will be hired "by arbitrary whim". Arneson does not see this as wrongfully discriminatory.

second best candidate because they are a friend, I might be guilty of nepotism but surely not of anything akin to racism or sexism).

I therefore conclude that Alexander's account of wrongful discrimination is a good one.

2.8 Conclusion: Racist/Sexist Policies

A question which has been left open by the discussion so far is the status of healthcare allocation policies, and legislation more generally. Can policies be racist/sexist such that they are intrinsically wrong? I suggest they can, to the extent that they are dependent upon their supporters' biases and unjustifiable judgments of lesser moral worth. The question of whether a policy or law is racist or not can come down to the motivations of the people who passed it. If it was an expression of their negative attitudes towards people with a given trait, and the policy or law wouldn't have said what it does without those supporters' biases, then the policy or law is racist. For example, suppose two bus companies place the same, rather heavy emphasis on the need for drivers to be strong. The policy is defensible but at the top end of what reasonable people would do in terms of emphasis on safety concerns. Other companies quite reasonably place somewhat less emphasis on safety and more on passenger comfort, punctuality etc. In London Transport 1, the policy is a result of the drivers in the above case switching from a demand for direct discrimination against women to a demand for strength tests which have the same effect. The motivation is clearly still to exclude women. At London Transport 2, the requirement has been imposed by a primarily female management board. Their reason is a genuine concern for safety. I suggest we would have reason to think the former policy is sexist while the latter policy is not.

Similarly, if healthcare allocators deny funding for patients with "the English gene", and their motivation is that English people have it, that policy might be racist. But if their motivations are all to do with cost-effectiveness and maximising benefit, it might not be.

And finally, returning to Harris's argument, it does not seem to be true that the underlying principles which debar racism and sexism apply to all traits equally. Harris mentions race, creed, colour gender, and age. These all have some claim to be HSD traits (traits which have been discriminated against in the past). As such, as argued, groups with those traits should be protected from further disadvantage. However, Harris also mentions life expectancy and quality of life. In themselves, these have little claim to be HSD traits and are therefore not covered by any protection we might offer under the HSD principle. They are only protected by the Alexander's principle debarring discrimination motivated by negative judgments. So, in order to show that healthcare allocation based on these traits is wrong, Harris would need to show that healthcare policy-makers are motivated by negative judgments about life expectancy and quality of life. But generally, they are not motivated in this way. Thus Harris cannot sustain his objections against policies which deprive patients of treatments because of their poor prognosis in terms of life expectancy or quality of life.

Nevertheless, Harris's line of argument could cause problems in one case: that of life expectancy. As noted above, given that age is an HSD trait, perhaps it could be argued that one of the considerations debarring discrimination by race and sex also debars discrimination by life expectancy, since that could sometimes involve indirect discrimination by age. However, there are countervailing considerations regarding age, which I will consider in chapters 6 and 7.

3. What benefit should healthcare allocators aim to produce?

3.1 Introduction

What kind of benefit should healthcare allocators aim to produce when they decide which treatments to fund? This is a foundational topic, in that benefit provides the "currency" of a number of plausible principles of healthcare allocation: benefit is the item to be distributed in the way specified by the principles. For example, perhaps policy-makers should aim to maximise the benefit they produce. But what should they maximise - preference satisfaction, or happiness, or opportunities, or something else? Clarity on this question will help clarify our obligations with respect to benefit maximisation. The same could apply to distributive principles, such as egalitarian principles: clarity regarding the benefits of healthcare could assist clarity regarding the metric of those principles. In summary, the rest of this thesis is about how to distribute the benefits of healthcare. This chapter is about what we distribute.

To outline my line of attack on this question, my assumption is that characterising the benefits produced by healthcare involves determining the value of different health states, so that we can then determine the value of moving from one health state to another. I approach this question about the relative value of health states by considering who we should ask about the value of health states. There are three groups one could potentially ask for health state preferences: patients (people in the state), public (here defined as people *not* in the state), and experts such as doctors.

I argue that usually, the most authoritative view on how good or bad a health state is will be the view of people in that state - patients - since usually, patients are in a better position to evaluate it appropriately. I consider an objection to this, that taxpayers, as the funders of a public healthcare system, are entitled to have their health state preferences respected (and have different treatments prioritised, on the grounds that they produce the health states valued by members of the public). But I find that even the public should respect patients' health state preferences, since patient preferences will be the public's preferences should they ever need healthcare.

However, there are exceptions to this "patient first" rule, these being cases in which the public's evaluation of health states should be preferred, or perhaps even the judgments of experts. In particular, this is so where patient preferences are influenced by evaluative adaptation or where patient preferences are irrational.

I conclude from these observations that anyone could be mistaken about how good or bad a health state is. This suggests that health state value claims are objectively true or false. So the traditional focus on the question of patients vs public misses the point. The aim should be to get the best judgment. It doesn't matter whose judgment it is as long as it is right.

In line with this, I then consider an account of Daniel Hausman's according to which health states are valuable or not partly depending on the extent to which they constrain the possibilities of living well and pursuing valuable objectives. Although I endorse this view in outline, I take issue with Hausman's claim that we can rely on general principles which describe how various health impairments combine to constrain the possibility of living well. I argue that as far as practically possible, we must evaluate health states individually for the way they constrain the possibilities of living well, since health conditions have sui generis impairments and because the impact of a combination of impairments is not deducible from the their impact individually.

3.2 Background: Current Practice (NICE)

As background, I will start by characterising the kind of benefit that healthcare policy-makers actually concern themselves with. I will consider the UK since, helpfully for analytical purposes, it has a quite formalised and precise approach to assessing benefit.[60] In the UK, the National Institute for Health and Care Excellence (NICE) stipulates that "health effects should be expressed in QALYs".[61] A QALY is a quality-adjusted life year: viz., a life year, adjusted for quality of life. For example, if a patient expects to live two years at 50% quality of life, they can expect one QALY. To summarise NICE's procedure for assessing cost-effectiveness, NICE compares the number of

---

[60] Although a similar approach is used in UK, Netherlands, Australia, Canada and New Zealand

[61] NICE (2013)

QALYs a patient can expect without treatment with the number they can expect with treatment and derives how many QALYs the treatment adds. NICE compares this with the cost of the treatment to arrive at an "incremental cost-effectiveness ratio" or ICER.

In more detail, how does NICE calculate how many QALYs a patient can expect? This question requires us to estimate how many life years a patient has in prospect and what the patient's quality of life will be during those years. The patient's expected life years can be straightforwardly derived from the diagnosis. The job of characterising quality of life is slightly more complicated. NICE requires a two step procedure:

> the measurement of changes in health-related quality of life should be reported directly from patients and the utility of these changes should be based on public preferences using a choice-based method.[62]

Step 1 involves describing the health-related quality of life associated with a condition (from here on I will simply refer to this as quality of life). Step 2 involves valuing that quality of life. NICE's procedure requires that the quality of life associated with a condition be described by patients but then valued by the public.

NICE's preferred descriptive measure is a survey instrument called the EQ-5D.[63] The EQ-5D has a descriptive system comprising the following 5 dimensions:

- mobility
- self-care
- usual activities
- pain/discomfort
- anxiety/depression[64]

Each dimension has 3 levels:

- no problems
- some problems
- extreme problems.

---

[62] NICE (2013)

[63] NICE (2013)

[64] Rabin et al (2011)

This gives 243 possible health states (3 x 3 x 3 x 3 x 3).

Once a condition is characterised by means of this instrument, NICE wants to know how good or bad the condition is, as far as the public is concerned. Perhaps the most common way of doing this is by means of a time trade-off survey, whereby we ask respondents how much life expectancy in full health they would sacrifice to avoid a certain condition, and look for the point at which they are equivocal between two health outcomes.[65] For example, if a respondent is unsure between 4 years in full health or ten years with paraplegia, then paraplegia is deemed to be 40% quality of life.[66]

This procedure will generate QALYs figures for the prognosis with and without treatment, or more likely, a QALY figure for the prognosis based on the existing treatment and another QALY figure for the prognosis based on a proposed new treatment, so NICE can assess the new treatment for its cost-effectiveness as a possible substitute for the old treatment. For example, suppose that patient X has a serious, life-threatening condition, and would benefit from a new drug. If he continues receiving standard treatment he will live for 1 year and his quality of life will be 0.4 (0 or below = worst possible health, 1= best possible health). If he receives the new drug he will live for 1 year 3 months (1.25 years), with a quality of life of 0.6. The new treatment is compared with standard care in terms of the QALYs gained:

Standard treatment: 1 (year's extra life) x 0.4 = 0.4 QALY

New treatment: 1.25 (1 year, 3 months extra life) x 0.6 = 0.75 QALY

Therefore, the new treatment leads to 0.35 additional QALYs (that is: 0.75 –0.4 QALY = 0.35 QALYs).

Now suppose that the cost of the new drug is assumed to be £10,000, and standard treatment costs £3000. The difference in treatment costs (£7000) is divided by the QALYs gained (0.75 - 0.4 = 0.35) to calculate the cost per incremental QALY. So the new treatment would cost £20,000 per QALY.[67] Generally, if a treatment costs more than the ICER threshold (£20,000-30,000

---

[65] Attema et al (2013)

[66] A second common method is the Standard Gamble: Ask respondents to choose between living with a condition or accepting a gamble with a 40% chance of cure or 60% chance of sudden death. A third common method is the Visual Analogue Scale: Ask respondents to rate a condition on a scale from 0 to 100, with 0 representing being dead and 100 representing perfect health.

[67] The figures used in this example are from NICE (2017)

per QALY), then it would not be considered cost effective unless there are special considerations in play.

3.3 Whose Preferences?

A feature of interest in this system is the fact that preferences are elicited from the public, not from patients or any other group. There is some controversy over the question of who to ask for preferences. NICE assumes that the public's values are the appropriate perspective, and it is not alone in this. For example, consider the methods for assessing cost-effectiveness recommended when the U.S. Department of Health and Human Services commissioned its influential review of cost-effectiveness in health and medicine. The panel suggested that:

> preferences from the general population rather than preferences of particular subgroups should be used in a reference case [cost-effectiveness analysis] ... patient preferences should only be used ... for studies designed to compare alternative therapies for a patient group in a setting where resources have already been allocated to the treatment of their condition.[68]

But contrary to the assumption of healthcare allocators in the UK and US, I aim to show that patient preferences are generally better than the public's preferences as a guide to healthcare allocation.

The question of who to ask is contentious because the results we get from the different groups can be very different. Paul Dolan & Daniel Kahneman note that

> In the first empirical study on this issue, Sackett and Torrance (1978) asked the public and home dialysis patients to value a remaining lifetime with chronic dialysis using the TTO method. The average utility from the general public was 0.39, as compared to 0.56 from dialysis patients. If it were possible to return dialysis patients to full health, then the gain would be 0.61 QALYs per year (1.0-0.39) if we used public utilities and 0.44 if we used patient utilities. ...Since then, whilst some studies have found little difference in public and patient valuations, most have produced similar results to those found by Sackett and Torrance, with

---

[68] Gold et al (1996)

> similarly important implications for the results from cost-per-QALY ratios. In a review of 39 studies, including their own, de Wit et al. (2000) found that 23 studies report patient values to be higher than public ones, 2 report public values to be higher, 11 report no difference and 3 report contradictory findings.[69]

I will now consider arguments regarding whose health state preferences should be respected: patients or public? I will first consider arguments favouring patient preferences, and then arguments favouring the public's evaluation of health states. The arguments for each group come in two kinds: 1. Arguments in terms of alternative distributive principles (such as arguments about the likely consequences of endorsing principles to do with benefit maximisation and reciprocity). 2. Arguments which assume benefit maximisation, and conflict over which group is best placed to assess benefit.

I conclude that no group's views can be relied upon in all cases as the basis of healthcare allocation decisions. Sometimes we should ask patients, sometimes public, sometimes perhaps even healthcare providers. Occasionally we should ask more than one group and triangulate between them. However I claim that, as a general rule, going on what patients value maximises benefit *and* delivers reciprocity.

3.4 First Defence Of Patient Preferences: Patients Know More About their Conditions

I start by considering arguments in favour of going on patient preferences, and start on these with epistemic reasons to favour patient preferences. More often than not, a patient's evaluation of their own health state will be worthier of respect than the evaluation of a member of the public who is not in that state, simply because the patient is more informed about the state. There are two main sources of such public ignorance: functional adaptation on the part of patients, and the public's unfamiliarity with the condition itself.

On functional adaptation, Brock points out that one reason disabled patients might place a higher value on their disabled state than the public might be that disabled patients adapt to their condition; they "improve their functional

---

[69] Dolan & Kahneman (2008), p. 223

performance through learning and skills development,"[70] so they can still do many of the things that non-disabled people do. If divergence is entirely explained this way, we can take the public as making a mistake. The reason their judgment diverges from that of patients is that they fail to realise how much they would be able to compensate for a given disability to improve their functional performance.

On the public's unfamiliarity with the condition, one cause of such unfamiliarity is that the survey instrument with which health state preferences are elicited over-simplifies the health states in question. The EQ-5D instrument for measuring quality of life is vulnerable to this sort of problem. For example, Brazier gives an account of a study in which survey respondents who were asked to rate their own experience of their health state directly gave it a higher rating than when they were asked to rate the same state characterised by means of the EQ-5D. The hypothesis is that the EQ-5D omits important aspects of a health state, or over-emphasises the negatives at the expense of the wider context:

> The descriptive system may be insensitive to subtle differences within a dimension or simply miss out important dimensions of health altogether. This has been well illustrated by the Insinga and Fryback study of VAS [Visual Analogue Scale] data from the UK MVH survey, where the difference between general population VAS ratings of hypothetical states and their rating of their own health state by VAS was probably due to the limitations of the EQ-5D descriptive system. Most of the difference was accounted for by the discrepancy between the respondent’s self-rating and their separate rating of the EQ-5D state that coincidentally described their own health. While improving the descriptive systems can reduce this problem, it can never be entirely overcome.[71]

To explain this sort of finding, Peeters found that divergence between public and patients is often best explained by what they call "focusing". When public and patients are presented with a condition characterised wholly in terms of the five dimensions of the EQ-5D questionnaire, there is some evidence that they focus wholly on those dimensions, i.e. on how bad life is with that condition; whereas when patients with the same condition are asked to rate

---

70 Brock (2004)

71 Brazier et al (2005) citing Insinga & Fryback (2003)

their own experience of it directly they rate it higher, because they do not just see it wholly in terms of the negative aspects characterised in the EQ-5D.[72]

However, focusing is not the only explanation of divergence between patients and public. We see this when we consider studies in which patients are asked to rate their own health state, but they are not asked to rate their own experience of the state directly. Instead, they are asked to rate a *description* of their health state based on (for example) the EQ-5D questionnaire, without being told that the state in question is their own. It is found that even in these cases, patients assign higher values to EQ-5D profiles describing their own health state than public assign to the same profiles. For example, De Wit writes of survey involving such "hypothetical" states:

> The conclusion reached was that evidence that patients assign different values to *hypothetical* health states than 'outsiders' is growing, compared to the 1989 review of Froberg and Kane. Studies that found differences reported higher values from patients in most cases. This was affirmed in the present study. The SG and TTO values were compared for three hypothetical health states from dialysis patients and students and it was found that in five out of the six valuation tasks, patients assigned higher values than students.[73]

Of the explanations of divergence we have considered so far, the most applicable to this case look like functional adaptation or value change.

It should be noted that patient values are not always higher than public values. Although Dolan & Kahneman note that in the vast majority of cases where patient and public values diverge, patient values of an impaired health state are higher than public ones, they also note two exceptions:

> the two studies where patient values are actually lower than public ones are in samples of menopausal women and women in childbirth.[74]

Similarly, Dolan & Metcalfe find that

> Through our TTO [time trade-off] preferences, we seem to imagine that 'some problems walking about' is about as bad as 'moderate anxiety or depression.' The general public in the United States (and in the United Kingdom, too) is willing to give up about 15% of their remaining life

---

[72] Peeters et al (2012)

[73] De Witt et al (2000), p. 115

[74] Dolan & Kahneman (2008), p. 223

> expectancy to avoid each of these states. In contrast, through assessments of SWB, we find that anxiety/depression has about 10 times as much impact as mobility. Interestingly, in our imaginations, being confined to bed or having extreme pain is worse for our well-being than having extreme anxiety. But from the experiences of people's lives, extreme anxiety is worse than being confined to bed or having extreme pain.[75]

What the menopause, childbirth, anxiety and depression have in common is that much of their disvalue arises from mental suffering, which is relatively invisible to those who do not go through the relevant experiences. This suggests that non-sufferers do not realise how bad the condition actually is.

In all these cases it seems that the divergence between patient valuations and public valuations is the public's ignorance of what the condition is really like. In cases where this is the explanation, patient valuations can be trusted.

3.5 Second Defence Of Patient Preferences: Principle of Benefit Maximisation

It also looks like most plausible principles of benefit maximisation will commit us to characterising benefit from the point of view of the patient (on the assumption that a general principle of benefit maximisation will have specific implications for the benefits we distribute in healthcare allocation[76]). However, I should clarify before continuing that I'm not yet committing to a principle of benefit maximisation. Rather I merely argue that any plausible principle of benefit maximisation we might want to adopt is likely to commit us to characterising benefits from the point of view of the patient.

So, as an example of this line of argument, John Brazier analyses benefits in terms of preference satisfaction, and writes:

> [welfare economics] asserts the supremacy of an individual's evaluation of their own well-being. This implies that it is the preferences of the losers and gainers from a public programme that should be elicited, and not a sample of the general population who will be unaffected by the change; this would seem to suggest that patient values should be used.[77]

---

[75] Dolan & Metcalfe (2012), p. 3

[76] In Chapter 5 on disability discrimination, I consider whether healthcare allocators should take of non-health benefits they might produce, concluding that they shouldn't.

[77] Brazier et al (2005), p. 204

Brazier is a health economist, but we needn't commit to classical welfare economics to endorse an argument of this form. A similar argument could be put in terms of obligations of beneficence. In summary, unless a condition is very common, most respondents to a survey of the general public will not have the condition. That means they do not stand to benefit from the treatments of interest. And the value of a benefit is its actual value to those who will actually benefit, not its hypothetical value for those who will actually not benefit. So, if there is a divergence of values between patients and public, a survey of patients will better predict its actual value (to patients) than a survey of the public. In more detail, the argument starts with the premiss that, all else being equal, policy-makers are under a pro tanto obligation to produce the most good they can with their resources. This requires us to ask, to the extent that a health state is good in the way we aim for as healthcare providers, who is it good for - the people in that state, or other people? It seems clear that we aim to produce health states that are good for the people in them. To illustrate with a particular theory of the good (a theory held by some health economists, though not by me), let's suppose for the sake of argument that a health state is good to the extent that it satisfies preferences. When we ask the public about their health state preferences, we ask about a hypothetical situation in which *they* are in one or other health state; we do not ask them what health states they would like to see produced in actual patients. Thus whichever patients we treat, we will not satisfy the public's expressed preferences; if a member of the public said they would prefer to be in health state A and not health state B, we will not satisfy that preference by putting *patients* into state A rather than state B. Having said that, I accept it is possible that certain members of the public will have preferences regarding what health states should be prioritised as healthcare outcomes, even though they themselves will not be the beneficiaries of those outcomes. But in terms of the kind of benefit we aim for as healthcare providers, we won't *benefit* those individuals by satisfying those preferences. So, if benefit is defined in terms of preference satisfaction, healthcare allocators must consider the extent to which healthcare states are good from the point of view of patients, not from the point of view of the public.

I conclude that usually we should go on patient preferences. This is because patients are always in a better position to know what a health state is like, and usually they are also in a better position to evaluate it. But I emphasise this is only usually; there are exceptions, as I will consider later.

3.6 Objection to Deriving Analysis of Benefit from Principle of Benefit Maximisation

Meanwhile, I must consider a proponent of public preferences who objects that in deriving our analysis of benefit from the principle of benefit maximisation in this way, we have ignored another relevant principle. I will consider the objection before rejecting it. The objection invokes a principle of reciprocity, according to which we should go on the public's evaluation of health states, on the grounds that taxpayers, as the funders of a public healthcare system, are entitled to have their preferences respected. For example, David Gauthier argues:

> Providing services against which one would not choose to insure oneself (given average resources and a general knowledge of health care needs without specific knowledge of one's own condition) is not only not required by fair access, but, in diverting resources from preferred uses, is positively unfair.[78]

David Hadorn, commenting on this, argues that treatments must represent good value to the people paying the bills, viz., premium payers. Patient preferences do not take account of cost and therefore will not be a good indicator of good value in these terms. "Permitting patients unlimited access to care based on post-illness preferences would too often result in the provision of marginally beneficial care."[79] The unfairness on Hadorn's account is that premium payers would be paying premiums for services they do not want because they are poor value (as far as the payers are concerned), and then finding that the services they want are not provided because there is no budget left for them.

Having said that, perhaps Hadorn had the US system in mind, which was largely funded via private insurance when he wrote this. However, similar defences have been offered of being guided by the public's evaluation of health states in the context of publicly funded systems. For example, Brazier writes:

> A related argument is the insurance perspective, where resource allocation decisions in healthcare are akin to the decisions about which services should be covered by an insurance package. Public funding can

---

78 Gauthier (1992) quoted in Hadorn (1991)

79 Hadorn (1991)

> essentially be seen as public insurance and so it is the ex ante the public's evaluation of health states that should be used to value health states.[80]

Hadorn and Brazier both rely on a similar premiss concerning the rationale for a publicly funded health service. They assume that publicly funded healthcare acts as a kind of mutual insurance fund. On this model, we all make affordably small, regular contributions to a mutual fund in return for the ability to make occasional big withdrawals as needed. This is a quasi-contractual model, or a model based on reciprocity. Hadorn and Brazier's argument is that "He who pays the piper calls the tune". Taxpayers, as the funders of the health service, are entitled to specify what services it should provide.[81] In contrast, the Brazier-inspired argument above ignores people's status as part-funders of the healthcare system; patients are seen merely as prospective beneficiaries of healthcare and healthcare policy-makers are seen as moral agents under a general obligation of (say) beneficence, assessing how they can produce most benefit with the resources available to them.

However, I will now argue that this argument fails. Even if we grant that public preferences should be sovereign in healthcare allocation policy, it is rational for the public to respect patient preferences, since patient preferences will be the public's preferences should they ever need healthcare. To make out this argument in a bit more detail, it is helpful to start by considering the question of what attitude we should take to our own future preferences. Suppose I can reasonably expect my preferences to change over time, for example as a result of getting older. For example, I will not want to party as much. Should I want those future preferences to be satisfied? As a general rule, I will have reason to want my future self's preferences to be satisfied. For example, I have reason to expect that my future self will be older and wiser so

---

[80] Brazier et al (2005), p. 204

[81] Perhaps the "mutual insurance" rationale will seem contestable. It might seem to picture a seemingly noble and altruistic institution as a grubby self-interested scheme, whose participants need pay no attention to needy non-payers. So, to defend this conception of a publicly funded health service: consider the case of two people who both have some unpleasant disease such as eczema or persistent backache, a disease which doesn't stop them working or pursuing most normal activities, but which does significantly reduce their quality of life. If we had to choose who to fund treatment for, I suggest that taxpayers might have a stronger reason to fund treatment for the fellow scheme member than for the other person, simply on grounds of reciprocity. Even setting aside contractual considerations, we can say they paid their dues and they are entitled to a return on them. This suggests that a publicly funded health service can be seen, at least in part, as a mutual protection scheme.

my preferences will be worth respecting.[82] I suggest that many cases of patient/public divergence in terms of health state preferences are like this. Usually I can expect that if I get into a particular health state, my preferences regarding that health state will go the same way as other patients in that state, and that the reason for this change will be that I become better informed about the health state. As such, I have reason to want my future self's preferences regarding that health state to be satisfied rather than my own preferences regarding that health state. So even if publicly funded health service can be seen as a mutual insurance scheme, and therefore should satisfy the preferences of its "members", the preferences that should be satisfied are the preferences its "members" *would* have were they in the health states in question, rather than the preferences they have now. Generally, that means going on the preferences of actual patients.

However, I must consider an objection to this conclusion that can be derived from a point of Daniel Brock's. He offers as one explanation of divergence what he calls adjustment, that is "altering one’s life plans to give greater importance to activities in which performance is not diminished by disability."[83] He comments:

> neither the nondisabled nor the disabled need have made any mistake in their different evaluations of quality of life with that disability. They arrive at different evaluations of the quality of life with that disability because they use different evaluative standpoints as a result of the disabled person’s adaptation, coping, and adjustment. Disabled persons who have undergone this process can look back and see that before they became disabled they too would have evaluated the quality of life with that disability as nondisabled people now do. But this provides no basis for concluding that their pre-disability evaluation of the quality of life with that disability was mistaken, and so in turn no basis for discounting or discarding it because mistaken. The problem that I call the perspectives problem is that the nondisabled and the disabled evaluate the quality of life with the disability from two different evaluative perspectives, neither of which is mistaken.

---

[82] Admittedly there may be times when I shouldn't, as when I have reason to believe I will be corrupted by the wealth I am due to inherit, or if I expect to get a health condition to which I will adapt in Sen's suboptimal way; but these cases are exceptional

[83] Brock (2004), p. 293

The point I take from this is that if the divergence between public and patients is not explained by the public's lack of information but by their different values, then there are no grounds to disparage the public's preferences. So it might look like both public and patients are equally entitled to have their values respected, since neither has made a mistake. Patient preferences and public preferences are on equal footing. This may seem to undermine our reason for going on patient preferences.

I concede that Brock's analysis will be applicable in some cases. However it will not be true of all cases; in fact, it characterises a minority of possible cases. Consider the case of someone who has acquired paraplegia and who rates their quality of life as 50%. Presumably before they got paraplegia, they would have given it a lower rating than that, like everyone else. So now they are paraplegic, how will they assess their previous evaluation of paraplegia? I suggest that evaluation will seem like a mistake. We have no reason to question that retrospective judgment. That individual has known both states, so as of now they are better informed. Also, suppose that this individual is then cured; I doubt they would retract the judgment they made when paraplegic that their quality of life was 50%. So this is a counterexample to Brock.

However, values change for different reasons and sometimes there will be no such grounds for saying that the public is mistaken. This is particularly where a patient abandons goals that have been frustrated by their health impairment, such as switching from rugby to art as a result of paraplegia. This could be accompanied by a change of values, with the individual seeing their art as a motivating and valuable project. In such a case, we needn't think of the old values or the new values as mistaken. But even in such cases, as a rule, we should still favour patient preferences - not because they are better, but simply because they are the preferences the public will have if they need healthcare. For example, consider someone who is pursuing a rugby career. Upon taking an extremely advanced and reliable personality test, he learns that if he was ever to get a spine injury that stopped him from playing rugby, he would discover that he had great talent as an artist. The personality test also reveals that after five years, he would end up with completely different goals and values, and would not want to return to rugby even if his paraplegia was cured. In considering healthcare priorities with respect to paraplegic patients, should this individual defer to the preferences of his possible paraplegic self (according to which paraplegia is, say, 50% quality of life) or should he insist on his own current preferences (according to which paraplegia is worse than death)? I contend he still has reason to defer to the preferences of his possible paraplegic self. However this is not because those preferences are better informed, but simply that they are not to be disparaged (they are not irrational, for example), and they are the preferences he would have at the time he needed healthcare. To simplify, suppose I am the rugby player. Call the actual rugby player "me-now" and his possible paraplegic self "paraplegic me". If I become paraplegic, me-now won't be around any more; the prospective beneficiary of healthcare will be paraplegic me. If paraplegic me would prefer health state A to B even though me-now prefers B to A, to insist on me-now's preferences would be to frustrate the preferences of the only person who can benefit from whatever interventions are on offer. As a rule, it would seem irrational to frustrate one's future self's preferences.[84] Thus generally, the public should defer to patient priorities where any difference of priorities is explicable in terms of divergent values.[85]

Of course it might be seen as problematic that adaptive preferences are effectively forced on the patient by their condition. As such shouldn't the preferences be seen as non-autonomous and therefore not to be respected? In line with this, Jon Elster argues that wants should not be taken as given:

> [we should] inquire into their rationality or autonomy. These, in the general case, are properties that cannot be immediately read off the wants themselves. ... Rationality in the broad sense depends on the way in

---

[84] Even if one would rather be dead than paraplegic, I suggest that is not a reason to frustrate one's possible paraplegic self's desire to stay alive. Instead, knowledge that one's desires would change in this respect would be reason to think that paraplegia wasn't as bad as it seemed. (Thanks to Carl Fox for the query that prompted this point)

[85] It should be noted that this is in the usual case where patients acquire their distinctive values as a result of having the condition. An unusual case is possible in which values are not the result of health impairments, but a partial cause of them. For example, suppose that all the sports played by a given population are dangerous, such that they produce lots of injuries. As a result, a large proportion of patients are sports lovers. However, not everyone is a sports lover; in fact, a large majority of the general population values intellectual activities and not sport. In this case, the majority of the public does not have the same reason to respect patient preferences, since they will not form the same preferences as patients if they need healthcare, even if they get the same conditions as the sporty patients. In this case, considerations of reciprocity might pull against considerations of benefit maximisation. To see this, suppose a condition produces both a physical impairment and a mental impairment, and so suppose that sporty patients with the condition want the physical impairment treated and non-sporty patients want the mental impairment treated. In such a case, the public could appeal to considerations of reciprocity to argue for prioritising treatment of the mental impairment even if a slight majority of patients want treatment of the physical impairment. In contrast, benefit-maximising considerations support treating the physical impairment. Thus in such a case, there can be a tension between the principle of reciprocity and the principle of benefit maximisation. However, in reality such cases are unlikely.

> which the states are actually formed. Two individuals may be exactly alike in their beliefs and wants, and yet we might assess them differently from the point of view of rationality, judgment and autonomy."[86]

But I respond that the provenance of patients' adaptive preferences does not necessarily give us reason to disparage them. Suppose a patient knows her life will go better with preferences for A rather than B, because the patient's condition means she can have can have A but not B, and therefore the agent forms preferences for A rather than B. Suppose an "ideal observer" version of the agent would have done the same thing, with full information and time to reflect. Then those preferences should be respected on the grounds that the agent's life will thereby go better. Of course, we shouldn't just dismiss all Amartya Sen's worries about adaptive preferences, but there are important differences between this case and the paradigm cases which prompted Sen's analysis. In particular, the hardships that Sen observed were brought about by unjust social arrangements. In such cases, policy-makers have reason to remove the injustice rather than respecting the preferences. In contrast, we can't usually cure a disease by removing a social injustice that caused it. Also, I tentatively suggest that in some of the situations that Sen observed, people might have convinced themselves they had preferences which actually they didn’t. Give them a choice and they would immediately take the better option, despite not having shown any interest in it previously. If expressed preferences aren't real, we needn't respect them.

But does that not lead to seemingly unfair consequences in the disability case? Suppose that if we go on patient preferences rather than public preferences, curing disability will produce less preference satisfaction, because the patients' functional adaptations and evaluative adaptions mean their life with the disability is not as bad as the public thinks. Then my proposal entails that we should place lower priority on curing disability. It seems unfair that as a result of making the best of their condition, patients then get stuck with that condition. But this "unfairness" pervades policy-making. We all made choices early in life as result of accidents to do with our early circumstances, and no doubt formed somewhat adaptive preferences favouring the opportunities that life presented to us. Now, policy-makers must respect the resulting preferences. It would be expensive and counterproductive to try and turn back the clock and reinstate the preferences we "should" have had in ideal

---

86 Elster (1985), p. 140

circumstances. Disabled patients are in the same boat as the rest of us, in this respect.

3.7 A Defence of the Public's Evaluation of Health States

I now consider other reasons for going on public preferences, where the public is defined as non-patients. I have already considered one argument, based on considerations of reciprocity, but rejected it. However, considerations of benefit maximisation can sometimes favour going on the public's evaluation of health states, on the grounds that patient preferences should be considered unreliable. I will consider three cases which might trigger such concerns: adaptive preferences; irrational preferences; and suicidal people. I will argue that there are some cases in which we should favour public preferences, because of the concerns stated above, but these situations are exceptional.

The first case is evaluative adaptation. Sometimes, going on patient preferences looks like it could fall foul of Amartya Sen's worry that people evaluatively adapt to problems in their life, so their preferences are not a good guide to their welfare. It looks like many patients evaluatively adapt to their health impairments. To see this, note that most patients' evaluation of their own health state is higher than the public's evaluation of the same health state (for example, see De Wit's finding, quoted above, that out of 39 studies, 23 reported patient values to be higher than public ones, whereas only 2 reported public values to be higher). There might be explanations of this that throw patient's high evaluation of their health state into question. For example, Menzel quotes Sen:

> Utilitarian ethics is guilty of an "overdependence on what people ‘manage to desire"’ that is "neglectful of the claims of those who are too subdued or broken to have the courage to desire much. A thoroughly deprived person, leading a very reduced life, might not appear to be badly off in terms of the mental metric of desire and its fulfilment, if the hardship is accepted. In situations of long-standing deprivation, the victims do not go on grieving and lamenting all the time. The extent of a person’s deprivation, then, may not at all show up in the metric of desire

> fulfilment."[87]

Sen's point helps us see the dangers of accepting a patient's positive evaluation of their health state. However, it should be noted that contrary to this, Menzel holds that there is a disanalogy between Sen's case and that of patients:

> Adaptation often involves genuinely successful achievements and shrewd control over the trajectory of a person's inner life. In these cases, the adapted person is anything but ''broken'', and hardly ''subdued''. If what is still seen as deprivation is handled by a person in terms of challenge and ''achievement'', does not a metric of desire fulfilment regain its integrity? Thus, while Sen's argument from entrenched deprivation should give us pause about too readily or generally using adapted patients' quality-of-life ratings, it does not justify an across-the-board rejection of values shaped by adaptation.[88]

The suggestion here is that the patient has to be in quite an extreme situation before we question their positive evaluation of their health state. But I object that a patient does not need to be broken or subdued for their evaluation to be questionable. Lowered expectations may be enough to give us pause, as with a patient who has stopped wanting to go out because they can't. They find their indoors life OK, because they've forgotten what they used to do outdoors and they've found other things to keep them occupied indoors. But we might have reason to think that such a patient would get a lot of benefit from going out and we can therefore reject their relatively positive evaluation, even if they are neither broken nor subdued. They get more pleasure than they used to from being indoors, but without their health impairment they could get much richer pleasures from the things they could do outdoors.

Thus there may be more circumstances in which we would be justified in rejecting a patient's positive evaluation than Menzel suggests, and going on public preferences.

My second case is patients whose preferences can be called into question on grounds of irrationality. Now, before I can consider this category of preferences, it must first be acknowledged that philosophers such as Hume questioned the idea that preferences can be irrational. For example, he said:

---

87 Menzel et al (2002), p. 2154

88 Menzel et al (2002), p. 2154

> It is not contrary to reason to prefer the destruction of the whole world to the scratching of my finger.[89]

But contrary to Hume, some preferences seem clearly irrational, or at least discreditable to the point where it is legitimate for others to go against them. In some cases, perhaps this is because, by intervening, we will change the patient's preferences, and they will end up with more and deeper preferences getting satisfied even though the none of the preferences they had before will be satisfied. In other cases, it may be because some preferences involve the patient neglecting their ethical responsibilities. For example, consider the preferences of alcoholics or drug addicts or gambling addicts. Some gamblers express no desire to stop. So, if we were guided by patients preferences alone we would not treat them. Nevertheless we often treat them; family and friends may put pressure on them to get treated, effectively against their will. Now, the proponent of patient preferences might object that this is because we should go on such individuals' *informed* preferences. The argument would be that after being cured, these patients will think of themselves as better off, and at that stage they will be better informed, having experienced life with and without the addiction, and it is this informed preference that justifies our imposition of treatment. But in response I contend that many gamblers will already be quite well informed about what life is like without addiction; they may remember quite clearly what their life was like before they became addicted, but addiction has changed their sense of priorities such that important things like job and family look unimportant. Furthermore, the "information" presented to a gambler during treatment may not change their health state values. Yet we will not respect *these* better informed preferences following treatment; instead we will say that the gambler was not cured, and that their preferences continue to be distorted by their continuing addiction. So, contrary to the proponent of patient preferences, even the patient's informed view loses out to public opinion regarding health state values.

Having said that, cases of this kind are not always clear-cut. Consider Body Dysmorphic Disorder, which the NHS defines as "an anxiety disorder that causes a person to have a distorted view of how they look and to spend a lot of time worrying about their appearance."[90] Perhaps Body Dysmorphia cases lie on a spectrum in terms of whose preferences we respect. At one extreme, as suggested by the designation as an "anxiety disorder", these cases should be

---

[89] Hume (1739/1978), section 2.3.3.6

[90] NHS (2016)

seen as manifestations of a psychological condition rather than a physical condition. For example, this will be the usual diagnosis when the patient wants to remove a leg. In that case, the patient's preferences will not be seen as a reliable guide to the kind of treatment that is required. They will be viewed as requiring therapy to change their preferences and beliefs, rather than surgery to remove the leg. But at the other extreme it might be right to treat such cases as physical disorders. For example, suppose that psychological therapy does not change the patient's view. Also suppose that the bodily change they seek is not obviously harmful; this is a female with a completely flat chest who wants breast enlargement. Her existing breast size is not causing problems and breast enlargement will not be beneficial in respect of anything but the patient's mental state. So there are considerations that could justify an analysis of the patient's condition as a psychological condition, as in the leg removal case. But the overall balance of considerations might favour going along with patient preferences and funding breast enlargement, thereby treating it as a physical condition. However, this does not undermine my point that in some cases, we will treat the patient's preferences as irrational and insist on going with the public's view of health state values.

My third case consist of suicidal people. Sometimes we will stop a suicidal person from acting on their wishes, even if they genuinely think death would be better than life. And perhaps sometimes, such interventions are justified, as when if we think there is a chance that life will get better for the suicidal individual. Similarly, euthanasia is banned in most countries, partly on the grounds that the patient's expressed and well-informed wishes may not reflect their true interests.

Thus sometimes, we should go on the public's evaluation of health states; in particular, where patient preferences are influenced by evaluative adaptation or where patient preferences are irrational or have the wrong kind of provenance.

3.8 Hausman's account

The main lesson I wish to draw from the above is that anyone could be mistaken about how good or bad a health state is. The public could be wrong, such as when they are uninformed about what it is actually like to be in that state. But also, even people in the state could be mistaken about the value of their own health state, such as when their expectations have been lowered by their health state.

The possibility of mistakes suggests that health state value claims are objectively true or false. This may be so even if health state values are different for different people (similarly, it be objectively true that eating wholegrain bread is good for me, but also objectively true that it is bad for you, because you have a gluten allergy).

So, if health state value claims are objectively true or false, that suggests that the traditional focus on the question of patients vs public misses the point. The question isn't whether one or other groups is the default source of preferences simply in virtue of their position, for example as healthcare beneficiaries or healthcare funders. The question is which health states are most valuable. That means the aim should be to get the best judgment. It doesn't matter whose judgment it is as long as it is right.

I will therefore now explore theories that work on the assumption that health state value judgments are objectively true or false. Hausman is one such theorist. The focus of my discussion now moves from the question of who to ask to the question of the nature of the benefit we should aim to deliver in healthcare.

As background for his argument, Hausman writes:

> The general cure that philosophers have offered to the problems that arise when preferences reflect false beliefs and cognitive deficiencies is to link value (and well-being in particular) to rational or informed preferences rather than to manifest preferences. The same cure is available in the case of health-state evaluation: let informed preferences rather than manifest preferences be the standard by which health states are compared. To measure informed preferences, provide survey respondents with information about the character and consequences of the health states. Give them time to reflect. Frame the questions in multiple ways. Challenge them to give reasons.[91]

However Hausman eventually goes further than this, developing a quite systematised theory of health state values. As an illustration of the way he wants to go, Hausman argues that we should aim to understand the features

---

[91] Hausman (2006), p. 262

that make health states valuable and derive the value of health states from that analysis:

> Suppose, for example, that people cared about health only insofar as health affects longevity. In that case, even if value depends on preferences, the best way to evaluate health states would be by determining their consequences for longevity rather than by measuring preferences for health states.[92]

Essentially, Hausman does not think we need to spend a lot of time understanding people's preferences for individual health states. Once we understand the features that make health states valuable, we can derive the value of any given health state by performing a non-evaluative assessment of the extent to which each it exhibits the valuable features. Ultimately, Hausman agues that the value of a health state comes down to whether it constrains the possibilities of living well and pursuing valuable objectives. Here he gives longevity as an example: some health states would deprive us of longevity and thereby constrain the possibilities of living well and pursuing valuable objectives. His claim is that instead of evaluating every health state, we can analyse the extent to which health states make these outcomes possible.

Hausman offers the following argument for this view:

> a preference for health state S over another health state S' does not "happen" to a person in the way one's mouth starts watering when one sees a rich chocolate cake. One does not wake up one morning with a preference for being blind rather than deaf, as one wakes up hungry or with a tune running through one's head. One makes up one's mind about which health states are better and which are worse. If asked, "Why is S better for you than S'?" one responds by explaining how or in what regard S is better for oneself, not by insisting on one's preference. ... Except in the limiting case of basic preferences, each of us constructs his or her preference ranking on the basis of judgments of value, which in turn depend on reasons. Those concerned to evaluate health states should be concerned with these reasons.[93]

Similarly, Hausman contrasts two imaginary worlds, "Moneyland" where people care only about money and believe that health states should be

---

[92] Hausman (2006), p. 258

[93] Hausman (2006), p. 267

evaluated in terms of the expected incomes they permit, and "Tasteville," where people have only basic preferences among health states, for which they can give no reasons. He asks:

> Why ask Moneylanders what they think the consequences of health states for income are (which is what health economists are doing when they measure the preferences of Moneylanders) rather than tackling questions about consequences directly?[94]

Thus although Hausman earlier entertained the claim that informed preferences were an improvement on uninformed preferences, Hausman now rejects even informed preferences as a guide for policy-makers. He argues that as soon as we recognise the superiority of informed preferences over uninformed preferences, we have reason to look beyond preferences for the source of value. He quotes Scanlon:

> The introduction of the adjective 'informed,' which looks like a small qualification, in fact represents a significant departure. Informed desires are desires which are responsive to the relevant features of their objects.[95]

Hausman then comments:

> Scanlon is right to point out that an informed preference view shifts attention away from whether preferences are satisfied to the unspecified features of health states that make them better or worse in some regard.[96]

So, moving on to specifics, Hausman's theory is as follows:

> Health states should be evaluated in terms of how severely they limit the range of valuable lives individuals can live. Rather than asking respondents ‘Do you prefer H1 to H2?’ or ‘Are people better off in H1 or in H2? or simply ‘Is H1 better than H2?’ the question to ask is ‘Does H1 constrain the possibilities of living well and pursuing valuable objectives more than H2 does?’ The evaluative relation ‘G’ in terms of which health states should be compared and ultimately quantified is not preference, consequences for well-being, nor some unspecified ‘better than’ relation. It is rather something like ‘capability enhancement’ or, put negatively, ‘capability constriction’.[97]

---

[94] Hausman (2006), p. 266

[95] Scanlon (2003) quoted in Hausman(2006), p. 263

[96] Hausman(2006), p. 263

[97] Hausman (2010), p. 287

Hausman also suggests a new health state classification in terms of just two dimensions: activity limitations and health-related feelings:

> [these are] the two aspects of health that most directly bear on the question of how health limits capabilities. In making specific activities impossible or difficult, health states can rule out certain kinds of lives and projects. In addition, suffering and distress, both physical and mental, interfere directly with achieving many good lives. The features of health states that are relevant from a public perspective can be summarized by an activity-limitation/feeling (a/f) pair. For policy purposes, health-state evaluation can focus on just two dimensions along which health states can be located: activity limits and feelings. Although it would be possible to construct questionnaires to determine directly activity limitations and feelings, what I have in mind instead is mapping health states, as defined in a more complicated health-stated classification system such as the EQ-5D ... into the equivalence classes defined by the simpler classification in terms of feelings and activity limitations. This task is largely non-evaluative. The a/f pair that typically corresponds to which array of functional limitations is a complicated matter of fact.[98]

As an illustration, Hausman offers the following activity/feeling classification system:[99]

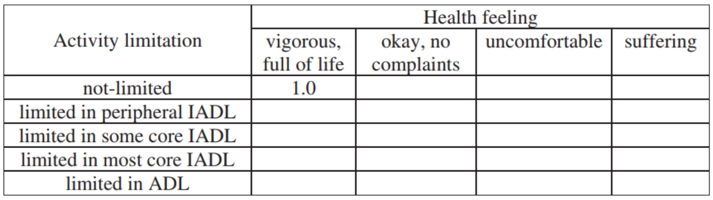

Activity limitation	Health feeling			
	vigorous, full of life	okay, no complaints	uncomfortable	suffering
not-limited	1.0			
limited in peripheral IADL				
limited in some core IADL				
limited in most core IADL				
limited in ADL				

Fig 1: Hausman's illustrative activity/feeling classification system

'ADL' stands for 'activities of daily living', and IADL stands for 'instrumental activities of daily living'. Each row is defined to be a more severe limitation than the row above. Hence, for example, the third row should be defined more precisely as 'limitations in peripheral IADL without limitations in core IADL or in ADL'.

Hausman emphasises that classification of health states into the above matrix of activity/feeling pairs is a factual matter:

> It is important to distinguish the classification of the activity limitations

---

[98] Hausman (2010), p. 288

[99] Hausman (2010), p. 292

> health states cause from the evaluation of a/f pairs in terms of the extent to which they limit the range of lives and projects open to individuals. The classification, unlike the evaluation, is primarily a factual matter.[100]

As a result of this classification, a health state might get assigned a factual description to the effect that the feeling is "uncomfortable" and activities are "limited in ADL". The health state thus characterised must now be evaluated on a scale of 0-1:

> The task is to assign numbers to a/f pairs – that is, to the cells in [the above figure]. Full health (the top left cell) has the value 1 and death has the value 0. The numbers assigned to health states are supposed to quantify how health affects capabilities.[101]

The numbers assigned to each a/f pair

> represent how severely a/f pairs limit how people can live and what significant projects they can pursue.[102]

Although these numbers represent evaluations, Hausman emphasises that they do not represent evaluations of the health states themselves:

> the problem evaluators face is not to figure out their preferences – their all-things-considered ranking of health states. Their problem is different and much more narrowly defined. Their task is to rank and ultimately to assign numbers to a/f pairs, where these numbers represent how severely a/f pairs limit how people can live and what significant projects they can pursue.[103]

However, Hausman recognises that

> evaluative disagreements are bound to remain. People can invoke well-known ideals governing what kinds of lives and activities are most valuable to argue that some a/f pairs make available a wider range of better lives for more people. It would however be overly optimistic to suppose that consensus will result. What then? ... the only way left to cope with disagreement is procedural, which in the case of roughly democratic societies involves deliberation coupled with some mechanism for decision making while deliberation proceeds. Health-state evaluation

---

[100] Hausman (2010), p. 288

[101] Hausman (2010), p. 293

[102] Hausman (2010), p. 288

[103] Hausman (2010), p. 288

> should, I suggest, depend on deliberative groups coupled with public debate, rather than on surveys of individual judgments.[104]

To summarise, Hausman proposes that there is a sequence of empirically determinable connections running from a health state, to activity-limitation/feeling pairs, to available lives/projects. For example, being paraplegic (health state) determines that one lacks mobility to a certain degree (reducing one's score for activity-limitation in the a/f pair), and this in turn prevents one participating in certain sports (narrowing the range of possible lives and projects). On Hausman's account, no evaluation of health states is required to justify the claim that the health state should get a lower a/f score ("the problem evaluators face is not to figure out their preferences"), or the claim that the lower a/f score means fewer lives and projects are possible. However, that narrowing of lives and projects can be evaluated, and as a result the a/f pair can also be given an evaluative score based on its implications for possible lives and projects. Our (deliberative) evaluation of the a/f pair can be seen as a sort of "thick" evaluation which combines the factual judgment regarding impact on lives/projects with an attitude towards those facts. The extent to which a/f pairs limit lives/projects is a matter of fact, but our evaluation of those limitations on lives/projects need *not* be seen as a matter of fact.

This completes my exposition of Hausman. However, as an addendum, it should be noted that a couple of things are missing from his account. First is the question of how he would handle life expectancy. I take it that for any given a/f pair, it makes a difference whether the patient can expect to live for only a month in that state or for 70 years. This dimension of health needs to be included in the health outcome descriptions to be assessed by his deliberative groups.

Second, Hausman does not say what living well consists of. This seems like a significant omission, considering that his deliberative groups will be asked to consider the extent to which health impairments constrain the possibilities of living well and pursuing valuable objectives. However, in setting this question, Hausman says he is not interested in the consequences of a health state for wellbeing, and it looks like this explains the omission (since I assume that to live well is to have wellbeing). It also looks like his motivation for this

---

[104] Hausman (2010), p. 288

lack of interest can be found in an earlier argument that the role of government is limited:

> Individuals appraise actions and states of affairs, including health states, from the perspective of their own specific objectives. ... The state, on the other hand, evaluates policies without committing itself to any personal objectives. Its goal is to expand and secure the range of alternatives that are accessible to individuals. ... Government serves as a referee, a protector, a facilitator, and an insurer, not as a big brother. ... The constituents of a meaningful life, such as close companionship, are not in general things that government should provide. Individuals must provide them for themselves. The role of government is to provide individuals with the means and opportunities to pursue for themselves as wide a range as possible of these goods.[105]

Back in the sphere of health, Hausman concludes from this that

> while from a personal perspective what matters most about health states is their bearing on how successfully an individual functions in the way of life the individual has chosen, the criterion of evaluation from a public perspective should be how severely health states limit the range of alternative lives and pursuits that are open to individuals. From a personal perspective, ... Bad health diminishes subjective well-being or prevents individuals from achieving their ends. From a public perspective, the significance of bad health lies not in ultimate outcomes, but in the extent to which it diminishes capabilities.[106]

But the question of whether a health state constrains the ability to live well will raise more questions than Hausman wants to answer here. If we don't know what it means to say someone lived well, how can we say whether or not a given health state constrains the possibilities of living well?

Perhaps Hausman would respond that as a capability theorist, he can stay neutral on this question: his deliberators will aim to state capabilities that are valuable regardless of your view of the good life. But it is not possible to specify the most important capabilities without first identifying the most important aspects of the good life for the prospective beneficiaries of the resulting policy. Why should we expect people from a sports-loving nation

---

105 Hausman (2010), p. 286

106 Hausman (2010), p. 287

like Norway[107] to prioritise the same capabilities as people from a book-loving nation like India?[108] Perhaps in response, Hausman will concede that his deliberators will need to understand what the good life consists of, but argue that his they will be perfectly able to apply their judgment to this question on a case by case basis, applying their intuitive concept of wellbeing, so there is no need to give any account of the concept. Well, doubtless they could make a fist of it, but they might do even better with more conceptual clarity. On this basis, before I move on to raise some queries about Hausman's account, I will briefly consider what living well consists of. I will interpret this as a question about wellbeing, on the grounds that to live well is to have wellbeing, and I will therefore consider the main accounts of wellbeing that have been offered.

There are three popular theories of wellbeing, which define wellbeing in terms of preference satisfaction, subjective contentment, and an objective list of goods. In what follows, I argue that no single theory gives a complete account of wellbeing, and that we should adopt a pluralist theory according to which all three can give us perspectives on wellbeing.

To take the first category, preference-satisfaction theories, we have already seen a problem, which is that a patient's preferences, even their informed preferences, cannot always be relied upon as a guide to their wellbeing.

Also, we may rationally care about things that do not involve our welfare. For example, Stephen Darwall writes:

> There are many things I rationally take an interest in, such as the survival of the planet and the happiness of my children long after I am dead, that will make no contribution to my welfare. A person may have rational interests that go well beyond what is for her good or in her interest.[109]

Partly in response to such issues, Dolan and Kahneman propose a theory of the second kind, based on subjective contentment: they say that healthcare should be allocated according to how good patients' hedonic experience of the outcomes is.[110] But it does not seem that subjective contentment is the only thing that makes life good. If I have reason to believe that a certain question in philosophy or science is important, it may not be irrational to pursue the answer, just for the sake of learning and revealing the truth, even if I do not

---

[107] Business Insider (2014)

[108] NOP World (2005)

[109] Darwall, (2002) quoted in Adler (2012)

[110] Dolan and Kahneman (2008)

expect the pursuit to make me happy. If I then find the answer, I might be said to have had a good life even if I was not as subjectively content as I could have been. Also, as a basis for healthcare allocation, this assumption could lead to intuitively unacceptable discrimination. For example, consider that some diseases, such as cancer, tend to cause high rates of anxiety and depression. The side effects of treatments may even exacerbate these problems. If healthcare allocation was based on subjective contentment, it would be difficult to avoid the conclusion that these patients should be de-prioritised for funding. Intuitively, this would be unfair.

The third main approach is the objective list approach, based on a list of things that are held to make life good, regardless of individual preferences. This is particularly associated with the capabilities approach. For example, the capability theorist Martha Nussbaum suggests that the objective goods are: life; bodily health; bodily integrity; senses, imagination, and thought; emotions; practical reason; affiliation; other species; play; and control over one's environment.[111] But a common objection to objective list theories is that a single group of people (those who assemble the list) can't know what's good for everyone else in the world.[112] Frances Stewart thinks Nussbaum's list "reflects the values of a typical 21st century American liberal rather than a set of timeless universal values or a contemporary global overlapping consensus."[113] There are also other, related worries. Some items, such as "senses, imagination, and thought", look more like categories than specific goods. They are not specific enough to be any use to policy-makers. But when we try to be more concrete, we will find that one community's good things involving imagination and thought are not the same as another's. Also, policy-makers will need to prioritise the goods, and again, priorities will vary from area to area, whether they involve sports or books or something else. So although objective lists might have value, e.g. because of worries about adaptive or irrational preferences, they cannot be a complete guide to policy. Policy-makers need to understand more about the communities they legislate for, whether in terms of preferences or sources of contentment or something else.

Thus my conclusion is that no single theory of wellbeing is a completely satisfying account. Nevertheless each theory represents aspects of wellbeing.

---

111 Nussbaum (2011)

112 For example, see Okin (2003), Stewart (2001) and Wells (2018)

113 Stewart (2001)

Someone can be said to be well off because they are content, or because they have what they desire. We can say that someone would be better off if they could feel emotions more, or played more, or had more control over their life, even if they don't want it and perhaps even if it wouldn't immediately make them happier (for example, initially it might be uncomfortable to have greater control). So my response to these concerns is to argue that we should adopt a pluralist theory of wellbeing, according to which all these perspectives have the potential to represent aspects of someone's wellbeing.

This raises the question, how do these conflicting considerations trade off against each other in determining the precise level of someone's wellbeing? To answer this question, I refer to David Ross and his influential view as to how plural moral considerations trade off against each other in determining the right thing to do:

> there is no principle by which we can draw the conclusion that it is on the whole right or on the whole wrong. ... we have more or less probable opinions which are not logically justified conclusions from the general principles that are recognized as self-evident".[114]

I contend that something similar is true of wellbeing: to decide whether a life exhibits wellbeing, we must apply our judgment to determine how the different considerations or perspectives on wellbeing balance against each other. There is no algorithm. On this basis, I endorse Hausman's proposal that the best procedure for determining this question is "deliberative groups coupled with public debate".

However, I would add two qualifications. The first is that as a general rule, the view of patients with the conditions in question should carry a lot of weight in the discussions, since, as argued above, we have reason to respect their values and knowledge.

My second qualification is that Hausman's deliberative groups will have another task, this being the delicate task of assessing whether patient preferences can be trusted. Are the preferences influenced by evaluative adaptation; are they irrational; do they have the wrong kind of provenance? This will be a matter of judgment, as much as the substantive task of health state valuations.

---

114 Ross (1930, 2002), p31

3.9 Objections to Hausman

Moving on to challenges to Hausman's account, I will now argue that Hausman over-simplifies the way our overall evaluations of outcomes derive from their non-evaluative features. To summarise, firstly, we cannot adequately characterise all health outcomes with just three descriptive parameters, viz., the two dimensions of his a/f pairs plus life expectancy. And even if Hausman adds further descriptive parameters, he cannot capture the value of all health outcomes by asking his deliberators to evaluate just a few outcomes and then deriving an algorithm which implies the value of other outcomes. So I argue that we'll need to ask Hausman's deliberators to evaluate a wide range of specific health states, not just a few.

In detail, I highlight two main issues with Hausman's account. The first involves Hausman's characterisation of the task of

> mapping health states, as defined in a more complicated health-stated classification system such as the EQ5D ... into the equivalence classes defined by the simpler classification in terms of feelings and activity limitations.[115]

Hausman says that this task will be "largely non-evaluative." But this task will require decisions about how the five dimensions of the EQ-5D trade off against each other, and this will require value judgments. For example, the facts alone will not tell us whether the loss of two legs limits our activities to the same degree as the loss of a leg and an arm (and thus whether these two impairments merit the same a/f score). Each impairment limits different activities, and it is not obvious what non-evaluative facts about these different activities would determine that they should or should not get the same a/f score. We need evaluations. For example, if we accept that health state values derive from their impact on possible lives, we must consider the possible lives that each impairment precludes and then consider the relative importance of those possible lives.

The second objection is that Hausman's two-dimensional descriptive system is far too simple. We won't be able to assign a fixed evaluative score to every location on the matrix (a/f rating), because the a/f system does not use specific enough descriptors to let us do that. Two health states might be assigned feeling ratings of "uncomfortable" and yet they might restrict the possibilities

---

[115] Hausman (2010), p. 288

of living a good life in radically different ways. "Uncomfortable" anxiety might debar good lives A and B and allow good lives C and D while uncomfortable pain allows A and B but debars C and D. If these two feeling states are associated with the same activity rating, say, "not limited", then Hausman will assign the same "factual" a/f rating to them (they will occupy the same place in the matrix). In consequence we will have to assign both states the same evaluative score, on the assumption that both limit our ability to live good lives to the same extent. But it's not hard to imagine that given full information about these two health states, at the evaluation stage we might rank, say, anxiety lower than pain because of the particular lives it debars.

Part of the problem is that Hausman stipulates that a health state's a/f rating must be a "factual" rating. This leaves it open that two states with the same rating may merit different evaluations. One solution would be to stipulate that a/f ratings are not factual; for example, Hausman could stipulate that health states are to be evaluated directly in terms of the way they restrict our ability to live good lives. But it is not practical to evaluate all health states in such terms; at some point we have to extrapolate from specific cases, using general principles about how health state features debar good lives. Also, different deliberators will be involved in the assessment of each condition, including people without the condition, and they will need a description of the condition which does not bias them one way or the other as to its disvalue. It is noteworthy that existing health technology assessment systems all share the same requirement for an initial, non-evaluative characterisation of each condition (such as in terms of the EQ-5D) before the condition is evaluated on a quality of life scale.

As a factual description, an a/f rating is at a disadvantage to an EQ-5D profile because, being based on fewer dimensions, it tells us less about a health impairment, so it tells us less about which good lives it debars and therefore how bad the impairment is. Ideally, to tell how much a health impairment debars good lives, we need to evaluate it individually, based on experience of the disease (viz., a patient evaluation), or a disease-specific characterisation. Failing that, we need a reasonably specific characterisation of the ways it impairs health. The EQ-5D, having a wider descriptive vocabulary (five dimensions), allows a more specific characterisation of the nature of impairments than Hausman's scheme, and is therefore to be preferred in terms of this desideratum.

Additional problems arise once we start combining feeling impairments with activity impairments. For example, compare two health states. The first involves an activity limitation of "limited in some core IADL" combined with pain rated "suffering" on the feeling scale. The second involves the same activity limitation combined with anxiety, also rated "suffering". These two health states might restrict the possibilities of living a good life differently even though they have the same a/f rating (they are at the same location in Hausman's matrix). The reason is that the two feeling states might interact with the mobility impairment differently. Perhaps being mobile helps with extreme pain, by distracting the sufferer from the pain, so a loss of mobility is worse for someone in such pain. In contrast, certain forms of anxiety might mean the sufferer does not want to leave the house, so a loss of mobility has less of an impact.

So it may be when we evaluate two health impairments individually in terms of how they debar good lives (such as anxiety and mobility), those two separate evaluations will not necessarily help us predict how bad the combination of the two health states would be. Our ratings of anxiety and pain in isolation (no other impairments) might suggest that anxiety and pain have the same disvalue (they get the same a/f ratings). This would suggest that when each of them was combined with an activity limitation, the resulting two "combination" health states would be as bad as each other - but in fact, pain plus an activity limitation could be worse than anxiety plus the same activity limitation.

Such "unpredictable combinations" are not uncommon. For example, there are topical creams that cause an unpleasant burning sensation in one place to distract from pain somewhere else (such as capsaicin creams). Evidence suggests these are effective in relieving pain.[116] An unpleasant burning sensation is normally bad, but here is good.

Again, it is generally bad to drink something that tastes so awful it makes you nauseous. But if you believe it's a medicine that could cure you, and the taste enhances the placebo effect, the awfulness itself could be good-making. A fictional passage vividly conveys how this placebo effect can work:

> Ever since [ship physician] Stephen Maturin had grown rich with their first prize [about 1790] he had constantly laid in great quantities of asafetida, castoreum, and other substances, to make his medicines more

---

[116] Ellison et al (1997), Epstein et al (1994)

> revolting in taste, smell and texture than any other in the fleet; and he found it answered - his hardy patients *knew* with their entire beings that they were being physicked.[117]

From a more theoretical perspective, Kappel & Sandoe present the following example to criticise what they call "feature-independence":

> Why should the features relevant for utility always bear the same value? I might be tall, and enjoy some utility from that. But it seems that the utility of the tallness of a person is increased if for example he is a basketball player. In general the utility of any feature would seem to depend on the other features of that person's life, notably the desires of the person.[118]

Other domains of value provide a useful analogy. Jonathan Dancy says: "I might value salt less than I value sugar, but I value meat with salt much more than I value meat with sugar. I might value friendship more than I value honesty, but I still prefer (that is, value more highly) an honest assessment of my prospects of good health to a friendly one."[119] Dancy says this runs counter to what he calls the "Atomic Principle", which is a constraint on rational choice saying that "if we value A+B more highly than we value A+C, we necessarily value B more highly than we value C". Dancy's examples (and the other examples above) instead support holism about reasons, the thesis that "a feature that is a reason in one case may be no reason at all, or an opposite reason, in another".[120]

The EQ-5D can avoid some of these problems with unpredictable combinations because it distinguishes impairments that Hausman's system does not. For example, the EQ-5D assesses pain on a different dimension from anxiety, so it is better placed to pick up any interaction between those two dimensions of health.[121]

An additional problem for Hausman is that many diseases induce entirely sui generis quality of life impairments which are not seen in other diseases. As

---

117 O'Brian (1990) quoted in Moerman (2002). This is a fictional case.

[118] Kappel & Sandøe (1992), p. 299

119 Dancy (2004), p199

120 Dancy (2004), p73

[121] However, I do not claim that this greater comprehensiveness is guaranteed to pick up all such interactions, and supplementary investigations may be needed where they are suspected (for example, based on disease-specific survey instruments). I merely claim that the EQ-5D is less likely to fall into this kind of trap, whereas Hausman's system will find it hard to pick up any such interactions.

an analogy, consider the observation of the narrator of Anna Karenina that "Happy families are all alike; every unhappy family is unhappy in its own way."[122] With many diseases, to fully capture the disvalue of the impairments in question, you need disease-specific measures as well as generic measures, or instead of them. As a case in point, Jull et al compared an instrument called the Charing Cross Venous Ulcer Questionnaire (CXVUQ) with two generic instruments, the SF36 and EQ-5D. They were assessed as instruments for tracking health-related quality of life in venous ulcer patients. Ulcer healing is the outcome of interest to physicians in these cases. It is expected that health-related quality of life will change with ulcer healing. The purpose of this investigation was to evaluate the CXVUQ alongside the SF36 and EQ-5D for their responsiveness to such changes in health-related quality of life. The study found that CXVUQ was more responsive than the generic instruments to changes in ulcer status over a 12 week period.

The problem highlighted here is that there are diseases such that, if we try to describe them using descriptors which are not specific to them, we will not succeed in capturing the impairments they cause. So, we can best capture a variety of impairments by using a wide range of quite specific descriptors. Hausman proposes just two, quite generic descriptors, so he will not be able to describe what's uniquely bad about many impairments. The EQ-5D has a greater descriptive vocabulary (five dimensions instead of Hausman's two), so it is better placed to capture the variety of impairments we might see with different diseases. Having said that, even the EQ-5D will not capture the state of an ulcer as well as the CXVUQ. But it will do so better than Hausman's two dimensions (plus the EQ-5D is a general measure, applicable to a wide range of diseases, so it is better than the CXVUQ for comparing different diseases).

Having said that, I acknowledge that the use of a wide range of dimensions is not guaranteed to avoid these problems completely, given that it will only ever be possible to deploy a finite number of dimensions in a generic measure of health. We must be aware of the possibility that the instrument with which we characterise health states may not capture them adequately, and be ready to supplement the generic instrument with a disease-specific instrument when this is recognised, using the findings from the disease-specific instrument to enrich the findings from the generic measure.

---

[122] Tolstoy (1877)

In summary, Hausman suggests we can simplify to extremes, but in consequence he loses accuracy. For example, it does not seem he would necessarily see any need to elicit preferences regarding every disease, since we could understand what people's important functions are without eliciting that many preferences. Instead he seems confident that we could do a non-evaluative assessment of how each disease limits capabilities, described in terms of his two dimensions, and then map that onto a general matrix of evaluations regarding capability limitations. The problem is that this will fail to match the judgments we would reach if we assessed each disease individually, because of sui generis impairments and because the impact of a combination of impairments characterised in term of Hausman's dimensions is not necessarily deducible from the their impact individually. This pushes us towards eliciting preferences regarding each disease.

However, there is clearly a balance to be struck between getting preferences regarding every single health state of interest, which is not feasible in practical terms, and alternatively being guided by a quite general/generic descriptive system, which could lead to error regarding the value of specific health states. The question of which approach between these two extremes is correct may have different answers in different circumstances. For example, policy-makers may end up at different points on the spectrum according to what budget they have for running preference surveys. But regardless of where on the spectrum a policy-making body falls, given that everyone must rely on general principles to an extent, all should be alert to the risk that an evaluation of a health state derived from general descriptors may fail to reflect its real value. Where there is any reason for doubt, they should investigate further, for example by conducting a small scale qualitative survey, and then if necessary a larger scale quantitative survey to quantify the importance of the considerations that are thereby revealed.

3.10 Conclusion

To conclude, healthcare allocators produce the right kind of benefit when they help patients get into more valuable health states. I endorse Hausman's view that health states are valuable or not partly depending on the extent to which they constrain the possibilities of living well and pursuing valuable objectives. We must understand wellbeing in a pluralist way which involves considerations to do with preference satisfaction, subjective contentment, and goods which anyone rational can see to be good.

The way these different goods trade off in different circumstances will vary in unpredictable ways. So to know how good or bad a health state is, we must apply our judgment to determine how the different types of wellbeing balance against each other.

Usually, the most authoritative view on how good or bad a health state is will be the view of people in that state, viz., patients. This is because patients are always in a better position to know what a health state is like, and usually they are also in a better position to evaluate it appropriately. But sometimes, we should be guided by the public's informed evaluation of health states, or perhaps even the judgments of experts; in particular, this is so where patient preferences are influenced by evaluative adaptation or where patient preferences are irrational or have the wrong kind of provenance.

Ideally the process of understanding how good or bad a health state is must involve an evaluation of each health state individually; contrary to Hausman, the impact of a health state on someone's wellbeing cannot be determined algorithmically from principles derived from a few cases. However, in practical terms, it is unlikely that it will be possible for healthcare policy-makers to evaluate every health state. Thus there is a balance to be struck between getting preferences regarding every single health state of interest and alternatively being guided by general principles.

Given the constrained resources for doing health state valuations in most healthcare systems, it will often be defensible to rely on an algorithmic valuation method of assessing health outcomes, somewhat in the manner of QALYs. However, there would be three key differences between the QALYs proposed here and NICE's QALYs. First, NICE base their QALY calculations on public preferences; I have argued that often, patient preferences would be better. Second, NICE's QALYs are based on health state preferences; I endorse Hausman's proposal of a deliberative assessment of whether health states constrain the possibilities of living well. Third, in applying such an algorithm we need to be alert to the possibility of interaction between impairments in terms of their impact on the overall value of a health state, and ready to apply disease-specific methods of valuing a health state where these are noticed.

4. Should healthcare policy-makers relax their normal criteria in the case of rare diseases, when deciding whether to licence and fund treatments for them?

In the EU, a rare disease is defined as one that affects fewer than 5 people in 10,000.[123] However, because there are so many rare diseases, in combination they affect more people than this might suggest. It is estimated that in the EU, 5-8,000 distinct rare diseases affect 6-8% of the population.[124] In this chapter I offer a new argument to show that policy-makers have reason to make exceptional provisions to help rare disease patients get treatments for their diseases, at least in wealthy societies.

After giving some background, I start by considering an objection of Chris McCabe's, who argues that all else being equal, healthcare policy-makers should produce as much health benefit as possible. He points out that orphan drugs policies incur a cost in terms of the principle of benefit maximisation, but we lack any other principle of equity we can appeal to justify this cost. He therefore suggests they cannot be justified.

I respond to this with two arguments. One is a quite general argument, derived from Taurek and Lawlor, which aims to undermine McCabe's assumption that the only way of defending policies favouring rare disease patients is by placing a higher value on the health gains of rare disease patients. However, this argument does not vindicate privileging orphan drugs specifically.

The second argument is a new, more targeted defence of orphan drugs policies. I argue that if healthcare allocation policy is driven purely by benefit-maximising principles, unmitigated by orphan drugs policies, rare disease

---

[123] European Commission (2013)

[124] European Commission (2013)

patients will stand very little chance of benefiting from future treatments, because of a problem to do with economies of scale. In effect, they will be permanently excluded from access to new treatments. A consideration of egalitarianism or inclusiveness gives us reason to introduce orphan drugs policies. Common disease patients with expensive treatments do not have the same complaint, because they do not have the same reason to think they will be permanently excluded from future treatments.

4.1 Background

Orphan drugs are drugs for orphan diseases, viz., "diseases so rare that sponsors are reluctant to develop them under usual marketing conditions".[125] The diseases are orphans because they would be neglected without special provisions. Orphan drugs tend to be much more expensive than medicines for common diseases. This is because it does not cost less to develop a medicine for a rare disease than for a common disease, but there are fewer patients to recover the development costs from, so the eventual cost per patient is higher (see Figure 1).[126] As a result, in the absence of special orphan drugs policies, it is harder to make money even with higher prices, and so there is less incentive for pharmaceutical manufacturers to develop treatments for these diseases. Currently there are less than 400 approved treatments for rare diseases.[127] In 2010 it was estimated that less than 5% of rare diseases had drugs in development for them.[128]

---

[125] Orphanet, 2016

[126] Of course the costs of R&D will not be the only factor affecting orphan drug prices. Another factor will be profit margins, and in fact McCabe et al (2006) accuses drug companies of making big profits from orphan drugs. However, the evidence of low investment prior to the introduction of orphan drugs policies suggests that the "economies of scale" problem was a real one for orphan drug R&D, and this chapter is premised on the assumption that it was.

[127] Rare Genomics Institute (2015)

[128] Fishman (2012) quoting National Institute of Health (2010)

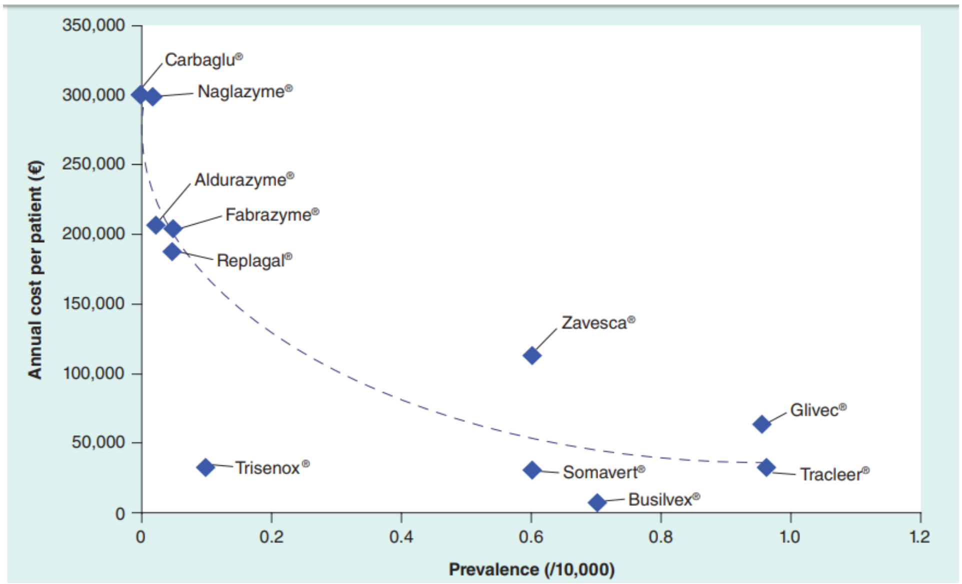

Figure 1. Increasing acquisition cost per patient with decreasing prevalence[129]

Healthcare policy-makers in many jurisdictions have responded to this by implementing orphan drug policies, relaxing their criteria for licensing and funding such medicines. For example, they allow pharmaceutical companies a longer period of exclusivity before competitors are allowed into the market. In the US, new orphan drugs are granted market exclusivity for 7 years whereas other drugs only get 5 years.[130] In healthcare allocation, some countries set aside specific budgets. For example, Scotland has a £80 million New Medicines Fund for rare or end-of-life conditions.[131] Other countries relax their assessment criteria. France accepts lower quality evidence if the total budget impact for the indication of a licensed drug is less than €30 million, accepting the evidence that was submitted to gain marketing authorisation, which may be based on smaller phase II trials and a literature review.[132] Similarly in Germany, additional benefit of orphan drugs with revenues not exceeding €50 million in the past year is automatically considered as proven

---

[129] Taken from Schlander et al (2014)

[130] "Generally, the term of a new patent is 20 years from the date on which the application for the patent was filed in the United States. ... Exclusivity is exclusive marketing rights granted by the FDA upon approval of a drug and can run concurrently with a patent or not" (FDA, 2015). The problem with patents is that drug developers have to file for a patent several years before a new drug is launched, while they are developing and trialling the drug. So there may not be much patent protection left by the time they launch. This is where the FDA's extra market exclusivity can help.

[131] The Scottish Government. (2015)

[132] Tordrup (2014)

through marketing authorization.[133] Germany will also accept evidence from "surrogate endpoints" in these cases (for example, biomarkers that indicate the progress of a disease has slowed, as opposed to, say, direct evidence of increased life expectancy).[134]

Such orphan drug policies have had an effect. In 2007, it was observed that "In the 24 years since the Orphan Drug Act was passed in the United States, 282 drugs and biologic products came to market under the legislation. In contrast, in the 8 to 10 years preceding the act, only ten treatments for rare diseases had been approved by the Food and Drug Administration and brought to market".[135] This represents a tenfold increase in the pace of new launches.

Thus orphan drugs policies have achieved their aims. And intuitively these aims seem reasonable. However, some writers have objected to orphan drugs policies.[136] The objection is basically that in funding expensive orphan drugs, we lose the opportunity to deliver more benefit elsewhere, or to benefit more patients. I will now consider this objection.

4.2 Objections to Orphan Drugs Policies

One of the most vehement and persistent objectors to orphan drugs policies is the health economist Chris McCabe. He says:

> Consider two groups of people who have similar diseases (J and K). ... imagine that the cost of the orphan drug for J is higher than the treatment for K. Suppose the cost of treating one case of J is £1000, the cost of treating one case of K is £100, and the budget is £1000. Then the real choice posed by orphan status is between treating 1 person with J or 10 people with K. To argue that the patient with J should get treatment implies that that health gain of people with J should be valued 10 times higher than that of people with K. The idea that decisions should be made based on valuing health outcome more highly for no other reason than rarity of the condition seems unsustainable and incompatible with other

---

[133] Rémuzat et al (2014)

[134] Tordrup (2014)

[135] Drummond, M. F., Wilson, D. A., Kanavos, P., Ubel, P., & Rovira, J. (2007). Assessing the economic challenges posed by orphan drugs. International journal of technology assessment in health care, 23(01), 36-42.

[136] For example, Juth (2014), McCabe (2005), and McCabe (2006)

> equity principles and theories of justice. Why should a person's health be valued less simply because the condition is not rare?[137]

I read McCabe as starting here with the reasonable assumption that policies must be justified in terms of morally relevant principles. He then assumes that one morally relevant principle is that healthcare policy-makers should produce as much health benefit as possible. It is therefore problematic that orphan drugs policies reduce the total amount of benefit the healthcare system produces. The only way of justifying orphan drugs policies in terms of benefit maximisation would be by taking health gains for rare disease patients to be worth more than health gains for other patients. This is not defensible.

The failure to maximise benefits is therefore an issue with orphan drugs policies. However, McCabe is *not* committed to benefit maximisation as the sole measure of the best policy. He canvasses various other possible objectives and equity principles, such as "equality of health outcomes, equality of resource use, or allocation of resources in proportion to the severity of the individual's ill health".[138] But he then argues that none of them will help justify orphan drugs policies: "rare diseases may not be particularly advantaged or disadvantaged by alternative objectives and principles of equity."[139]

In summary, McCabe argues that by instituting orphan drugs policies, we incur a cost in terms of the principle of benefit maximisation, but we lack any other principle of equity we can appeal to justify this cost. So we invoke an arbitrary, unjustifiable criterion (rarity) and thereby breach principles of equity.

McCabe can be seen as putting a complaint on the part of common disease patients. Consider a group of common disease patients whose treatments are quite expensive. However their treatments would be funded were it not for orphan drugs policies. Common disease patients can complain that as a result of orphan drugs policies, there is no budget left for those treatments.

The question I address in this chapter is whether orphan drugs policies can be vindicated in the face of McCabe's criticism and the complaints that could be expected from unfunded common disease patients whose treatments are as

---

137 McCabe et al (2005), p. 1018

138 McCabe et al (2005), p. 1017

139 McCabe et al (2005), p. 1017

cost-effective as those enjoyed by the rare disease patients benefiting from orphan drugs policies.

4.3 An Initial Response to McCabe: Egalitarianism, Pluralism

I will offer two responses to McCabe. The first is a quite general argument, derived from Taurek and Lawlor, which aims to undermine McCabe's assumption that the only way of defending policies favouring rare disease patients is by placing a higher value on the health gains of rare disease patients. However, this does not vindicate privileging orphan drugs specifically. The second argument is a new, more targeted defence of orphan drugs policies.

The first response addresses a key assumption in McCabe's argument. The assumption is that healthcare policy-makers should maximise the aggregated benefits they produce, unless specific equity principles imply otherwise. So for example, if the benefits per patient are equal, and there are no other relevant considerations favouring some patients over others, then we save the greater number. McCabe infers that since there are no specific considerations favouring rare disease patients, and since orphan drugs policies fail to maximise aggregated benefits, they cannot be justified. One way of defending orphan drugs policies against McCabe's argument would be to deny that it makes sense to speak of maximising benefits, at least in the normal case where the prospective beneficiaries are different people between whom policy-makers must choose. A philosopher who would take this line is John Taurek. I will now consider his arguments. My conclusion will be that although they fail, Taurek points to an important consideration which may be of some assistance to rare disease patients.

Taurek argues that in situations where we must decide who to save, we should not consider "the relative numbers of people involved as something in itself of significance in determining our course of action".[140] Instead, Taurek says he would choose who to save in a way that expressed "equal concern and respect for each person."[141] In some circumstances, this might mean flipping a coin, for example. However, Taurek realises this is counterintuitive. How does

---

[140] Taurek (1977), p. 293

[141] Taurek (1977), p. 303

he explain this stance? He presents a number of arguments to challenge our intuition that we should try to save as many people as possible. One of his arguments is that to count the numbers would be to treat people just like valuable objects:

> If six objects are threatened by fire and I am in a position to retrieve the five in this room or the one in that room, but unable to get out all six, I would decide what to do in just the way I am told I should when it is human beings who are threatened. Each object will have a certain value in my eyes. If it happens that all six are of equal value, I will naturally preserve the many rather than the one. Why? Because the five objects are together five times more valuable in my eyes than the one.[142]

The key point here is that each object has value *to him* such that five will have five times more value to him than one. However, if he is asked to decide which of two groups of strangers to rescue, his decision is not determined by their value as objects:

> I empathize with them. My concern for what happens to them is grounded chiefly in the realization that each of them is, as I would be in his place, terribly concerned about what happens to him. It is not my way to think of them as each having a certain objective value, determined however it is we determine the objective value of things.[143]

Here Taurek observes that, of course, each prospective rescuee values their life. Taurek now argues that this is the only value of each life to each rescuee is the value that each person's life has to that individual.

> Five individuals each losing his life does not add up to anyone's experiencing a loss five times greater than the loss suffered by any one of the five ... Each person's potential loss has the same significance to me, only as a loss to that person alone. Because, by hypothesis, I have an equal concern for each person involved, I am moved to give each of them an equal chance to be spared his loss ...My way of thinking about these trade-off situations consists, essentially, in seriously considering what will be lost or suffered by this one person if I do not prevent it, and in comparing the significance of that for him with what would be lost or suffered by anyone else if I do not prevent it. This reflects a refusal to take seriously in these situations any notion of the sum of two persons'

---

142 Taurek (1977), p. 306

143 Taurek (1977), p. 306

> separate losses.[144]

So, each rescuee's life has value to the person living it, but not to the prospective rescuer. As Taurek puts it, "Each person's potential loss has the same significance to me, only as a loss to that person alone." *More generally, the value which each life has to the person living it will not be apparent to anyone but the person living it*. So, what demands a response from the rescuer is not the value of each life (since that value is only apparent to the rescuee). Instead, the rescuer merely recognises that each life has value *to the person living it,* and this is what gives the prospective rescuer reason to save each person (in Taurek's terms, the rescuer feels concern). And since the value of each life is not apparent to the rescuer, the rescuer can't add up the value of the lives in the way an owner of valuable objects can add up their value. The problem this poses for the benefit maximiser is that in order for our talk of "maximising benefits" to make sense in this case, we would have to be able to add up the value of different people's lives. Since we can't do that, talk of maximising benefits in this case makes no sense. And since maximising benefits was our purported reason for saving the greater number, we have no basis for saving the greater number.

This leads Taurek to his counterintuitive conclusion: we should not save the greater number merely because they are the greater number.

We see a similar argumentive strategy in his discussion of pain. He says:

> I would like to combat the apparent tendency of some people to react to the thought of each of fifty individuals suffering a pain of some given intensity in the same way as they might to the thought of some individual suffering a pain many or fifty times more intense. I cannot but think that some such tendency is at work in the minds of those who attribute significance to the numbers in these trade-off situations.[145]

He goes on to argue:

> The discomfort of each of a large number of individuals experiencing a minor headache does not add up to anyone's experiencing a migraine."[146] He imagines members of a group asking an individual to suffer an episode of extreme pain so that they can each avoid an episode of moderate pain. Taurek has them pleading: "Think of the awful sum of

---

144 Taurek (1977), p. 306

145 Taurek (1977), pp.18-19

146 Taurek (1977), p. 308

> pain that is in the balance here![147]

The problem with this plea is that pain must be suffered by someone. No-one suffers 50 people's pain. So there is no such sum of pain.

This leads him to his conclusion regarding the betterness or worseness of the outcome:

> I cannot understand how I am supposed to add up their separate pains and attach significance to that alleged sum in a way that would be inappropriate were any of those involved to do it.[148]

Taurek makes a similar point regarding how outcomes must be valued in life-saving cases:

> The claim that one ought to save the many instead of the few was made to rest on the claim that, other things being equal, it is a worse thing that these five persons should die than that this one should. It is this evaluative judgement that I cannot accept. I do not wish to say in this situation that it is a worse thing were these five persons to die and David to live than it is or would be were David to die and these five to continue living. I do not wish to say this unless I am prepared to qualify it by explaining to whom or for whom or relative to what purpose it is or would be a worse thing.[149]

As in the pain case, the problem with the claim that "it is a worse thing were these five persons to die and David to live", says Taurek, is that it only makes sense once it is indexed to a particular individual or purpose.[150] If we are to capture everything good or bad about an outcome in a single evaluation, we would need to aggregate the goodness and badness of different people's harms and benefits. However, since, on Taurek's view, the value of different people's respective harms and benefits can only be determined from their individual perspectives, and there is no overarching evaluative framework subsuming these individual frameworks, such aggregation will be impossible.

Taurek concludes as follows:

> His loss means something to me only, or chiefly, because of what it means to him. It is the loss to the individual that matters to me, not the

---

[147] Taurek (1977), p. 309

[148] Taurek (1977), p. 309

[149] Taurek (1977), pp. 303-4

[150] As seen in the last sentence of the quoted passage.

> loss of the individual. However, should any one of these five lose his life, his loss is no greater a loss to him because, as it happens, four others (or forty-nine others) lose theirs as well. ... Five individuals each losing his life does not add up to anyone's experiencing a loss five times greater than the loss suffered by any one of the five.[151]

Perhaps rare disease patients can exploit this to offer a line of defence against McCabe. McCabe argues that, absent other considerations, we must maximise benefits, and points out that orphan drugs policies fail to do this. However, in line with Taurek, rare disease patients can respond that the idea of an objective quantity of benefit gained by different groups of patients, conceived independently of the perspective of any one patient, makes no sense. So McCabe cannot argue that healthcare policy-makers should maximise the aggregated benefits they produce. Thus the health benefits produced by devoting a given budget to a few rare disease patients are neither greater in aggregate nor smaller in aggregate than the health benefits we would produce by devoting the same budget to a larger number of common disease patients. So, if there are no other considerations favouring either group of patients, patients in both groups must be understood as having an equal claim to treatment, and our policy must reflect this. Perhaps one way would be Taurek's coin toss.

However, Parfit raises an objection to Taurek, and succeeds in showing that in fact we *can* make sense of 50 people's pain being a greater sum of pain than one person's pain:

> Consider first pains that are felt by one person. I might decide that fifty minor headaches would be worse than a single migraine. If I had to endure the fifty headaches, I would suffer more. In other words, my "sum of suffering" would be greater. Such comparisons are, even in principle, rough. There is only partial comparability. But that does not make the comparisons senseless. And this use of the phrase "sum of suffering" would, I believe, be understood by Taurek. At any rate, he says nothing against it.
>
> Suppose, next, that each of fifty headaches would be had by a different person. If these headaches were about as bad, they would again together involve about as much suffering. The "sum of suffering" would be about as great. This is not a different use of this phrase. It is the same use. Since

---

[151] Taurek (1977), pp. 307

> he understands this use when applied within one life, Taurek thereby understands it when applied to different lives.[152]

Parfitt's argument shows that contrary to Taurek, we can make sense of aggregated harms and benefits.[153]

However this is not all there is to say about Taurek's arguments. I take Taurek's paper to have two lines of argument; the negative line we've seen where he aims to show that the idea of comparing the aggregated harms and benefits of different groups makes no sense, and a positive line where he aims to show the value of offering prospective beneficiaries of our help an equal chance. He spends less time on this so it is not often noticed, but for example, he says that tossing a coin

> would seem to best express my equal concern and respect for each person.[154]

Lawlor argues that although Taurek has not characterised a consideration here that will be decisive in every case, nevertheless Taurek's principle is plausible if it is reframed as a pro tanto consideration:

---

[152] Parfit (1978)

[153] Parfit does also go on to consider whether Taurek can be understood as offering a second line of argument, granting that different people's pains can be summed but denying that a greater sum of pain is morally worse than a lesser sum: "If one person is harmed, that would be just as bad as if any number are each equally harmed" (Parfit, 1978, p. 295). I will not consider this line of enquiry of Parfitt's, since, as argued above, I take Taurek's argument to be entirely dependent on denying that different people's pains can be summed; I take his further arguments about moral value of consequences to be dependent on that initial denial that different people's harms can be summed. In addition, I suggest it would be difficult for Taurek to pursue Parfitt's second line of argument, given his other commitments. To see this, note that Taurek does not deny that he has a reason to prevent harms to others if he can; he never considers the option of rescuing no-one. For example, he says of his original case that "we were to imagine that I must choose between sparing David the loss 'of his life and sparing five others the loss of their lives" - here, the word "must" suggests that Taurek does not see saving no-one as a permissible option. Perhaps more to the point, the interest of his argument for the rest of us is the question of "who to help", not "whether to help". So, what is Taurek's reason for helping? It must be that the outcome of helping is better for the beneficiaries than the outcome of not helping. This reasoning makes sense on Taurek's account because he can say "to whom or for whom or relative to what purpose [the outcome] is or would be a worse thing" (my interpolation). So, we can set aside any scepticism as to whether we should help at all. Taurek's accepts he has some reasons to help someone; he merely denies that we have specific reason to help the greater number. But Taurek's acceptance of such reasons makes it difficult for him to pursue the kind of argument Parfit has him pursuing here, the second line of argument based on accepting that different people's pains can be summed but denying that he has reason to prevent the greater harm. If we can make sense of aggregated harms, it would seem arbitrary to say we have reason to produce the better outcome in respect of harms that will be suffered by a single individual, but not in respect of such aggregated harms.

[154] Taurek (1977), p. 303

> Typically, people have argued as if there was a choice to be made: either numbers matter, in which case we should save the greater number, or numbers don't matter, but rather there is moral value in giving each person an equal chance of survival, and therefore we should toss a coin. My claim is that we do not have to make a choice in this way. Rather, numbers do matter, but it does not follow that we should always save the greater number. And likewise, there is moral value in giving each person an equal chance of survival, but it does not follow that we should always toss a coin.[155]

Lawlor's argument for this starts with intuitions:

> When considering the case of one versus two, or even one versus four or five, people will often think this presents a real dilemma. As soon as we consider bigger disparities, one versus a million for example, all sense of there being a dilemma falls away.[156]

To make this point more strongly, Lawlor goes on to offer a case in which we face a choice between saving 1,000,000 or 1,000,001.[157] Here, it seems intuitively compelling that a coin toss is in order. However, all else being equal, number-counters such as McCabe seem compelled to say we must save the 1,000,001. But it seems intuitively clear that both parties should be given an equal chance. Now, moving to the case at hand, suppose the 1,000,000 patients consist of (say) ten groups of rare disease patients and the 1,000,001 patients consist of one group of common disease patients. On this basis, contrary to McCabe, the rare disease patients can make a case for being given equal chances without claiming that their prospective health gain is worth more than that of the common disease patients. Our intuitions regarding such cases are best explained by an egalitarian principle according to which everyone is entitled to some of the beneficial resources that society makes available, even if this sometimes means fewer people benefitting from those resources.

So Lawlor offers a pluralist[158] position according to which "there is some value to saving the greatest number" but also "there is some value to giving each person an equal chance of survival".[159] In some cases these two

---

[155] Lawlor (2006), p. 160

156 Lawlor (2006), p. 160

157 Lawlor (2006), p. 160

158 Perhaps the canonical expression of moral pluralism is Chapter 2 of Ross (1930, 2002)

159 Lawlor (2006), p. 161

considerations may conflict, creating a dilemma. Nevertheless, in many cases one consideration may clearly outweigh the other, as in the above case of 1,000,000 vs 1,000,001, where Lawlor suggests that

> the moral value of giving each person an equal chance of survival *can* outweigh the moral good of saving the biggest group in this particular case.[160]

Lawlor contrasts the above case with one where with a case in which you can save one, or you can save two, and argues as follows regarding such cases:

> Someone might argue that in both cases it is only one extra life that is at stake in either case, so why should we think that we should save the greater number in one case, but toss a coin in the other. The answer is simple. In the case where we can save one, or we can save two, the extra life we can save is weighed against the moral value of giving one person – the lone individual – some chance of survival. In the case where we can save 1,000,000 or we can save 1,000,001, however, the one extra life we can save is weighed against the moral value of giving a million people a chance of survival. This explains why many will have the intuition that we should toss a coin in one case, but not the other.[161]

So perhaps orphan drugs policies can now be defended on the basis of this modified version of Taurek's arguments. Although, contra Taurek, we can make sense of the idea of comparing the benefits accruing to different groups of patients, and although this is a relevant pro tanto consideration, there are also other relevant pro tanto considerations. Taurek has highlighted one: we should show "equal concern and respect for each person"[162] by giving rare disease patients *some* share of healthcare resources, even if though this means fewer people benefitting from those resources. Orphan drugs policies are one way of respecting this consideration, whether it be by hypothecating budgets or by relaxing the cost-effectiveness threshold.

However, the conflicting values we have been considering may justify seemingly contrasting policies in different circumstances. Orphan drugs policies needn't favour rare disease patients in all cases. This provides a defence against McCabe claim, considered above, that we should treat the 10 person with K rather than the one person with J, and his argument that this

---

160 Lawlor (2006), p. 160

161 Lawlor (2006), p. 160-161

162 Taurek (1977), p. 303

tells against orphan drugs policies. From a pluralist point of view, there is no inconsistency in saying that value A outweighs B in one situation, while value B outweighs A in another, where A is exhibited to a lesser degree; McCabe might be right that we should be swayed by benefit-maximising considerations in his case, and yet consistently with this the pluralist can say that in another case (such as 1,000,000 vs 1,000,001), we should be swayed by egalitarian considerations rather than benefit-maximising considerations, since here the egalitarian policy does not cost as much in terms of benefit. Thus a moderate orphan drugs policy can deliver the verdict that McCabe expects, which is to save the ten people with K instead of the 1 with J, whilst also prescribing that in the right circumstances (for example, if the choice is between 9 with J and 10 with K, we should save the smaller group.

4.4 Return of McCabe: General egalitarian considerations cannot justify discrimination between rare disease patients and common disease patients.

However, although I endorse Lawlor's pluralism of benefit maximisation and Taurek-style egalitarianism, this combination of pro tanto principles does not entirely vindicate orphan drugs policies. Although the Taurek-derived principle reduces the relative importance of cost-effectiveness comparisons in our policy-making, it does not help us differentiate between different groups whose cost-effectiveness is the same. This is what is needed if we are to justify orphan drugs policies, since, all else being equal, orphan drugs policies will discriminate between drugs whose poor cost-effectiveness results from rarity and drugs with the same cost-effectiveness resulting from other factors. Specifically, McCabe can point out that although the Taurek-derived principle might give us reason to give some chance of funding to rare disease patients, with the result that orphan drugs might get funded instead of more cost-effective treatments for common disease patients, it does not give us reason to favour rare disease patients over common disease patients whose cost-effectiveness is the same. To see this, consider two groups of 1,000 patients. Group A consists of 10 subgroups of 100 rare disease patients whose cost-effectiveness is (say) double the threshold. Group B consists of a single group of 1000 common disease patients whose cost-effectiveness is also double the threshold (suppose this is because their manufacturing costs are high). According to the Taurek-derived principle espoused above, both these groups are equally entitled to a chance of treatment. However, under orphan drugs policies, the rare disease patients could get funded while the common disease patients do not. Thus the arguments above do not help to justify all the decisions which could result from orphan drugs policies. As we have seen, on the basis of Taurek-style arguments, rare disease patients can argue that they have as strong a claim to healthcare resources as common disease patients; but they cannot use Taurek-style arguments to defend "privileging" rare disease patients over common disease patients, which is what orphan drugs policies do.

McCabe rebuts certain defences of orphan drugs policies using just this sort of argument. Specifically, he addresses a defence of orphan drugs policies based on the argument that many rare disease patients lack alternative treatments:

> Hughes et al recount another frequently cited argument for special treatment – 'ensuring access to treatment where no other treatment exists.' Like 'gravity' this is not a defining characteristic of an orphan drug, but it is a frequently cited argument for their special status in licensing and reimbursement. Not being unique to orphan drugs, it cannot be a justification for their special status.[163]

McCabe here assumes that a good defence of orphan drugs policies must be based on a feature which is "unique to orphan drugs,"[164] pointing out that the feature picked out by the defence in question (a lack of alternative treatments) is not specific to the beneficiaries of orphan drugs policies but is also exhibited by some patients who lose out by orphan drugs policies.[165] I believe that in appealing to this uniqueness criterion, McCabe has in mind cases such as the above, where orphan drugs policies privilege rare diseases over common diseases. The problem is that without such uniqueness, orphan drugs policies look arbitrary. For example, the attempted justification of orphan drugs policies in terms of a lack of alternative treatments makes orphan drugs policies look arbitrary in privileging rare disease patients lacking alternative

---

[163] McCabe et al (2006), p. 341

[164] McCabe et al (2006), p. 341

[165] I don't believe McCabe intends to insist that the feature picked out by a defence of orphan drugs policies be entirely unique to orphan drugs; in fact it might be helpful if the feature is found in cases outside healthcare where similar special provisions are justified, such as the GP clinics and fire stations considered earlier. Such non-healthcare analogies help bolster the case that a real consideration is at stake. But McCabe does mean to insist that the feature is unique in the context of healthcare allocation, for the reasons given here.

treatments over other patients lacking alternative treatments. Given that, for example, my groups A and B both share the feature in question, the attempted justification does not justify this privileging of one group over the other.[166] A similar problem will recur in the case of all arguments expressed in terms of features not specific to rare disease patients (by which I mean features shared by other patients who lose out by orphan drugs policies). In the absence of special considerations,[167] such arguments will not justify the way orphan drugs policies privilege rare disease patients with the feature over other patients with the feature.[168]

In summary, McCabe relies on a theoretical adequacy condition to the effect that a good justification of orphan drugs policies must be in terms of a feature which is specific to rare disease patients and not shared by other patients, since orphan drugs policies will prioritise rare disease patients over those other patients. Many have tried to assemble such a justification, but I am not aware of any that succeed. Other writers such as McCabe and Gosain have surveyed the literature and concluded that no extant justification succeeds.[169]

---

[166] It also creates a puzzle about policy-makers' motivation - if a patient's lack of alternative treatments is a relevant consideration, why not articulate the policy in terms of that feature, rather than such a poor proxy as rarity?

[167] Admittedly, it might be possible to argue that rarity is a reasonable proxy for the feature in question, and that pragmatic considerations justify the injustice of any unfair distinctions made by orphan drugs policies. But in the absence of such pragmatic arguments, McCabe's uniqueness criterion seems reasonable. And generally, I suggest such pragmatic arguments will be difficult to sustain. For example, consider the justification for orphan drugs policies in terms of rare disease patients' lack of alternative treatments. Actually, it is quite easy to determine which conditions have treatments licensed for them, so it would not be hard to formulate a policy which picks out diseases lacking alternative treatments. In fact, some jurisdictions have done so. For example, Cookson notes that the Australian Pharmaceutical Benefits Advisory Committee specifies three requirements for the application of the rule of rescue to national pharmaceutical coverage decisions, the first of which is that "No alternative exists in Australia to treat patients with the medical condition meeting the criteria of the requested restriction." (Cookson 2008). So the fact that we have orphan drugs policies rather than policies based on the lack of alternative treatments can't be explained by the difficulty of formulating a policy based on lack of alternative treatments.

[168] This is true even if the other patients with the purportedly morally relevant feature themselves merit "special" treatment. For example, suppose the other patients with the feature are disabled patients who stand to gain less benefit from life-saving treatment. As argued in Chapter 5 on disability discrimination, it may be unfair that benefit maximising considerations tell against such patients, so perhaps we should not apply benefit maximising considerations in the same way in their case. But this would not mean we could offer a good justification of orphan drugs policies in terms of a feature exhibited equally among rare disease patients and disabled patients. The problem would be that orphan drugs policies discriminate between rare disease patients with the feature and disabled patients with the feature, and an argument in terms of this feature would not justify this discrimination.

[169] McCabe (2005) and Gosain (2015)

However, I do not have space to go through all the attempted defences here. Instead I shall simply present what I consider to be the correct defence.

4.5 A New Defence of Orphan Drugs Policies: Rare disease patients should not be permanently excluded from a key benefit of society.

I now aim to show that we should reject McCabe's conclusion that orphan drugs policies are "incompatible with ... equity principles". Specifically I aim to show that we can have good reasons to institute orphan drugs policies, such that contrary to McCabe, they are neither arbitrary nor unjustifiable. We have seen that the mere fact that orphan drugs are expensive does not give us any reason to favour them over equally expensive common disease drugs. But I will argue that in another respect, rare disease patients are in a different position to common disease patients, in that without special measures, they stand very little chance of benefiting from future treatments. A consideration of egalitarianism or inclusiveness gives us reason to adjust our policy-making criteria to reflect this. I contend that my argument satisfies McCabe's demand for a feature that is specific to rare disease patients. My starting point for this argument is to have rare disease patients pointing out that if healthcare allocation policy is driven purely by benefit-maximising principles, unmitigated by orphan drugs policies, then they do not get an equal chance of benefiting from future treatments. In fact, they will be permanently denied future treatments. In contrast, common disease patients do not have the same complaint, because they do not have the same reason to think they will be permanently excluded. I will put the argument first by means of an analogy, and then elaborate it by identifying grounds for a reasonable complaint on the part of rare disease patients in the absence of orphan drugs policies.

The analogy is with decisions about where to locate GP clinics and fire stations and schools. In such decisions, we do not only aim to produce the greatest benefit (or benefit the greatest number). If we did, all the facilities would be built in locations accessible to large populations, viz., towns and cities. However, a proportion of facilities are built in rural areas, at some cost in terms of benefit maximisation. We are similarly concerned to achieve close to universal coverage for postal services, phone networks, power networks, and so on.

Again, to forestall a complaint of McCabe's, rural residents do not need to claim that a higher value should be placed on the benefits they will derive from such services than on the potential benefits of such services to others. Rural residents merely need to argue that policy-makers should try and distribute benefits more equally to all members of society, even if this means delivering slightly less benefit in total. Perhaps the underlying concern is to do with reciprocity; rural residents have accepted all the liabilities of membership of society, such as the liability to pay tax when they earn enough. They can expect a share of the key benefits of society in return.

However, notice that we do not demand universal coverage at any price. Some remote areas must do without fire stations or electricity (for example, if it is very expensive to install the facilities or if the population is very small). So we compromise between benefit maximisation and broad accessibility. Such cases are best explained by a pluralist theory, according to which a principle of benefit maximisation and an egalitarian principle are in tension with one other. As in the case of the egalitarian considerations identified by Taurek, the trade-off may go different ways in different cases.

Rare disease patients can argue that similar conflicting principles apply in their case. They can start by arguing that benefit maximisation should not be our only aim in healthcare allocation. We should also aim at including different patients groups. An egalitarian principle or principle of inclusiveness prescribes that the important benefits of society should be shared out amongst all its members. If we do nothing to respect such a principle of inclusiveness, patients with rare diseases have a reasonable complaint. The reasonable complaint is that whereas anyone with a currently untreatable common disease can hold out some hope that a treatment will be developed at some point for them or their offspring, and then funded, there is very little hope for people with currently untreatable rare diseases in a system with no orphan drug policies. Because of their small patient numbers, even if a treatment is developed it will never be able to compete in terms of cost-effectiveness with treatments for common diseases. Furthermore, in a system without any orphan drug policies, that hurdle is a permanent, nearly insurmountable hurdle. As I will argue later, among all the patient groups competing for funding, rare disease patients are the only well-defined[170] group that are

---

[170] By "well-defined" I mean a group that can be simply and unambiguously identified in legislation. I will grant later that there may be small groups of other patients who also face similar hurdles, but it would be difficult to pick most of these other groups out without encouraging gaming on the part of pharma companies.

known to face this hurdle as a group. Few if any patient groups with common diseases can offer similar evidence that they face any such permanent hurdle. Even if they do not have a treatment now, for all we know they stand a chance of getting one in future. In virtue of this difference between rare disease patients and other patients, my defence of orphan drugs policies satisfies McCabe's uniqueness criterion. He missed this feature because he focuses on one-off cases, rather than looking at the long-term impact of the policies he defends.

This permanent exclusion of a social group from one of the critical benefits of society is a pro tanto concern.

4.6 Objection: Permanent Exclusion Is Not A Relevant Consideration

But common disease patients could question whether an expectation of *permanent* exclusion carries all the weight it is given here. Can't anyone who is denied funding complain that the policy-making criteria are somewhat unfair in respect of that denial - not just patient groups who get no chance of funding under the policy-making criteria? So doesn't the fire station analogy help them equally?

In answer, I start by observing that if we have limited budgets, someone has to lose out. If everyone who lost out had a complaint to the effect that the policy-making criteria were unfair to them, simply because they were denied funding, then all possible healthcare allocation criteria would be unfair. I will assume that contrary to this, it is possible to articulate fair healthcare allocation criteria. So the question is, who must lose out and why? To answer this, I offer a new fire station analogy, which helps demonstrate that patients don't have grounds for a complaint just because a decision goes against them; they must show that the grounds of the decision are such that such decisions would never go their way. I previously asserted that the position of rare disease patients in the absence of orphan drugs policies is analogous to that of rural dwellers who would always lose out on decisions about fire station locations under a benefit-maximising regime. Now, in contrast to the position of rural dwellers, consider the case of city dwellers who happen to lose out on a fire station decision because a site in urban area A happened to be cheaper than another site in urban area B where they live, or because there was simply no site available in area B. Urban traffic congestion being what it is, we can suppose that the journey time from the fire station to urban area B is much greater than to urban area A, putting people in urban area B at greater risk of death by fire than people in urban area A. Suppose that, as far as policy-makers could tell in advance of making this specific decision, the policy-making criteria didn’t give people in urban area B a significantly lower chance of getting a fire station. In that case, people in urban area B don't have the same kind of case for special consideration that rural dwellers had, since, at the point where the policy-making criteria are formulated, but before they are applied, it is not foreseeable that the policy-making criteria will give them no chance of getting a local facility. It is the way the policy-making criteria guarantee from the start that people in rural areas will be excluded that is intuitively unfair. In contrast, the way people in urban area B happened to get excluded on a particular occasion is not unfair, because it could not have been predicted by looking at the policy-making criteria. To at least that extent, all of us must take our chances as to whether we will gain or lose from public policy-making principles. I further suggest that the position of common disease patients is analogous to the position of such city-dwellers rather than to that of the rural dwellers. Consider a group of common disease patients which loses out across a sequence of allocation decisions. Assume it would not have been *foreseeable* that they were going to always lose out, from an examination of the benefit-maximising criteria before they were implemented. In that case the policy-making criteria did not relevantly "guarantee" that those patients would lose out. To that extent the policy-making criteria could not be said to be unfair, at least not in the way that the exclusion of rare disease patients under a benefit-maximising policy could be said to be unfair. The common disease patients are the victims of acceptably tough luck rather than of an unfair policy.

The upshot of this is that, contrary to McCabe's argument, we have identified a principle of equity which supports instituting orphan drugs policies, so they are not "incompatible with other equity principles". Furthermore, again contrary to McCabe, this consideration does not rely on valuing a person's health less because their condition is not rare; any given change in health status is valued the same in rare disease patients and common disease patients, but considerations other than the value of the health change give us reason to prioritise health benefits in rare disease patients. Finally, the consideration we have identified is specific to rare disease patients, and so we have satisfied McCabe's uniqueness criterion.

Having established that permanent exclusion is a relevant consideration in healthcare allocation as well as other areas such as fire station allocation, it is a suitable moment to question whether McCabe's simple paired comparisons make good analogies with the cases at issue. For example, consider McCabe's case where we have to choose between 1 person with disease J or 10 people with disease K. I would argue that such "one-off" decisions are not a fair test of the fairness of orphan drugs policies. In such a case, one naturally thinks of a choice that is unconnected with all the other choices we would ordinarily make in a healthcare system. The consideration I have highlighted (a principle of inclusiveness as applied to policy-making principles) is not obvious in such contexts. However, in the real world of policy-making, we adopt sets of principles, which determine the outcome of a great number of decisions over an extended period. The considerations I have highlighted are more obvious here. The case for orphan drugs derives from considering the impact of these principles and these extended sets of decisions; in particular, the fact that a pure benefit-maximising framework would preclude virtually all R&D on rare diseases, as we saw before orphan drugs policies were introduced. Yes, it is good to save more lives and generate more benefits of other kinds; but there are other considerations too, and when we consider the effect of a pure benefit-maximising framework across multiple decisions, it becomes apparent that the rarity of a patient's condition may be one such consideration. McCabe's simple cases of decisions made in isolation are not a good analogy with decisions made in this wider policy-making context. Having said that, the principle of inclusiveness is applicable even McCabe's in one-off cases; favouring the larger group in the one-off case is unfair if our reasons are such as to guarantee that rare disease patients will lose out. However, one-off cases do not trigger our intuitions in the same way as sequences of cases based on the same criteria, because we don't see any long-term exclusion in the one-off case.

4.7 Objection: Some Common Disease Patients Are Permanently Excluded

However, some common disease patients may seem to have another objection, which accepts that permanent exclusion is a relevant consideration. Consider common disease patients whose cost-effectiveness is higher than the threshold, and who therefore do not get funded because no special provisions are made for them, while rare disease patients with the same cost-effectiveness get funded as a result of orphan drugs policies. Some of these common disease patients might appear to face difficulties as great as those faced by rare disease patients. For example, some cancers have a cunning ability to evolve in response to treatment, a bit like the way some infectious bacteria can evolve to develop resistance to antibiotics. Perhaps the scientific difficulties on these diseases are such that it will never be possible to develop a treatment whose cost-effectiveness falls below the threshold. If orphan drugs policies are justifiable based on the permanent exclusion that rare disease patients can expect without orphan drugs policies, then these patients are equally entitled to special provisions. In line with McCabe's argument, privileging rare diseases over these permanently expensive common diseases looks arbitrary.

My response is to start by conceding that if we somehow knew that the scientific difficulty of developing treatments for a certain group of common disease patients were extreme and permanent, such that they faced the same insurmountable cost-effectiveness hurdle as rare disease patients, then we might have reason to institute similar special provisions for them. The trouble is, we have don't know this. We do not have reliable estimates of the costs we can expect to incur in doing R&D on different diseases. Generally you only get to know the R&D costs by doing the R&D. The whole point of scientific research is that it is a journey into the unknown, full of surprises. So the above case of cancers that evolve in response to treatment is somewhat unusual (and even in that case, it is not wildly overoptimistic to expect that someone will eventually find a way of dealing with the problem). Broadly, we do not have strong grounds for picking out patients who are disadvantaged by a benefit-maximising policy on grounds of the scientific difficulty of their disease. This is not to say that pharma companies have no idea about how easy or difficult the R&D would be on the diseases they have under review. Presumably they have a view on this question, and maybe even a view about how much it would cost to achieve R&D success on each possible project. However, I take it they are uncertain about these estimates; or at least, they should be. The vast majority of R&D projects end in failure; this would not be so if our knowledge of R&D project chances was certain. In consequence, pharma companies will be uncertain about which R&D projects would be more expensive than others.

In contrast, we can be sure that most rare disease patients are disadvantaged by the low number of other patients with each condition, such that average R&D costs per rare disease patient will be higher than they would have been had we addressed the same scientific difficulties on behalf of larger populations of patients.

In addition, a couple of practical difficulties are worth noting. First, it would be difficult to avoid incentives targeted on scientific difficulties going to unintended recipients. It is relatively straightforward to formulate orphan drugs policies that only benefit the intended diseases. However, scientifically difficult diseases are difficult to characterise in general terms suited to policy-making, so it would be much more difficult to formulate special provisions for scientifically difficult diseases that avoid gaming on the part of pharma companies such that benefits go to diseases that are actually not that difficult, relatively speaking. Second, our state of knowledge regarding the respective difficulties influences the efficacy of special provisions - pharma companies will respond to orphan drugs policies more than to incentives based on scientific difficulties because they are more confident about the nature and scale of the difficulties addressed by orphan drugs policies.

Bearing these points in mind, consider the following two generalisations:

> Generalisation 1. Without orphan drugs policies, rare disease patients will not see successful R&D projects on their diseases, but with orphan drugs policies they will see many more.
>
> Generalisation 2. Without similar special provisions, common disease patients with scientifically difficult diseases will not see successful R&D projects on their diseases, but with special provisions they will see many more.

Policy-makers must rely on generalisations such as these in accepting or rejecting policies.[171] I also suggest we can be much more confident about the first generalisation than the second. This means we can justify orphan drugs policies in terms which are not available in the case of scientifically difficult diseases.

---

[171] Even if the generalisations in question are only "for the most part" generalisations. Very few laws avoid unintended consequences. Realistically, the measure of a good law must be that unintended consequences are minimal; it can't be that they are entirely absent.

4.8 Conclusion

We saw earlier that McCabe argued that by discriminating against common disease patients simply on the grounds of the commonness of their disease, we invoke an arbitrary criterion and thereby breach principles of equity. Contrary to this, I have shown that rarity can be a morally relevant consideration. I have argued that orphan drugs policies can be justified by appeal to a pro tanto reason to mitigate permanent exclusion of a group from the benefits distributed by the state.

In line with this I have offered two responses to the argument from McCabe I quoted at the beginning. First, in the case he offers where we must choose between 1 person with disease J or 10 people with disease K, I can deliver the verdict he expects, which is to fund the 10 people with K. But in the context of a pluralist theory, this is consistent with the view that in *some* circumstances, the funds should go to the smaller group (say, if the choice is between 9 with J and 10 with K). Second, I have argued that McCabe's simple cases of decisions made in isolation are not a good analogy with decisions made in a wider policy-making context. It is only when we consider the effect of a pure benefit-maximising framework across multiple decisions that the unfairness becomes manifest, in that it becomes apparent that it gives rare disease patients almost no chance of benefiting, in contrast to other groups, showing that the rarity of a patient's condition may be a relevant consideration.

However, on the pluralist view I am defending, this does not answer the question of whether to institute orphan drugs policies in any given case, or which policies to institute. The considerations favouring orphan drug policies are only pro tanto considerations. Other, strong considerations can come into conflict with them: in particular, the "efficiency" consideration that we ought to maximise either the benefits we produce or the number of people we benefit. Nothing I have said so far entails a specific conclusion regarding what we should do when the considerations favouring orphan drugs policies conflict with benefit-maximising considerations, as they will. Where there are good reasons for instituting orphan drugs policies, but also good reasons not to, and neither set of reasons is decisive, orphan drugs policies will be permissible but not obligatory, and it will be permissible for different healthcare systems to take a different views on how these conflicting considerations should be traded off against each other.

Thus I have shown McCabe's conclusion to be false. In cases where orphan drugs policies are permissible, they are not arbitrary. It is not unreasonable for policy-makers to sometimes place more weight on egalitarian considerations than on considerations of benefit maximisation.

5. How to Avoid Unfair Discrimination Against Disabled Patients in Healthcare Resource Allocation

In this chapter I address the potential for benefit-maximising healthcare allocation policies to discriminate[172] against disabled patients. To take an example case from an earlier paper of mine,[173] consider two groups of patients, the only difference between them being that one group is disabled and the other not. For simplicity, suppose that both groups have life-threatening conditions such that they can only expect to live a week without treatment, but the prognosis for both if they are given life-saving operations is 40 years, during which they will have their preoperative health-related quality of life (hereafter "quality of life").

For example, in the UK, the potential for discrimination between the two groups of patients arises because the UK's unit of benefit, quality-adjusted life years (QALYs), are based on the value placed by the public on various health outcomes, and surveys show that the public places a very low value on disabled states. For example, one way of eliciting public preferences for health states is by using the EQ-5D survey:

> The [EQ-5D] descriptive system comprises the following 5 dimensions: mobility, self-care, usual activities, pain/discomfort and anxiety/depression. Each dimension has 3 levels: no problems, some problems, extreme problems. The respondent is asked to indicate his/her health state by ticking (or placing a cross) in the box against the most appropriate statement in each of the 5 dimensions.[174]

---

[172] In my usage, it is not analytic that discrimination is impermissible. For example, it may be permissible for an employer to discriminate between job candidates by intelligence where that is required for the job.

[173] Sinclair (2012)

[174] van Reenen & Janssen (2015)

The EQ-5D allows us to describe health states. Now we have to value them. So let us consider a health state valuation quoted in Dolan (1997).[175] Survey respondents were asked to value EQ-5D health state "32211". In this state, as interpreted against the five dimensions of the EQ-5D, the patient has the worst possible mobility in that they are confined to a bed, and the patient also has "some problems" with washing and dressing themselves and with performing their usual activities, although they do not have pain or discomfort and they are not anxious or depressed. The survey produced a quality of life figure of 0.152 for this health state. This valuation was elicited using the time trade-off method, so we can understand respondents as saying that they would be willing to sacrifice nearly 85% of their life expectancy to avoid the health state.

So let us suppose that the quality of life of the disabled patients in our example above would be deemed by the public to be just 15%, whereas that of the nondisabled patient is 100%. This means that even if the life-saving operations produce the same number of incremental life years for both groups of patients, they will only produce six QALYs for the patients with disabilities compared with 40 QALYs for the non-disabled patients. So we gain more QALYs by saving the lives of the non-disabled patients than by saving the lives of the patients with disabilities. This means that the cost-effectiveness ratio (cost per QALY) of our life-saving treatment will be higher for disabled patients than for non-disabled patients. Now suppose that as a general rule, treatments will only be funded if their cost-effectiveness ratio is below a certain threshold, as in the UK. That could mean that the treatment is funded for non-disabled patients but not for the disabled patients.

But not only would such discrimination be illegal in most jurisdictions, on the grounds that disability is a protected characteristic; it would also be intuitively wrong. No doubt disabled patients and members of the public alike would consider it to be a case of unfair discrimination.

I aim to vindicate such intuitions. In doing so, I investigate whether this potential for disability discrimination points to a deeper problem with cost-effectiveness analysis. I conclude by offering a new theoretical basis for valuing health outcomes for the purposes of healthcare allocation.[176] In this approach, a key role is played by the causal provenance of the health outcomes that are expected with and without treatment. I will consider an

---

[175] Dolan (1997), p1105 table 3

[176] I thereby abandon my previous approach based on Patient Preference Theory (Sinclair, 2012).

account from Frances Kamm, who sees the potential that causation could play in an adequate theory.[177] However, I argue that Kamm's theory ultimately fails to realise the potential of this approach. I then propose my new, causation-based theory. The core of my defence of this is an argument that healthcare policy-makers should only focus on certain benefits of their interventions and not others. The account implies that disabilities unrelated to the condition being treated do not detract from the benefits of treatment. My case for this account depends on something like Kamm's "separate spheres" principle. Unfortunately, I find that Kamm has not offered a strong argument in defence of this principle, so it looks ad hoc. But I develop a new argument which gives us independent reason for adopting a separate spheres principle.

5.1 An Outline of the Difficulty

I will set the scene by highlighting some relevant intuitions regarding disability. I believe the following claims would be widely endorsed:

Intuition 1. The prevention intuition: It is permissible (perhaps sometimes obligatory) for governments to spend money on public health programmes and safety campaigns with a view to reducing the incidence of disability; for example, osteoporosis prevention programmes, or campaigns to reduce the incidence of deafness in noisy workplaces.

Intuition 2. The disability cure intuition: It is permissible (perhaps sometimes obligatory) for governments to spend money curing disabilities.

Intuition 3. The life-saving intuition: The disabled are as entitled to life-saving operations as the non-disabled.

Benefit maximising considerations in isolation cannot vindicate these intuitions, whatever we assume about people's quality of life. If we assume disabled people's quality of life is the same as that of non-disabled people (as perhaps we would sometimes find if we went on patient evaluations), we must renounce intuitions 1 and 2. But going on the public's health state valuations, we must assume it is worse. In that case, we must controvert intuition 3, with the consequence that we could end up depriving disabled patients of a life-saving treatment that is given to non-disabled patients. I

---

[177] Bognar (2010) also considers a causation-based account. I would argue it is vulnerable to counterexamples, but I do not have space to consider it here.

suggest there is a category of cases in which such "benefit-maximising" decisions would be particularly counterintuitive. These are cases where the disability in question is completely independent of the disease or treatment under consideration. Consider a patient with disability resulting from one disease who requires life-saving treatment for another disease, which threatens to kill her soon. Here, the disability is intuitively a background feature or incidental feature of the case and is therefore not a feature which healthcare allocators should take into account. Most people would find it completely unacceptable to refuse this patient her life-saving treatment in favour of a non-disabled patient purely on the grounds of the difference in their disability status.

I will briefly consider a couple of attempts to avoid one or other of these counterintuitive consequences. I will then consider a line of enquiry proposed by Frances Kamm. I will argue that this does not succeed as it stands but that it merits more detailed consideration; in fact, it ultimately inspires the theory I propose.

One attempt to avoid the counterintuitive consequences acknowledges that benefit maximisation is merely a pro tanto consideration and concedes that other considerations carry weight. In particular, the fact that a group has suffered historical disadvantage can sometimes seem to be a relevant consideration. On this basis, perhaps disabled people sometimes merit special provisions, either to compensate them for past injustices or to make sure they feel fully included despite their disadvantages, thereby ensuring social cohesion. Such considerations might sometimes give us reason to favour disabled patients for life-saving operations even if we cannot save as many lives that way as we could by directing the same budget elsewhere. Analogously, take an ethnic minority who have equally suffered historical discrimination. Suppose a gene disorder is much more common in that ethnic group than in other ethnic groups. As a result, a treatment is less effective for them than other groups, so their cost-effectiveness is poorer, so a healthcare allocation decision threatens to go against them. The proponent of a pro tanto consideration to do with benefit maximisation could concede that considerations to do with historical injustice and social cohesion might justify ignoring the seemingly poorer cost-effectiveness associated with that gene disorder. The same kind of argument could be applied in the case of disabled patients.

However, although historical injustice will often be a relevant consideration, I doubt it can explain all cases in which disability discrimination seems wrong in healthcare allocation. Even if disabled people had not suffered historical discrimination and did not feel socially excluded in other respects, they might still have grounds to complain if funding decisions kept going against them because their disability counted against them in a pure benefit-maximising policy-making framework.

Furthermore, this analysis would imply that there is a kind of compromise here between benefit maximisation and other considerations. But intuitively, there is no compromise; in fact, many would want to say that if the disability is a mere incidental or a background condition, unconnected with the condition or treatment, then it is not relevant at all to the question of who gets treated. Thus a theory which counts the disability as a "cost" thereby incurs a cost itself, in terms of its intuitive fit. (And there will be a similar problem with any other pluralist account which pits benefit maximisation against some other consideration).

This motivates a search for theories which avoid the potential for discrimination in such cases. But what factor could we incorporate into such theories to distinguish the cases of interest? The answer I propose to investigate is cached in the above example. What distinguishes the most counterintuitive forms of discrimination, I suggest, is that the disability in question is a separate condition, causally, from the condition to be treated. It neither causes it nor is caused by it. Perhaps an account based on causal factors could eliminate the problem in an important category of cases.

5.2 Kamm's Account

I will now consider a set of arguments from Frances Kamm, who has written on this topic in a number of papers. In Kamm (2009), Kamm considers various theoretical issues bearing on "real-life cases in which lifesaving resources are scarce."[178] Specifically she examines

> the role of quality and quantity of life in nondiscriminatorily allocating lifesaving and non-lifesaving resources between disabled and

---

[178] Kamm (2009), 148

> nondisabled individuals.[179]

Kamm canvasses a set of principles which she thinks together can help explain why it would be wrong to take account of people's disabilities in healthcare allocation. Each successive principle is discussed for its potential to correct counterintuitive implications of the preceding principles. I will consider four main discussions in this sequence:[180]

(1) Principle of Irrelevant Goods based on Sufficiently Good Only Option Argument

(2) Principle of Irrelevant Goods based on Equal Respect Argument/Irrelevant Identity

(3) Causative Principle

(4) Principle of Irrelevant Identity (Kamm calls on this principle again to deal with one final problem)

5.2.1. Principle of Irrelevant Goods based on Sufficiently Good Only Option Argument

To start at the beginning, Kamm's Principle of Irrelevant Goods says that "sometimes the fact that we can produce an additional good if we choose to perform one act rather than another is morally irrelevant".[181] Perhaps Kamm's most compelling illustration of this principle at work is given in an earlier paper of hers:

> suppose, for example, we have a choice between saving A's life and saving B's, and alongside B is C who has a sore throat. Our drug that can save B's life can also in addition cure C's sore throat. ...I believe it would be wrong to deprive A of his 50% chance to be saved simply in order to get the extra utility of curing C's sore throat associated with saving B.[182]

In Kamm (2009) she also offers a second and third case in which the same principle is applicable. In her second case, she suggests we should not favour A over B just because A stands to gain 15 years whereas B only stands to gain 10 years. She tentatively suggests that the extra five years for B is "a morally

---

[179] Kamm (2009), 148

[180] The fourth discussion to be covered here, in section 2.4, reverts to a previously mentioned principle, the Principle of Irrelevant Identity, but applies it differently.

[181] Kamm (2009), p. 149

[182] Kamm (1993), p. 101

irrelevant good."[183] In her third case, she says that if we have to choose between saving one fully healthy person and another who has lost a hand, the missing hand should not influence our choice. Kamm says "I think it is morally wrong to decide whom to aid on this ground."[184] Since the latter case involves a disability (a missing hand), the Principle of Irrelevant Goods looks well-placed to explain why we may not discriminate against the disabled in healthcare allocation.

But it is not enough for Kamm that a principle has intuitively satisfying consequences for specific cases. If this was her only theoretical adequacy condition, her principles could end up looking ad hoc. I take it that this is why Kamm does more than merely pointing out the intuitively satisfying consequences; she defends her principles with independent arguments.

Kamm offers three initial arguments in defence of her Principle of Irrelevant Goods, of which perhaps the most persuasive is the Sufficiently Good Only Option Argument. Ultimately this argument does not do the required work, but nevertheless it will be helpful to review it because it provides context to explain Kamm's eventual position. Kamm argues that if the alternative is death, one might want life years with a quality of life of 0.5 as much as life years with a quality of life of 1. Specifically, Kamm says:

> Suppose one can only have a life rated at .5 and not 1, and the alternative is 0 (death), which is very bad. One may reasonably want .5 as much as one would want 1 if one could have it. So, for example, given that .5 is all that one can have and 0 is very bad, one might reasonably do as much to achieve .5 (e.g., spend as much money, suffer as much) as one would do to achieve 1 if one could have it.[185]

To explain this, suppose, for example, that one is faced with the choice between death and a quality of life of 0.5. Then the strength of one's desire for 0.5 might be 99%, where it is not possible to want something more strongly than 100%. Now suppose that instead, one is faced with a choice between a quality of life of 1 and death. Here Kamm is saying that the strength of one's desire might also be 99% (the same), even though the quality of life on offer is 1 rather than 0.5.

---

183 Kamm (2009), p. 149

184 Kamm (2009), p. 159

185 Kamm (2009), p. 163

So the right way to value the outcome of life-saving treatments is to ask: how hard would someone fight to have this outcome rather than death? I suggest that the alternative to death would have to be pretty bad before someone's will to fight for that alternative was at all weakened. Most quality of life impairments would be trivial concerns in the context of such a choice.[186] Thus when we are allocating life-saving treatments, where the alternative for patients is indeed death, most quality of life impairments do not detract at all from the value of the outcome for the purposes of determining entitlement to those treatments.

This argument relies on the intuitively compelling thought that the value of an outcome depends what the available alternatives are. In doing so the argument invokes the idea of what we might call choice frames, where a choice frame is the range of options one considers in making a decision. The standard measure of quality of life is based on the public assessing the value of impaired health states versus the alternative of perfect health. The Sufficiently Good Only Option Argument says that this is the wrong choice frame when we are assessing the value of alternative health outcomes for the purposes of deciding which life-saving treatments to fund. Instead, we should assess the value of the impaired health state as compared with the alternative of death.

When we map health states against values, the two choice frames produce two differently shaped curves. When one axis is quality of life and the other axis is the value of life compared with full health, the two scales are related somewhat tautologically, given that quality of life is standardly measured in terms of life years of full health, e.g. via time trade-off questions. This means that, for example, a 10% increase on one scale corresponds with a 10% increase on the other scale, so the graph has a straight line.

---

[186] One way of operationalising this criterion might be to consider what the "average" citizen would do if s/he needed a life-saving treatment which would cost all their money, when there is also something else they would dearly like to spend that cash on, such as education for a daughter or a treatment that would cure a spouse's quality of life impairment (thanks to Rob Lawlor for this example). The options would be to accept the expensive life-saving treatment or refuse it in order to fund one of these alternative options. Perhaps if someone had 90% or 80% quality of life, refusal would not be a popular option, but if one had only 20% or 30% quality of life perhaps one of these alternatives would be preferred.

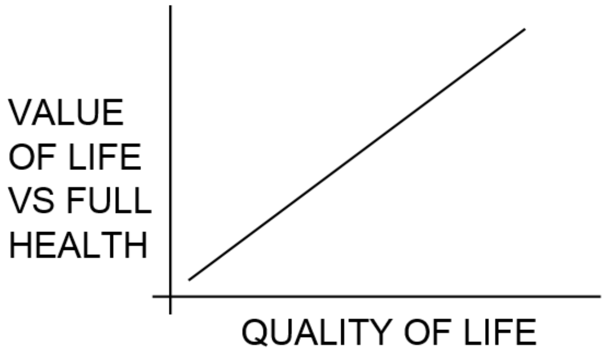

Fig 2: Value of life vs full health. The straight line reflects that quality of life is standardly measured in terms of life years, e.g. via time trade-off questions. This guarantees that equal intervals on the quality of life scale will correspond to equal intervals on the value of life scale.

Now suppose we take various points lying at equal intervals along the quality of life axis (for example, 0.1, 0.2 etc), and instead map those points against the value of life compared with death. We will find that although at very low quality of life there may be a rising trend (people with better quality of life will fight harder to avoid death), at a certain point the curve will flatten out. This is how we would represent Kamm's point that "one might reasonably do as much to achieve .5 ... as one would do to achieve 1".[187]

Fig 3: Value of life vs death. Although at very low quality of life there may be a rising trend, at a certain point the curve will flatten out to reflect that fact that people with a quality of life above some threshold will fight equally hard to stay alive.

---

[187] Kamm (2009), p. 163

To summarise, standard methods of assessing benefit are based on valuing each outcome (with and without treatment) as an alternative to full health. This yardstick indicates that (for example) life with one hand is worse than life with two hands. But Kamm's defence of the Principle of Irrelevant Goods is that this is not the relevant choice frame for the purposes of allocating life-saving treatments. In that context, the goodness of having two hands rather than one is an irrelevant good. Instead the relevant choice frame is where each health outcome is valued as an alternative to death. In this choice frame, having one hand and 100% quality of life are, to all intents and purposes, equally valuable, in that they are equally worth pursuing. And perhaps the same applies to worse quality of life too; for example, perhaps even paraplegia and 100% quality of life are equally worth pursuing, when the alternative is death.

5.2.2. Principle of Irrelevant Goods based on Equal Respect Argument/Irrelevant Identity

However, Kamm worries that when understood and defended this way, the Principle of Irrelevant Goods may have too much scope, implying that we must ignore the difference between saving one patient for 5 years and another for 50 years:

> suppose one person can be saved to live for five years and another for fifty years, and everything else is equal between them. Five years is a very significant good, and given that it is someone's only option, she might well reasonably do everything to get it that someone who could live for fifty years would do to get that. If there is a moral difference between taking account of expected length of life (both in the future and in how long someone will have lived if not aided) and taking account of disability or some other quality-of-life factor (both in the future and in the life someone will have had if not aided), then another argument apart from those we have considered is needed to justify this.[188]

She thinks the problem with relying on the Sufficiently Good Only Option argument to support the Principle of Irrelevant Goods is that the principle will then imply "that we should treat sufficiently good only options that involve *quantity* of life in the same way as we treat sufficiently good only options that

---

[188] Kamm (2009), p. 168

involve *quality* of life."[189] But intuitively, if a one treatment produces more life expectancy than another, this is not an irrelevant good. So it turns out that the explanation of irrelevant goods in terms of the Sufficiently Good Only Option argument has counterintuitive consequences. This leads Kamm to seek an alternative principle, one that will show that quality of life can be an irrelevant good but not that life years are an irrelevant good. The assumption is that such a principle will permit discrimination by life years but not by quality of life.[190]

So, Kamm changes the argument with which she defends the Principle of Irrelevant Goods. She now aims to defend it in terms of a Principle of Irrelevant Identity (the Principle of Irrelevant Identity is held to entail the Principle of Irrelevant Goods).[191] The Principle of Irrelevant Identity says that identity is irrelevant for purposes of allocation. This is important because our identity is partly constituted by aspects of our quality of life. Kamm defines our quality of life as constituted by various "synchronic" properties. To define synchronic properties, Kamm distinguishes two notions of the worth of life: "the notion that involves evaluating lives as better or worse, where all properties of the life are included in this evaluation" and "the notion of the worth of life, or going on living, to someone".[192] The first of these is partly based on the quality of the person's life, "thought of as a set of synchronic properties that modify any period of her life."[193] In contrast, the second notion is not conceptually tied to quality of life in the same way: going on living may be as valuable to someone with unfavourable synchronic properties/poor quality of life as to anyone else. Kamm suggests that our synchronic properties are constitutive of our identity, so they should be ignored for the purposes of healthcare allocation.

However, once again Kamm again does not regard it as sufficient justification for a principle that it happens to have intuitively satisfying

---

[189] Kamm (2009), p. 167

[190] Kamm (2009), p. 168

[191] Note that she also offers various other arguments in defence of the Principle of Irrelevant Goods. She sets them aside as they have counterintuitive consequences or do not support her target conclusion regarding disability discrimination. To give a complete list, the arguments she offers are the Major Part argument, Moral Importance argument, Sufficiently Good Only Option argument, and Principle of Irrelevant Goods/Equal Respect argument. These are each intended to support the Principle of Irrelevant Goods (see Kamm, 2009: p. 169). In what follows, I now consider Principle of Irrelevant Goods/Equal Respect argument.

[192] Kamm (2009), p. 165

[193] Kamm (2009), p. 165

consequences; she needs an independent argument for it. She appeals to the equal respect argument in support of the Principle of Irrelevant Identity:

> Each person is entitled to equal respect and (at least for purposes of an impartial distribution of scarce resources) equal concern. That may mean that (at least certain) synchronic properties, such as whether one is or will be paralyzed, even assuming that they significantly affect quality of life, should not bear on selection for scarce resources. If these synchronic properties are appropriately thought of as determining one's identity, one might say that equal respect makes identity irrelevant for purposes of allocation. (Call this the Principle of Irrelevant Identity.) However, taking into account, for example, how long a person can live if he gets a scarce resource is not treating someone differently because of the type of person he is or will be qualitatively; the latter (it is being suggested) is done only if we consider someone's synchronic properties (properties that determine the character of his time alive). It is, theoretically, compatible with each synchronic type that a person could be, that he could be that type for longer or shorter amounts of time.[194]

The Principle of Irrelevant Identity and its supporting argument in terms of equal respect supplants the earlier Sufficiently Good Only Option Argument as Kamm's justification for the Principle of Irrelevant Goods, since the Principle of Irrelevant Identity implies that we must ignore a patient's prior disability in valuing the outcome of a life-saving treatment without counterintuitively implying that we should ignore all life expectancy differences.

Finally it should be noted that Kamm thinks that, although the demands of equal respect require us to ignore many quality of life impairments, even quite serious ones, equal respect does not require us to ignore *all* quality of life impairments. In extreme cases, where quality of life is very poor, we do not disrespect someone by failing to extend their life:

> Counting the qualitative differences that the people themselves bring, I have suggested, is not consistent with respect for different types of persons, except when the difference results in a life below a certain minimum (e.g., a life not worth living, or not worth doing a lot to save).[195]

---

[194] Kamm (2009), p. 168

[195] Kamm (2009), p. 177

The qualification at the end is interesting, and somewhat difficult to explain. Why is it consistent with respect to refuse to save a life not worth living? My suggested interpretation is that generally, to respect a person, we must respect every aspect of them, including their (minor) impairments. One way of expressing such respect is to give them an equal claim to beneficial treatments. We recognise their worth as human beings by helping them survive, or relatedly by giving them the same chance to survive as other humans. However this gesture only expresses respect because the treatments in question are beneficial. In contrast, we would not express any respect by offering something worthless. But a treatment which does not extend life would be worthless, and perhaps we could say that a treatment which extends a life not worth living is similarly worthless.

However this gloss does not explain why it might be consistent with respect to refuse to save a life that is not worth doing "a lot" to save. The suggestion seems to be that we may refuse to treat a patient whose quality of life is poor, even if their life is not as bad as death. The problem is that in such a case, a life-extending treatment might still be beneficial, though the benefit would be small. Would it not be as disrespectful to refuse such a patient as any other patient? Perhaps we could say that the smaller the benefit we refuse to give, the less disrespectful the refusal is. Perhaps in such circumstances, the obligation to maximise the benefits we produce outweighs the duty of respect. The problem with this is it reverts to make the strength of a patient's claim dependent on how much benefit they will get, which is what Kamm was trying to avoid. Perhaps in all consistency she must abandon this somewhat nuanced picture; if quality of life impairments are not such as to make a patient's life worse than death, they are as entitled to treatment as any other patient.

5.2.3. Causative Principle

But now, Kamm worries that the Principle of Irrelevant Goods implies that

> we should not prefer saving a paraplegic who, as a side effect of lifesaving treatment, will also be able to walk again to saving a paraplegic who will remain a paraplegic.[196]

Similarly, it implies that

---

[196] Kamm (2009), p. 173

> if one unparalyzed person would become paraplegic as a side effect of lifesaving treatment but another would not, this should make no difference in whom we choose to save.[197]

These implications are problematic, as they suggest it would be impermissible to take account of treatment outcomes more generally in healthcare allocation. So, the Principle of Irrelevant Goods looks too strong - it says some quality of life impairments that are in place before a patient receives treatment are too minor to influence our decision about who receives treatment. But the principle then implies that the same impairments are also too minor to take into account when a treatment helps cure them or avoid prevent them subsequent to treatment. This is a counterintuitive consequence.

Kamm thinks these cases raise the following possibility:

> (1) Sometimes a sizable extra synchronic good (or bad) that we can produce in the outcome, if we treat one person rather than another, should be morally relevant in deciding whom to help with a lifesaving resource. (2) Yet, if candidates for treatment who present themselves have this difference in good (or bad) between them, and this is why it shows up in the outcome, the extra good (or bad) should be morally irrelevant in deciding whom we help with a lifesaving resource. The Principle of Irrelevant Goods cannot account for the simultaneous truth of (1) and (2).[198]

Kamm's response is to supplement the account given so far with an additional principle. Kamm says we need:

> a principle that will explain why the fact that a person is and will be disabled to a certain degree should sometimes be irrelevant in deciding whose life to save, but the presence of the same disability in an outcome can sometimes be morally relevant and used in a non-discriminatory fashion in deciding whose life to save.[199]

Kamm's therefore introduces the Causative Principle, which says:

> "We may decide whom to help based, in part, on the synchronic difference we can make to a person’s situation, not on the difference he brings to the situation ... the Causative Principle is concerned with the

---

197 Kamm (2009), p. 173

198 Kamm (2009), p. 173

199 Kamm (2009), p. 173

> differential effectiveness of our treatment in producing nondisability and our being entitled (though not obligated) to bring about a better outcome by using our skills in this way ... The Causative Principle tells us to ignore this difference in outcome when it arises in any other way, whether because the disability inheres in the person or will arise because of what inheres in him, or even will arise from causes outside of him other than our treatment."[200]

Kamm's justification of the Causative Principle is similar to her justification of the Principle of Irrelevant Identity:

> when outcomes are affected by who a person is and/or by what we do, counting only what we do is consistent with the account of respect for different types of persons given above ... Counting the qualitative differences that the people themselves bring, I have suggested, is not consistent with respect for different types of persons, except when the difference results in a life below a certain minimum (e.g., a life not worth living, or not worth doing a lot to save)[201]

To clarify, Kamm should not be seen as abandoning the Principle of Irrelevant Identity here, but as supplementing it with the Causative Principle. As per the Principle of Irrelevant Identity, she still holds that respect involves ignoring the qualitative differences that people bring. But a qualitative difference we produce is not a difference that people bring, and so we may take account of it.

5.2.4. Principle of Irrelevant Identity - again

However, Kamm now recognises one final problem for her account of the principles involved in non-discriminatory healthcare allocation. Kamm asks us to imagine a case in which, "if we save person I, we will also cure his paralysis, whereas if we save J, we will just save his life, there being no paralysis to cure."[202] Kamm observes that in the terms of the causative principle, "we produce a significantly larger difference if we treat person I than if we treat J."[203] As a result the Causative Principle says that it is permissible to save the paralyzed person rather than the unparalyzed. Kamm

---

[200] Kamm (2009), p. 176

[201] Kamm (2009), p. 177

[202] Kamm (2009), p. 183

[203] Kamm (2009), p. 183

thinks this is "the wrong conclusion" and that "there is no good reason for favouring one over the other."[204]

To solve this problem, Kamm does not need a new principle; she invokes The Principle of Irrelevant Identity again. This principle implies that to favour the paralysed person in the above case would be to take account of the identity of the prospective patients in disrespectful ways. For example, it would involve holding against the person who is nondisabled the fact that he is nondisabled, giving him less respect and concern than the other patient. In the following, she describes how the Causative Principle combined with the Principle of Irrelevant Identity issue in different prescriptions according to whether the outcomes expected in two individuals are different or the same:

> When the synchronic outcomes expected in different individuals are different, to pay attention to anything but the causative difference we make would be to make the difference in them affect our decision of whom to aid. Hence, we can abstract from who they are by attending to the causative component. Or, alternatively, we can imaginatively add the good property that one party is missing into his outcome ... But when the synchronic outcome we expect in different individuals is the same, the decision to attend only to the causative difference we make results in the differences in who they are playing a role in our decision regarding whom to treat.[205]

Thus Kamm thinks that she can still deal with the earlier case, recommending that we should give the life-saving treatment to the paraplegic who will be able to walk again instead of to the paraplegic who will remain a paraplegic. In that case we discriminated on the basis of the different outcomes. But she can also deal with the latest case; her account does not recommend either patient (neither the unparalysed patient whose life we can save, nor the paralysed patient, whose paralysis we can cure and whose life we can save). In this case, the outcome that will be enjoyed by the two patients is the same, so we have no grounds to discriminate between them. We must ignore the fact that one patient could go from being paralysed to unparalysed, since that would involve taking account of who they are when they come to us. We must simply look at the outcomes.

---

204 Kamm (2009), p. 183

205 Kamm (2009), p. 184-185

5.2.5. Kamm's Conclusion

Kamm then summarises her final position as follows:

> In evaluating outcomes, the difference a resource can make to quantity of life in one person rather than another may be relevant. A difference in quality of life in one person's outcome rather than another's may be relevant when a large difference in quality between persons would be produced by our efforts (rather than when we save a person whose much better quality of life is not due to us) or when someone's unchangeable quality is like, or causally linked to, what we aim to treat.[206]

Thus Kamm is more permissive regarding discrimination by life expectancy than discrimination by quality of life. On Kamm's account, life expectancy differences are always potentially relevant, but quality of life differences are only relevant when we produce those differences or when there is a likeness or causal link with the impairments we're treating or when a patient's quality of life falls below a certain minimum. To put it another way, there is a big contrast between the moral status of quality of life factors and life expectancy factors when they are dimensions of a patient's prognosis without treatment, but there is less of a contrast when they are dimensions of treatment benefit. With regards to prognosis without treatment, the principle of irrelevant identity debars discrimination by quality of life impairments but not life expectancy impairments,[207] since Kamm holds that quality of life impairments are part of a patient's identity whereas life expectancy impairments are not. In contrast, when we are considering the benefits of treatment, life expectancy benefits and by quality of life benefits will often both be relevant. This is because Kamm does not view quality of life benefits produced by a treatment as part of the patient's identity; the patient does not bring those benefits to the clinic.

5.3 Commentary on Kamm's Account

Although I believe parts of Kamm's account are right, particularly the parts involving causation, I also contend her account faces several difficulties.

---

[206] Kamm (2009), p. 196

[207] Thus the argument from equal respect should be seen as supplanting the earlier arguments, such as the sufficiently-good-only-option argument, since those earlier arguments would debar most discrimination by life expectancy.

First, we can question why the equal respect argument picks out quality of life factors and not life expectancy factors. Other writers who demand equality demand a more wholesale rejection of discrimination than this. For example, in Chapter 2 on discrimination, we have seen Harris rejecting discrimination by "life expectancy, or quality of life so long as that quality of life is worth having for the person whose life it is."[208] As far as I can tell, no-one other than Kamm defends an egalitarianism which implies that discrimination by one aspect of a patient's prognosis is OK but by another aspect is not. This is not an objection in itself, but it suggests this is not an easy distinction to defend.[209] Kamm does not spend long on this, but she says:

> If these synchronic properties are appropriately thought of as determining one's identity, one might say that equal respect makes identity irrelevant for purposes of allocation.[210]

But the conclusion that identity is irrelevant does not follow from the premise that synchronic properties determine one's identity. Even if we grant Kamm that we owe each other equal respect, we need an account which shows what kinds of characteristic we may discriminate by whilst showing equal respect. To consider some of Kamm's options on this question, consider the most radical view that says that all discrimination by patient characteristics is wrong. But I have already argued against this in Chapter 2. Kamm would also reject this position, as in her conclusion when she claims that: "A difference in quality of life in one person's outcome rather than another's may be relevant when a large difference in quality between persons would be produced by our efforts,"[211] and that it may be relevant if we can increase life expectancy by more in one patient than another.

The natural thing for Kamm to say is that equal respect demands equal treatment in respect of any characteristic where people with the characteristic find it to be an important part of their identity. The grounds for this could be that such discrimination would be demoralising, conveying a sense that society disvalues characteristics which are centrally important to people's

---

208 Harris (2005)

209 Kerstein makes a somewhat similar point about the lack of a defence of the special status of quality of life factors over life expectancy factors: "Kamm does not offer an explanation of where, morally speaking, the difference lies between privileging someone for lifesaving treatment on the basis of expected lifespan and privileging someone for such treatment on the basis of blood type." (Kerstein, 2015)

210 Kamm (2009), p. 168

211 Kamm (2009), p. 196

sense of self-worth. But there are at least two problems with this line of defence. The first is that patients do not always self-identify with their quality of life impairments in this way. People in pain do not generally identify themselves as essentially pained. Similarly, a common refrain of people with disabilities is that others need to "see beyond/through the disability."[212] One individual with a psychiatric disability is quoted as saying "I am a normal person, I just happen to get sick sometimes."[213] The second problem is that this defence does not support Kamm's distinction between quality of life factors and life expectancy factors. If I am a 20 year-old with a disease that causes muscle ache and will kill me at 25, why is the fact that I get muscle ache constitutive of my identity and not the fact that I will die at 25? If anything, the fact that I will die at 25 will shape my goals and values and personality more than the fact that I get muscle aches, and as such it would seem to be more part of my identity. So Kamm cannot justify debarring discrimination by pain or disability on the grounds that most people with those impairments self-identify with the impairments.

Thus it looks difficult for Kamm to explain why an obligation of mutual respect requires us to ignore quality of life factors but not life expectancy factors. As a result, Kamm's principles are in danger of looking ad hoc. Of course, Kamm could simply state her Principle of Irrelevant Identity in terms that pick out the right cases, and then assert that this extensional adequacy is sufficient grounds for endorsing it. But the principle needs a supplementary defence against the accusation that it picks out the right cases by means of some contingent and morally irrelevant feature which just happens to appear in those cases. Analogously, no doubt it would be possible to come up with a "quasi theory" of wrongful discrimination by generating some enormously complicated generalisation that uniquely picked out all the cases of wrongful discrimination. But stating an "ethical principle" based on such a generalisation is unlikely to help *explain* why the actions it picks out are wrong, and it would therefore lack interest. Also, this method would guarantee that our moral theory will always vindicate the intuitions with which we start, whereas it is implausible that our intuitions are always right in the moral domain when we are so regularly wrong in other domains.

So, *some* further defence of an explanatory principle is needed beyond the claim that it captures all the intuitively right cases. What else we could

---

[212] For example, see Department of Communities, Queensland Government (2007)

[213] Anne, quoted in Department of Communities, Queensland Government (2007)

demand? There are at least two things. We could demand that our explanatory principles be defensible on the basis of analogies, explaining distinctions we would make in a variety of cases other than the ones at issue. The applicability in uncontested cases provides evidence that the principle is valid. Alternatively we could demand that our explanatory principles be intuitively compelling (a principle that is extensionally adequate but explains our moral beliefs in terms of seemingly bad-making features will not be an intuitively compelling explanation of why doing the right thing might somehow be good).

As noted earlier, Kamm herself seems to respect this demand for something more than extensional adequacy in the way she defends, for example, the Principle of Irrelevant Goods and the Causative Principle: she offers arguments for them and doesn't merely show that they have intuitively satisfying consequences.

Kamm has also endorsed this demand more explicitly. For example, after discussing whether distance matters in the duty to rescue, she says:

> I have dealt with whether and how intuitively we think distance matters. Now, we should consider whether we could justify these intuitions. For I accept that intuitive support is not enough to justify a principle of morality. We must find morally significant ideas underlying the intuitions for the principle to be justified.[214]

So Kamm is not as sceptical about such explanations as she is sometimes held to be.[215] Elsewhere, while she admits to not spending as much time on this explanatory task as other writers, she doesn't dismiss the value of such explanatory work:

> I do not necessarily intuitively find these underlying factors morally significant, though I may. ... When I do not have an intuitive response to the found factor, I only claim that it accounts for my intuitions, and if my intuitions are morally correct, then these distinctions which generate them, perhaps surprisingly, have moral import. So, it may be that we do have to draw distinctions not drawn before in order to account for the appearances. ... Next, we might go deeper to investigate whether there are any ideas that we think have some intuitive moral significance with which these found distinctions are connected. I said in the introduction to

---

[214] Kamm (2000), p. 674

[215] For example, Norcross (2008)

> my book that I would do much less of this sort of thing, focusing rather on considering my intuitions and what differences in the cases seem to generate them.[216]

So, if I'm right that it is difficult to explain why equal respect requires us to ignore synchronic properties, I take it that Kamm would share the view that this is an issue.

Second, a problem with the principle of irrelevant identity is that it does not debar discrimination by intuitively irrelevant differences of life expectancy. For example, suppose that two patients will die in a month without treatment. In one case treatment is expected to add ten years and in the other case ten years and one month. Intuitively that difference should not influence our decision. If a missing hand is too minor to be relevant in a life-saving case, this difference in life expectancy is irrelevant for the same reason; it is too trivial, in context. But neither the principle of irrelevant identity nor the causative principle oblige us to ignore this difference; in fact, Kamm specifically introduces the causative principle to permit discrimination by such aspects of the outcome.

Finally, there is a problem with the scope of Kamm's Causative Principle. As background, consider a pair of cases introduced by Kamm, in which she argues that discrimination is permissible: "one unparalyzed person would become paraplegic as a side effect of lifesaving treatment but another would not".[217] Kamm comments on this case that "Sometimes a sizable extra synchronic good (or bad) that we can produce in the outcome, if we treat one person rather than another, should be morally relevant in deciding whom to help with a lifesaving resource,"[218] and she offers the Causative Principle in response.

Now compare Kamm's cases with a third case: that of a patient whose disease will kill them without treatment or cause paraplegia if we give them life-saving treatment. Kamm's Causative Principle and her Principle of Irrelevant Identity debar taking the paraplegia into account when deciding who to treat. For example, The Causative Principle tells us to ignore this kind of difference in outcome when it arises in any way other than the exercise of our skills, "whether because the disability inheres in the person or will arise

---

[216] Kamm (2008)

[217] Kamm (2009), p. 173

[218] Kamm (2009), p. 173

because of what inheres in him, or even will arise from causes outside of him other than our treatment."[219] But this discrimination between Kamm's patient and mine is counterintuitive. It is implausible that we are permitted to discriminate against Kamm's unparalyzed person whose lifesaving treatment would cause paraplegia but not my patient whose disease would cause paraplegia if we save them. If it is relevant that one candidate for treatment would become paraplegic as a result of treatment while another would not, it is equally relevant that one candidate for treatment would become paraplegic as a result of the disease we are treating, while the other would not. In my case as in Kamm's, there is a close causal connection between the paraplegia and the alternative outcomes we are considering; we are not dragging in aspects of the patients circumstances completely unconnected with the condition on which we are focused. Disabled people should have no complaint about discrimination in my case, since they do not stand to lose out by it; their disability will still not be taken into account. And those who stand to lose by discrimination in my kind of case have no more complaint about this than about any of the other prognosis-based discrimination which forms part of a defensible system of assessment by cost-effectiveness.

5.4 A New Three-Part Account: New Causal Account, Plus Considerations to Do with Egalitarianism and Historical Disadvantage.

For these reasons, I contend that Kamm has not explained what is wrong with disability discrimination in healthcare allocation. To summarise, I will now argue that the following pro tanto factors are relevant in determining whether discrimination against disabled patients on grounds related to their disability is permissible or not:

Principle 1. An assessment of the net benefit produced by the treatment, in terms of both quality of life and life expectancy. To assess this, we compare the health impairments predicted to result from the untreated disease with the health impairments predicted to result from the disease and its treatment if patients are treated.[220] We ignore symptoms of conditions unrelated to the

---

[219] Kamm (2009), p. 176

[220] I set aside for a moment the question of whether quality of life and life years are commensurable or whether the assessment of the various dimensions of each health outcome

condition we are treating, as well as the costs and benefits of treating those unrelated conditions, *unless* those causally unrelated costs and benefits involve dimensions of health impaired by the disease/treatment we're considering treating (for example, the condition we're considering treating paralyses patients' leg muscles. The treatment cures the paralysis in all patients except those with a comorbidity which causes paralysis in the same muscles by a different causal pathway). In such a case the unrelated disease/treatment should be deemed to affect the efficacy of the treatment in question.[221]

Principle 2. A consideration of egalitarianism which implies that all have some entitlement to healthcare even if their prognosis/cost-effectiveness is poor (regardless whether it is poor in respect of life expectancy or quality of life).

Principle 3. Considerations to do with historical disadvantage, social cohesion and perceived social inclusiveness can give us reason to relax the criteria in the case of members of the disadvantaged groups, such as disabled patients.

Under Principle 1, what sets my account apart from Kamm's is the concept of causation I have employed in the benefit-maximising criterion. For Kamm, causation is about "the synchronic difference we can make to a person’s situation".[222] By this she means the changes produced by our treatment. But this leaves her with a problem. She considers a case where, "if we save person I, we will also cure his paralysis, whereas if we save J, we will just save his life, there being no paralysis to cure".[223] Kamm's Causative Principle implies we should favour the patient in whom we can cure paralysis, but Kamm thinks "this is the wrong conclusion and that there is no good reason for favoring one over the other".[224] I think she is right. She tries to deal with this case by resorting to the Principle of Irrelevant Identity, arguing that we must ignore a patient's prior disability status for the sake of equal respect. However, I have noted reasons to question whether that principle is valid. My account needs no such principle, since according to my account, patients' health states prior

---

must be done in a more implicit, case-by-case way. I merely argue that one way or another, the positive and negative effects of treatment must be considered in this way.

[221] This is somewhat like Kamm's Treatment Similarity Principle, which I have not considered in detail.

[222] Kamm (2009), p. 176

[223] Kamm (2009), p. 183

[224] Kamm (2009), p. 183

to treatment have no bearing on our assessment of their treatments; we simply compare the outcomes with and without the treatment of interest. For example, in Kamm's case, we can say that the outcome of treatment option 1 (treating I) is that I lives and J dies; whereas the outcome of treatment option 2 (treating J) is that I dies and J lives. When we compare these two outcomes we get precisely the result that Kamm wants, which is that all else being equal, there is nothing to tell between them. The curing of the paralysis doesn't threaten to enter the assessment, on this principle, so there's no need to block it with anything like Kamm's Principle of Irrelevant identity.

I suspect the reason Kamm introduces her Causative Principle, telling us to focus only on the synchronic difference we can make to a person’s situation, is to exclude the effects of other diseases. I agree this is a desideratum. My account achieves the same effect by requiring us to focus on the effects of the target disease and our treatment, and not the effects of unrelated diseases.[225] (So the difference between us is that my account requires us to take account of the effects of the disease we're treating, whereas Kamm's doesn't).

Principle 2 is intended to deal with Kamm's cases of intuitively irrelevant goods, such as sore throats and missing hands in life-saving cases. I outlined how egalitarian considerations played their role as pro tanto considerations in Chapter 4 on rare diseases. The egalitarianism I defended there implies that if an impairment is trivial it should be ignored, whichever kind of impairment it is.

Principle 3 highlights that we may sometimes take account of contingent facts such as historical discrimination. One question that Kamm does not address is whether disability merits special treatment compared with other quality of life impairments, such as severe but non-disabling pain or anxiety. She seems to assume that disability can be dealt with in the same breath as other quality of life impairments. On the contrary, sometimes disability seems to merit special treatment. For example, compare two diseases. Disease A shortens life and cause severe but non-disabling pain. Disease B shortens life and causes a severe disability. Both have the same life expectancy and the same quality of life. We have life-extending treatments for both diseases which would add the same number of years. I suggest it might be intuitively acceptable to take account of the severe pain in the case of disease A (reducing the chances of funding) whilst ignoring the disability in the case of disease B.

---

[225] Of course, in line with my earlier demand, this principle needs to be justified in terms of something more than extensional adequacy. address this demand below.

If so, I suggest the grounds for this distinction would be contingent facts about the actual circumstances of disabled patients, such as historical discrimination. The motivation might also be to enhance disabled patients' sense that they are socially included and enhance social cohesion. So the justification would be grounded in contingent historical circumstances and not any feature that can be identified a priori. Sometimes special status is needed for such contingent reasons. The UK list of protected characteristics is quite a motley list, including items such as gender reassignment, pregnancy, religion and belief. It doesn't look like these characteristics share any feature which would explain their inclusion on the list other than their association with historical discrimination. Given this, if people had not been historically discriminated against on grounds of these traits, policy-makers would have had no reason for conjoining them in a list of protected characteristics. Conversely, if people started discriminating against people with blue eyes, such that people with blue eyes started suffer terrible disadvantage, then blue eyes would be added to the list of protected traits.

5.5 Objections and Replies

I now address a couple of objections to my account. The first objection relates to my stipulation under Principle 1 that that we should only take account of the effects of the target disease and its treatment - not the effects of unrelated diseases. The objection is that the stipulation is ad hoc. If we are interested in benefits, why not the overall net benefit that comes about as a result of our intervention? Isn't it arbitrary to ignore some benefits?

To justify my restriction, I will offer a quite general argument that healthcare policy-makers should only focus on certain benefits of their interventions and not others. Now, as it happens, Kamm has presented a very similar argument, known as the "separate spheres" argument. I find the conclusion of this argument intuitively compelling, but I do not find that the conclusion is well-defended. So the main job that needs doing is to find a better justification of the position she wants to defend. In the following, I articulate a new defence of a separate spheres principle, and then consider the implications. I find that one implication is that that we should only take account of the effects of the target disease and its treatment, and not the effects of unrelated diseases.

I start by summarising Kamm's argument from separate spheres. Kamm starts with an example where we must choose between saving a great philosopher or a not-so-great non-philosopher. Kamm argues that it would be impermissible to save the great philosopher instead of the non-philosopher for the sake of the valuable discoveries we expect from the philosopher (even if those patients who would lose out to the philosopher would also recognise the value of those prospective discoveries).[226] She says:

> in the sphere of lifesaving ... we cannot say that we will have satisfied the requirement of the sphere, namely to save life, if we save one life and go on to produce a cultural achievement as well. The good of lifesaving rather than this extra utility must continue to be our paramount objective. Under a theory of separate spheres or specialised aims, health resources are to be used to produce adequate conscious life and/or what people are willing to live for if it is life expectancies than objectively adequate.[227]

I find this principle intuitively compelling. I suggest, however, that Kamm's defence of it needs supplementing. In Kamm's argument the principle is supported merely by examples which illustrate how counterintuitive it would be to go against it. Kamm does not do enough to explain *why* we should respect it. Her argument is as follows:

> the goal of improving health or saving life is sufficient unto itself, and there is corruption in the achievement of this aim if the achievement of some other good is combined with it in the selection of persons.[228]

These are weak arguments. To see this, consider a policy-maker who holds that we should save the philosopher rather than the non-philosopher. The policy-maker grants that of course saving lives is valuable, but points out that the philosophical discoveries we could expect would also be very valuable; her grounds are therefore that saving the philosopher produces two benefits rather than one. Kamm replies with two assertions: 1. that "the goal of saving life is sufficient unto itself" and 2. that "there is corruption in the achievement of this aim if the achievement of some other good is combined with it." Unfortunately, the language here is somewhat bombastic and unclear (uncharacteristically of Kamm). When we try to clarify the assertions, we find that neither of them gives such a policy-maker any reason to change her mind. When Kamm says that "the goal of saving life is sufficient unto itself", she seems to say that the goal of saving life should be considered in isolation;

---

[226] Kamm (1993), p. 258

[227] Kamm (1993), p. 259

[228] Kamm (1993), p. 259

incidental benefits do not increase the value of a life-saving intervention. But this merely asserts the conclusion (the conclusion being that, when we're considering whose life to save, we should not take account of other benefits of our intervention). So this does not provide independent support for the conclusion. Our policy-maker would deny that the goal of saving life is sufficient unto itself, since that is a straightforward denial of what the policy-maker asserts. The second of Kamm's assertions is that the aim of saving lives will be corrupted if we combine other aims with it. The word "corruption" is loaded, suggesting policy-making will be debased or tainted if it takes account of other aims. These connotations come at the expense of clarity; we cannot assess this claim until we clarify exactly *how* policy-making is expected to become debased. The idea appears to be that if we take account of other aims, we may achieve less of the first aim of saving lives. For example, in the case at hand we might save the philosopher even though the philosopher is not expected to live as long as the non-philosopher. But again, this assertion does not give our policy-maker any reason to change her mind. The policy-maker will point out that we have to balance conflicting aims all the time in healthcare, as when we must choose between extending a life or improving the quality of a life. If an aim is important we should take account of it even if it conflicts with other aims. I take it Kamm would recognise this, since her analysis posits many conflicting considerations. Once she is taken to have accepted this, the point of her second assertion merely seems to be that saving lives is a legitimate aim of healthcare but facilitating philosophical discoveries is not. But again, this merely asserts the conclusion which Kamm wishes to defend.

Since Kamm has not provided a good defence of her separate spheres principle, I will attempt my own explanation. I will start by surveying some other policy areas to illustrate the influence of something like this principle. One example is income support: to decide whether to give me income support, the government might means test me, but they will not try to determine how happy the money will make me. To determine whether to send police officers to investigate the burglary at my house, the police will judge the seriousness of the crime and the chances of solving it, but not how much I need the stuff back.[229] To determine whether I get a university place, a university may forecast my ability to do well on the course, but they will not

---

[229] Except in extreme case, as when the stuff includes irreplaceable life-saving medicines (thanks to Rob Lawlor for this point)

generally try to forecast how much benefit I will derive from it afterwards, such as how well I will do in the job market afterwards. Thus each of these policy areas seems to have its own "separate spheres" constraint in operation. Policy-makers in each area seem to intuitively recognise a similar ethical consideration, providing prima facie evidence that some such principle is indeed applicable in policy-making.

In all these cases, government prioritises services that deliver *universally beneficial* resources or capabilities.[230] Generally these will be direct benefits of the service, viz., benefits which recipients derive from the service before they go on to pursue whatever personal projects or plans they have. The direct benefit of welfare is money; the direct benefit of a successful police investigation is the restoration of one's stuff; the direct benefit of a university course is knowledge. And the direct benefits of successful healthcare are, for example, an ability to walk or absence of pain. These benefits are universal in the sense that all who use each of these services can be expected to find it beneficial in these ways.[231] Generally, the benefits are universal because they are direct benefits of the service. In contrast, any further benefits (indirect benefits) that service users derive will vary from user to user according to their subsequent choices and circumstances. I suggest that unless such indirect benefits are universal, they should not enter into the question of who to allocate the service to. Citizens who need the service and stand to gain such universal benefits from it should stand an equal chance of getting the service. The universal benefits are the point of public provision. The further benefits depend on each citizen's own circumstances and predilections and projects and priorities; they are therefore each citizen's own private business.

To try and explain *why* it's important that benefits be universally beneficial, as background we saw in Chapter 3 that there are two accounts of the rationale for publicly funded healthcare. The first accounted for publicly funded healthcare as a kind of mutual insurance fund. On this model, we all make affordably small, regular contributions to a mutual fund in return for the ability to make occasional big withdrawals as needed. This is a quasi-contractual model, or a model based on reciprocity.[232] The second is in terms

---

[230] I defended a capability account of the benefits of healthcare in Chapter 3.

[231] If it's argued that the ability to walk is not intrinsically beneficial, my response is that it's nevertheless instrumentally beneficial, allowing all patients to do things they couldn't do otherwise

[232] Actually some systems of universal healthcare still rely on insurance funds. This does not necessarily mean private insurance funds and patchy coverage, as in the US. On the contrary,

of a general obligation to help people in need; governments are entitled to tax their citizens to fulfil citizens' obligations on their behalf.

I will try to explain why the benefits that drive policy should be *universally* beneficial for service users in terms of each of these two models. First, when the system is understood on the mutual insurance model, the reason is that citizens can only be expected to support a mutual insurance fund to the extent that it delivers the kind of benefits that they might need, *regardless of what they are doing with their life*. In contrast, it would be unreasonable to require participants in a mutual insurance scheme to contribute to a scheme that funds other people's idiosyncratic benefits. So, for a communally funded system to be justifiable to its funders it needs to be focused on universal benefits. We all want the capabilities that health gives us, but we do not all use those capabilities the same way. We'll all pay for a service that helps all of us maintain our eyesight, because we all want our eyesight. But it would be unreasonable to expect me to fund an eyesight restoration service if others were going to get ahead of me in the queue for the service because they got a lot of enjoyment from reading pulp fiction while I didn’t get much enjoyment from reading. Everyone has their own uses for eyesight and that gives everyone a reason to fund an eyesight restoration service if it is allocated on the basis of likelihood of benefit in terms of eyesight. But some people would have less reason to fund an eyesight restoration service if it was prioritised on the basis of uses that some patients had for eyesight while others didn’t. More generally, if government services were only beneficial for people living their life according to plan A rather than plan B, people's support for government would be conditional on whether they were planning to live the rest of their life according to plan A. For government to have universal support (such as people being willing to pay their taxes), it needs to provide services that are potentially beneficial regardless of your life plan. (Bearing in mind that I might have an interest in having such services available even if I don't need them now; I am willing to pay through my taxes to maintain heart surgery centres not because I need heart surgery right now but because I may need it and I want heart surgery centres in place now to cover that contingency).

---

in a Bismarckian system of social insurance, such as operates in Germany, insurance is provided by non-governmental bodies, but insurance coverage is mandatory, so that coverage is universal. Contributions are based on ability to pay, and those who can't pay are covered by government money. The insurers are often long-established bodies who started out covering particular trades. So, such Bismarckian systems seem to provide the same benefits as publicly funded systems, but if we asked for people's intuitions in a Bismarckian system, one would expect them to account for its value in terms of the mutual insurance account.

So, on a mutual insurance conception, patients and public can legitimately insist that healthcare be allocated according to universally beneficial capabilities that all service users stand to derive, and not other benefits. For example, I don't want to be denied healthcare just because I have fewer ambitions than others, with the consequence that I will not pursue such valuable projects as others. The mutually acceptability of a universal, publicly funded healthcare system depends on the allocation criteria being articulated in terms of universally valuable dimensions of health and not in terms of such patient-specific considerations.

What if the rationale for publicly funded healthcare is understood in terms of the second account, namely a general obligation to help people in need? I offer two arguments why we must ignore non-universal benefits when the rationale is understood this way. First, a lot of non-universal benefits do not satisfy important needs, since important needs tend to be universal. We all need life years and a reasonable health-related quality of life to do anything important. But none of us needs (in the relevant sense) to read pulp fiction. Suppose I love pulp fiction, but unfortunately there are no works of pulp fiction available locally. No doubt my subjective contentment will be lower than it could be, and my preferences will be unsatisfied; but still, I can read other stuff. Eyesight is an important need, whereas pulp fiction is not. Given constrained budgets, governments should focus on important needs only. In this context, works of pulp fiction are irrelevant utilities and as such they there is no general obligation on us to deliver them.

And second, healthcare policy-makers are not well-qualified to assess the non-health benefits of their services, nor to identify the best methods of producing those benefits. They are not expert in the relevant areas. In addition, it would create duplication and confusion if policy-makers in separate areas of policy simultaneously tried to influence the same outcomes. So, to avoid inefficiency and "mission creep" in each area of policy, it is better to have a strict rule that policy-makers in each area focus exclusively on factors in their area. These will be direct benefits (benefits which do not depend on patients' subsequent choices). Such direct benefits tend to be universal.[233]

---

[233] This consideration supplements the other considerations given here (so we have more than one reason for ignoring, say, enjoyment of pulp fiction as a benefit of eyesight restoration)

Thus on either conception of the rationale for publicly funded healthcare, a focus on universally defensible benefits is defensible.

My account can explain why, when we evaluate individual treatments for cost-effectiveness, we should focus only on health impairments produced by the disease and the benefits and side-effects[234] of treatment. For example, consider the case of antidepressants for patients whose mobility is restricted. Assume that the direct effects of antidepressants are as great for mobility-restricted patients as for others (for example, in terms of energy levels, ability to concentrate, and mood). In that case, any negative impact of immobility on quality of life should be ignored for purposes of assessing antidepressants. It may be that a patient's restricted mobility means she gets fewer indirect benefits from antidepressants (such as in terms of employment - perhaps patients who recover from depression are more able to get jobs, but patients who lack mobility stand less chance of succeeding in this). But healthcare should be allocated according to the benefits we all stand to get from it, regardless of our circumstances. In the context of decisions about antidepressants, a patient's mobility and poor job prospects are as irrelevant as her life plans or her temperamental grumpiness. Thus to generalise, we should only take account of the effects of the target disease and its treatment and not the effects of unrelated diseases.

The second objection is that my account does not preclude all disability discrimination. If a disability is caused by the disease to be treated, its impact on quality of life may enter into our considerations and thus could reduce the patient's chances of funding. This includes the case where the disability is not one of the symptoms to be treated, such as where we supply a life-saving treatment for a disease which both shortens life and causes a disability. Our assessment of a life-saving treatment for the disease may take account of the disability it causes even though the disability is not something we target with treatment. In this case, my account implies that disabled patients' entitlement to life-saving operations is weaker than that of an otherwise similar patient who is non-disabled.

This is admittedly somewhat counterintuitive. However, note that my account only has this counterintuitive implication in the case of severe

---

[234] Negative health effects

disabilities. If a disability is mild, we would ignore it on the general egalitarian grounds adumbrated in Principle 2 of my account.

Also note that historically disadvantaged groups are entitled to some protection on my account. We may have reason to ignore the quality of life issue on the grounds that we ought to enhance disabled patients' sense that they are socially included and enhance social cohesion. If disabled patients are denied life-saving treatment because of their disability, that would add to disabled people's sense that they are not given equality.

Nevertheless there may be some cases of severe disabilities where my account implies that discrimination is justified, such that a severely disabled patient loses out when another patient whose prior disability is not caused by their disease does not. I need to bite the bullet on this consequence of my account and admit that the discrimination in such cases may be permissible. To try and bring intuitions round, I could offer an argument from analogy: where a disease causes pain, it is intuitively permissible to take that into account when assessing a treatment for the disease even if the treatment is not intended to treat the pain. If such discrimination is permissible, that might help to make it intuitively more plausible that disability discrimination is permissible in similar circumstances.

Nevertheless this consequence of my account may still be somewhat counterintuitive. I explain the intuitions in question as a mistake, driven by awareness of the historical injustices that disabled people have suffered and perhaps by our tendency to over-generalise from cases where disability ought to be ignored for purposes of healthcare allocation. My rejection of these intuitions is justified in terms of my method, which is to pursue reflective equilibrium: the philosopher considers intuitions regarding specific cases, develops a general theory encompassing those cases, and then modifies either the theory or her judgments regarding individual cases until the implications of the theory match her judgments regarding specific cases. This method cannot always be expected to vindicate our intuitions regarding specific cases; sometimes it will require us to give up our intuitions for sake of a principle that explains a lot of other cases and is intuitively compelling in its own right. I contend that my principle is strong enough to justify the rejection of such countervailing intuitions.

5.6 Conclusion

I conclude that it is impermissible to discriminate against patients who have a disability caused by something other than the disease we are treating or its treatment, because, when we assess the benefits of a treatment, only the effects of the disease we are treating or its treatment are relevant. We must ignore symptoms of conditions unrelated to the condition we are treating, *unless* those causally unrelated costs and benefits involve dimensions of health impaired by the disease/treatment we are treating.

Even if a disability is caused by the disease we are treating or its treatment, we may have reason to ignore the disability in evaluating the outcome. We should consider the historical disadvantages suffered by disabled people and make extra efforts to avoid any sense of exclusion.

We may also have reason to ignore a disability caused by the disease we are treating or its treatment because it is relatively trivial, and egalitarian considerations outweigh considerations to do with benefit maximisation.

However, if a serious disability is caused by the disease we are treating or its treatment, and considerations to do with historical disadvantage do not outweigh considerations of benefit maximisation, we may take account of the disability in assessing the benefits of the treatment.

6. Is it permissible to discriminate by age in healthcare allocation?

Some critics of benefit-maximising healthcare allocation policies say such policies are ageist, in that they favour the young. Elderly patients generally have fewer life years in prospect than other patients, so they stand to gain less benefit from a one-off intervention than the young, which could mean that younger patients get funding for a treatment while the older patients are denied the same treatment.[235] But some who defend age-based discrimination[236] say that such discrimination should be taken even further, for example with a "fair innings" policy which favours the young even if they will gain no more life years than the elderly.[237]

I consider this issue in the following two chapters. In the current chapter I consider an argument from Norman Daniels, who is perhaps the best known proponent of age-based discrimination.[238] He argues that in certain circumstances, age-based discrimination could be permissible.

To outline my plan, I start by summarising Daniels' account. Daniels' distinctive claim is that age-based discrimination in healthcare allocation is not necessarily unfair discrimination, and his argument for this is that since we can all hope to age, a stable system that treats different age groups differently won't necessarily treat *people* unequally. Daniels argues that if such a system treats people of the same age the same, and it is in people's lifelong interests, then it may be just.

I then address three problems that the account faces. Daniels can deal with the first and second, but the third is somewhat trickier. The problem is that Daniels accounts for the things we ought to do in healthcare in terms of the self-interest of its potential users. But some of our obligations in healthcare do not arise from self-interest. I do not reject Daniels' account of our obligations

---

[235] For example, Harris (1987)

[236] As usual, I used the term "discrimination" in a morally neutral sense according to which certain acts of discrimination may be permissible.

[237] For example, Callahan (1995)

[238] Daniels (1988) is one of the most cited works favouring age-based allocation, but others include Callahan (1995); Fleck (2010); Lockwood (1988); Nord (2005); Shaw (1994); Wagland (2012a); Williams (1997)

in healthcare, but I argue that it is incomplete; he only characterises one of our reasons for supporting healthcare. This leads me into the next chapter, where I explore a fair innings principle, giving us reason to favour the young whether or not it is in our interests.

6.1 Summary of Daniels' Prudential Lifespan Account

The conclusion that Daniels aims to defend is that age-based discrimination in the allocation of life-extending resources can be just. Specifically:

> There are conditions under which a healthcare system that rationed life-extending resources by age would be the prudent choice and therefore the choice that constituted a just or fair distribution of resources between age groups.[239]

To start my exposition with a summary, Daniels' argues as follows:

> Step 1. Age-based discrimination on the part of lifelong institutions need not be wrong in the same way as racism/sexism.
>
> Step 2. Hypothetical deliberators making choices under a veil of ignorance reveal what's relevantly prudent for us (prudent for us as actual people in the real world).
>
> Step 3. Such hypothetical deliberators would choose certain lifelong allocation policies that discriminate by age.
>
> Step 4. Therefore the discriminatory policies in question are prudent for us.
>
> Step 5. With regard to the question of justice between age groups, if an allocation policy is prudent for us, then it is just.
>
> Step 6. Conclusion: The discriminatory policies in question are just.

I now explain this argument in more detail.

Step 1. Daniels starts by refuting an objection that age-based discrimination is wrong in the same way as sexism and racism. Specifically, he aims to show that age-based discrimination need not discriminate between persons and therefore a fortiori need not constitute wrongful discrimination. As background to his argument, he starts by observing that "Important principles

---

[239] Daniels (1988), p. 91

of distributive justice prohibit our using 'morally irrelevant' traits of individuals, like race, religion or sex, as a basis for differential treatment in the distribution of important social goods."[240]

Daniels now constructs an argument on behalf of his opponent, purporting to show that age-based discrimination is wrong in the same way as racism/sexism on the basis that it satisfies the same key conditions. On my reading, the argument has the following steps:

1. When we distribute important social goods such as healthcare entitlements, important principles of distributive justice prohibit our differentiating by morally irrelevant traits such as race, religion or sex, such that people end up with unequal shares of those basic social goods.

2. Age is a morally irrelevant trait, as is recognised in age discrimination legislation. For example, our laws ban age discrimination in employment, housing and other contexts just as they ban discrimination by race and gender in those contexts.

3. Differentiating by age in healthcare allocation always means that people end up with unequal entitlement to healthcare.

4. Therefore we may not differentiate by age in healthcare allocation.[241]

Daniels thereby makes clear that his opponent's objection against age-based discrimination in healthcare allocation depends on showing that such discrimination results in different people getting unequal healthcare entitlements.[242] Daniels now aims to show it need not generate such inequalities. On the assumption that we all age,[243] if a policy of age-based discrimination is part of a stable institution that lasts a long time, it is possible for that policy to treat everyone the same. Assuming that most of us will live till we're old, policy-makers can justify such a policy by saying that what you

---

[240] Daniels (1988), p. 40

[241] Daniels (1988), p. 40-41

[242] I talk in terms of healthcare entitlements rather than healthcare, since I contend that this is the morally relevant metric for egalitarians in healthcare allocation. Different people need different amounts of healthcare, so it does not seem objectionable if they therefore receive different amounts of healthcare. But it might trigger egalitarian concerns if different people were differently entitled to healthcare, such that one person with a need gets healthcare and another person with a similar need does not.

[243] Daniels says "I here abstract from the obvious problem of early death, and I continue to work with this simplification throughout what follows." (Daniels, 2008). It may seem that an age-biased system could disadvantage the young, which could trigger an objection. But as it happens, both Daniels and I propose a bias to the young, and those who die young are not particularly disadvantaged by this aspect of the system.

lose when you're young, you gain when you're old, or vice versa.[244] Thus "Differential treatment by age, over time, is not unequal treatment of persons, even if it is unequal treatment of age groups on each occasion, at each moment.[245] So, "even if age is a morally irrelevant trait, using it in certain distributive contexts will not generate an inequality in life prospects for morally irrelevant reasons - because it generates no inequality at all."[246] In this respect, Daniels says discrimination by age contrasts with discrimination by sex and race. Since most of us do not change sex or race, a pattern of discrimination that favours one sex or race over another will always disadvantage some persons, namely persons with the disparaged trait.

The key to this argument is that we are considering lifetime wellbeing and lifelong institutions. So someone's gains at one time can compensate for their losses at another. If everyone stands to gain and lose in the same way at the same times of their life for the same reasons, they have no reasonable complaint. When we are evaluating lifelong institutions, the lifelong perspective must be the relevant perspective. In the case at hand, a policy that seems to treat the elderly worse than others actually does not.[247] The argument is that, from the long term perspective, age discrimination has been as much in the interests of the elderly as everyone else, even though they are losing out now (bearing in mind we are defending the policy in an idealised world where the policy always been in place). In contrast, race-based discrimination offers no similar compensation for the race that loses out.

So there is a key difference between age discrimination and discrimination by race or sex. However, this does not close the question of whether age-based discrimination is permissible. Daniels only takes himself to have shown that his opponent's argument

> fails to show definitively that there is no distinct problem of justice

---

[244] This is easiest to defend with a policy that favours the young, which is the policy considered in this paper; then, no-one has a complaint if they don't get old enough to see the flip side of the policy. They only ever benefited from the policy and never lost out by it.

[245] Daniels (1988), p. 42. This idealisation ignores' start-up' problems that arise when we begin such policies, as Daniels clarifies in Daniels (1989).

[246] Daniels (1988), p. 41

[247] Regarding the claim that the elderly have been treated the same as the young, I believe Daniels intends this in terms of people's entitlements, which can be understood in terms of value of people's healthcare coverage, e.g. the premium they would have to pay a private insurer to get the healthcare entitlements they have at each age. Of course, different people use different amounts of healthcare; it may be that one person needs a lot of healthcare while another never uses healthcare. But two such people are still relevantly equal if they are *entitled* to the same healthcare at each age should they need it.

> between age groups. Our banal fact means we cannot rule out unequal treatment by age simply through a direct appeal to equality in the way we can rule out unequal treatment by race or sex. The basic question remains: Which unequal treatments of age groups are just or fair?[248]

Step 2. To develop an answer to this question, Daniels' now aims to show that, in relation to the question of allocation between age groups, hypothetical deliberators making choices under a veil of ignorance reveal what's relevantly prudent for us (as actual people in the real world).

He starts by setting aside *general* questions of justice governing how resources should be distributed between persons:

> We may appeal to prudence to solve the age-group problem only if we *frame* that problem. We must constrain prudential reasoning about the age-group problem by assuming that other principles of distributive justice already govern interpersonal distributions. These principles of overall justice define the overall budget that prudent deliberators must allocate over the lifespan. We must now consider what form of prudential reasoning is appropriate within this frame.

Thus Daniels does not address general questions of justice involving interpersonal transfers. Once these general questions of justice have been set aside and each person has a notional healthcare budget,[249] a relevant consideration for policy-makers will be what allocation of those budgets is in the best interests of the beneficiaries. Here it is relevant that Daniels is considering a "policy that is stable over time" (he ignores "'start up' problems that arise when we begin such policies"[250]). This makes a difference because:

> From the perspective of institutions that distribute basic goods over our lifespan, transfers between age groups are equivalent to transfers within a life.**[251]**

So the question becomes, when in our lives do we most need healthcare? In answering this, Daniels assumes that "a time-neutral concern for wellbeing over the lifespan is one of the demands of prudence itself,"[252] on the grounds

---

248 Daniels (1988), p. 42

[249] Or perhaps better, a healthcare *insurance* budget, entitling citizens to healthcare *should they need it*.

250 Daniels (1988), p. 42-43

251 Daniels (1988), p. 67

252 Daniels (1988), p. 57. Daniels here ignores the need to take account of uncertainty. By this I mean that, in assessing the pros and cons of a plan of action, it might be rational for us to

that "what counts is the quality of our experience, not when in our lives it takes place"[253]

So the relevant conception of people's best interests involves their lifetime wellbeing. Daniels suggests we cannot understand this by considering actual people's views and preferences. The problem is that "my conception of what is good in life changes".[254] People won't take whole life view of their interests; they will be biased by their desires and circumstances at the time we ask them.

> If we set fully informed consumers the task of budgeting their fair share of health care over their lifetime, we would have to restrict them to choices made at an early point in life. Otherwise, jumping from plan to plan would lead to them exceeding their fair share. But then we seem to bias plans in favour of what the young take to be prudent and ignore the prudence of the old.[255]

Daniels solves this problem by asking us to consider the preferences of hypothetical, rational, self-interested deliberators when making their choices "from behind a veil of ignorance that keeps them from knowing their age or their conception of what is good in life".[256] Daniels argues that the features of his hypothetical deliberation are such as to ensure that the deliberators' choices will reflect our whole life interests. First, regarding deliberators' ignorance of their age, Daniels argues that this avoids bias towards the needs of one specific age group. At any given age, if we are asked to make choices while knowing our age, we are likely to be biased in our preferences towards that age. But as noted, Daniels assumes that "a time neutral concern for wellbeing over the lifespan is one of the demands of prudence itself"[257] Second, regarding deliberators' ignorance of their conception of the good, Daniels argues this ensures that they are required to keep their options open, which Daniels argues ensures they cater to the possibility that our conception

---

discount prospective harms and benefits that fall far into the future, on the grounds that we may not be around to see them. However, Daniels does touch on this point elsewhere. For example, with respect to one decision, his hypothetical deliberators are expected to take account of the fact that most of us try to get the important things done early in life, since "Most people are well aware of their mortality and construct plans in which the tasks and rewards of early and middle years are integral to their success." (Daniels 1988, p. 90)

[253] Daniels (1988), p. 56

[254] Daniels (1988), p. 57

[255] Daniels (1988), p. 63-64

[256] Daniels (1988), pp. 64-65

[257] Daniels (1988), pp. 57

of the good changes over the course of a life.[258] And third, regarding deliberators' self-interested evaluative framework, this ensures that what appears desirable to them is actually prudent for us (or at least, most of us - see below). Perhaps the prudent deliberators could be seen as ideal observers, able to see better than we can what is good for us from a long term perspective. Daniels says his veil of ignorance "gives us a way of making the reasoning about lifetime wellbeing appropriately neutral."[259]

It is worth clarifying here that although Daniels places quite Rawlsian-looking constraints on his hypothetical deliberation, he emphasises that his reasons for doing so are not Rawlsian. Firstly, although, as noted, Daniels assumes that general questions of distributive justice have been resolved before his hypothetical deliberation takes place,[260] he is not committed to a Rawlsian answer to these general questions:

> My conditional claim does not depend on the acceptability of any particular general theory of justice, such as Rawls's contractarian theory.[261]

Secondly, Daniels' justification for his constraints on his hypothetical deliberators is not Rawlsian. Daniels suggests that Rawls' justification is to do with fairness.[262] In contrast, with regard to his own account, Daniels says "The reasons for appealing to this veiled form of prudence derive ... from the requirements of prudence alone".[263] He also says his prudent deliberators "cannot attempt to solve problems of justice which cross the boundaries between persons".[264] It is critical to understand that Daniels uses his hypothetical deliberation device for different purposes from Rawls; specifically, to reveal what's in our interests, from a whole life perspective.

Step 3. Daniels now argues that such hypothetical deliberators would choose certain lifelong allocation policies that discriminate by age. He considers two schemes for rationing a selected range of high cost treatments, Scheme A (A for ageist) and scheme L (L for lottery). One scheme is based on age

---

[258] Daniels (1988), pp. 54-60

[259] Daniels (1988), p. 64

[260] For example, he says "these deliberators must already know that the basic goods being distributed constitute a fair or just share, that is, that more general principles of distributive justice already solve problems of distribution between persons." (Daniels 1988, p. 67)

[261] Daniels (1988), p. 71

[262] Daniels (1988), p. 62

[263] Daniels (1988), p. 64

[264] Daniels (1988), p. 62

discrimination and the other is not. Scheme A says that no-one over 75 can have any of the treatments on the list. Scheme L distributes the treatments by lottery.[265] The pros and cons are that scheme L reduces the chance that the young will reach old age, but in return ensures that once anyone reaches the age of 75 they stand a higher chance of getting older. To eliminate benefit-maximising considerations, Daniels stipulates that "Both yield the same expected lifespan, but they do so differentially".[266] One option gives a high chance of a middling life expectancy (for example, everyone lives to 75 and then dies), the other a lower chance of a longer life expectancy (for example, half die at 55 and half die at 95). On expected utility theory, whereby options are assessed by weighting their possible outcomes by the probability of the outcomes, both options could yield the same expectation. Daniels argues that prudent deliberators would prefer an age rationing scheme to a lottery. To show this, he considers two alternative decision rules which they might adopt. One is Rawls' maximin rule, which tells us to make the worst outcome as good as possible.[267] Daniels thinks it is clear that the worst outcome is dying young, and scheme A minimises the likelihood of that outcome, making it the best scheme under this decision rule. Daniels next considers an alternative decision rule, "the standard rule" based on expected utility, requiring prudent deliberators to multiply the value of each possible outcome by its likelihood, and then to sum all the resulting probability-adjusted values. If utility is constituted by bare life years, the two schemes tie under this decision rule, because the average expected life expectancy under the two schemes is the same. Daniels thinks even this result is interesting, since it shows that age rationing is not imprudent for the deliberators, indicating that age discrimination is not necessarily unjust. However Daniels thinks this conclusion is conservative, since it does not take account of the fact that our later years are likely to be impaired by disabilities and illness. Once deliberators take this into account, scheme A is more prudent from the perspective of the standard rule. The deliberators will want to give life years to younger people who will enjoy them in full health rather than to older patients who will not.

---

[265] An alternative scheme L just offers one of the high cost treatments to everyone who needs it leaving the other treatments completely unfunded.

[266] Daniels (1988), pp. 87

[267] Daniels argues this might be appropriate for circumstances in which there is genuine uncertainty regarding the possible outcomes, rather than quantified probabilities; also for when the worst outcomes are so grave that they cannot merely be weighed against better outcomes. (p. 88)

Step 4. Daniels can now conclude from the above that the policies in question are prudent for us:

> My conclusion from these versions of the prudential argument is that there are conditions under which a healthcare system that rationed life-extending resources by age would be the prudent choice ...[268]

Step 5. Daniels now needs to show that, with regard to the question of distributive justice between age groups, if an allocation policy is prudent for us (as actual people in the real world), then it is just. We could not rely on any such conditional in other areas of distributive justice (those involving interpersonal transfers), but Daniels thinks we can rely on it with respect to the question of justice between age groups (where discrimination needn't be understood as involving interpersonal transfers). His argument is based on the same observation as his defence of age-based discrimination against accusations of wrongful discrimination, namely the observation that we all age:

> To find out what is just between age groups, we must seek principles to govern the design of the institutions that distribute goods to us over our lifespan. From the perspective of such institutions, transfers between age groups appear as transfers between the stages of a life, not between persons. The shift in perspective I have been urging thus means we should not seek the typical principles of justice which govern distribution between competing individuals or groups. Rather, we must seek principles governing allocations within a life. The facts that this shift is plausible in the case of age groups is what makes it a distinct problem of distributive justice. For it, but not for problems of distributive justice in general, prudence is a safe guide to justice.[269]

Step 6: Daniels infers from this that a policy of age-based discrimination could be just:

> There are conditions under which a healthcare system that rationed life-extending resources by age would be the prudent choice and therefore the choice that constituted a just or fair distribution of resources between age groups.[270]

---

268 Daniels (1988), p. 91

269 Daniels (1988), p. 47

270 Daniels (1988), p. 91

Before I go on to critique aspects of Daniels' account, it bears repeating at this point that although Daniels' hypothetical deliberation looks Rawlsian, his claim "does not depend on the acceptability of any particular general theory of justice, such as Rawls's contractarian theory."[271] And while Rawls' justifies his constraints on his hypothetical deliberators in terms of fairness (on Daniels' reading),[272] Daniels says of his own account that "The reasons for appealing to this veiled form of prudence derive ... from the requirements of prudence alone".[273] He explains:

> My use of Rawlsian devices does not depend on an appeal to is robust Kantian account of the nature of persons or to his claims that the choice situation is procedurally fair to such persons ... I have argued that prudence itself (under the standard assumptions) requires that individuals respect their own changes in their conception of what is good at each stage in life. Their concern for their own lifetime wellbeing will require them to abstract from full information in order to be neutral about each stage of their lives, at least when they are considering the design of institutions that affect them over the whole lifespan.[274]

In summary, it is important to be clear that Daniels uses his hypothetical deliberation device for very different purposes from Rawls; his purpose is to reveal what's in our interests, from a whole life perspective. This is why Daniels' account is best known as the "prudential lifespan account." Although Rawls' deliberators select social arrangements that are in their interests, Rawls wouldn't claim that these arrangements are necessarily in each of our interests in the real world, or even in the majority's interests.

6.2 Issues for Daniels

I now address a couple of problems that this account faces. Daniels can deal with the first, but the second is somewhat trickier.

First, it may be objected that Daniels' hypothetical deliberators are not a good guide to people's interests in a lifelong sense. For example, perhaps

---

[271] Daniels (1988), p. 71

[272] Daniels (1988), p. 62

[273] Daniels (1988), p. 64

[274] Daniels (1988), p. 62

hypothetical deliberators could not know what's good for an individual who has certain values without having those values themselves.

But if some such objection turns out to be fatal for Daniels' account of prudence, I believe Daniels' most distinctive and interesting claim and argument can still be defended. His distinctive claim is that age-based discrimination in healthcare allocation is not necessarily unfair discrimination, and his argument for this is that if a healthcare allocation system promotes the lifelong wellbeing of healthcare users, such discrimination could be justifiable. It is this claim and the supporting argument that represents Daniels' distinctive contribution to the debate on healthcare allocation. The argument depends on the assertion that policy-makers should take a whole life perspective on their constituents' wellbeing, but it does not depend on construing that whole life wellbeing in terms of the preferences of hypothetical deliberators. So if Daniels' account in terms of hypothetical deliberators is not defensible, Daniels can abandon that account and still make his distinctive claim based on a different account of whole life wellbeing. To demonstrate the claim, without appealing to the argument being challenged here, Daniels needs to show that either "a time neutral concern for wellbeing over the lifespan is one of the demands of prudence itself"[275], or that for some other reason, "institutions that distribute basic goods over our lifespan"[276] must ignore the bias that any single user might display towards their own age group. Once one of these assumptions is accepted, Daniels can make at least a prima facie case for a policy of age-based discrimination that treats people equally across their lifetimes, on the grounds that it does not discriminate between persons. So it will not be fatal to Daniels' case if he has to abandon his claim that prudence must be understood in terms of the choices of hypothetical deliberators, as long as he can hang onto one of these assumptions.

But second, I will argue that citizens' self-interested concerns do not explain everything a healthcare system ought to do. For example, consider people who already have diseases. Specifically, consider genetic diseases, such that there is no age at which sufferers don't have them (sufferers have these diseases their entire lives), and no age at which non-sufferers have them (they're not acquired). For someone with such a disease, it is in their interests that the healthcare system prioritises that disease, but for someone who

---

[275] Daniels (1988), pp. 57

[276] Daniels (2008), p. 67

doesn't, it is not. At any one time only a small minority of the population has any genetic disease, even when all the patients with any genetic disease are aggregated into a single group. To sharpen up the case, let's consider diseases that have their effects in the young, as many genetic diseases do (though I think other kinds of diseases are implicated by this line of argument too). Suppose it is not possible to formulate a policy which designates these diseases individually; for example, variants are appearing all the time, and although a sufferer might realise they have one of the diseases of this type, it is not possible to diagnose or treat it until it develops. So the only way of accommodating these patients in healthcare policy or legislation is by designating them as patients who would die young without treatment. Also suppose these diseases are expensive to treat. Finally suppose that people are likely to die young if they have one of these diseases, but they are unlikely to die young otherwise.

I suggest the case could be such that intuitively, the healthcare system should make special provision for these younger patients. But the interests of this minority diverge from the interests of the majority. The majority has no interest in relaxing the cost-effectiveness threshold for younger patients. This makes it difficult for Daniels to generate the intuitively satisfying result. The problem is that an intuitively satisfying policy involves interpersonal transfers from those who do not have such diseases to those who do, so Daniels' justification in terms of the interests of (most) healthcare users cannot be applied here (Daniels is explicit that his prudential lifespan account does not have implications regarding interpersonal transfers).

Perhaps in response, Daniels can modify the setup of his hypothetical deliberation such that his deliberators choose a healthcare allocation policy that accommodates those with such genetic diseases. For example, he could make his hypothetical contractors ignorant of what diseases they have. Also, to avoid deliberators being influenced by the low probability that they have one of the genetic diseases, Daniels could stipulate that they don't even know the prevalence of the diseases. This would parallel Rawls stipulation that his deliberators don't know the probability that that they will be among the worst-off, which led to them adopting a maximin rule, telling us to make the worst outcome as good as possible. In the case of Daniels' deliberation, this might be enough to ensure that those with short lives get priority.

But there are three problems with this solution. First, it would require Daniels to deprive his deliberators of basic epidemiological knowledge, which would make it difficult to make any sensible decisions about what kind of healthcare system to adopt. Second, if his deliberators know so little, it is difficult to see them playing the role Daniels had them playing as ideal observers who reveal the whole life prudential interests of the majority of the actual population (i.e. real citizens). In fact, if deliberators choose a policy which diverts healthcare cash to the unfortunate minority, they will be going against the interests of most of the actual population. This is not an argument against the policy (I will argue that such policies don't need justifying in terms of the interests of the actual population). But it means Daniels' original line of argument is no longer viable; he cannot defend funding healthcare for sufferers of these genetic diseases in terms of the interests of the actual population. And third, whereas he was able to give independent justification for his original setup in terms of getting hypothetical deliberators to see our whole life interests, he would now seem to be tweaking his setup to get the answers he wants, in which case the device of the hypothetical deliberation would no longer seem to be providing any independent support for his desired conclusions.

Another answer for Daniels would be to bite the bullet and say that, where different people's interests diverge, the democratic answer is for the system to serve the interests of the majority. Sufferers of these genetic diseases must go unfunded, despite their short life expectancy. But that is a counterintuitive conclusion. And I don't think it's necessary. We can just say that the majority has an ethical obligation to this minority. The obligation does not arise from self-interest, but self-interest is not the only reason we provide healthcare to each other. I argued in Chapter 3 that there are two accounts of the rationale for publicly funded healthcare. The first is in terms of a general obligation to help people in need; governments are entitled to tax their citizens to fulfil citizens' obligations on their behalf. The second accounted for publicly funded healthcare as a kind of mutual insurance fund. On this model, we all make affordably small, regular contributions to a mutual fund in return for the ability to make occasional big withdrawals as needed. I suggest that Daniels' defence of age-based policies is based on a view of healthcare in terms of the mutual insurance model; healthcare priorities are explained in terms of our interests. Although Daniels reaches an ethical conclusion, i.e. the conclusion that an ageist policy could be just, the justice of such a policy is ultimately founded on its being in the interests of the population it serves. I do not deny that this is a valid explanation of one of our reasons for supporting a publicly funded healthcare system, and one basis for healthcare policy-makers to determine the system's priorities. But perhaps there are some priorities we ought to support in healthcare even if they are not in our interests (perhaps these arise from Daniels' background general theory of justice).

6.3 Conclusion

I conclude that Daniels has characterised one reason we might have for favouring the young in healthcare allocation, this being that it is better for us in terms of our life interests. But this is only one pro tanto consideration, and we also may have other reasons for favouring the young.

In the next chapter, I explore whether there's an argument that we should favour the young even if it is not in the interests of the majority. In other words, a justification of favouring the young that involves intrapersonal transfers. I continue to take Daniels' whole life perspective on our interests. However, our whole life interests are deployed as a metric in a different principle which we do not see in Daniels, a principle from the egalitarian/prioritarian family. Specifically, I defend a fair innings principle according to which anyone expected to die young merits priority for life-extending treatment in order to help bring them up to society's average life expectancy. Thus I introduce distributive considerations which might not be recognised by his hypothetical deliberators, and sometimes these considerations may trump the consideration mentioned by Daniels. But I argue that we should view the age-related consideration highlighted by Daniels as just one pro tanto consideration among several. So this potential conflict between my principle and Daniels' principle should not be seen as problematic for my overall position. I present this fair innings principle as an additional pro tanto consideration in my pluralist theory.

7. Fair Innings: Should we favour the young over the old in healthcare allocation?

We have seen Daniels give one defence of age-based discrimination[277] in healthcare allocation. In this chapter I will give another. The starting point for my discussion is the existing principles that are defended in the literature, which say that we should treat "the young before the old."[278] This implies that the young should be favoured even if they will gain fewer life years than the elderly. This is based on intuitions that, for example, all else being equal, it seems right to give 5 years to a 40 year-old rather than 6 years to a 70 year-old, even though the older patient gets more benefit in terms of life years.

Such principles also seem to chime with ordinary intuitions. Aki Tsuchiya reviews a number of studies of the public's age-related preferences done in Japan, The Netherlands, USA, and Australia. The surveys consistently show that the majority of the public is clearly in favour of direct age-based discrimination.[279] Figure 4 maps the results from a number of studies, showing how the strength of a patient's claim to an extra life year is held to decline as patient age increases (each line has been derived from a separate study). Greg Bognar also reviews a number of studies suggesting that the public supports policies favouring younger people.[280]

---

[277] I will use "age-based discrimination" as a catch-all term for all discrimination by age, whether the patient's current age or their expected age at death. As usual, I use "discrimination" in a morally neutral way, such that a discriminatory act is not necessarily impermissible.

[278] Kappel & Sandoe (1992), p. 314

[279] Tsuchiya (1999).

[280] Bognar (2015), note 3.

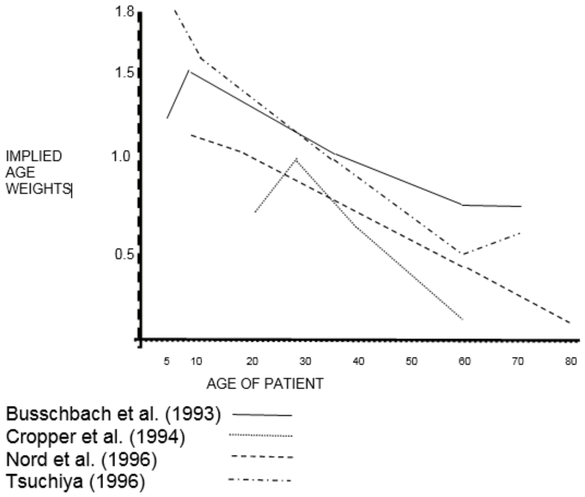

Fig 4. Public opinion regarding age weights[281]

Age-based discrimination also persists in healthcare despite the existence of laws against ageism. For example, increasing age is associated with "markedly decreased rates of histological verification, surgery and chemotherapy"[282] Also, "In the UK there has been a longstanding and effective policy of covert age-based rationing of dialysis"[283].

I aim to defend a form of age-based discrimination. Actually I consider two new principles in succession. First I consider a fair innings principle according to which anyone who, without treatment, is expected to die younger than society's average life expectancy merits priority for life-extending treatment in order to help bring them up to society's average. Then I defend a generalisation of this principle, concluding that anyone with a *quality-adjusted* life expectancy lower than society's average merits priority for any treatment to help bring them up to society's average quality-adjusted life expectancy.

To outline what follows, I start by clarifying some of the claims which fair proponents of age-based discrimination need *not* be committed to. For example, they need not be committed to discriminating directly by the patient's current age, as has traditionally been assumed. This could be seen as wrongfully ageist (and presumably for this reason would be illegal in many jurisdictions). But I will argue that, given the justifications that have been offered for age-based discrimination, the morally relevant consideration is not

---

[281] Graph constructed using data from Tsuchiya (1999)

[282] Austin & Russell (2003)

[283] Sayers & Nesbitt (1998)

the patient's current age, but the age at which they are expected to die (which is why the principle I defend is articulated in terms of when people are expected to die, not how long they have lived already). So the fair innings principles that have been defended miss the point, in that they talk in terms of the patient's current age rather than the age at which the patient is expected to die. I then consider a counterexample to age-based discrimination raised by John Harris. Addressing this case requires us to distinguish two types of principle: "straight line" principles according to which the strength of your claim to life-extending treatments declines at a steady rate according to your expected age at death, and threshold or "stepped" principles according to which your entitlement declines faster as your expected age at death rises past a certain age, such as 70. I adopt a threshold principle (a "fair innings" principle) which largely avoids the counterintuitive implications in Harris's case. Specifically, I defend a view according to which the fair innings in a given society is equal to that society's average life expectancy. However the real work is not characterising the principle, but defending it, so as to avoid being ad hoc. I do this with an innovative set of cases in which society's average life expectancy is varied to show its importance for intuitions. I also offer a new set of egalitarian arguments to defend the same principle. Finally, I consider whether a similar principle can be applied to QALYs or their equivalent, and not just life years. I consider a QALY-based proposal from Alan Williams, noticing that it has some counterintuitive consequences. However, I find that if Williams accepts some relatively painless constraints on his weighting factors, he can dodge the counterintuitive consequences. I therefore conclude by endorsing his proposal.

To clarify the difference between this chapter and the last one, Daniels' defence of age-based policies was based on our interests as potential users of the healthcare system. It is better for us that our healthcare entitlements increase or decrease at certain ages. Although Daniels reaches an ethical conclusion, namely the conclusion that an ageist policy could be just, the justice of such a policy is ultimately founded on its being in the interests of the individuals it serves, on Daniels' view. I now aim to show that this is not the only way of justifying age-based policies, because self-interest is not the only reason we provide healthcare to each other. According to the argument I now give, policy-makers are obliged to prioritise the needs of the worst-off in our society, and those expected to die young are among the worst-off. So we have moral reasons to give more weight to their healthcare needs.

7.1 Defences of Age-Based Discrimination

How can we justify age-based discrimination without appealing to self-interest? I will quickly consider a couple of arguments which avoid a Daniels-style appeal (at bottom) to citizens' self-interest. First, Daniel Callahan offers an argument based on fair opportunity. In outline, he argues:

> our ideal of old age should be achieving a life span that enables each of us to accomplish the ordinary scope of possibilities that life affords, recognising that this may encompass a range of time rather than pointing to a precise age. On the basis of that ideal, the aged would need only those resources which would allow them a solid chance to live that long and, once they had passed that stage, to finish out their years free of pain and avoidable suffering.[284]

Second, Klemens Kappel & Peter Sandoe offer an argument based on egalitarian considerations:

> Is it really from the point of view of justice the same whether we give the liver to the young person or to the old? We think that the answer must be, No. If we give the liver to the older person, he will get another ten years on top of those 60 that he has already got. And then he will end up with 70 life years. Whereas if we give the liver to the younger person he will only end up with a total of 30 life years. To give the liver to the older person rather than the younger is like giving money to the rich instead of the poor.[285]

They then offer one plausible distributive principle and show its implications for the case at hand. They think the best measure of whether people are relevantly equal is whether we would be indifferent between their lives if we had to choose. Of course, healthcare policy-makers can't make everyone completely equal in that sense, but by giving the liver to the young person, we get closer to equality between the young person and the old person than if we give the liver to the old person:

> To treat people equally can plausibly be spelled out to mean that we should distribute resources among them in a way so that when in turn imagining their respective situations we get as close as possible towards being indifferent between being in the shoes of one person or being in the

---

[284] Callahan (1995), p. 135

[285] Kappel & Sandoe (1992), p. 313

> shoes of the other. In the described situation with the young patient and the older patient and only one available liver transplant we would get as close as possible to being indifferent by giving the transplant to the young person ... Therefore, if we want to treat people equally, we should in the case under consideration (other things being equal) give the liver to the young patient.[286]

I will return to the question of how we justify age-based discrimination later, looking at the implications of our justification for the precise shape of the principle. But first I will address a basic objection.

7.2 Objection: Direct Discrimination by Age Is Impermissible

I now address the objection that the policies of age-based discrimination that have been defended are directly ageist, in that the age of patients is mentioned in the allocation criteria.[287] In virtue of this direct discrimination, such principles contrast with QALY-based discrimination where the discrimination against the elderly is only indirect.

Defending direct discrimination will generally be harder than a defence of indirect discrimination. The law is helpful in highlighting the morally relevant issues in this area (it isn't always so helpful, of course). The UK Equality Act says indirect discrimination may be legal if there is a legitimate reason for it.[288] For example, if it is not possible to do a job without travelling at short notice, an employer may require its employees to travel at short notice even though this may indirectly discriminate against primary child carers, who tend to be women. Similarly, allocating healthcare to maximise benefit can be justified even though this may sometimes involve indirect discrimination against the elderly. In contrast, the law would usually prohibit direct discrimination against women or the elderly, even if such discrimination served some legitimate requirement (such as short notice travel). I presume this is because, to the extent that organisations have legitimate requirements, there will be no good reason why they can't discriminate by people's

---

[286] Kappel & Sandoe (1992), p. 314

[287] For example, see Kappel and Sandøe (1992), discussed below. But direct discrimination is not a necessary feature of fair innings principles, as we will see.

[288] UK Citizens Advice (2017)

performance on those requirements rather than by proxy measures, allowing people with protected traits who satisfy the requirements to apply. In the case of historically disadvantaged traits such as gender and age, this has led to a presumption in law that direct discrimination by the traits is impermissible unless proven otherwise, to minimise the chance of historical attitudes and practices being perpetuated. In line with this legislative presumption, the ethical worry about any proposal to directly discriminate against the elderly in healthcare allocation would be that it is motivated by discredited attitudes. So there is an extra burden of proof on proponents of any such proposal to show that it serves some legitimate purpose.

However, I argue that policies of age-based discrimination need not be directly ageist; in fact, they should not be, given the underlying justifications that have been offered for them. To see this, consider Kappel & Sandoe's argument that: "To give the liver to the older person rather than the younger is like giving money to the rich instead of the poor"[289] and their principle that we should distribute resources in such a way that "we get as close as possible towards being indifferent between being in the shoes of one person or being in the shoes of the other." I understand the metric of Kappel & Sandoe's principle to be whole life welfare, or opportunity for welfare, on the assumption that we will be indifferent between whole lives if and only if we can expect the same welfare or opportunity for welfare in those whole lives. But on this picture, the fair innings principles that have been defended miss the point, in that they talk in terms of the patient's current age. Kappel and Sandoe themselves make this mistake. They say: "QALYs are not ageist enough. ... we shall argue that considerations of fairness speak in favour of treating the young before the old".[290] If the point is whole life welfare, and if the older patient is expected to die at a younger age than the younger patient, then the younger patient would be "richer" in the way they describe, and so contrary to their assumption, fair innings considerations might justify treating the old before the young.

To give a specific example, and to try to capture intuitions, suppose that we must choose which of two patients to treat. We know that we could extend the life of either by ten years. The first patient is 30 years old; she has multiple sclerosis; without treatment, she is expected to live until she is 70, and our treatment would enable her to live until she was 80. The second patient is 40 years old; he has motor neurone disease; without treatment, he is expected to

---

[289] Kappel & Sandoe (1992)

[290] Kappel & Sandoe (1992), p. 314

die at 50, and our treatment will enable him to live until he is 60. Let's suppose that other factors (such as quality of life factors) are equal. I suggest that we have some reason to prioritise the older patient, on the grounds that he will die younger without treatment. This tallies with Kappel and Sandoe's analogy: in their terms, the older patient is "poorer", and this gives us some reason to treat him. Even if a patient is younger, if we're sure they will get to live a long life even without our help, then it would seem unfair to prioritise them over an older patient who we know will die young as a result of our denying them life-extending treatment.

Also consider Callahan's argument that everyone should be able to "accomplish the ordinary scope of possibilities that life affords." Once again, if there is an age at which we can say people have had a reasonable chance to exploit life's possibilities, and if policy-makers are obliged to consider whether patients will have had some such chance, then the relevant age for policy-makers to consider is not the age at which patients present for treatment but the age at which they are expected to die.

I therefore suggest that our principle of age-based discrimination, if we are to adopt one, should be expressed not in terms of the patient's current age but in terms of the patient's expected age at death. I suggest that historically, debate on this question has been confused by the fact that the age and the expected age at death of patients often correlate. As a result these two factors have been conflated in some discussions. Specifically, patients who need life-extending treatment will often be expected to die soon without treatment; for example, as a general rule, if a patient is going to die young, she will learn about it only a few months or years before her expected age at death. So if she is expected to die in her 20s she will learn about it in her 20s or just before, but if she is expected to die in her 40s, she will learn about it in her 40s or just before. These factors are presumably where the historical focus on age has come from.

This helps us deal with the objection that fair innings principles are directly ageist. In fact, they need not be. For example, Kappel & Sandoe's distributive principle is only indirectly ageist, in that although it will often favour the young, age is not the metric of the principle. This makes the principle somewhat easier to defend. Assuming that, in line with UK law, discrimination is morally permissible as long as it is the consequence of pursuit of a legitimate aim, and assuming that the pursuit of equality of lifetime welfare is a legitimate aim, then Kappel & Sandoe can argue that the age discrimination implied by their principle is not morally wrong. In fact, a policy expressed in terms of expected age at death is as defensible against the "ageist" objection as QALY-based allocation, for roughly the same reason: it is indirectly ageist, but not directly ageist.

It might be objected against my proposal that we should go on the patient's current age because we can't be sure when younger patients will die. Just because a 30 year-old is healthy, we can't assume she will live until she is 70; she might catch a fatal infection at 40. In contrast, we can be sure a 70 year-old will have had 70 years of life before he dies. So, even accepting my principle, this might suggest we should go on the patient's current age, just as traditional fair innings principles have implied. In response, I accept we should try to take account of such "unpredictable" causes of death in our estimates of a young patient's life expectancy. But ordinary estimates of life expectancy do exactly that; they factor in the chance of the patient being hit at an early age by fatal illnesses, fatal accidents etc. This is why our expected age at death gets higher as we get older, because we have evaded some of the things that might have killed us when we were younger. So by going on ordinary estimates of life expectancy to determine who should get priority, we take account of the kind of factors the objector is pointing to.

Before addressing a more substantial objection from John Harris, we should clarify one other thing to do with the metric of the principle I will ultimately defend. Although initially I will focus on life years, by the end of the chapter I will extend this to defend an egalitarianism of quality-adjusted life expectancy. Now, the metric of this is valuable health outcomes, where the value of health outcomes is understood in terms of how they constrain the possibilities of living well (the same metric as I defended in Chapter 3 on how to understand benefit). I thereby defend a focus on opportunities for welfare rather than welfare itself, and a focus specifically on health-based capabilities rather than opportunities made available by people's non-health circumstances. This account positions me somewhat in the vicinity of Callahan with his metric of "the possibilities that life affords". But it puts me at more distance from Kappel & Sandoe, at least if we understand their metric as whole life welfare. Nevertheless we can grant that for the rational patient, the ultimate point of healthcare would be not just health but whole life welfare. So, why don't we just take whole life welfare as our egalitarian metric? There are two restrictions in my account, a restriction to opportunities for welfare rather than welfare itself, and a restriction to health-related capabilities rather than other non-health-based opportunities. I take each restriction in turn. First, why do we focus on opportunities for welfare rather than welfare itself? Here I reiterate Hausman's conclusion, given in Chapter 3: "Government serves as a referee, a protector, a facilitator, and an insurer, not as a big brother".[291] I offered several arguments in defence of the point that governments should not try to provide welfare directly: it would create moral hazard; it would detract from citizens' autonomy; it would be democratically unwieldy; and it would be more expensive than letting individuals manage their own welfare. Second, why do we only focus on health-related capabilities rather than other non-health-based opportunities? In Chapter 5 on disability discrimination, I gave an argument in terms of separate spheres. One of the rationales for a publicly funded health system is that it is a mutual insurance system. On this model, Government should deliver *universally beneficial* resources or capabilities, because citizens can only be expected to support a mutual insurance fund to the extent that it delivers the kind of benefits that they might need, *regardless of what they are doing with their life*. We all want the capabilities that health gives us, but we do not all use those capabilities the same way.

7.3 Objection: Problems for a Non-Threshold Principle of Age-Based Discrimination: Counterintuitive Discrimination Between Adjacent Age Groups

However, now we face a more serious objection from John Harris, which requires us to consider exactly what our distributive principle should say. But before I go into detail, I should note in passing that Harris has a confusing sequence of positions on age-based discrimination; he starts by opposing it, then accepts a qualified form of it, and then in later works rejects it again. In the following I will consider his early rejection; I will also later consider an argument he offers in defence of a revised principle.[292] So, to start with his early objection to age-based discrimination, he says:

> Imagine a group of people all of an age, say a class of students all in their mid-20s. If fire trapped all in the lecture theatre and only twenty could be rescued in time should the rescuers shout "Youngest first!"? Suppose they had time to debate the question or had been debating it "academically"

---

[291] Hausman (2010)

[292] I will ignore his later arguments, which are dealt with adequately elsewhere; for example, in Kappel and Sandoe (1994).

> before the fire? It would surely seem invidious to deny some what all value so dearly merely because of an accident of birth?[293]

This case threatens to undermine the intuitions which motivated age-based discrimination in the first place. It's all very well to discriminate in favour of someone who will die at 20 over someone who will die at 90. But as Harris suggests, if those who will die without our help are close in age, the intuitive basis for discriminating in favour of the younger ones is much weaker.

Having said, that there are a number of confounding factors in this case which make Harris's "anti-ageist" conclusion more difficult to draw than he thinks. Nevertheless there is a way of reconstructing the case so as to avoid these. I will now run through the confounding factors, before reconstructing the case accordingly. I will then respond to it with a modified principle of age-based discrimination.

The first confounding factor is that, even if there is no intuitive distinction between, say, 25 year-olds and 26 year-olds in Harris's situation, this is not a counterexample against age-based discrimination. In summary, the proponent of age-based discrimination can argue that this difference is so small, other considerations could easily outweigh this one. As an analogy, consider that in Chapter 4 on rare diseases, Lawlor offered a pluralist position according to which "there is some value to saving the greatest number" but also "there is some value to giving each person an equal chance of survival" (the consideration that Taurek pointed to). Although the amount of benefit we can produce is a pro tanto consideration, another is that we should express our "equal concern and respect for each person" by giving all patients *some* share of healthcare resources, even if they will not gain as much benefit from it as others. If benefit maximising considerations are weak, then egalitarian considerations may take precedence. This does not mean that we may never try to maximise benefit. Similarly here, it could be said that although age-related considerations give us some reason to save the younger students, that reason is very weak (especially since these students face a sudden and unexpected death, an important consideration in its own right as argued in Chapter 8; surely this is a much more important consideration than the fact that they are of slightly different ages). The upshot is that perhaps we have a

---

[293] Harris (1985), p. 89. It should be noted that this counterexample is specifically targeted against "straight line" principles, according to which the strength of our claim to life-extending treatments declines at a steady rate according to how long we can expect to have lived at our death.

stronger reason to give each student an equal chance of survival than to discriminate by age. So the fair innings theorist can agree with Harris that we should give each person an equal chance of survival in this case, without conceding that we may never take account of someone's expected age at death.

It is also relevant that Harris asks us to imagine the rescuers asking certain people to hold back and let others go forward. In a situation of life and death, this is too much to ask of most people. Perhaps ought implies can, so if people's instincts prevent them from sacrificing themselves for others, perhaps we must say that there is no obligation to do so. In contrast, when we refuse a cancer treatment to a patient because it is way too expensive for the sake of a month's extra life, we do not *ask* the patient to forgo it; we just refuse to offer it. So there is no similar demand for self-sacrifice.

So our intuitive resistance to selection by age in Harris's case can be explained by other factors. But perhaps Harris can eliminate these confounding factors and get a similar result. In fact perhaps he can return to the case of healthcare allocation. Suppose healthcare allocators must choose between treating a group of 20 year-olds who without treatment will die in 5 years time (at the age of 25) and another group of 20 year-olds who without treatment will die in 15 years time (at the age of 35).[294] Suppose that treating either group will add 20 years to their respective life expectancies. The age-based principle we have been considering implies that in any case where, without treatment, one group of patients is expected to die younger than another group of patients, then that first group should be prioritised for life-extending treatment. This would imply treating the group who would die younger without treatment, viz., those who would die in 5 years time. But perhaps this seems somewhat counterintuitive; intuitively, there is little difference between these two groups in terms of the strength of their respective claims to extra life years. Although there is a significant difference in the life expectancy of each group without treatment, the important consideration is that both groups will die very young without treatment. So our age-based principle has counterintuitive implications. And this cannot be explained with the same confounding factors as Harris's original case. Are there other ways we can save my age-related principle here? The only contender for a supplementary consideration to help explain our intuitions is

---

[294] I hypothesise cases with these longer time scales to eliminate suddenness and urgency, which introduce their own considerations, as I show in the next chapter.

a Taurek-style egalitarianism. But if such egalitarian considerations outweigh our age-related principle in this case, then our age-related principle does not look strong enough to do any work in the other cases where it is needed; it looks too weak to vindicate some of the intuitions that motivated it. For example, it would not support the intuition that patients who will die at 60 should be treated ahead of those who will die at 80. So, if instead we can modify the principle to deal with this case in a more direct manner, we can hope for a more intuitively satisfying principle.

7.4 A Modified Principle Of Age-Based Discrimination: Threshold Principle

In response to this difficulty, I will now consider a modification of our principle of age-based discrimination which may help address some of Harris's cases, or at least reduce the number of problematic cases. The modification is to introduce a threshold age, such as 70, at which entitlement declines faster than at other ages (all else being equal). To compare the alternative principles, lets consider some graphs mapping a patient's expected age of death against the strength of their claim to an extra life year.[295] On the original proposal, there was no threshold, so the graph has a straight line which slopes down from left to right, representing the gradual decline of entitlement with expected age at death:

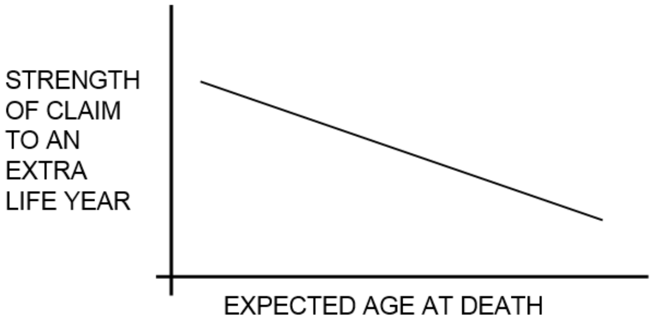

Fig 5. Straight line principle

---

[295] Throughout this chapter, in line with my pluralist framework, I presuppose that benefit maximisation is a pro tanto consideration. Here, in mapping patient age against the strength of a patient's claim to an extra life year, I have assumed that we should prioritise treatments that produce more life years, all else being equal. Later I will consider a proposal from Williams which incorporates quality of life considerations, relating quality-adjusted life expectancy to strength of claim to a QALY.

Here, each extra year of life expectancy is treated as if it is worth less than the previous one. For example, a treatment that takes someone's expected age at death from 30 to 31 is treated as if it produces less benefit than one that takes someone's expected age at death from 20 to 21. This principle has the problem with Harris's counterexamples involving youngsters of a similar age. So instead of this, consider an alternative principle according to which entitlement declines more suddenly at a certain age, such as 70, than at others:

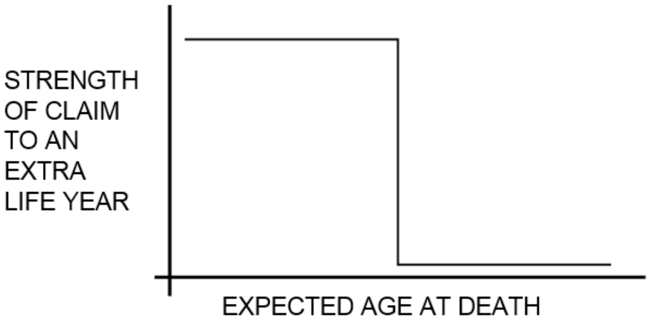

Fig 6. Sharp threshold principle ("fair innings")

This step-shaped line represents a threshold principle, according to which principle, we should prioritise those whose expected age at death is below the threshold over those above the threshold. Assuming the step is at age 70, this principle would say that if your expected age at death is less than 70, any life years added by treatment get up-weighted; a year gets treated as if it is more than a year. But if your expected age at death is greater than 70, no weighting is applied. (And if a treatment takes your expected age at death from, say, 68 to 72, then your first two years will get up-weighted and the second two years will not). If we must choose between Harris's students in their 20s, this line implies no difference in weighting; every added life year gets the same weight, whether you're in your early 20s or late 20s.

However, the above graph implies a sudden cut-off around the threshold age, according to which a treatment adding, say, 3 months life expectancy would get much less weight for a patient expected to die at 70 than for a patient expected to die at 69 years and 9 months. It would be counterintuitive to discriminate so harshly between two patients whose expected ages at death are so close. Also, given our overall justification for age-based discrimination, there may be nothing to justify such a sudden cut-off in the morally relevant facts associated with different ages at death. On these grounds we might want to adopt a compromise proposal, according to which the strength of one's claim always declines with expected age at death, but the decline accelerates over a certain period:

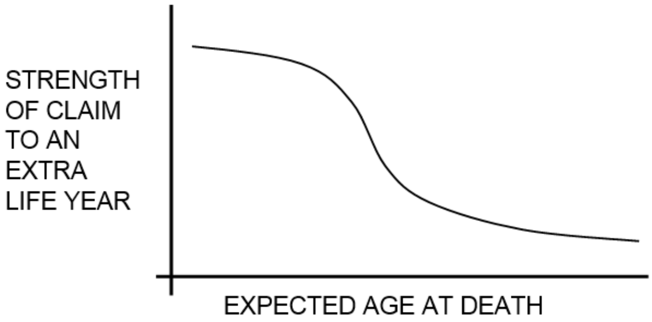

Fig 7. Graduated threshold principle

A response to Harris based on this principle would say that, although patients expected to die at 25 have a somewhat weaker claim than patients expected to die at 35, this consideration is easily outweighed by other considerations such the countervailing considerations mentioned above (because the strength of different students' age-related claims is not very different). But if the expected ages of death are around the threshold age (say, some 5 years above and some 5 years below), then their age differences would carry more weight (the claim of the older ones would be significantly weaker than that of the younger ones). In such a case, age-related considerations are more likely to outweigh other considerations.

The threshold principles represented by the above two graphs reflects much of the talk about age-based considerations, which does indeed involve a threshold. For example, in setting up a case, Kamm asks us to suppose that a "reasonable" life is seventy years long.[296] Similarly, in defending his threshold-based fair innings theory, Callahan cites "The biblical idea of a full life as about three score and ten years."[297] And in the same vein, on Harris's understanding, the fair innings argument runs as follows:

> The fair innings argument takes the view that there is some span of years

---

[296] Kamm (1993), p. 239

[297] Callahan (1995), p. 68

> that we consider a reasonable life [for a person to have had], a fair innings. Let's say that a fair share of life is the traditional three score and ten, seventy years. Anyone who does not reach 70 suffers, on this view, the injustice of being cut off in their prime. They have missed out on a reasonable share of life: they have been short-changed. Those, however, who do make 70 suffer no such injustice, they have not lost out but rather must consider any additional years a sort of bonus beyond that which could reasonably be hoped for. The fair innings argument requires that everyone be given an equal chance to have a fair innings, to reach the appropriate threshold but, having reached it, they have received their entitlement. The rest of their life is a sort of bonus which may be cancelled when this is necessary to help others reach the threshold.[298]

The consequence of this "threshold" theory of age-based discrimination would be that it is easier to justify discrimination between one group of patients below the threshold and another group above the threshold than between two groups who are both below the threshold or two groups who are both above the threshold. Let us return for a moment to the Harris-inspired case we considered earlier, involving a gap of ten years between two groups' expected age at death. The proponent of a threshold principle of age-based discrimination can say that age-based considerations give us only quite weak reasons to discriminate between two groups of young patients separated by this gap; as such, age-based considerations are easily outweighed by other considerations. Such a theorist can also say that age-based considerations give us little reason to discriminate between two groups of older patients separated by the same gap, who are all expected to die above the threshold.[299] In contrast, a threshold principle gives us stronger reasons to discriminate in the mixed case where some are below the threshold and some above. Specifically, if, without treatment, one group of patients will die younger than the threshold (say, at 65) whereas the other group will die older than the threshold (say, at 75), then all else being equal, this principle gives us a relatively strong reason to prioritise the patients who will die younger.

Thus a "threshold" fair innings theory may get the best of both worlds, avoiding counterintuitive discrimination between patients who are expected

---

298 Harris (1985), p. 91

[299] But of course, if there is a bigger gap between two groups' expected age at death, then we might get stronger reasons to discriminate between two groups even if both are above the threshold or both below the threshold

to die very young (e.g. younger patients) while licensing discrimination in favour of patients who without treatment would not make the threshold over patients who will.

However, if a threshold principle is merely cited as a means of getting the intuitively acceptable result in specific cases, it will look ad hoc. In particular, we can ask why the line is flat (or close to flat) on either side of society's average, and yet vertical (or closer to vertical) at a certain point? As Harris puts it:

> the very arguments that support the setting of the threshold at an age which might plausibly be considered to be a reasonable lifespan, equally support the setting of the threshold at any age at all. ... After all, what's fair about the fair innings argument is precisely that each individual should have an equal chance of enjoying the benefits of a reasonable lifespan. The younger patient can argue that from where she's standing, to age of 40 looks like much more reasonable a span than that of 30, and that she should be given a chance to benefit from those ten extra years. ... This argument generalised becomes a reason for always preferring to save younger rather than older people, whatever the age difference, and makes the original anti-ageist argument begin to look again the more attractive line to take.[300]

Harris is saying that there are no grounds for a steeper drop at 70 than at any other age. Younger patients are in as good a position as older patients to argue that they should be favoured over someone (say) ten years older.

So the question we must address is, why do age-related considerations give us no reason to favour students in their early 20s over students in their late 20s, while giving us reason to favour people in their 20s over people in their 90s? Additional argument is needed to give us independent grounds for thinking that there is a threshold, and to help *explain* intuitions. Kappel and Sandoe's argument does not support a threshold; their principle favours the younger patient regardless of the age of the patients in question. I will therefore consider Callahan's defence of a threshold theory, before rejecting it and offering my own.

---

[300] Harris (1985), p. 91-92

7.5 Rejected Explanation of Threshold: Callahan

Callahan defends a sharp threshold principle, according to which the strength of our claim to life-extending treatment drops to zero after a certain age. Callahan's position is that:

> Government has a duty, based on our collective social obligations, to help people live out a natural life span, but not actively to help extend life medically beyond that point. ... Beyond the point of a natural life span, government should provide only the means necessary for the relief of suffering, not life-extending technology.[301]

Callahan offers an argument for this based on fair opportunity. In outline, he argues:

> our ideal of old age should be achieving a life span that enables each of us to accomplish the ordinary scope of possibilities that life affords, recognising that this may encompass a range of time rather than pointing to a precise age. On the basis of that ideal, the aged would need only those resources which would allow them a solid chance to live that long and, once they had passed that stage, to finish out their years free of pain and avoidable suffering.[302]

In detail, he argues that if we can agree a notion of a "tolerable death", then we will have "the basis for a correlative idea of a natural life span and thus, perhaps, the foundation for an appropriate goal for medicine in its approach to aging".[303] Callahan's definition of a "tolerable death" is as follows:

> the individual event of death at that stage in a lifespan when (a) one's life possibilities have on the whole been accomplished; (b) one's moral obligations to those for whom one has had responsibility have been discharged; and (c) one's death will not seem to others an offense to sense or sensibility, or tempt others to despair and rage at the finitude of human existence.[304]

These stipulations need clarification. Firstly, to clarify what it means to say "one's life possibilities have on the whole been accomplished ", Callahan says

> Life affords us a number of opportunities. These include work, love, the

---

[301] Callahan (1995). The three passages quoted here are taken from p. 135, p. 137 and p. 138.

[302] Callahan (1995), p. 135

[303] Callahan (1995), p. 66

[304] Callahan (1995), p. 66

> procreating and raising of a family, life with others, the pursuit of moral and other ideas the experience of beauty, travel, and knowledge, among others. By old age - and here I mean even by the age of 65 - most of us will have had a chance to experience those goods; and will certainly experience them by our late 70s or early 80s. It is not that life will cease, after those ages, to offer us some new opportunities ... But what we will have accomplished by old age is the having of the opportunities themselves, and to some relatively full degree.[305]

In saying this, Callahan acknowledges that some of these opportunities may not come to certain individuals before old age, but he argues that in that case those opportunities are unlikely to come in old age either. In other cases, further opportunities will keep coming in old age, but:

> we will *on the whole* already have had ample time to know the pleasures of such things. ... No amount of time would make it possible to do everything possible. The biblical idea of a full life as about three score and ten years must have had behind it a perception of that kind rather than a purely biological observation.[306]

Regarding the second part of his definition, Callahan says:

> When I speak of having discharged one's moral obligations, I have in mind primarily family obligations, particularly to one's children. Obligations to children are very special and inescapable. The death of parents when children are still wholly dependent on them is easily and rightly seen as particularly sad and wrong, and that is so even if others can step in and assume the parental role.[307]

Regarding the third part of the definition, Callahan says:

> the death of an elderly person who has lived a rich and full life is not, in any society, accounted as an evil, as if symptomatic of a deranged and cruel universe. ... [the main reason normally being] that the elderly have lived a full life, have done what they could, and thus are not victims of the malevolence of the forces either of divinity nor of nature.[308]

On the basis of the above, Callahan says:

---

[305] Callahan (1995), pp. 66-67

[306] Callahan (1995), p. 67

[307] Callahan (1995), p. 69

[308] Callahan (1995), p. 71

> I have proposed that our ideal of old age should be achieving a life span that enables each of us to accomplish the ordinary scope of possibilities that life affords ... I will, therefore, define need in the old as primarily to achieve a natural life span and thereafter to have their suffering relieved.[309]

How should we assess this argument, as a defence of a threshold? One response might be that different people get their important things such as family and career done at different times of life. Many leave it till they're older. There is no sudden cut-off age and therefore no argument for a sharp threshold based on the fact that we all get the important things done by a certain age. But Callahan can respond that the relevant question for his account is the question of when people get their chances, not when they take them. One might have had a chance to have a family or career without actually having had a family or career.

But when have we had our chances? Callahan's age threshold for withholding life-extending treatments comes at the point where we have had various important opportunities "to some relatively full degree". A lot depends on what counts as having had an opportunity to this degree. Here, Callahan is stuck on the horns of a dilemma. He can either give an easily-satisfied definition such that we have had our opportunities in the relevant sense when we are quite young, in which case the threshold comes at a counterintuitively young age. Or he can give a hard-to-satisfy definition such that we haven't had our opportunities until the age at which everyone has fulfilled their desire for family or career, in which case the threshold comes at the end of our natural lifespan where it has no practical implications. On the other hand, it is difficult for Callahan to defend a threshold between these two extremes, because any such intermediate threshold will seem ad hoc, motivated by the need to satisfy intuitions rather than by independent considerations.

Let's consider the two extreme options in more detail. First, Callahan could give an easily-satisfied definition of opportunity, such that we have had our opportunities in the relevant sense at some point well short of our natural lifespan. The definition could be in quite narrow, health-related terms, such that you have the opportunity to have a family or pursue a career provided you have the physical ability to do those things. Then, to the extent that the

---

[309] Callahan (1995), p. 135

healthcare system can give you that physical ability and does so, it has fulfilled its obligation to give you the opportunities which Callahan's principle says it should. The problem with this account is that it means the threshold is counterintuitively young. In the UK people can legally start a family at 16, so the kids could be leaving home when the parents are in their mid 30s. Similarly, if reaching one's career potential means going as far as one will go, most of us could fulfil our career potential by the age of 30 or 40, *if* we were focused and worked hard enough.[310] Similar points apply to the other areas that Callahan mentions. So, assuming that "opportunity" is defined in narrow, health-related terms, it seems that the vast majority of people have had the opportunity to carry out their most important plans by the time they are 35, perhaps 40 in some cases.

The problem for Callahan is that intuitively, we have not had a fair innings by the age of 40. For example, it would be very counterintuitive to favour a 35 year-old over a 45 year-old for life-saving healthcare on the grounds that the 45 year-old has had all the chances anyone could expect, but then to say there is no difference between a 45 year-old and a 75 year-old in terms of the chances they’ve had. To the extent that we have intuitions about a fair innings threshold, I suggest it is age 70 or higher. Callahan shares this intuition (again, see his reference to the his reference mentioned above to the "The biblical idea of a full life as about three score and ten years"). Various theorists other than Callahan have the same intuitions. Actually, few theorists explicitly endorse a specific fair innings figure, but it is suggestive that almost all theorists use 70 as their example when they put their arguments (for example, Kamm and Harris).[311]

So now we can consider Callahan's second option, which is to switch to the hard-to-satisfy definition such that we haven't had our opportunities until the age at which everyone has fulfilled their desire for family or career. This seems to be Callahan's preference, e.g. his reference to the "The biblical idea of

---

[310] For example, there are 29 CEOs of US listed companies aged 35 and under (Kirdahy, 2007). The UK's Higher Education Statistics Agency says there were 25 professors under the age of 30 working in the UK in 2013-14 (HESA, 2015 as quoted in Grove, 2015). To emphasise, I'm only considering opportunities of the kind healthcare can appropriately aim to give us, such as the opportunities deriving from able-bodiedness; not all the other opportunities that young CEOs had, deriving from their other abilities and circumstances.

[311] For one of her examples, Kamm asks us to suppose that a "reasonable" life is seventy years long (Kamm 1993, p. 239). Harris asks us to suppose that "a fair share of life is the traditional three score and ten, seventy years" (Harris 1985, p. 91).

a full life as about three score and ten years".[312] But if we wish to give sufficient opportunity to all, he may have trouble defending any threshold short of our maximum natural lifespan, since some people die having spent their entire lives searching for love or for work that suits them. But giving such people every last chance by moving the threshold to the end of our natural lifespan does not match fair innings intuitions, whereby the threshold comes at some point well before death (many live beyond Callahan's "three score and ten years"). So this justification has counterintuitive consequences. It also means that the fair innings principle is irrelevant for policy-making purposes.

At first sight it may seem like Callahan avoids this consequence, saying that "even by the age of 65 - most of us will have had a chance to experience those goods".[313] The word "most" here suggests he does not propose giving these opportunities to all. In virtue of this qualification, can Callahan defend a more intuitively acceptable threshold at the age of 65 or thereabout? Well, first he must address three problems.

First, what proportion of the population is an adequate proportion, on his account? I suggest it will be difficult for him to state a proportion that can be defended in non-ad hoc terms, viz., terms that do not presuppose that our intuitions regarding the threshold are correct. To simplify, we could say something like "The threshold ought to be at age 70, so let's set the proportion at 80%, since that implies a threshold at age 70". This is ad hoc. The account is motivated by our intuitions and therefore cannot offer to explain them.

The second problem with this account is that it excludes the minority who do not get the relevant opportunities by the time they reach the threshold age. This minority of older people has not yet had a chance to do those things. Further life years could allow them to do them. If some of the young need their life years for the sake of doing the important things, and if this need gives us a compelling reason to allocate healthcare to them, then we have the same reason to allocate healthcare to the minority who don't manage to get their important things done in youth and middle age. If we are obliged to satisfy a certain type of need in one case then surely we are equally obliged to satisfy the same need in other cases.

---

[312] Callahan (1995), p. 68

[313] Callahan (1995), pp. 66-67

The third problem is faced by any account on which our reason for giving healthcare is to give people various important opportunities such as opportunities for career and family. The problem is that this could lead to unfair discrimination against people who do not have such opportunities. For example, in terms of the opportunity to have family, we would have less reason to give healthcare to infertile people. We would also have reason to treat men until they die but less reason to treat women until they die, since men are fertile into their old age but women are not (creating a double jeopardy for women - loss of entitlement to healthcare as well as loss of fertility). In terms of the opportunity to have a career, we would have less reason to give healthcare to the disabled who cannot work. Of course we might have reason to treat people's infertility or their disabilities, but if we cannot cure these conditions, then Callahan gives us little reason to extend the lives of the affected individuals. These would seem like unacceptable consequences.

So, to conclude on Callahan, in setting the threshold at the point where we have had various important opportunities "to some relatively full degree", Callahan has three options, each with its problems as the basis of a sharp threshold. On one option the threshold age will be counterintuitively young, to reflect that on a "physical possibility" conception of opportunity, most of us have the opportunity to pursue family and/or career when we are young. Alternatively the threshold will come at the end of our natural lives, to accommodate those who take a while to get going. This does not match fair innings intuitions, whereby the threshold comes at some point before death. If in response, Callahan tries to set the threshold at some point between these extremes, his solution will look ad hoc, driven by the need to satisfy intuitions rather than by independently plausible considerations. It will therefore lack explanatory interest.

7.6 New Proposed Explanation of Threshold

So we have not yet found a defensible way of setting the threshold and justifying it. I will now offer my own defence of a threshold principle. I will defend a view according to which the fair innings in a given society is equal to that society's average life expectancy. The argument starts with an appeal to ordinary intuitions, before moving on to state a principle in terms of which one might aim to vindicate those intuitions.

I start by arguing for the egalitarian significance of society's average life expectancy (where this is understood as the average age at which people are expected to die). I suggest this is important because it is where everyone would be if life years were redistributed such that everyone had the same number of life years. As such it gives us a kind of target; it is the place we want everyone to be, as egalitarians.[314]

In support of this point, I will offer some cases which suggest that society's average life expectancy is intuitively important. Specifically, these examples show that society's average life expectancy is an important benchmark for fair innings intuitions, in that if society's average life expectancy varies, our intuitions as to who merits priority also vary. If we must choose between saving some patients who are expected to die at 40 and some patients who are expected to die at 50, our intuitions vary according to the average life expectancy of the society they live in.

Consider the decision whether to fund a group of 20 year-olds who will die at 40 without treatment or another group of 20 year-olds who will die at 50 without treatment. Suppose that in each society we must choose between giving one extra year to the group who will otherwise die at 40 or two extra years to the group who will otherwise die at 50. Compare three societies these patients might belong to, each with different average life expectancies. So in one society, healthy 20 year-olds can expect to live healthily till age 30, in another they can expect to live healthily until 48 and in another they can expect to live healthily until they are 110. For vividness and to bring out the issues, also suppose that deaths are clustered: 95% of the population in each society die within one year either side of the average age of death.

Start with the society where healthy 20 year-olds live till 48 on average. In that society, I suggest there is an intuitive reason to give the one year to the patients who will otherwise die at 40. 95% of the community get to live till 48, yet these patients will die at 40. In contrast, the patients who will otherwise die at 50 will live beyond the age at which 95% of their contemporaries will die. So even without treatment, they are privileged. So although it is relevant that the patients who will otherwise die at 50 would get extra benefit from their treatment, I contend that is a strong intuitive pull towards giving the extra year to the younger patients.

---

[314] Albeit we also wouldn't wish harm on anyone, so we wouldn't take life years from anyone. But we might think that a world in which everyone now living had the same life expectancy would be better than the world we're in.

Now consider the society in which healthy 20 year-olds can expect to live healthily until age 110. From the perspective of this society, both the patients who will otherwise die at 40 and those who will otherwise die at 50 will seem severely underprivileged by comparison with their contemporaries. In fact, in the grand scheme of things, both groups are roughly equally underprivileged. In such circumstances, I suggest intuitions do not favour the patients who will otherwise die at 40 as strongly. In consequence, in the overall trade-off between age-related considerations and our concern to do most good, our intuitions might be swayed by the fact that the patients who will otherwise die at 50 get extra benefit from their treatment. Many would judge that those patients should get the treatment.

Finally consider the society where healthy 20 year-olds can only expect to live till age 30. Intuitively, I suggest that from the perspective of this society, both the patients who will otherwise die at 40 and those who will otherwise die at 50 will seem privileged. Both will live well beyond the age at which most of their contemporaries will die. They are the equivalent of people who expect to live into their 90s and 100s in the present day developed world, in terms of the life they will have had compared with the average. For this reason, once again, I don't think intuitions favour the younger patients as strongly. In the context of this society's average life expectancy, the comparison between the patients who will otherwise die at 40 and those who will die at 50 does not make those who will die at 40 seem as badly off as they would seem in the first society we considered. In this context, I suggest our intuitions driven by age-related considerations are weaker. In fact, I suggest many would favour giving the treatment to the patients who will otherwise die at 50, because they will get more benefit.

To summarise, if both groups have life expectancies lower than society's average life expectancy, or both have higher than average, then the younger group's claim is weaker by comparison with that of the older group than in the case where the life expectancy of each group falls on either side of society's average life expectancy.

Having said that, in appealing to these cases, I have tried to avoid relying on any claim about exactly what policy is correct in each society. I merely aim to show that intuitively, the relative weight of two salient considerations has the potential to vary depending on another, "background" consideration. The first salient consideration is that the patients expected to die at 50 will get more benefit than those expected to die at 40. The second salient consideration is that those expected to die at 40 are expected to die younger; this is what our age-related principle will latch onto, if it makes a difference.[315] These two considerations are held constant across the three cases. The background consideration, which varies between the cases, is society's average life expectancy.

These intuitions tally with the stepped principle I defended above, with the vertical part of the line at society's average life expectancy. For example, if we must choose between two groups, and the life expectancy of both is below society's average, this principle does not favour either, and benefit maximising considerations might take precedence. But if one group is below and one above, these age-related considerations favour the group below. If we wish to draw a graph representing the principle which seems to be driving intuitions, it might look something like this:

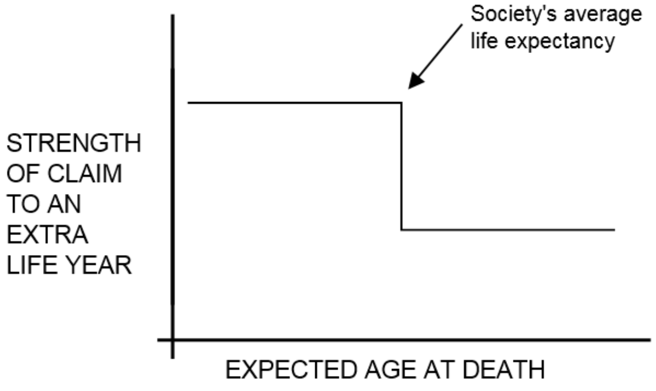

Fig 8. Intuitive strength of claim vs society's average life expectancy: stepped

As usual, the bottom axis represents expected age at death. The principle represented by this line can then be combined with a benefit-maximising principle to capture the intuitions described above. All else being equal, an allocation based on this line will not differentiate between two patients who are both below the average, nor between two patients who are both above it. As a result, other considerations may make all the difference. For example, in

---

[315] To keep things simple, I have not considered the more complicated cases where treatment would take a patient's life expectancy from below the average to above the average. In all three of these cases, I have assumed that if a patient's life expectancy without treatment is lower than the average for their society then their life expectancy with treatment will also be lower than average; similarly, if a patient's life expectancy is higher than average without treatment, their life expectancy with treatment will also be higher than average.

the above cases, the other considerations are benefit maximising considerations: our intervention therefore goes to the patient who will benefit the most, this being the patient with greater life expectancy. But the principle represented in the graph *does* distinguish between one patient who falls above society's average and another who falls below. In such a case, benefit-maximising considerations have less influence on the overall decision, and it may be that patients whose life expectancy is below society's average may get the treatment despite benefiting less than other patients.

Alternatively, the line could be as follows:

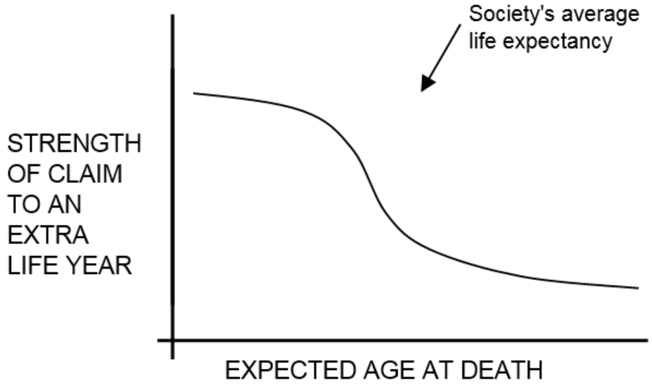

Fig 9. Intuitive strength of claim vs society's average life expectancy: graduated

Here, there is a slope in the line either side of society's average, but the slope gets steeper at the average. This represents the way age-related considerations *always* favour patients with low life expectancy over patients with higher life expectancy, regardless where the two patients stand with respect to society's average. Specifically, even if two groups of patients are *both* below society's average, the worse off group may still get treatment despite benefiting less. However, the line between the life expectancies of two such groups of patients only slopes gently, so this is less likely; generally, the patients that benefit more will get treatment. But as before, if the patients fall either side of society's average, the line between them is much steeper, and so benefit-maximising considerations have less influence; the worst-off patients are more likely to get priority, even if they benefit significantly less than the better off patients.

I contend that one of the principles represented in the above two graphs can explain the intuitions regarding the three societies described above (I would tentatively defend the second, graduated principle over the first, step-shaped principle; if a fair innings principle is applicable at all, surely there will be *somewhat* greater priority for the worst-off, even over others who are also below society's average). More generally, these principles can explain threshold-shaped intuitions according to which patients whose life expectancy is *below* the age threshold all have roughly the same entitlement to life-saving treatments, and patients whose life expectancy is *above* the age threshold also all have roughly the same entitlement to life-saving treatments - but patients with life expectancy above the age threshold have sharply less entitlement than patients below the age threshold.[316]

But as noted, to defend an allocation scheme, it is not enough to appeal to intuitions about specific cases. Those intuitions must be explained in terms of a compelling principle. So, what ethically compelling principles can explain our responses in the above cases? Specifically, why are the line segments either side of society's average life expectancy flatter and why is the line so steep at the average?

It is tempting to say that there is a sufficientarian principle at work, but that does not work. To explain, the line representing a sufficientarian principle would indeed have a threshold, to represent the way the strength of someone's claim suddenly drops at the point where they have enough. For example, sufficientarianism regarding food might say that someone has a strong claim to food whether they have a bit less than they need to keep them healthy or a lot less; but as soon as they have enough food to keep them healthy they have no claim to further food. The graph representing this principle would look a lot like the threshold graphs above; it would have a steep bit representing the boundary between the need for more food and the absence of such need. But this sufficientarian explanation of the above lines does not work, since the location of the steep bit in the above lines depends on society's average life expectancy: if we changed the cases such that society's average life expectancy was higher or lower, the threshold would move accordingly. In contrast, assuming we need life years, this need does not change according to society's average life expectancy (just as the question of

---

[316] Because, in the case of life-saving treatments for diseases which threaten to kill patients soon, the age of the patient correlates with their expected age at death, so patient age is a good proxy for the factor which egalitarians are concerned with, which is expected age at death. In contrast, this egalitarian principle would not be well-placed to explain similar intuitions regarding treatments for diseases which threaten to kill their sufferers a long time in the future, if the intuitions took exactly the same shape for younger vs older patients, since in such cases it is possible for younger patients to have a higher expected age at death than older patients.

whether someone has enough food does not depend on their society's average food intake). So the steep bit of the above lines does not represent the boundary between need and its absence, and these lines do not represent a sufficientarian principle.

Here is my attempt at explaining why the flatter line segments are flat (or close to it), and why the steep bit is so steep (in comparison). Or in ordinary language, I will try to explain why this principle does not discriminate much between groups below society's average life expectancy, nor much between groups above society's average, but nevertheless discriminates between groups above and below the average. I start by reiterating my earlier point: society's average life expectancy is important, because it is where everyone would be if life years were redistributed such that everyone had the same number of life years. As such it gives us a kind of target.[317] This means our attitude towards those below the average is different to our attitude to those above it, as follows. (For simplicity I will speak as if the line either side of the average is flat, although it is probably not; it is simply *flatter* than the line running through the average).

I start by considering two groups who both have a life expectancy above the average. In this case, the issue is that whoever we help, we will be moving them further away from the point where they would be if society was equal. From that perspective, there is not much difference between helping someone who is only a bit above average and someone who is well above average. Either intervention takes us away from where we want society to be. So age-related considerations will not imply priority for one or the other. This is reflected in the flat line to the right of the threshold.

On the other hand, if everyone we can help has a life expectancy *below* the average, we are equally concerned to bring all of them up, however far they are below the average (as egalitarians/prioritarians; again, I set aside other considerations that may play a role). We feel a comparable degree of concern wherever someone falls below the average. The important thing is that not everyone is at the average. So, if we have to choose between bringing up someone who is a long way below the average and bringing up someone who is closer to the average, we don't care much more about bringing up the person who is a long way below, if both can be moved an equal distance towards the average. Both moves would work equally well at bringing the

---

[317] The principle I characterise here is therefore somewhat vulnerable to the levelling down objection, which I will return to later.

distribution of life years towards the shape we want, so both interventions are equally desirable. Thus age-related considerations will not imply priority for one or the other. This is reflected in the flat line to the left of the threshold. And as we saw in the cases I considered above, when age-related considerations do not tell strongly either way, this will often leave the field clear for other considerations to be decisive, as per the above case where our intuitions appear to be swayed by benefit maximising considerations. (Admittedly it is counterintuitive that people who are a long way below society's average get no priority over those who are only a bit below. If we are at all concerned about inequalities, we'll surely be concerned about these inequalities too. This suggests that in fact, pace Harris, the line to the left of society's average is not entirely flat after all. The line is flat enough to mean that we will not usually discriminate between people in their 20s, but steep enough to mean that we will usually discriminate between 20 year-olds and 60 year-olds. In summary, the graduated slope looks more compelling than the stepped slope)

For these reasons, whether we have to choose between two groups below the average or between two groups above the average, age-related considerations do not give us much reason to discriminate between the two groups.

In contrast, if one group is below the average and one group is above, age-related considerations give us much more reason to discriminate. We care about moving those below the average towards the average; we don't care about moving those above the average away from the average; so we have reason to favour the group below the average.

Perhaps the key feature distinguishing this threshold principle from other varieties of egalitarianism and prioritarianism is the way it benchmarks the prospective beneficiaries of our intervention against society's average. This governs the location of the steep bit of the line. In contrast, other egalitarian principles involve a comparison solely between the prospective beneficiaries of our intervention.[318]

---

[318] Is my principle egalitarian or prioritarian? It refers to society's average life expectancy and values benefits more or less highly according to whether they move people towards the average. Thus the principle I defend places a certain value on equality as an outcome, which might suggest it is egalitarian. However, the principle only implies that we should value some benefits over others; it does not imply that we should value harms. This makes it a prioritarian principle, in line with Parfit's original distinction. He defined the Priority View as saying that "Benefiting people matters more the worse off these people are." (Parfit, 1997). The scope of the Priority View is benefits, which get weighted more or less according to who receives them. It

This makes it important to have a non-arbitrary way of defining the reference group, viz., the society whose average determines the location of the steep bit of the line. The obvious reference group would be the society we are making policy for.[319] But consider the case of regional policy-makers in a national system. Suppose they must choose between two patient groups. The better off group stands to gain slightly more from their treatment. Treatments for both groups cost the same and all else is equal. Suppose the life expectancy of both groups is below the average for the region. If policy-makers consider only their region, then my principle suggests treatment should go to the better off treatable group. But if their region is quite well off, local average life expectancy may be higher than national average life expectancy, such that the life expectancy of the two groups falls either side of the national average life expectancy. This would indicate that treatment should go to the worst-off treatable group. For my principle to have clear implications for this case, it needs to be clear which of the regional average or the national average is the benchmark average. I tentatively suggest that a difference might be made by the extent to which the local community is networked into the national community. So, to the extent that a region is independent in terms of its economy and policy-making, to that extent perhaps the local average is the relevant average (one might argue that societies need equality, and so policy-makers are responsible for producing equality in their society). On this view, the identity of the relevant reference group may be somewhat indeterminate; perhaps both averages are relevant to an extent.

Assuming this issue can be dealt with, I argue that a threshold fair innings principle of age-discrimination can be defended and can explain why we want to discriminate by life expectancy in many cases, but not in Harris-type cases.

7.7 Williams' Account: Quality-Adjusted Life Expectancy

Can we extend the logic of fair innings theory from life years to health-related quality of life[320] impairments? I now briefly consider an account

---

does not apply to goods which people already have, and it therefore avoids the Levelling Down Objection which Parfit raised in his 1997 paper (the objection that some forms of egalitarianism imply that there is something good about a harm which brings someone privileged down towards society's average, even if no-one else benefits)

[319] This fits existing practice; it is my impression that egalitarianism is more often applied within societies than between societies.

[320] Hereinafter "quality of life"

offered by Alan Williams which offers to do exactly that. Williams would prioritise an individual whose expected lifetime QALYs are reduced whether as a result of short life expectancy[321] or quality of life impairments. The account is interesting as a quite specific operationalisation of an age-based discrimination principle, being precisely stated in terms of QALY weights, enabling us to see the potential pitfalls that lie in wait for any principle of this kind. It should be noted that this account does not have the shape I proposed above; in particular, it lacks the "step" at society's average life expectancy. But there is nothing in principle to stop it being amended along those lines. Williams merely offers his account as an "exploration of certain equity issues", and does not assert that his principle is the sole defensible way of addressing such issues.[322] In assessing the account, I find it is at risk of counterintuitive consequences, but that these can be avoided with some care in setting the weights to be assigned to patients with severe impairments. I therefore conclude by endorsing a version of the account. On the principle I endorse, anyone with a quality-adjusted life expectancy lower than society's average merits priority for any treatment to help bring them up to society's average quality-adjusted life expectancy.

I start by introducing Williams' account. He considers Harris's conclusion that

> people who had achieved old age or who were closely approaching it would not have their lives further prolonged when this could only be achieved at the cost of the lives of those who were not nearing old age.[323]

Williams comments that this raises three questions:

> (a) what is a relevant personal characteristic by which to classify people for such policy purposes?
>
> (b) how are we going to measure health?

---

[321] The phrase "life expectancy" is ambiguous. It can denote the total number of years a patient is expected to have lived by the time they die, or the years remaining to them at their current age. I will use it to denote the former, if I must use it. But in general I will try to use other constructions, such as "patients who are expected to die young" or "patients whose lives will have been short" (to substitute for the first meaning) or "life years remaining" or "life years in prospect" (to substitute for the second meaning).

[322] Williams says he is motivated by two desires: the desire to "relate discussions of social justice, as typically conducted by non-economists, more closely to ways of thinking that are natural to economists" and the desire to "impose some quantitative rigour upon the assertions made by non-economists about what is equitable".

[323] Williams (1997), quoting Harris (1985)

> (c) how are we going to measure a health inequality?[324]

Williams then says:

> With this information before us, we then have to decide just how averse we are to any inequality so described, i.e. what sacrifices in the original efficiency maximand we would be prepared to accept to achieve a specified reduction in the policy-relevant inequality.[325]

Williams takes social class as the focus of his analysis, but he says this is only "an example of how policy analysis of health inequalities might proceed."[326]

Williams proposes that we should adopt a quantitative approach, noting that:

> On the whole, debates about equity are not cast in quantitative terms. ... If the nature and implications of particular positions are to be clarified in a policy-relevant way, this discussion has to move on to seek quantification of what are otherwise merely vaguely appealing but ambiguous slogans ... At present, although reassurance is frequently offered that equity considerations have been taken into account, there is no way of establishing what bearing, if any, those principles actually had upon the outcomes. ... Quantification thus has potential for clarification, for performance measurement, for accountability and for policy analysis and reappraisal.[327]

He examines how social class correlates with differences in life expectancy:

> If we take life expectancy at birth as defining a 'fair innings' within any society, then in the UK the differences in male survival rates between the professional and managerial groups (social classes 1 and 2) on the one hand and the semi-skilled and unskilled manual workers (social classes 4 and 5) on the other, has been estimated to differ by about 5 years (72.5 compared with 67.7). ... The equalization of life chances in terms of life expectancy seems to require some changes in public policy, though not wholly confined to health care. But limited though the contribution of health care may be, it could be exploited more fully by weighting additional life years gained from the various health care activities

---

[324] Williams (1997), p. 119

[325] Williams (1997), p. 119

[326] Williams (1997), p. 120

[327] Williams (1997), p. 120

> according to the social class of the potential recipient.[328]

Williams then proposes that,

> if it is to capture the full flavour of this kind of thinking, the concept of a 'fair innings' needs to be extended beyond simple life expectancy to embrace quality-adjusted life expectancy. Otherwise it will not be possible to reflect the view that a lifetime of poor quality health entitles people to special consideration in the current allocation of health care, even if their life expectancy is normal.[329]

since

> It is clear from these data (shown in Fig. 3) that surviving members of social classes IV and V have noticeably worse health than their contemporaries in social classes I and II, especially once they are past the age of 40 years.[330]

The consequence of this is that

> When these data are combined with the differences in survival rates we find that the quality-adjusted life expectancy at birth of someone in social classes 1 and 2 is nearly 66 QALYs, but for someone in social classes 4 and 5 it is only about 57 QALYs. To achieve the mean value of about 61.5 QALYs (a 'fair innings' for a pure egalitarian) they would need to live to be 65 and 71 years old, respectively, a feat achieved by about 76% of social classes 1 and 2, but by only 46% of social classes 4 and 5.[331]

he then sets the scene for his example of a trade-off formula by asking

> how big a sacrifice in the overall health of the population would you be prepared to accept in order to eliminate the disparities in health between [two social groups]?[332]

Then, returning to social class differentials between males in the UK, he asks us to

> suppose people were prepared to sacrifice 6 months of life expectancy at birth, in order to eliminate the disparity of 5 years.[333]

---

[328] Williams (1997), p. 121

[329] Williams (1997), p. 121

[330] Williams (1997), p. 121

[331] Williams (1997), p. 121

[332] Williams (1997), p. 123

[333] Williams (1997), p. 123

This means

> we should attach twice the weight to improving the life expectancy at birth of people in SC4 and 5 as we do to doing so for SC1 and 2.[334]

7.8 Objections to Williams

I now consider two objections to Williams. First, it looks like Williams' account could have some counterintuitive consequences. To show this, I will talk in terms of two "broad dimensions" of health: quality of life and life expectancy. The problem is that Williams' account implies that priority carries over from impairments in one broad dimension to benefits in the other. For example, Williams' account implies that of two 20 year-olds expected to die at 50, the patient in severe pain should get priority for life-extending treatment, all else being equal: the pain reduces the patients quality-adjusted life expectancy, and extra life years would bring it back up towards society's average. I suggest that this implication is highly counterintuitive (I am willing to grant that, as we saw in Chapter 5 on disability discrimination, a patient's pain might sometimes be *neutral* with respect to their claim to life-saving treatment, but it is counterintuitive that it could *increase* their entitlement). Similarly, of two patients with a headache, Williams' account implies that the one who is expected to die at 50 gets priority for painkillers over the one who is expected to die at 70. Again, this is counterintuitive.

These problems arise because Williams' account implies that priority carries across different broad health dimensions (in contrast, the account I have been considering so far only implies that short life expectancy implies priority for extra life years; I have not been considering the possibility that short life expectancy might imply also priority for painkillers). Contrary to Williams, I suggest that intuitively, priority does not carry across different broad health dimensions. Quality of life impairments don't give priority for life-extending treatment, and patients with short life expectancy don't get priority for quality of life improvements. (Note that these intuitions don't go against anything I say above. Intuitively, priority does carry *within* broad health dimensions. Priority carries from quality of life impairments to quality of life benefits, as well as from life expectancy impairments to life expectancy benefits. For example, a patient in severe pain should get priority for pain relief over a

---

334 Williams (1997), p. 124

patient with only moderate pain. And as argued, a patient expected to die young should get priority for life-extending treatment over a patient expected to end up having more life years than society's average.)

The above counterintuitive implications look like a problem for Williams. However, he can avoid the problem and ensure his account aligns with our intuitions by taking care over the weights he applies to reflect the severity of patients' impairments. The solution needs to be based on the fact that when a patient's severe impairment is in a different broad health dimension from the treatment benefit, there will be a tension between the factors determining the overall strength of that patient's claim to treatment. In summary, egalitarian considerations will favour the patient, but benefit maximising considerations will go against the patient. Williams can exploit this tension. Specifically, although Williams' account entails that the patient's treatment benefits must be up-weighted to reflect the patient's severe impairments, this need not lead to that patient getting any priority. The patient's impairments are in a different broad dimension of health, and as a result the impairments will detract from the benefit we attribute to the treatment. This negative impact on the amount of benefit will counteract the up-weighting factor we apply to the benefit. As an example, consider patient A with severe pain and a slightly impaired life expectancy. Compare patient B with the same slightly impaired life expectancy but with unimpaired quality of life. Suppose we must choose which patient receives a life-extending treatment. Although Williams'' account entails that patient A's severe pain means that that a positive weighting factor must be applied to their treatment benefits, the pain also means that patient A derives fewer unweighted QALYs from the life expectancy gains than patient B. So Williams can argue that, as long as the weightings deriving from the quality of life impairment are quite low, patient A will get the same number of weighted QALYs as patient B, or fewer. This is in line with intuitions.

In contrast, if a patient's treatment benefits are in the same broad health dimension as the patient's severe impairment, there will be a synergy between the factors determining that patient's claim to treatment, and no "background" condition detracting from the benefits of treatment. For example, suppose that without treatment, patient A is expected to die younger than society's average age at death and patient B is expected to die older than society's average age at death. Also suppose both are expected to gain 5 years from life-extending treatment. Neither has any quality of life impairments and there is no other relevant difference between them. Then both will get the same number of unweighted QALYs from that treatment. Williams' account will then assign patient A higher weighting factor because of their shorter life expectancy. This will ensure patient A gets the treatment, in line with the fair innings theory I offered above.

In summary, the implications of Williams' account are highly dependent on the weights he assigns to reflects a patient's impairments. If the weights are quite light, such that the amount of benefit a patient is expected to derive from a treatment is still the most important factor determining their chances of getting that treatment, then Williams can avoid implying that a health impairment in one broad dimension gives priority for a patient who stands to benefit in another broad health dimension. Nevertheless, the resulting account will still make a difference in the right cases: even when the weights are light, when a patient's treatment benefit is in the same broad dimension as a severe impairment, the patient will get priority over another patient who stands to gain the same benefit but who is not severely impaired.

The step-shaped principle I defended above can be operationalised under the same constraint. For example, the flat line on the left can represent that patients whose quality-adjusted life expectancy is below society's average will have their treatment benefits up-weighted by a factor of only 1.1 (a small factor), while the flat line on the right can represent that those whose quality-adjusted life expectancy is above society's average will not have any weight applied to their QALYs.

The second problem for Williams is that his account entails that the elderly get lower priority for quality of life improvements than the young. It would seem harsh to give pain relief to someone young but not someone older in the same degree of pain, simply because the older patient had exceeded society's quality-adjusted life expectancy. In response, as a pluralist I need to admit that sometimes, if everything else is equal, a difference in quality-adjusted life expectancy might tip the balance in favour of the younger patient. But perhaps this is not as counterintuitive as it seems. This would only happen in a case where we can choose only one patient, and there are no other morally relevant differences between them, or at least no difference big enough to compensate for the difference in quality-adjusted life expectancy. In such a case, perhaps it is not so counterintuitive that decision could be swung by a difference in quality-adjusted life expectancy.

7.9 Conclusion

I conclude that if a patient's quality-adjusted life expectancy is below society's average, the benefits they can expect from treatment should be up-weighted by a small factor.

8. Are there rule of rescue obligations in healthcare allocation?

8.1 Introduction

In this chapter I will consider whether a "rule of rescue" obligation is ever applicable in healthcare allocation. By this I mean (roughly) an obligation to help an individual whose life is imminently at risk, where the intervention is relatively costly and therefore does not maximise the expected benefit we can produce with the resources at our disposal. Outside healthcare, such an obligation sometimes seems applicable. For example, consider the Chilean government's 2010 rescue of the miners stuck down the Copiapó mine, or the Australian government's 1997 rescue of the lone yachtsman Tony Bullimore, lost in the southern ocean after his boat had capsized. Suppose that on one of these occasions, a minister had announced that he would not authorise the rescue because the cost per life year was estimated to be 20% above the normal threshold for healthcare interventions, and she intended to use the budget to extend more lives via healthcare interventions. I suggest this would have prompted strong protests from the public, indicating that intuitively at least, a strong obligation was applicable despite the relatively poor cost-effectiveness of the rescues.

I here investigate whether our rule of rescue intuitions can be vindicated. I argue that our rule of rescue intuitions can be explained in terms of a plurality of considerations, which together have intuitively satisfying implications for the cases which trigger such intuitions. This has various implications for healthcare policy, particularly in relation to life-extending treatments.

After clarifying some ambiguities in the rule of rescue debate, I start the investigation proper by considering evidence that some healthcare policy-making is driven by rule of rescue intuitions. I then consider the most popular account of our rule of rescue obligations in terms of the identifiability of the prospective victims (for example, the prospective rescuer knows that a particular person is going to die). I conclude that identifiability is not a morally relevant factor. However, it is closely associated with another feature involving the distribution of risk, and this is more defensible as a morally relevant feature. Specifically, our behaviour in many rule of rescue cases can be explained in terms of an intuitive obligation to prioritise life-extending interventions for people who face a high risk of an early death. I will argue that this intuitive obligation in turn can be explained in terms of another principle I defended in the last chapter, viz., a principle saying that we should prioritise life-extending interventions for people whose quality-adjusted life expectancy is below society's average.

However, this account is incomplete. This principle does not explain all of our rule of rescue intuitions. In particular, it does not explain why we prioritise life-extending treatment for those who are given very little notice of their death rather than those who are given more notice. We are averse to letting people die suddenly. I give a new account that explains this in terms of the value of people having time to reconcile themselves to their death, sort out their affairs, say their goodbyes to friends and family. This account explains why our rule of rescue intuitions are stronger in cases involving sudden death than in cases involving sudden quality of life impairments: there is relatively little benefit in being warned about an imminent quality of life impairment sooner rather than later; whereas in the case of one's imminent death, it makes a huge difference to be given some warning.

So the value of an extra month or year depends on the context: an extra month is worth a lot for someone who has just been diagnosed and who is only expected to live a month to live without treatment. But it is worth less for someone is expected to live two years without treatment.

In summary, our rule of rescue intuitions can't be explained by a single unitary rule of rescue principle. Rather they are explained by a plurality of considerations. In healthcare, these sometimes give us a reason to relax our normal assessment criteria in terms of cost and benefit maximisation.

Finally I evaluate NICE's end of life premium in the light of rule of rescue obligations. I propose various ways NICE's end of life premium falls short when seen as an attempt to capture rule of rescue considerations. For example, NICE defines its timings in terms of how long the patient can expect to live as at the time of the treatment decision, whereas my principle implies that the timings should be specified in terms of time from diagnosis.

My principle also implies that in prioritising preventive measures, the priority should go to heart attack prevention rather than preventing slow growth cancers which get detected early, other things being equal, on the grounds that heart attacks kill suddenly, whereas a patient who learns early about a slow growth cancer has notice of their death.

8.2 Clarifying the Rule of Rescue Debate

Before we can make progress on the rule of rescue debate, we must clarify the topic. The debate is hampered by a couple of obscurities. First, there is ambiguity in the very term with which the debate is denoted. As Mark Sheehan says,

> is unclear how the 'Rule' is to be understood. From one perspective it appears to be a rule of morality but from another it describes a fact of human psychology.[335]

In line with this, we can distinguish two rules of rescue. In relation to a certain set of conditions X involving an individual facing imminent death, we can assert:

1. A descriptive rule of rescue: In conditions X, people will ordinarily go to great lengths to rescue the individual.
2. A normative rule of rescue: In conditions X, one *should* go to great lengths to rescue the individual.

Albert Jonsen, who coined the term "rule of rescue" as applied to this debate, used it primarily to denote an aspect of human psychology. Jonsen finds that in applying the "felicific calculus"[336] to healthcare technologies, he has encountered a "barrier" consisting of the "the imperative to rescue endangered life". He describes the dilemmas his policy-makers faced:

> Should we encourage the development of the artificial heart? Of course we must, it is said, because it rescues the doomed from certain death.[337]

and so

> We reach a conclusion contrary to the utilitarian principle: We benefit a few at cost to many.[338]

and Jonsen says

---

[335] Sheehan (2007), p. 352

[336] Jonsen here uses a phrase originally employed in Bentham (1789)

[337] Jonsen (1986), p. 172

[338] Jonsen (1986), p. 174

> I call this barrier the rule of rescue.[339]

We must take Jonsen as using the term "rule of rescue" to denote a psychological fact, since he does not commit himself on whether or not this barrier represents a genuine moral norm; he leaves the normative question quite open:

> I am not claiming that this impasse is either salutary or malign; I am not suggesting that we are better off living within this limit or breaking through it. I merely report my experience with serious, conscientious efforts to discern the utility of medical technology.[340]

Similarly, David Hadorn also sees the rule of rescue at work in public opinion regarding healthcare priorities:

> there is a fact about the human psyche that will inevitably trump the utilitarian rationality that is implicit in cost-effectiveness analysis: people cannot stand idly by when an identified person's life is visibly threatened if effective rescue measures are available.[341]

Hadorn goes on to adopt Jonsen's "rule of rescue" term for this fact about the human psyche. Similarly, McKie & Richardson use "rule of rescue" to denote an impulse, leaving it open whether the impulse might be ethically justifiable.[342]

However, other writers use the term to invoke the normative rule, such as Bettina Schöne-Seifert[343] and Hughes & Walker.[344]

This possibility of different usages need not be a problem as long as one is clear in one's usage. I will use the terms "rescue disposition", "rescue behaviour" and "rescue intuitions" in discussing the descriptive rule of rescue, and "rule of rescue obligation" to invoke the (purported) normative rule of rescue.

The second obscurity is a lack of clarity over what is being discussed. The main problem is a lack of clarity regarding the circumstances in which rule of

---

[339] Jonsen (1986), p. 174

[340] Jonsen (1986), p. 174

[341] Hadorn (1991), p. 2219

[342] McKie & Richardson (2003)

[343] Schöne-Seifert (2009)

[344] Hughes & Walker (2009)

rescue behaviour tends to be triggered, or in which our purported moral obligation is triggered. All are agreed that these circumstances include that someone is facing imminent death. But this is not sufficient. A 100 year-old on her death bed after a long illness does not ordinarily trigger rule of rescue behaviour and would not ordinarily be held to trigger rule of rescue obligations (we will not up-weight any benefits she stands to get from treatment). So any plausible account must incorporate additional "trigger circumstances", viz., situational features which are held to trigger the behaviour or obligation. All accounts do indeed offer additional conditions. For example, the following have been defended as the relevant features (I focus here on descriptive conditions, but most of these have been offered as possible normative conditions too):

Identifiability (a particular individual is known to be at risk)

A reasonable chance of effective rescue

Unusual situation (exceptional for the prospective rescuer)

Shock (the circumstances are extreme or unusual, such as to cause shock)

Acceptable risks or costs to the rescuers[345]

This lack of clarity regarding the trigger circumstances is a difficulty, since there is *also* a fundamental disagreement over whether there is any such thing as a rule of rescue obligation. This makes for a confusing combination of disagreements - a fundamental disagreement over the moral status of rescue acts *and* disagreement over which acts are at issue, or at least the circumstances in which the behaviour or purported obligation is triggered. Most other normative debates start with a better-defined focus. I suggest this is why most papers on rule of rescue do not open with a statement of the general principle to be considered or a general characterisation of the acts of interest, but instead open with a set of paradigm cases such as those I opened with: the miners stuck down a mine or yachtsman lost at sea. This helps to locate the topic whilst avoiding controversial commitments as to the type of circumstances in which the behaviour or obligation is triggered. This leads the writer into a casuistic attempt to draw out the morally relevant features of the paradigm cases, features which can then be identified in other situations where the writer can argue that similar behaviour is called for.

---

[345] Permutations of these have been defended by, for example, McKie & Richardson (2003), Orr & Wolff (2014), Schöne-Seifert (2009)

This lack of clarity in the focus of the rule of rescue debate has consequences for the prospects for the debate. If it is demonstrated that a particular circumstance which was held to be morally relevant is actually not morally relevant, that will not close off the question of whether there is a rule of rescue obligation, since it might be possible to characterise paradigm rule of rescue situations in terms of other features which *are* morally relevant. Furthermore, a morally relevant feature might not even have been noticed yet as a feature common to rule of rescue situations, let alone defended in terms of its normative status. I will argue our rule of rescue intuitions are explained by a plurality of considerations rather than a single unitary rule of rescue principle. The nature of the topic makes it difficult to be sure that an account is comprehensive, but I aim at least to characterise all the considerations that apply in the context of healthcare allocation.

However, despite the difficulty of characterising the conditions in which a rule of rescue obligation is triggered, all are agreed that rule of rescue behaviour involves a relative disregard for costs (whether opportunity costs or financial costs). So, for the purpose of this chapter, I will define "rule of rescue" behaviour as any behaviour involving an intervention to help an individual whose life is imminently at risk, where the intervention is relatively costly and therefore does not maximise the expected benefit we can produce with the resources at our disposal.

I will follow the classic procedure of identifying salient features of paradigm rule of rescue cases as the basis of my account of the obligations that justify this behaviour.

8.3 Evidence That Some Healthcare Debates Are Driven by Rule of Rescue Intuitions

I begin my investigation by considering some cases where healthcare policy debates seem to have been driven by rule of rescue intuitions. The first case involves the occasions when healthcare policy-makers refuse funding for a new, life-extending cancer drug on grounds of expense. The emotive response is reminiscent of the kind of response one might expect to the imaginary government minister in the Chilean miners case or the Tony Bullimore case. For example, prior to the introduction of its end of life premium in 2009, NICE regularly came up against public opposition to its strict application of benefit maximising criteria. In 2005, after NICE refused funding for the breast cancer drug Herceptin, Health Secretary Patricia Hewitt questioned a ruling by Stoke local health bosses not to fund the drug for a patient. After mounting pressure they reversed the decision.[346] Other local health bosses also fell into line with the Stoke decision, and started funding Herceptin. I take this was largely in response to a sense that the objectors had public opinion behind them. I also construe public opinion as an intuition that the refusal of funding was wrong. Cases like this motivated the introduction of NICE's end of life premium not long after, in 2009, whereby NICE relaxed the cost-effectiveness threshold for life-extending treatments for patients expected to live less than two years without treatment.

A well-known case from the US illustrates similar intuitions at work. This is the case of the first Oregon healthcare plan of the early 1990s. According to a widely-cited analysis by Hadorn, the plan foundered because it was guided entirely by cost-benefit comparisons:

> Specific examples taken from a single page of the 161-page list illustrate the problem. Surgical treatment for ectopic pregnancy and for appendicitis are rated just below, or as less important than, dental caps for "pulp or near pulp exposure" and splints for temporomandibular joint disorder, respectively. This priority order occurred despite the fact that the former surgical procedures are virtually 100% effective in treating otherwise generally fatal conditions, while the latter conditions are minor and may resolve even without treatment. This counterintuitive preference order did not occur as a result of faulty data, as was suggested by OHSC, or by chance, but as an inevitable consequence of the application of cost-effectiveness analysis.[347]

As noted, in Hadorn's view the problem was not faulty data, nor faulty analysis. Hadorn judges that the estimates of cost-effectiveness for various treatments appear reasonable, but:

> These reasonable estimates did not translate into reasonable (relative) priority ratings, however. Although both surgical procedures for appendectomy and ectopic pregnancy were correctly estimated to entail a far higher level and duration of benefit than either of the two minor treatments, the relatively high costs of surgery effectively neutralized

---

346 Triggle (2005)

347 Hadorn (1991)

> these outcome considerations, producing nearly identical priority ratings for all four treatments.[348]

But although Hadorn thinks the ranking of treatments was based on a reasonable analysis of their cost-effectiveness, still, he argues that the attempt to apply cost-effectiveness analysis rigorously foundered in the face of a very human propensity:

> people cannot stand idly by when an identified person's life is visibly threatened if effective rescue measures are available.[349]

Once again, diseases which threaten imminent death seem to have a special status in terms of ordinary intuitions, and these intuitions are strong; policy-makers can't ignore them.

There are also various other cases in which a similar principle seems to be at work. For example, consider the privileged status of treatment vis-à-vis prevention. One estimate was that for a given healthcare budget, you could get 10 QALYs from dialysis of kidney patients or 1197 QALYs from stop-smoking campaigns.[350] Yet we give dialysis to everyone who needs it, no doubt taking money from public health campaigns that would save more lives. Once again, it seems that patients at risk of imminent death merit special status, intuitively.

8.4 Can we Explain Rule of Rescue Obligations in Terms of the Identifiability of the Prospective Victim?

I now propose to investigate whether we can vindicate rule of rescue intuitions. The most popular candidate for a morally relevant feature in rule of rescue cases is the identifiability of the prospective beneficiaries of the rescue. Before examining the arguments, it should be noted that identifiability comes in more than one form. Our conclusions in this debate could depend on what form of identifiability we take to be invoked in an argument. A prospective rescuee is clearly identifiable in one sense if someone else can name them. But this is a somewhat narrow construal; not all rescuees are known by name to their rescuers. Alternatively perhaps a prospective rescuee is relevantly

---

348 Hadorn (1991)

349 Hadorn (1991)

350 Harris (1987) quoting BBC1 (1986)

identifiable if someone else can point to them and say "That person is at risk of imminent death". But again, prospective rescuees in most paradigm rule of rescue cases were not visible to their rescuers prior to the rescue. But perhaps in the same vein, a prospective rescuee is relevantly identifiable if someone else knows that there is a particular person who uniquely satisfies a certain, non-tautologous[351] description and who is at risk of imminent death, and the description picks out the prospective rescuee. So I am identifiable in this sense if I am stuck down a mine facing imminent death and you know that there is someone stuck down that mine facing imminent death, even if you don't know my name and can't see me. This is the most inclusive construal of the ones I have considered, in that it applies in all the cases the previous ones applied in. I would expect it to be the one on which most accounts in terms of identifiability implicitly rely.

Many descriptive papers characterising the human disposition to perform a rescue in paradigm rule of rescue cases take identifiability to be the key feature of the initial conditions in which rescue behaviour is triggered. For example, Hadorn's characterisation of the behaviour in question is that "people cannot stand idly by when an identified person's life is visibly threatened"[352]

But the question is whether we can explain why identifiability is morally relevant. Anyone who wishes to say that identifiability is morally relevant needs to say who the prospective victim is identifiable to. One option is that identifiability consists of being known to the public at large. McKie et al offer a defence of identifiability understood in these terms, though initially they acknowledge reasons to be sceptical about the moral relevance of any such feature:

> the RR typically comes into play when the life of an "identified" individual is "visibly threatened" and rescue measures are available. In the case of anonymous individuals, by contrast, there is no opportunity for "human nature" or the "human psyche" to override the "utilitarian rationality" of [cost-effectiveness analysis]. But being identified does not seem to be a morally relevant criterion for discrimination. After all, those

---

[351] To explain "non-tautologous": someone facing imminent death is not relevantly identifiable just because I know that there is someone facing imminent death who is facing imminent death.

[352] Hadorn (1991). Other writers who take identifiability to be the condition that triggers rule of rescue behaviour include McKie & Richardson (2003); Cookson et al 2008; Hope (2004).

> anonymous individuals who quietly die from preventable cancers in hospital wards are no less real than the trapped miner or lost sailor whose life is "visibly threatened", and their families will grieve no less for their deaths. ... Why should those who are lucky enough or manipulative enough to attract media attention be thought to have a special claim on resources?[353]

However, McKie et al end up defending the relevance of "public identifiability" after all, on the basis of the possibility that, as Hadorn points out,

> the media coverage given to an individual denied access to lifesaving treatment would result in "continual societal discomfort and shame"[354]

McKie et al then offer a utilitarian argument for countering such feelings in society. They say it may be

> desirable to promote the self-perception of being a humane society by accepting the inevitability of some media-based priorities. Utilitarians can justify this by pointing out that in this way social welfare may be maximised overall.[355]

However, I contend that this defence fails. There are two objections. First, when healthcare policy-makers refuse funding for a drug, going on opinions expressed in the media, I see little evidence that society feels shame. I take it this is because, to the extent that people think a decision to refuse funding for a treatment is wrong, they blame the policy-makers for that decision.

And second, even if the public feels shame, it would be wrong to pander to that shame by favouring high-profile patients with a weak claim over the lower profile patients who actually ought to get treated. Some kinds of "social welfare" do not count in favour of the policies that produce them. As an analogy, consider the case of mental health. For many years it was routine for those working in psychiatry to complain about the relative lack of investment in research into mental disorder compared with other conditions such as cancer.[356] In 2001 the WHO concluded that "Mental and behavioural disorders are estimated to account for 12 per cent of the global burden of disease, yet the mental health budgets of the majority of countries constituted less than 1 per

---

[353] McKie & Richardson (2003), p. 2408

[354] McKie & Richardson (2003) quoting Hadorn, 1991a, p. 2219

[355] McKie & Richardson (2003), p. 2147

[356] Wahlberg & Rose (2015)

cent of their total health expenditures."[357] The result was the development of a powerful movement to address 'the grand challenge in global mental health."[358] Intuitively this is a positive development. Treatment of physical impairments may produce more "social welfare" than treatment of mental health problems, because physical impairments are more visible, but that doesn't mean we should prioritise physical impairments. The public's uninformed sympathies can sometimes be disparaged as a basis for policy. Returning to the case at hand, the public's pleasure at seeing patients with higher profile diseases get priority for treatment may result from the public's ignorance of some of the relevant facts (such as lower profile deaths of other patients), or from a failure to think carefully about the issues. If so, the public's pleasure does not count in favour of prioritising those patients. We must do the right thing if we can, and a morally wrong policy can't be justified by the mere fact that public opinion favours it (unless unusual circumstances make it impossible to do the right thing). Instead, policy-makers should explain their policies, for instance by highlighting the lower profile patients with a stronger claim to treatment. If despite thorough education, the public still feels shame in respect of the high profile patients with a weaker claim, then perhaps the best explanation is that the shame is justified; perhaps the public has connived at a wrong in the case at issue. The wrong must either be that healthcare resources were distributed wrongly, or that there were not enough resources, in which case the public must accept higher taxes.

Another possible construal is identifiability to healthcare professionals and/or policy-makers.[359] But we still need to explain *why* such identifiability is important. One writer who has tried to explain it is Sheehan. Sheehan first distinguishes between agent-relative obligations and agent neutral obligations. For example, although I may have an agent-neutral obligation to look after other people's children in certain circumstances, this agent-neutral obligation may sometimes be trumped by an agent-relative obligation to look after my own children, even if other people's children would derive greater benefit from the things I do for my children. He then argues that:

> the general empirical fact about our human proclivities may be indicative of an agent-relative kind of obligation. That is, the fact that we are strongly inclined to use large amounts of resources (or otherwise risk

---

357 WHO (2011)

358 Collins et al., 2011

359

> great cost to ourselves) in order to save an identifiable individual suggests that we have a prima facie agent-relative obligation to those in need of rescue. Importantly, in some cases this obligation may override our agent-neutral obligations to do the most good with the resources we have. The thought is that we stand in a special relationship, perhaps a relationship of circumstance, to those in need of rescue and as such have a prima facie obligation to save them. Since this obligation looks to be an agent-relative obligation, there may be circumstances in which it is more pressing than an agent-neutral obligation, say, to use resources to do the most good.[360]

Sheehan takes this to be a vindication of our ordinary practice:

> the idea is that good explanations function as justifications for that of which they are good explanations. So if an explanation does a good job of capturing a distinction within our ordinary practice, we are justified in thinking that the distinction and its elements exist. The agent-relative/agent-neutral distinction functions in this way - it is intended to register a significant distinction within the ordinary obligations that we take ourselves to have. If the distinction between agent-relative and agent-neutral obligations is a good one and so genuinely captures a relevant feature of our practice, then we can take ourselves to have these kinds of obligations. That is, we are justified in thinking that we have these kinds of obligations.[361]

Sheehan then runs through some of the circumstances which might cause people to prioritise a prospective victim. He observes that we might feel an obligation to make special efforts to rescue someone who served us in the local supermarket, or people who live near us. He suggests that circumstances like these might create agent-relative obligations to rescue such people, should they need it. Sheehan concludes that "it is not identifiability that matters in the rule of rescue but the way in which things like identifiability, proximity and acquaintance contribute to our agent-relative obligations."[362]

The problem with his argument is that Sheehan simply takes our intuitions as evidence of normative considerations without further argument. Jecker objects to Sheehan as follows:

---

[360] Sheehan (2007), p. 359

[361] Sheehan (2007), p. 359

[362] Sheehan (2007), p. 361

> [Sheehan's analysis] moves too quickly from a purely descriptive observation of our tendency to want to help people in peril to a clearly normative conclusion concerning our ethical obligation to help. To show why this move is so problematic, consider an analogous case. Humans have raped and plundered for all of human history. Yet it hardly follows that there is a prima facie obligation to respond to lust with rape, or to rage with plunder. We need instead to intercede before we make the move from feeling to action by reflecting on the ethical valuations that support the putative obligation.[363]

Jecker's objection is that the mere existence of a disposition does not itself provide sufficient grounds for an assumption that the resulting behaviour can be justified. Something else is needed.

However, Sheehan offers what seems like a response to this when he says his argument is put:

> on the assumption that much (or at least enough) of our ordinary practice is in order ... I suspect that those who would deny the plausibility of the assumption would also deny the overall project of the paper.[364]

Sheehan is basically arguing that in order to do normative theorising the way he is trying to do it, he has to assume that "our ordinary practice is in order". I take it he means that he, along with most other ethicists, needs to assume that some of our intuitions are true as a springboard for philosophical enquiry.

There is something in this as a characterisation of philosophical procedure. But to continue the argument on Jecker's behalf, although I grant Sheehan that for the purpose of moral theorising it is necessary to assume that a reasonable number of our intuitions are "in order", it is not necessary to assume that *all* our intuitions are in order. In pursuit of reflective equilibrium we articulate intuitively compelling general principles and then consider whether those principles have intuitively satisfying implications for specific cases. Sometimes, we may find a principle's implications for specific cases unacceptable and therefore reject the principle; but on other occasions, we may accept a principle with one or two counterintuitive implications because it is otherwise strongly supported (for example, it is intuitively compelling in its own right and has intuitively satisfying implications for a wide variety of

---

[363] Jecker (2013)

[364] Sheehan (2007), p. 365, note 25

other cases). Thus, contrary to Sheehan, we may sometimes find that our best moral theory implies that certain aspects of ordinary behaviour are impermissible, despite the overall assumption that a reasonable number of our intuitions are "in order".

So Sheehan needs to do more work than he does here to vindicate rule of rescue behaviour. To see what extra work is required, note that two key claims come up in Sheehan's account:

1. The claim that we have distinct agent-relative obligations and agent-neutral obligations.
2. The claim that rule of rescue behaviour is obligatory in the case of, say, a prospective victim who served the prospective rescuer in a supermarket or a prospective victim who lives near the prospective rescuer.

Sheehan starts by defending the claim that we have distinct agent-relative obligations and agent-neutral obligations. He then infers from this (and our ordinary intuitions) that rule of rescue behaviour can be justified in terms of an agent-relative obligation in the types of case he mentions. But it is unclear how he can infer the conclusion from the premises. The distinction between agent relative and agent neutral obligations is just a conceptual scheme. Sheehan's argument that the distinction is sustainable does nothing to show that one or other concept is applicable in rule of rescue cases. In particular, his claim that we have agent-relative obligations to our children does not show that we have agent-relative obligations in every other case where we ordinarily take ourselves to have them. To draw any conclusions about our rule of rescue obligations from the case of obligations to our children, Sheehan would need to argue from analogy. For this he would need the cases to share some morally relevant feature. He does not do enough work to show that they share any such feature. If the supposed shared feature is that we have agent-relative obligations in both cases, this would presuppose the conclusion that Sheehan is trying to draw, viz., the conclusion that we have an agent-relative obligation in the situations which generally trigger rule of rescue behaviour.

So Sheehan's defence fails. I also suggest there are counterexamples to the idea that identifiability carries weight. For example, consider the case of the 100 year-old on her death bed after a long illness. There is no intuitive obligation in such a case to perform a rescue or relax our normal cost-effectiveness criteria for life-extending treatments. Yet the patient is identifiable to potential rescuers.

Also, consider a case where patients are sometimes identifiable to policy-makers: the case of panels that review Individual Funding Requests. IFRs are requests submitted by hospital consultants to Clinical Commissioning Groups on behalf of patients who have an exceptional characteristic such as a comorbidity, so they don't fit on one of the standard care pathways, but the consultant thinks a treatment will be beneficial for them. For IFR panels, patients are identified. In some cases, patients are allowed to attend the panels. In other cases, they are identified by name. Even if they are not, they are at least picked out individually to the same extent as, say, miners stuck down a mine might be for a politician deciding whether to launch an operation to get them out, in that they constitute known individuals subject to a known risk. Suppose that one or other of these forms of identification is available to panel. Still, intuitively it would seem quite reasonable for the panel to apply the same cost-effectiveness criteria as are applied in standard commissioning decisions applying to "statistical" patients. To the extent that data is available,[365] it would not seem impermissible to consider what it implies for the standard cost-effectiveness criteria, especially if the panel are allocating money from the same overall budget as in standard commissioning decisions. In fact, it would seem quite unfair to other patients if IFR decision-makers *didn't* apply the same criteria as are applied to other patients, to the extent that this is possible.[366]

To make this point vivid, suppose that policy-makers approve an IFR for a given type of case, and then decide they need to formulate a standard care pathway for that type of case. They formulate criteria based on the same broad cost-effectiveness criteria as applied in all other therapy areas. So, when the new pathway is published the week after the IFR case, a hospital consultant tells a clinically identical patient that according to the newly published criteria, the patient will not get funding for the drug they need. The patient knows that the first patient was clinically identical, and was funded, and therefore asks why he was refused. When the hospital consultant seeks clarification from policy-makers, they reply that they met the first patient but not the second one, and so different criteria were applied. This would seem very unfair. So, even when patients are identifiable to policy-makers, it doesn't

---

[365] IFRs tend to involve small patient populations for which there is less data available.

[366] Admittedly, the data is often not robust, because of the small patient populations. But there is usually *some* data, and if there is enough data to draw some very approximate conclusions about cost-effectiveness, it would not seem impermissible to do so, albeit allowing for large margins of error.

seem that policy-makers are subject to any rule of rescue obligations which might be said to be triggered by such identifiability.

I conclude that identifiability is not a morally relevant feature.

8.5 First Morally Relevant Factor: Distribution of Risk (Williams' Egalitarianism of Quality-Adjusted Life Expectancy)

Rather than trying to defend the moral relevance of identifiability itself, Karen Jenni and George Loewenstein focus on a closely associated feature. They explain the "identifiable victim effect" in the public preference for - or greater toleration of - wide rather than narrow distributions of risk.[367] For example, studies have shown that people are more concerned about the risks of a vaccination program if only 10 per cent of the population are susceptible to adverse side effects, than if the whole population is susceptible, even if the number of people who will die will be around one thousand in either case.[368] Jenni and Loewenstein speculate that the "identifiable victim effect" may be explained by this antagonism towards concentrated rather than shared distributions of risk.

Jenni and Loewenstein end up concluding that public opinion is indefensible, on the grounds that someone's level of risk is not determinate; your assessment of someone's risk will depend on what group you assign them too.[369] But we can vindicate the public's views in terms of a quite

---

[367] Jenni and Loewenstein (1997). Thanks to Rob Lawlor for pointing me to this way of vindicating the concern with identifiability in terms of a concern with risk.

[368] Ritov & Baron, 1990

[369] They explain public opinion in terms of reference groups. For example, on their account, the public considers rescue to be obligatory in paradigm rescue situations because "n out of n people will die if action is not taken" (Jenni & Loewenstein, 1997). As McKie et al comment, "The crucial question is, of course, what is the morally relevant reference group? When people’s behaviour conforms to the RR the relevant reference group becomes the identifiable person. But why is it the identifiable person, rather than those anonymous individuals who could (individually and collectively) benefit, perhaps even more, from the expenditure of the same resources?" (McKie et al, 2003). Jenni and Loewenstein go on to object that "Given that reference group size is often a matter of framing - a reference group of arbitrary size can be specified for virtually any hazard - a blanket endorsement of a policy that treats fatalities differently based on what proportion of the reference group they compose is normatively dubious." For this reason, Jenni and Loewenstein argue that public opinion regarding distribution of risk cannot be vindicated.

ordinary, defensible conception of probability of harm.[370] This provides the materials for a principle that says that people who are likely to die soon should be prioritised over others because we should equalise people's chance of dying imminently if we can. For example, suppose we urgently need to vaccinate a town of 20,000 to protect against a plague that will otherwise kill them all. We have the choice of two vaccines. The first will bring a 50% risk of death to a subgroup consisting of 2,000 people. The second will bring a 5% risk of death to all 20,000 people. Thus both vaccines can be expected to cause around 1,000 deaths. Jenni and Loewenstein's evidence suggests that public opinion would favour the second.

I aim to vindicate the intuitions the public could be expected to have about this case. More broadly I aim to defend a principle that implies that shared distributions of lifetime risk are better than concentrated distributions of lifetime risk. I do not believe I need to offer any new principle to do so. It is implied by a principle I have already defended, Williams' egalitarianism of quality-adjusted life expectancy, which I introduced in Chapter 7 on fair innings. We can see the preference for shared distributions of risk over concentrated distributions of risk as motivated by a kind of egalitarianism of risk. Williams can be seen as generalising this egalitarianism of risk to all health outcomes. I will summarise the implications of a Williams-style principle for the vaccines case. We have seen that he thinks not only that interventions are better to the extent they produce valuable benefits, but also that benefits have more value to the extent they increase equality. So, in the vaccine case, both vaccines produce the same benefit, but one vaccine gives a small group of people lower quality-adjusted life expectancy than everyone else, so Williams' principle would down-weight the benefits of that vaccine to reflect the associated inequality. In contrast, the other vaccine ensures that everyone's quality-adjusted life expectancy is reduced equally, so there is no more inequality with the vaccine than without, and as a result the benefits from that vaccine do not get down-weighted so much. Assuming that both vaccines produce the same total amount of benefit, Williams' account will imply that the vaccine that reduces everyone's quality-adjusted life expectancy equally is the better option.

---

[370] I avoid the problem with arbitrarily selected reference groups. Many accounts of probability allow us to conceive probabilities in terms not involving reference groups, such as Popper's account of probability (Popper, 1957).

I will explain how this could work with some example figures, in order to show how a Williams-style prioritarianism could explain some of our intuitions regarding distribution of risk. For each intervention to be assessed, we calculate the total benefit associated with the outcome and we then weight that benefit according to the level of inequality[371] (the more inequality an outcome has, the greater the reduction we apply to the benefits associated with that outcome). Now, suppose we are making decisions for a society of 20,000 people. For ease of calculation, suppose that each person stands to live for ten years in full health if they survive, so they get ten QALYs. That means the vaccine which puts all 20,000 people's lives at risk is expected to produce 190k QALYs (95% of the population having ten years in full health). However, there will be some existing social inequality, which the vaccine will not eliminate. This existing inequality requires us to apply a negative weighting factor to the QALYs. Let's suppose for the sake of argument that this factor is 10%. So we reduce the unweighted QALYs by 10%, which means a reduction from 190k to 171k QALYs. On the other hand, the vaccine which puts 2,000 people's lives at risk is expected to produce 190k QALYs but also to increase inequality. Those 2,000 people will end up with a quality-adjusted life expectancy much lower than society's average, because they get a 50% chance of dying immediately, while the rest of their society has a roughly zero chance of dying immediately.[372] The result will be that inequality of quality-adjusted life expectancy produced by this vaccine is much greater than is produced by the other vaccine. So the QALYs associated with this outcome will get reduced by more. Let's suppose for the sake of argument that the weighting factor in this case consists of a reduction of 20%, which means a reduction from 190k to 152k QALYs. So the vaccine which puts all 20,000 people's lives at risk produces the greatest weighted benefit. It produces the same unweighted benefit, but that unweighted benefit does not get reduced as much because of the lower inequality in the outcome. Thus a Williams-style prioritarianism would favour the vaccine which distributes the risk throughout the population, in line with the public's intuitions.

---

[371] Various measures of inequality are available, such as Gini (1912). I leave the question of which measure is most suitable until another occasion.

[372] I have assumed that the 2,000 are a random selection of the population. The only way equality could be increased in this scenario is if the 2,000 whose life expectancy is reduced by the vaccine happen to be the 2,000 with the longest life expectancy without the vaccine, and their life expectancy without the vaccine is at least double that of the other people.

Similarly, if a minister argued that the money allocated to saving the Chilean miners would be better spent on preventive public health measures, because such measures would prevent many more deaths, the response based on Williams' account would be that the miners' merit priority on the grounds that their quality-adjusted life expectancy is lower than society's average life expectancy. The population who could be helped with the preventive measures face a far lower chance of immediate death than the miners, so that their quality-adjusted life expectancy is much closer to society's average life expectancy. So although we do not produce as much benefit by rescuing the miners as we could with the preventive measures, Williams' principle entails that the expected benefits for the miners should be up-weighted, and that could justify directing the budget towards the miners rather than towards the preventive campaign. Thus Williams helps vindicate ordinary intuitions with respect to the miners.

However, Bettina Schöne-Seifert has raised a counterexample against this sort of principle. If we prioritise people known to be high risk, then:

> Rather than screening people susceptible for a potentially fatal disease and treat them early, effectively and at low cost one would wait until later - only to treat the very same patients at higher suffering, with higher risk and higher cost.[373]

Such a policy looks plainly irrational. However, there is a response to this. We should look for the people who are *actually* at risk, whether they are known to us or not. There are individuals whose physiology or environment or lifestyle puts them at risk, even if their high level of risk is not immediately obvious. Screening is justified because it enables us to identify those high risk individuals, who would otherwise be unjustly neglected despite their high risk.

So the lesson we can draw from all this is that identifiability was an imperfect proxy for what really matters, viz., underlying risk. But we can also explain why identifiability seemed to matter. In paradigm rule of rescue cases, the risk to which individuals are exposed is manifested in dramatic and vivid ways that are very motivating. This explains our sense of urgency in such cases. Nevertheless the morally relevant features of such cases are also present in cases where the marks of risk are less obvious.

---

[373] Schöne-Seifert (2009), p. 424

However, there are some cases in which Williams' principle does not imply any priority for patients facing imminent death. In particular, consider older patients whose quality-adjusted life expectancy is above society's average because they have had a long life with good quality of life, such as a 100 year-old on her death bed after a long illness. Williams principle does not imply that her expected benefits from treatment should be up-weighted, since by hypothesis, her quality-adjusted life expectancy is above society's average. Is this acceptable? I suggest it is.[374] Such patients would not ordinarily trigger rule of rescue behaviour and do not intuitively merit higher priority for life-extending treatment than other patients.[375] Nevertheless, there are situations in which older people merit the same priority as younger people. I will consider these now.

8.6 Second Morally Relevant Factor: That the Negative Outcome Is Sudden

Williams' account does not appear able to explain all of our rule of rescue-related intuitions. In particular, he cannot explain why we prioritise life-extending treatment for those who are given very little notice of their death rather than those who are given more notice. For example, consider the choice between some 40 year-olds who have known since they were 20 that they were going to die in a month, and some 40 year-olds who have only just learned that they will die in a month of the same disease. I suggest that if we have the choice between giving 10 extra weeks to the first group or 9 extra

---

[374] For this reason I do not believe we should defend a principle in terms of current risk rather than lifetime risk, i.e. a principle implying that shared distributions of current risk are better than concentrated distributions of current risk. Such a principle would require us to up-weight any benefits that the 100 year-old on her death bed could expect from treatment, because she has a much greater current risk of dying than most other people.

[375] Perhaps if an older person faced an extraordinary death we would ignore their age (for example, if an older person was lost at sea or stuck down a mine). But I suggest such cases introduce additional considerations. One is that the prospective victim is isolated from friends and family; as I will argue in a moment, it is important that people have a chance to say goodbye to friends and family, and get their affairs in order, once they know they are about to die. Paradigm rule of rescue cases make this difficult. This consideration could carry weight even with older people, in certain ordinary healthcare cases. In addition, the prospective death in some rule of rescue cases seems cruel and unusual, and we have reason to prevent such deaths even among older people. However, this consideration will not often carry over to the healthcare context, since most patient deaths are not unusual enough.

weeks to the second group, we have reason to give the 9 extra weeks to the second group, even though their health gain is smaller.

Thus there seems to be an intuitive distinction between those who are given very little notice of their death and those who are given more notice. A Williams-style egalitarianism of quality-adjusted life expectancy cannot account for this intuitive distinction, since both groups have the same quality-adjusted life expectancy.

How to explain the intuitive distinction? The first group does have certain considerations on their side, to do with the Williams-style egalitarianism we investigated in the last chapter. But the second group has those considerations on their side and others as well; something to do with the suddenness of their imminent death. In summary, it looks like sometimes we are concerned not with how long a patient has to live, but with how much notice they've had of their expected death. In this I am supported by public opinion. Consider the following finding:

> The study results suggest that [NICE's current end of life] policy may be insufficient in two ways. First, whilst it is concerned with patients' remaining life expectancy, the supplementary advice does not distinguish between sudden and non-sudden disease progression. Findings from the pilot, coupled with an examination of the reasons given by respondents in the tick-box tasks, suggest that for many people the preference for prioritising the treatment of end-of-life patients is driven by concern about how much time the patients will have had to prepare for death.[376]

But as things stand this principle is puzzling. Why do we place weight on death that comes suddenly? Without an explanation, a principle in terms of the suddenness of someone's imminent death will look ad hoc. But an explanation is available, which is that we recognise the value of people having time to reconcile themselves to their death, sort out their affairs, say their goodbyes to friends and family, do that one important thing they've always wanted to do, go through whatever rituals help make sense of their life so far and their imminent death. This is not just for practical reasons, for instance it wouldn't be much compensation for a sudden death that one was one good terms with family and had just organised one's affairs and said goodbye to everyone in order to go on a long business trip. For many of us there is a

---

[376] Shah et al (2014), p. 397

specific value in closing one's life story in the right way, putting a full stop on it, in a dignified way.

8.7 Third Morally Relevant Factor: That the Expected Outcome Is Death

Another feature of our rule of rescue intuitions should be noted: they are stronger in cases involving death than other cases. They seem to be much weaker in the case of most quality of life impairments. For example, consider a modified version of the above case. We must choose between some 50 year-olds who have known since they were 20 that they were going to get rheumatoid arthritis in a month, and some patients of the same age (50) who have only just learned that they will also get rheumatoid arthritis in a month. Intuitively, there is little reason to favour those who've only just learned about the imminent onset of the disease over those who have known about it for 30 years. If those who knew about it already were expected to get a bit more benefit, that could swing the decision their way (however, there might be an exception in the case of severely disabling impairments; I cover these below).

At first sight this distinction between imminent death and imminent quality of life impairments is puzzling. In assessing the benefits of a treatment we assess its effects on both life expectancy and quality of life. Why don't rule of rescue intuitions also treat both kinds of impairment equally?

My explanation is that there is relatively little benefit in being warned about an imminent quality of life impairment sooner rather than later; whereas in the case of one's imminent death, it makes a huge difference to be given some warning. For example, I contend there is a huge difference between the situation of a 30 year-old patient who learns he will die in a year and the situation of a 31 year-old patient who learns he will die in a week - even though they will both die at the same age. In contrast, there is much less difference between the situation of a 30 year-old patient who learns that their illness will cause lifelong disability in a year and a 31 year-old patient who learns that the same illness will give them the same disability immediately. Of course there is *some* difference between the two patients in the latter case; the patient who is given warning of their quality of life impairment can work out a bucket list of things which will be more difficult or less enjoyable after they get the impairment, and then get some of those things done. But this is also an advantage enjoyed by the patient given warning about their death over the patient given no warning of theirs. And the advantage enjoyed by the patient warned of their death is much greater than the advantage enjoyed by the patient warned of their quality of life impairment, since we can't do any of the important things after we die, but we can do many of the important things after we get most quality of life impairments. In addition to this point, as argued above, the idea of sorting out one's affairs and making the right kind of exit from one's life is very important to most of us. This is a big advantage for the patient given warning of their death over the patient given no warning of theirs. There is no equivalent advantage for the patient given warning about their quality of life impairment over the patient given no warning of theirs. You do not need time to say goodbyes if you are going to be in a lot of pain in a month.

Against my claim that people would want notice of their death, it might be objected that having notice of one's death will make a person gloomy and depressed, and so worse off than someone who doesn’t know it is coming.[377] However, most people want to know they are dying, presumably for the reasons I have given.[378]

This restriction to life years means that, for example, if we face a choice between one group of 20 year-olds facing a lifelong headache starting tomorrow vs a group of 18 year-olds facing lifelong headaches in two years time, my account does not imply any priority for the 20 year-olds just because they lack notice of their headaches. If we could produce slightly more benefit for the 18 year-olds, all else being equal, my account implies priority for them. This may seem counterintuitive. If so, some explanations are available. For one, the comparison I have given assumes that both outcomes are equally certain. But in fact, in respect of the 18 year-olds expecting headaches in two years time, there would be greater uncertainty about whether they’ll actually need treatment: maybe they will not get the headaches after all; maybe they will die; maybe their condition will change; maybe we'll develop an alternative treatment for their condition.

It should also be noted that, to the extent that a patient's pain is disabling, to that extent my account could imply that such a patient should get priority. For example, suppose a patient faces a state of extreme disability such that they won't be able to get any of the important things done that they would want to

---

377 Thanks to Carl Fox for this point.

378 Smith (2014)

do before they die, such as communicating with family and friends. Then, on my account, we might have the same case for prioritising that patient as if they were facing death.

8.8 Conclusion: Implications For Health Policy

I have explained our rule of rescue intuitions in terms of a plurality of principles, which have various implications for healthcare. Where a patient has a very poor quality-adjusted life expectancy, a Williams-style egalitarianism of quality-adjusted life expectancy justifies us relaxing our cost-effectiveness criteria and prioritising that patient. This justifies much rule of rescue-type behaviour. In addition, I have argued that we should ensure that people have notice of their death so they have time to prepare. This gives us reason to prioritise life-extending treatment for those who are not expected to live long after the diagnosis of the condition that will otherwise kill them.

In many cases, these two considerations will be mutually supportive. But in many cases, rescue-type behaviour may be solely motivated by my principle that people should be given reasonable notice of their death. How does my principle compare with existing policies? NICE's end of life premium looks like it might be motivated by somewhat related considerations. I will evaluate the end of life premium as if it is an attempt to capture rule of rescue considerations, ignoring for the sake of argument the possibility that it was not motivated by such considerations.

NICE's end of life premium gives priority to patients satisfying the following conditions:

1. The treatment is indicated for patients with a short life expectancy, normally less than 24 months and;

2. There is sufficient evidence to indicate that the treatment offers an extension to life, normally of at least an additional 3 months, compared to current NHS treatment, and;

3. The treatment is licensed or otherwise indicated, for small patient populations.[379]

---

379 National institute for Health and Clinical Excellence (2009)

When the above condition are met, NICE's Appraisal Committee is advised to consider:

- The impact of giving greater weight to QALYs achieved in the later stages of terminal diseases, using the assumption that the extended survival period is experienced at the full quality of life anticipated for a healthy individual of the same age, and;

- The magnitude of the additional weight that would need to be assigned to the QALY benefits in this patient group for the cost-effectiveness of the technology to fall within the current threshold range.

Thus policy-makers are given quite wide discretion in how they respond to the above condition. Nevertheless the "direction of travel" is clear. Firstly, within limits they can to give greater weight to QALYs gained in these circumstances (in that they can act as if more QALYs are gained than actually are). Secondly, the QALY calculation can be based on an assumption full quality of life, so if patients satisfying the above condition have poor quality of life, that will not negatively impact the deemed cost-effectiveness of their treatments.

This policy has had an impact:

> NICE has applied EoL flexibilities in 25 TAs since the guidance was introduced. Of those, 18 have resulted in NICE recommending use ... what really matters is the cost-effectiveness threshold used when [Appraisal Committees] consider these treatments. The magic number, based on an average across all positive recommendations, seems to be around £49,000 per QALY.[380]

This compares with the normal threshold of £20,000 - £30,000. Thus the effect of this policy has been to relax the cost-effectiveness threshold for life-extending drugs.

How does this policy look from the perspective of the principles I have defended above? I will focus on the first two conditions in NICE's policy, and I will compare these with my principle stipulating priority for life-extending treatments for patients who lack reasonable notice of their death (the purpose of this being to state my account more clearly; Williams' principle does not need such work as it is already quite precisely operationalised).

---

380 Barham (2016)

In making my principle more precise the first question would be: how should the life expectancy terms be specified? NICE defines its timings in terms of how long the patient is expected to live as at the time of the treatment decision. As seen from the perspective of rule of rescue obligations, this must be seen as a mistake. For example, it means that a 100 year-old on her death bed after a long illness qualifies under the policy. Much as any death is to be regretted, I suggest that such a patient does not merit special priority for life-extending treatment. More generally, if I am right that part of the point of a policy like this is to ensure that people have reasonable notice of their death and time to sort out their affairs, someone who has known for many years that they can expect to die this year should not be treated the same as someone who has only just learnt. To preclude such cases, the timings would need to be specified in terms of time from diagnosis. The question is, how long from diagnosis can the patient expect to live without treatment and how long they can expect to live with treatment? On the basis of medical advice, or commonly known information about ageing, an elderly patient should have realised for some time before their death that their death was imminent.

The next question is, how long does someone need to sort out their affairs? Of course the longer the better, and it will be difficult to generalise, but I would hypothesise that after a few months, a law of diminishing returns kicks in. Without wishing to sound harsh, after a year, even someone with a complicated life should have been able to sort their affairs. So NICE's stipulation that anyone with less than two years to live qualifies for special treatment would seem to be slightly on the generous side, from the perspective of rule of rescue obligations. But of course, there are no sharp lines on this question, and anyway empirical research would be required to answer it properly.

Having made this decision there will be another, closely related decision, which is the question of how many extra years produced by treatment qualify for the premium. The same reasoning implies the same limit: whatever time is needed to sort out one's affairs, that should be both the maximum life expectancy without treatment to qualify for the premium, and also the maximum life expectancy *with* treatment to qualify for the premium. There should be no weighting for life expectancy improvements above the period deemed necessary to sort out one's affairs. If a treatment produces more years, of course they count as a benefit, but they are not specifically needed for someone to sort out their affairs, and therefore should not qualify for a premium which is specifically motivated by that need.

So the value of an extra month or year depends on the context: an extra month is worth a lot for someone who has just been diagnosed and who is only expected to live a month to live without treatment. But it is worth less for someone is expected to live two years without treatment.

Now consider NICE's stipulation that to qualify for the premium, a treatment needs to produce at least an extra three months. But this misses out perhaps the most valuable period that a terminal patient could be given. If I've been given a week to live and a treatment promises to add another week on top, that’s a huge difference. The requirement should only be that a treatment adds *something*.[381] Also in relation to this point, I would suggest that the weightings stipulated in the policy should be graduated. The first few month or two after diagnosis are extremely valuable, both for allowing the patient to come to terms with their death and for allowing them to sort out their affairs. Thereafter, returns diminish. I would propose that if a patient is expected to live less than, say, three more months, the weightings for additional life expectancy within that period should be high. Thereafter, the weightings should get lower.

Perhaps imminent death is not the only thing that should trigger such measures. Consider illnesses that lead to degeneration of mental capacities, or a coma that lasts the rest of one's life, such that one could not communicate with one's family.[382]

In addition, my principle should also be understood as implying priority, not only for those clearly at risk, but also for those who are at non-obviously at risk, where such individuals can be picked up with screening.

Finally, in terms of preventive interventions for healthy patients, my principle implies that the priority should go to heart attack prevention rather than preventing slow growth cancers which get detected early, other things being equal.[383] Heart attacks kill suddenly, whereas a patient who learns early about a slow growth cancer has notice of their death. Perhaps this implication favouring heart attack prevention is somewhat counterintuitive, but not unacceptably so. It is an acceptable cost for an account which has otherwise

---

[381] Of course if it only adds an hour, there will be very little benefit to weight, so even the weighted benefit will be minimal.

382 Thanks to Rob Lawlor for this point.

383 But of course, other things might not be equal. Having a slow cancer, and living with the knowledge you are about to die, might be a consideration that runs counter to the considerations I've outlined. Thanks to Rob Lawlor for this point.

intuitively satisfying implications. Intuitions are not infallible. The account's intuitively satisfying implications in other cases give us reason to rethink our intuitions in this case and/or see them as mistaken.

9. How should we trade off conflicting considerations in healthcare allocation?

I have reviewed some of the factors which ought to influence us in determining where to allocate healthcare. The questions remains: what to do. It won't always be easy to say which policies are obligatory or permissible, since the relevant considerations may conflict with each other. For example, we may not be able to benefit a patient with a severe disease as much as a less needy patient, so the concern to maximise the benefits we produce may conflict with the concern to help the neediest patient. In this chapter I will address how to trade off such conflicting considerations.

I consider whether we should rely on general trade-off algorithms, or instead deliberate over the best policy case by case. I consider arguments on each side. I will argue that we should rely on algorithms where we have good evidence that they are reliable indicators of the best policy. In line with this, I consider a process called Multi-Criteria Decision Analysis (MCDA), which involves getting policy-makers and stakeholders to deliberate over criteria and their relative weights with a view to producing an algorithmic policy-making model. I argue that on the evidence available about MCDA, the process of agreeing an algorithm can produce better quality deliberation, so even the opponents of algorithms have some reason to support them.

9.1 Should We Use Trade-Off Algorithms?

There is a dispute over the best way to take account of conflicting considerations in order to reach a policy decision in each specific case. Alan Williams and Richard Cookson, both health economists, defend an approach based on general trade-off algorithms, i.e. general principles which state how conflicting considerations are to be traded off across a great variety of cases. In contrast, the philosopher Stephen Toulmin argues that we should judge how conflicting considerations trade off in individual cases without trying to state a general trade-off principle.[384] I will argue that we should pursue Williams & Cookson's strategy of stating general trade-off principles, but without assuming it will always be possible.

The question is whether or not to rely on trade-off algorithms in policy-making, and I will start by showing what such an approach entails in the context of healthcare allocation by means of a couple of examples. The examples consist of a pure efficiency trade-off principle (the QALY system itself), and an example of an equity-efficiency trade-off principle covering trade-offs between patient need and benefit maximisation.

The QALY system itself is an example of a trade-off algorithm. The QALY system is a way of valuing health outcomes relative to each other. As such it facilitates trade-offs between such outcomes, licencing policy-makers to say that, for example, ten years with tetraplegia is worth more than five years with full quality of life. Moreover, QALYs do not merely licence trade-offs between quality of life and life years; they imply a vast matrix of value comparisons between different levels of quality of life.

For my second example of how general algorithms can prescribe value trade-offs, I consider a principle that weighs equity against efficiency. Consider a trade-off algorithm regarding rarity, to the effect that if one disease is a hundred times rarer than another then, all else being equal, the cost-effectiveness threshold for the rarer disease is doubled (based on a smooth gradient of increasing weights). Assume this policy is triggered at the current threshold for rarity, so that if a disease affects 5 people in 10,000, no special provisions apply, but if a disease affects fewer than that, increasing weights apply with decreasing prevalence. This algorithm specifies how much we should be willing to sacrifice on one dimension of assessment (the obligation to maximise benefit) in order to deliver on another dimension (the obligation to make our policy-making criteria somewhat inclusive of rare disease patients). It prescribes trade-offs between the two dimensions of assessment in a wide range of cases. For example, in the case of a rare disease patient whose disease only affects 0.05 people in 10,000, the principle implies that even if this patient does not stand to gain as much from their treatment as another patient whose disease is not deemed rare, they may nevertheless have their treatment approved for funding while the common disease patient's treatment is refused, because for the purposes of assessment we double the life years

---

[384] Actually Toulmin does not address healthcare allocation specifically, but his argument is a general argument which applies across a variety of areas of policy-making.

added by the rare disease patient's treatment, or other benefits they get from treatment.

9.2 Toulmin: Argument Against Algorithms

Should we rely on such algorithms? Toulmin says we shouldn't. Specifically, Toulmin objects to what he characterises as "the revival of a tyrannical absolutism in recent discussions about social and personal ethics"[385] whereby absolutists deny

> all real scope for personal judgment in ethics, insisting instead on strict construction in the law, on unfeeling consistency in public administration, and-above all-on the "inerrancy" of moral principles.[386]

Toulmin starts with a history of Rome[387] to illuminate the alternative approaches he wants to discuss. He observes that:

> For the first three hundred years of Roman history, the legal system made no explicit use of the concept of rules. The College of Pontiffs acted as the city's judges, and individual pontiffs gave their adjudications on the cases submitted to them. But they were not required to cite any general rules as justifications for their decisions. Indeed, they were not required to give reasons at all.[388]

In answer to the question of how the law could operate "in the absence of rules, reasons", he observes

> Initially Rome was a small and relatively homogeneous community, whose members shared a correspondingly homogeneous tradition of ideas about justice and fairness, property and propriety... the disputes that the pontiffs adjudicated were typically ones about which the traditional consensus was ambiguous; the balance of rights and obligations between the parties required the judgment call of a trusted and disinterested arbitrator. In these marginal cases all that the arbitrator

---

[385] Toulmin (1981), p. 31

[386] Toulmin (1981), p. 33

[387] Actually I don't think Toulmin would insist that his "history" is necessarily 100% accurate; at least, he doesn't need to. We can read his story as a kind of quasi-history which helps him characterise the social dynamics that shape policy-making systems (just as we can understand Locke's "state of nature" as a useful fiction which helps him explain the forces that drive us into political communities)

[388] Toulmin (1981), p. 33

> may be able to say is, "Having taken all the circumstances into account, I find that on this particular occasion it would, all in all, be more reasonable to tilt the scale to A rather than to B." This ruling will rest, not on the application of general legal rules, but rather on the exercise of judicial discrimination in assessing the balance of particulars.[389]

However, Toulmin notes:

> This state of affairs did not last. Long before the first Imperial codification, Roman law began to develop the full apparatus of "rules" with which we ourselves are familiar.[390]

Toulmin believe that a number of factors contributed to this development, including that:

> the case load increased beyond what the pontiffs themselves could manage. Junior judges, who did not possess the same implicit trust as the pontiffs, were brought in to resolve disputes; so the consistency of their rulings had to be "regularized."[391]

Also that

> Rome acquired an empire, and foreign peoples came under the city's authority. Their systems of customary law had to be put into harmony with the Roman system, and this could be done only by establishing a concordance between the "rules" of different systems.[392]

Toulmin then notes that:

> Life in late-twentieth-century industrial societies clearly has more in common with life in Imperial Rome than it has with the Rome of Horatius ... The way we live now, people have come to value uniformity above responsiveness, to focus on law at the expense of equity, and to confuse "the rule of law" with a law of rules.[393]

In other words, we have moved to an 'ethics of strangers' instead of an 'ethics of intimates':

> in the ethics of strangers, respect for rules is all, and the opportunities for discretion are few. In the ethics of intimacy, discretion is all, and the

---

389 Toulmin (1981), p. 33

390 Toulmin (1981), p. 33

391 Toulmin (1981), p. 33

392 Toulmin (1981), p. 33

393 Toulmin (1981), p. 34

> relevance of strict rules is minimal.[394]

Toulmin tells a couple of stories to illustrate the problems with this development. One is the story of the abortion debate in America, where

> much of the public rhetoric increasingly came to turn on "matters of principle." As a result, the abortion debate became less temperate, less discriminating, and above all less resolvable. Too often, in subsequent years, the issue has boiled down to pure head-butting: an embryo's unqualified "right to life" being pitted against a woman's equally unqualified "right to choose." Those who have insisted on dealing with the issue at the level of high theory thus guarantee that the only possible practical outcome is deadlock.[395]

Toulmin also tells the story of a handicapped young woman who had difficulties with the local Social Security office:

> Her Social Security payments were not sufficient to cover her rent and food, so she started an answering service, which she operated through the telephone at her bedside. The income from this service-though itself less than a living wage-made all the difference to her. When the local Social Security office heard about this extra income, however, they reduced her benefits accordingly; in addition, they ordered her to repay some of the money she had been receiving. (Apparently, they regarded her as a case of "welfare fraud.") The television reporter added two final statements. Since the report had been filmed, he told us, the young woman, in despair, had taken her own life. To this he added his personal comment that "there should be a rule to prevent this kind of thing from happening.[396]

But Toulmin says rules were the problem here, and argues:

> the delivery of social services has become ever more routinized, centralized, and subject to bureaucratic routine. It should not take horror stories, like that of the handicapped young woman's answering service, to make us think again about the whole project of delivering human services through a bureaucracy.[397]

---

[394] Toulmin (1981), p. 35

[395] Toulmin (1981), p. 32

[396] Toulmin (1981), p. 32

[397] Toulmin (1981), p. 36

One consequence of this development is an unreasonable demand for equal treatment on the part of members of the general public, in their role as service users:

> Since we are unwilling to grant discretion to civil servants for fear that it will be abused, we leave ourselves with no measure for judging administrators' performance other than equality. As Kaufman remarks, "If people in one region discover that they are treated differently from people in other regions under the same program, they are apt to be resentful and uncooperative."[398]

Hence there arises a "general concern for uniform application of policy," which can be guaranteed only by making the rulebook even more inflexible.

In search of a solution, Toulmin refers back to Aristotle, who insisted there could be no rigorous "theory" of ethics:

> Practical reasoning in ethics, as elsewhere, is a matter of judgment, of weighing different considerations against one another, never a matter of formal theoretical deduction from strict or self-evident axioms.[399]

Toulmin tells the story of a commission he worked on

> In almost every case they came close to agreement even about quite detailed recommendations - at least for so long as their discussions proceeded taxonomically, taking one difficult class of cases at a time and comparing it in detail with other clearer and easier classes of cases. Even when the Commission's recommendations were not unanimous, the discussions in no way resembled Babel: the commissioners were never in any doubt what it was that they were not quite unanimous about. Babel set in only afterwards. When the eleven individual commissioners asked themselves what "principles" underlay and supposedly justified their adhesion to the consensus, each of them answered in his or her own way: the Catholics appealed to Catholic principles, the humanists to humanist principles, and so on. They could agree; they could agree what they were agreeing about; but, apparently, they could not agree why they agreed about it. This experience prompted me to wonder what this final "appeal to principles" really achieved.[400]

---

[398] Toulmin (1981), p. 36

[399] Toulmin (1981), p. 37

[400] Toulmin (1981), p. 32

I take it that the moral Toulmin wants to draw from this tale is that in deliberations about what to do about a case, general principles are unhelpful. The commissioners had very constructive discussions without reference to principles, but as soon as they started discussing principles the quality of discussion deteriorated. Toulmin's previous stories about the abortion debate and disabled woman add weight to this picture, showing that general principles are positively unhelpful as guides to policy. On the basis of these two points, Toulmin concludes that we should do without general principles as far as possible, and deliberate about cases individually. He agues that in their proceedings, the commissioners

> were reviving the older, Aristotelian procedures of the casuists and rabbinical scholars, who understood all along that in ethics, as in law, the best we can achieve in practice is for good-hearted, clear-headed people to triangulate their way across the complex terrain of moral life and problems. So, starting from the paradigmatic cases that we do understand - what in the simplest situations harm is, and fairness, and cruelty, and generosity - we must simply work our way, one step at a time, to the more complex and perplexing cases in which extremely delicate balances may have to be struck. For example, we must decide on just what conditions, if any, it would be acceptable to inject a sample group of five-year-old children with an experimental vaccine from which countless other children should benefit even though the risks fall on those few individuals alone.[401]

The lesson he draws from his arguments and his experience on the commission is that:

> we need to be reminded that equity requires not the imposition of uniformity or equality on all relevant cases, but rather reasonableness or responsiveness (epieikeia) in applying general rules to individual cases. Equity means doing justice with discretion around, in the interstices of, and in areas of conflict between our laws, rules, principles, and other general formulas. It means being responsive to the limits of all such formulas, to the special circumstances in which one can properly make exceptions, and to the trade-offs required where different formulas conflict. The degree to which such marginal judgments can be regularized

---

[401] Toulmin (1981), p. 37

> or routinized remains limited today, just as it was in early Rome.[402]

He concludes:

> we do need to recognize that a morality based entirely on general rules and principles is tyrannical and disproportioned, and that only those who make equitable allowances for subtle individual differences have a proper feeling for the deeper demands of ethics.[403]

And he says of the casuists that they

> grasped the essential, Aristotelian point about applied ethics: it cannot get along on a diet of general principles alone. It requires a detailed taxonomy of particular, detailed types of cases and situations.[404]

9.3 Williams & Cookson: Defence of Algorithms

I now contrast Toulmin's position with that of Williams & Cookson. They observe a similar sequence of developments in policy-making systems, but they take the opposite position to Toulmin, seeing them as positive developments.[405] Somewhat like Toulmin, Williams & Cookson base their argument on a quasi-historical account of the development of legal and policy-making processes. They observe four stages.[406] Each stage suffers a defect which is repaired in the subsequent stage, so he can characterise these historical developments as a sequence of improvements. In the first stage, policy decisions are based on intuitions regarding individual cases. But they say:

> The problem with decision making based on intuition - that is, stage one - is that it may give rise to suspicions of special pleading and favoritism. Public sector decisions are supposed to be made in an impartial and fair-minded way. However, all such decisions take place against a backdrop of political lobbying by vested interests. The "losers" from any decision will naturally be suspicious that the "winners" have exerted undue influence - and will challenge the policy maker to justify the decision.

---

402 Toulmin (1981), p. 34

403 Toulmin (1981), p. 38

404 Toulmin (1981), p. 38

405 However, Williams is not responding to Toulmin directly.

406 Williams & Cookson (2006), p. 4

> One justification is simply that the correct process was followed. But that is unsatisfactory in modern democracies, where increasingly well-educated citizens expect to be told why decisions were made - and on the basis of what information and evidence - rather than simply being asked to trust the wisdom of the decision maker or the collective wisdom of the decision-making process.[407]

As an aside, alongside Williams & Cookson's worry about "favoritism", in the healthcare domain perhaps we should also note a related worry that policy-makers will be influenced by the eloquent representative of one or other stakeholder group, who swings a decision to their favoured patient group with some finely judged rhetoric, whilst the next patient group, with equal grounds for funding, gets denied as a result of lacking such an effective representative. We might not want to describe this as favouritism, but it is equally a bias towards a particular group driven by irrelevant factors which could not be justified in terms of policy-making criteria (such as how articulate a patient group representative is).

Now, as it turns out, for Williams & Cookson, each subsequent stage turns out to have the same problems as stage 1, such as scope for favouritism, albeit the scope is narrowed each time. For example, in stage 2:

> Under case law, decisions are made and justified on the basis of principles distilled from precedents set by previous decisions. The decision maker’s discretion is restricted: decisions that are clearly inconsistent with precedent are ruled out unless extremely good reasons can be deployed to justify an exception.[408]

How do precedents restrict favouritism? First, unless a given patient group or their representatives get in early, their room for manoeuvre will be restricted. If, in an earlier decision, an impartial discussion led to a decision in favour of patient group A over patient group B in circumstances C, then in subsequent decisions, a proponent of patient group B will have trouble pushing for patient group B in those same circumstances C, because of the requirement to respect the precedent. Second, even if the proponent of patient group B gets in on an early decision and swings it in B's direction, a demand that subsequent decisions be justified in terms of precedent is more likely to expose any questionable justifications that had an influence. Policy-makers

---

[407] Williams & Cookson (2006), p. 4

[408] Williams & Cookson (2006), p. 4

will ask how each new case is importantly like that first one. It will not be enough to respond merely that the previous decision went B's way. Proponents of B will be asked to identify a morally relevant feature that the current case shares with the last one. If the previous decision favouring B was based on bias or prejudice to B, that is then likely to be exposed in the form of a weak answer to this question.

Again, however, Williams & Cookson think that "suspicions of favoritism may arise in particular cases". This time, the cause of the problem is likely to be that there are:

> several apparently inconsistent precedents from which the decision maker is free to pick and choose.[409]

Thus policy-makers will be forced in the direction of stage three: codification. At this stage, principles are generalized to cover all foreseeable cases. The demand for explicit justifications makes it even more difficult to argue vaguely that "patient group B seem to me to merit priority", and if the explicit considerations are clearly one-sided, other stakeholders will see this and complain. Transparency enhances accountability.

But still

> trade-offs will arise between different principles. So the decision maker will have to use discretion in determining what weight to give to different principles in different cases.

And so again, suspicions of favoritism may arise. So, finally, Williams & Cookson have us moving to stage four, at which point we're using fully-fledged algorithms:

> a formula that specifies what weight to attach to different general principles, and to the different specific values invoked by those principles, in the case of trade-offs.[410]

Some might be worried that the move to stage 4 leaves our democratic representatives with less room to exercise their judgment, but Williams & Cookson argue that there need be no loss of democratic legitimacy in the move from stage 1 to stage 4. They say:

> The "formulaic approach" of stage four is sometimes contrasted with the

---

[409] Williams & Cookson (2006), p. 4

[410] Williams & Cookson (2006), p. 4

> "deliberative approach" of stages one to three. But this distinction can be misleading. In comparing the four stages, the question is not whether deliberation, nor yet how much deliberation, but rather at what stage and within what parameters deliberation takes place. ... As we progress through the four stages, however, more of the deliberation about social value judgments takes place outside the context of the particular decision in hand. The move toward a formula can be embedded within a democratic process.[411]

In support of Williams & Cookson's position, a range of other problems with uncodified trade-off principles can be noted. I will consider five. Briefly, the first three are: lack of transparency; lack of consistency; and interest group capture. As Marsh et al summarise it:

> many resource allocation decisions are still based on non-transparent choices that reflect competing interests of governments, donors and other stakeholders ... frequently, decision-making is inconsistent and unstructured. Important criteria such as budget impact, equity and disease severity have not always been taken into consideration.[412]

On transparency, the problem is that stakeholders cannot see why decisions were reached. This is a problem because policy-makers must explain their decisions to the people who lose out.[413] It is not enough for policy-makers to say that on balance, their intervention seemed less important. On consistency, the problem is that it is difficult to ensure consistency when the principles driving decisions are not stated. This is a problem because inconsistency is unjust. If patient A gets funding for her treatment, then patient B whose relevant characteristics are entirely similar should also get funding. On interest group capture, deliberative processes without clear decision-making principles are open to manipulation by influential groups with vested interests. The more open and flexible the decision-making principles are, the more of a danger this will represent. Again, this leads to inconsistency.

A fourth problem is that (at least on my pluralist account) ethics introduces multiple criteria, and the evidence clearly shows that we are bad at juggling

---

[411] Williams & Cookson (2006), p. 4

[412] Marsh et al (2017)

[413] I take it that a policy-making system is transparent if stakeholders know they can get explanations of the reasons for policies should they want them. If stakeholders nevertheless do not seek such explanations, this does not necessarily impugn the system's transparency.

multiple criteria. As Devlin et al put it (in discussing Multiple Criteria Decision Analysis, which I will defend in a moment):

> The preferred options identified by MCDA are likely to out-perform the use of intuitive judgement alone. The decision theory and psychology literature abounds with examples of the various biases and heuristics that are evident when individuals are confronted with complex decisions (Kahneman 2003; Gilovich et al 2002). This is because the consideration of multiple criteria is cognitively demanding – arguably especially so when decisions are made in a committee context.[414]

Our failings when handling multiple criteria are most visible in the case of empirical questions, since in that case we have a universally accepted "gold standard" against which the results of deliberation can be checked, this being the facts as subsequently confirmed by the accumulation of further evidence. Meehl is one of the most famous investigators of this problem. As Grove and Meehl summarise it, in the context of clinical psychology:

> Given a data set about an individual or a group (e.g., interviewer ratings, life history or demographic facts, test results, self-descriptions), there are two modes of data combination for a predictive or diagnostic purpose. The clinical method relies on human judgment that is based on informal contemplation and, sometimes, discussion with others (e.g., case conferences). The mechanical method involves a formal, algorithmic, objective procedure (e.g., equation) to reach the decision. Empirical comparisons of the accuracy of the two methods (136 studies over a wide range of predictands) show that the mechanical method is almost invariably equal to or superior to the clinical method.[415]

And to emphasise the relevance of this, Phillips points out that Meehl "identified integration of multiple pieces of data as the problem, not the judgments about the pieces."[416]

There is no reason to think that our deliberations on policy and ethical questions are immune from these problems. If anything, we can expect them to be worse: we get less chance to learn from our mistakes in ethics, because our mistakes are generally not subsequently made as obvious to us as our empirical mistakes. This gives us reason to look for alternative methods.

---

414 Devlin & Sussex (2011), p. 19

415 Grove & Meehl (1996)

416 Phillips (2017)

9.4 Reconciling Toulmin and Williams/Cookson

However, the proponents of algorithmic policy-making don't have all the arguments on their side. I have been defending a form of moral pluralism, according to which multiple independent values are in play. Given their independence of each other, it would be surprising if their conflicts could always be captured in simple formulae. Toulmin is not the only one to see this. Ross, perhaps the most famous proponent of moral pluralism, also emphasises the difficulty and uncertainty of cases where we must weigh conflicting prima facie duties. In reaching one's final conclusion, he thinks we must take "a moral risk",[417] since there are no principles governing such weighing:

> The only possible premises would be the general principles stating their prima facie rightness or wrongness qua having the different characteristics they do have ... there is no principle by which we can draw the conclusion that it is on the whole right or on the whole wrong. ... we have more or less probable opinions which are not logically justified conclusions from the general principles that are recognized as self-evident.[418]

Thus I would accept that some, perhaps many, trade-offs will evade capture by simple generalisations. For example, perhaps no general formula will capture how quality of life considerations trade off against benefit maximising considerations. In order to commensurate quality of life issues against, say, benefit maximising considerations, we need a single, all-purpose measure which can capture the disvalue of every possible quality of life impairment. But many diseases cause sui generis quality of life problems which are not seen in other diseases. This could mean there is no standard way of valuing quality of life impairments which tells us how much priority each impairment merits. For example, suppose we value each possible quality of life impairment against life years using the time trade-off method. There may be no standard formula which maps the resulting valuation figures to priority ratings (e.g. weightings applied to QALYs or their equivalent) in a way that reflects the priority that each impairment really merits. But this does not

---

[417] Ross (1930, 2002), p30

[418] Ross (1930, 2002), p31

preclude the possibility of true generalisations which characterise how other factors trade off against each other, such as how fair innings considerations trade off against benefit maximisation. Such a principle might say that patients who are expected to die young without treatment get a specified weighting on their incremental life years according to how far their life expectancy falls short of society's average.[419] To be able to state some such true generalisation would represent a philosophical insight. The question is whether there is some such generalisation which is true for all the circumstances where it implies a specific policy. If there is, we should seek such generalisations.

But how could we *know* when such a generalisation is true? One way would be to take a bunch of wise policy-makers (experienced, trusted by knowledgeable stakeholders, etc) and set them up in epistemically ideal conditions, with loads of information and time to reflect (in other words, a process that would be unaffordable in ordinary policy-making; that's our motivation for testing the algorithm). Get them to make a number of decisions. Compare the decisions with the algorithm. If the two methods tally for a significant number of cases, we can start to be confident in the algorithm. Now get ordinary policy-makers to make the same decisions using ordinary, time-constrained and cash-constrained methods. If the algorithm makes fewer serious mistakes than the ordinary policy-makers, we have a prima facie case for the algorithm.

So, one response to this question is to conclude that even if Toulmin and Ross are right that with respect to many policy issues there are no true generalisations regarding how different considerations should be traded off, nevertheless if some policy issues *can* be analysed on the basis of such true generalisations, then we should state those generalisations and use them in policy-making. If we can articulate true generalisations in these corners of the ethical world, and state trade-off rules for policy-makers, the above evidence suggests that we will be more likely to get policy decisions right, and more

---

[419] As an example, we might say that if you're only expected to live until the age of 40, the incremental life years delivered by a treatment get roughly doubled, and if you're only expected to live until 20, the incremental life years get roughly tripled. A formula that would deliver these results would say that incremental life years should be multiplied by the square of 2 minus life expectancy, with life expectancy measured on a scale from 0 to 1, and 1 equating to the population mean life expectancy. So if your expected length of life is 40 years (0.5), then any incremental life years you gain from an intervention are multiplied by 2.25, which is the square of (2 - 0.5). In contrast, if your expected length of life is 20 years (0.25) then the incremental life years are multiplied by 3.1.

able to explain them to stakeholders with a legitimate interest in understanding them.

How far could this go? For illumination, it is worth briefly considering the possibility of a "black box" policy-making system. My hypothesised system is based on artificial intelligence. We have fed it a large number of cases previously considered by policy-makers, together with the resulting policy, and it has inferred a set of policy-making rules.[420] Using these rules, given an input of a policy question and relevant facts, it can output a policy. Based on the above test (involving unaffordably elaborate policy-making processes), its output has been found to match the carefully considered decisions of our finest policy-making minds across many thousands of decisions, at least within certain narrowly specified domains of policy-making. To clarify, the black box I envisage is not merely one that answers empirical questions such as prognosis with and without treatment; it answers policy questions. Of course we are a long way off having such a reliable system (a system that accurately reproduces our policies). But supposing one became available in future, might we be justified in relying on its prescriptions in the domains where it had demonstrated its "competence"?

One response is that currently, we delegate responsibility for these important decisions to people. The public delegates responsibility to politicians and they delegate responsibility to various officials. It would be a derogation of responsibility for those officials to then pass the buck to a machine. But suppose a proper decision is to be taken about whether the machine should be used (for example, suppose a referendum is to be held). Would it then be wise to decide in favour?

One argument against machine-based policy-making would be that there is always the chance of unusual case where precedents are a poor guide, such as a case involving disease-specific impairments which the algorithm has not "seen" before, or an unusual set of circumstances where the right trade-offs do not follow any previously seen pattern. But against this possibility we must balance the likelihood that policy-makers will continue making their usual mistakes as a result of relying on heuristics and being subject to myriad cognitive biases. Suppose the evidence suggests that the number of people unjustly denied treatment by policy-makers deliberating in the ordinary way exceeds the number that would be unjustly deprived by algorithmic policy-

---

[420] Not necessarily rules we would recognise as such, a point I consider in a moment.

making. Not to mention that the money spent on bringing policy-makers together could be spent on saving lives. Then we might have reason to reconsider.

However, one problem with such a black box is that it would be difficult to check whether its prescriptions were right. Compare the ordinary case where a policy-making committee comes up with a policy but the rest of us find the policy counterintuitive. We can then ask what the committee's reasons were for the policy, and object if the justification is inadequate. But in the case of a black box which has developed its policy-making "model" based on artificial intelligence techniques, there might be no way for us to understand why it came up with its policies. We could not engage in anything like "debate" with the system, challenging its assumptions.

Furthermore, the problem might be that not only can we not understand the system's reasons for its recommendations, but that the system may not be driven by anything like ordinary-style reasons at all. Using artificial intelligence techniques, the system can work up whatever weird and wonderful principles fit the previous cases we fed it. So its policy-making model may not be driven by anything like reasons of the ordinary kind. This means that if we implement its recommendations, there is no guarantee that our policies could be justified in terms of reasons of the ordinary kind. So the policy-making process would lack transparency. This would also mean the policy-making process lacked accountability of the right kind. If my fellow citizens deny me life-saving treatment, surely I am entitled to know why, and to insist that the reasons be good ones. With a policy derived from a black box, policy-makers could not give me good reasons. It is not a good reason that officials programmed an artificial intelligence system according to best practice. Rather, good reasons are based on features of the case at hand. Such reasons may not be available, not just because the system was too complex to understand, but because the system was not driven by reasons of the ordinary kind (no doubt there are many different ways of characterising the situations in which a given policy is right, and we can suppose the computer has found a way that we don't understand; when experts analyse the computations that lead the machine to its conclusions, they find the computations involve no recognisable morally relevant features). Thus neither the machine nor its users can justify the machine's recommendations in the ordinary way. In view of the right of stakeholders to know why policies went against them, this seems like a fatal objection against "black box" policy-making.

Thus if a policy-making algorithm is to stand any chance of acceptability, its policy-making parameters must be based on situational features which are recognisably reason-giving. Similarly, we must also be able to understand why the model trades off conflicting considerations as it does. If a particular set of trade-offs is counterintuitive, we need to be able to see why the system went the way it did. For example, suppose that although the machine beats ordinary policy-makers in terms of matching the judgments of our finest policy-making minds, occasionally it diverges even from them. In such cases, I suggest we could not simply assume without further argument that the mistake was made by the machine rather than by our finest policy-making minds. How to resolve the matter? In ordinary cases of such divergence, we need to engage in debate to determine who is right. The same facility would be needed with a policy-making machine. This means we need its decisions to be based on recognisable reasons.

The consequence should be that policy-makers are able to defend the policies derived from the system in the ordinary way, and take responsibility for the policies. The machine contributes to policy-making, but does not relieve policy-makers of their ultimate responsibility for the policies.

Of course, all this assumes that such a machine is possible. My argument is only that *if* such a machine was possible, we might have grounds for accepting the output of its algorithm.

However, even if such a machine was possible, I suspect there would still be resistance to such a process from Toulmin and others of his persuasion. Some will worry that such a process will lull policy-makers into false sense of security, so they rely on the system without thinking about its prescriptions carefully enough. In addition if we ever need policy-makers to make decisions, they will more prone to mistakes, because their skills will have deteriorated through lack of use.[421] My response is that the hypothesised process beats policy-makers working alone, in terms of "hit rate" of good policies. So there is no more reason to worry about mistaken policies with the machine than without it. So although it is true that policy-makers should always be vigilant, it is not clear why their sense of security should be seen as particularly problematic in this case. Others might point out that intuitively, we must understand the reasons for our views, and not just take our moral views on other people's moral testimony. In the hypothesised situation we

---

[421] Thanks to Carl Fox for this point.

might appear to be taking moral testimony from a machine. But again, this should not be a concern in the case of a system where policy-makers are required to understand the system's reasons for its recommendations.

No doubt there will be other objections. At base, I believe the concern motivating some of this resistance is to do with autonomy and dignity. For similar reasons, hospital surgeons resisted surgical checklists long after they were shown to reduce errors.[422] I expect worries about driverless cars will persist long after we have incontrovertible evidence of their superior safety. Many psychologists still resist Meehl's conclusions favouring the "mechanical method" of diagnosis despite the large amount of evidence in support of them. I suggest what motivates the objectors in such cases is that, like the rest of us, they find self-respect in doing their own practical reasoning. People find it demeaning to let a machine or other automated processes lead them in such matters, especially on evaluative questions. In some cases there may also be a concern about the risk of down-skilling jobs which are currently interesting and well-paid.

In summary, there are two broad reasons to be worried about algorithms: 1. Pessimism about the prospect of ever having an AI system that delivers the right policy in every situation, and 2. Respect for the ordinary human desire for autonomy. I will not try and resolve the dispute between AI proponents and AI doubters. Instead, I will try and finesse it. I will consider policy-makers' practical experiences of a process known as Multi-Criteria Decision Analysis (MCDA). MCDA is actually a collection of related approaches, but they share in common an attempt to get policy-makers and stakeholders to deliberate over criteria and their relative weights with a view to producing an algorithmic policy-making model. In theory, the model could then be used to generate policies with no further input from policy-makers. However, the point of MCDA in practice is not generally to use the algorithm in such a "blind" way; rather the value of the process is seen in the way it shapes the deliberation, making it more precise and revealing points of divergence between stakeholders. This prompts stakeholders to understand why they diverge and thereby reach a precisely stated consensus about the relative importance of various criteria. As Garau & Devlin say:

> decision-makers may fear that MCDA locks them in to a mechanistic

---

[422] The World Health Organization's Surgical Safety Checklist is an operating theatre checklist, and although surgeons resisted it, evidence suggests that it has had a positive impact on safety awareness and operating theatres mistakes worldwide (Haugen, 2013).

> approach - or that their judgments may be made redundant. However, as emphasised by many researchers in the field, MCDA instead aims "to help the decision-makers by structuring the information to support their deliberative process". (Thokala et al 2016). Further, "MCDA is not simply a technical process. Its successful implementation depends crucially on effective design of social processes by which the analysis is structured and conducted". (Dodgson et al 2009).[423]

The impact of this discipline on the discussion can be imagined from the following snatch of conversation which has the flavour of real life:

> Expressing those judgments in the explicit, quantified form of value functions and criterion weights is at first difficult for many experts, but I have found that gentle assistance is usually welcomed and after a brief period becomes almost second nature, with those quick to adapt, challenging others to "stop waffling and give us a number".[424]

As a result, disagreement is revealed and hopefully resolved:

> It is common for experts to disagree, but assigning numbers often reveals differences in perspective that reflect each person's past experience. Sharing that experience in a workshop setting provides a degree of learning that enables a model to be constructed.[425]

So, the worry on the part of the proponents of deliberation is that MCDA replaces deliberation. But we see here the possibility that an MCDA process will actually enhance the quality of deliberation, by bringing out divergence regarding criteria weights into the open and triggering a discussion of why the divergence exists. Instead of supressing deliberation, MCDA encourages it, by providing new, stimulating material for it. Yet we still get the transparency and consistency that proponents of MCDA demand. Either decisions go by the eventual model, transparently; or divergence from the model has to be explained, transparently, rather than exceptions getting nodded through on a hunch.

Given the precision that is demanded of policy-makers involved in MCDA, one might expect MCDA to expose a divergence of views. But seemingly, MCDA can sometimes help produce convergence. This is perhaps indicative that where there is one best answer, MCDA is well-placed to identify it:

---

[423] Garau & Devlin (2017). p. 294

[424] Phillips (2017)

[425] Phillips (2017)

> Applying these principles for creating an MCDA model can result in a model that truly represents the collective expertise of a group of experts and that could be replicated with another group of competent experts. This was demonstrated by a replication in 2013 for the European union of the 2010 UK drug harms study (van Amsterdam et al, 2015). The correlation between the final weighted preference values in the two studies was 0.993, quite unexpected from two different groups of experts using the same criteria after changing about 10% of the scores and reassessing al the weights.

So it looks like MCDA could be a valuable process even if the output will not be used to generate policies algorithmically.

However there is a practical question about how MCDA would work in the case of health technology assessment, where (for example) policy-makers evaluate a new cancer drug to decide whether to fund it. The classic MCDA process requires considering several policy options at once and developing policy-making criteria in the process of considering the options. But in health technology assessment, the policy options (technologies) are not considered alongside each other but in separate discussions which take place over the years as new technologies are released. The criteria are set beforehand, to ensure that all technologies are evaluated against the same criteria. The problem is that there does not seem to be room in this process for the to-and-fro between specific policy options and general policy-making criteria that characterises MCDA, because of the way different decisions are made successively rather than simultaneously. Once a decision is made regarding a technology, it is politically difficult to reverse it; and consistency requires us to then apply the same criteria to subsequent technologies, since otherwise, patients and technology manufacturers could reasonably complain that they have not been treated fairly. Thus it seems difficult for consideration of each successive new technology to influence the criteria in the way required by MCDA.

I suggest there is nevertheless room to adopt an MCDA-style approach in the context of health technology assessment. The first requirement would be that consideration of a new technology had the potential to change the assessment criteria. Each new technology would be considered as a potential contributor to the debate over criteria and not merely evaluated against pre-set criteria. Admittedly this could create difficulties. If consideration of a new technology resulted in modifications to the criteria, and the modified criteria entailed that a previously funded technology should no longer be funded, that could lead to protests from patients, which could be disruptive and costly. One way of softening the blow would be to allow existing patients to continue getting the treatment but to stop any new patients from getting it. (On the other hand, If the modified criteria entailed that an older technology that had been refused funding should be funded after all, one would not expect protests from patients).

However, these issues are mitigated by the fact that in a long-established process, it would be rare for a new technology to trigger a revolution in policy-making criteria. Existing principles would have been designed around many other cases, so any change made in the policy-making principles to accommodate a new technology would generally have multiple counterintuitive implications for other technologies. To explain, the existing principles can be idealised as arising from a kind of pursuit of reflective equilibrium between the general policy-making principles and intuitively correct decisions regarding individual technologies. For example, suppose that 100 technologies have been considered; 70 of them were approved for funding, 30 of them were refused. In order to generate these decisions, policy-makers adopted and then repeatedly revised their policy-making principles. The current principles are now such that a maximal number of these previous technology decisions are intuitively correct. So any tweak to the existing principles will mean accepting an increased number of counterintuitive implications with regard to those previously assessed technologies, such as a greater number of approvals for treatments that don't intuitively merit it, or a greater number of refusals for treatments that intuitively do merit funding. So the more decisions policy-makers have made previously, the greater the cost of a subsequent tweak to the principles will be, in terms of intuitively incorrect implications regarding previously assessed technologies.[426] Nevertheless it is not impossible that an individual technology decision will result in the principles being tweaked.

In summary, the distinctive features of an MCDA-based health technology appraisal process would be as follows:

1. Multiple criteria are taken into account, as identified in an initial round of decision-making.

---

[426] And also regarding new technologies going forward, since presumably a large number of previous decisions will mean the principles are more able to handle whatever is thrown at them

> 2. The relative importance of different criteria is quantified precisely. For example, for patients who are expected to die young without treatment, a precise weighting is applied to life-extending treatments depending on their expected age at death.
>
> 3. Technologies are given an overall score based on the weighted scores for each criterion, and technologies that exceed the threshold get funded.
>
> 4. But consideration of a specific new technology has the potential to bring policy-makers to change the relative weights of criteria, or even introduce new criteria. In that case, the new criteria and weights replace the old ones, and are applicable across the board - in principle, even to older technologies. (However it should be noted that political realities will restrict the application of new criteria to old technologies; it is not easy to withdraw funding for a technology for which funding was granted previously and which is now used by patients).

I suggest this ought to satisfy Toulmin, when he says:

> When Pascal attacked the Jesuit casuists for being too ready to make allowances in favor of penitents who were rich or highborn, he no doubt had a point. But when he used this point as a reason for completely rejecting the case method in ethics, he set the bad example that is so often followed today: assuming that we must withdraw discretion entirely when it is abused and impose rigid rules in its place, instead of inquiring how we could adjust matters so that necessary discretion would continue to be exercised in an equitable and discriminating manner.[427]

With MCDA used as a tool to support higher quality deliberation, we could claim to have satisfied Toulmin's demand that "necessary discretion would continue to be exercised in an equitable and discriminating manner". As noted, in the proposed approach, policy-makers will have complete discretion in that they will not be bound by any algorithm they might have adopted. Nevertheless in considering modifications to it, they must be cognisant of the implications for all the other technologies which fall into their remit.

In contrast we seem to have moved away from Williams & Cookson's desideratum that "more of the deliberation about social value judgments takes place outside the context of the particular decision in hand", because in accordance with point 4 of the MCDA process description above, there might

---

427 Toulmin (1981), p. 39

be regular deliberation about social value judgments within discussions of specific technologies. Nevertheless, I suggest we still have a system which matches the account at his final stage 4, viz., a system with "a formula that specifies what weight to attach to different general principles, and to the different specific values invoked by those principles". So every policy-maker has to formulate their arguments in terms of precise criteria weights, justify those weights, and get them agreed by other policy-makers. Policy-makers no longer have the discretion they had at Williams & Cookson's stage 3, the discretion to determine what weight to give to different principles in different cases. If it is difficult for policy-makers to be biased in Williams & Cookson's system, where policy-making criteria were set in stone prior to discussions of specific policy options, then I suggest it will be nearly as difficult with my proposed MCDA process. What prevents bias in Williams & Cookson's stage 4 system is that criteria are set once and for all for all technologies, and must therefore be justified in terms of their implications for all possible technologies. But this constraint is equally applicable in my proposed system. The fact that the discussion of criteria may be triggered by discussion of a specific technology does not soften the constraint that the discussion must be in entirely general terms. It would be difficult for policy-makers to be biased to a particular patient group (say, by weighting a criterion which happens to favour that group) without that having lots of counterintuitive implications for other cases. This will generally cause the resulting proposal to be rejected. So at least in terms of the formal constraints, there is no more room for bias in my system than in Williams & Cookson's. Now, having said that, it must be admitted that in my system there is room for the eloquent patient representative to arouse policy-makers' emotions, such that a criterion that they would ordinarily give little weight to gets more weight than it deserves. This is a weakness of the system. But this danger can be mitigated if we ask policy-makers to look out for it. And given that we cannot entirely eliminate the danger without also reigniting the problems that Toulmin talks about, such as policy-makers prevented from recognising new considerations raised by new technologies, I suggest the costs and benefits of my system ought to be acceptable even for Williams & Cookson.

9.5 Conclusion

I conclude that both deliberation and an algorithmic approach to policy-making have their distinctive advantages. If a healthcare technology assessment department is well-funded and staffed with wise policy-makers who will be given plenty of information and time to make their decisions, then a deliberative process will no doubt be the best route.

But otherwise, algorithmic approaches have the potential to be better than they are given credit for in the philosophical literature. In particular, if a policy-making system is stretched, or it is vulnerable to pressure from stakeholders such as patient groups and pharma companies, then there will be a strong case for relying on algorithmic methods for at least some areas of policy. But this conclusion comes with the caveat that algorithms should not be black boxes. They should be based on ordinary, reason-giving features of the cases to be assessed. Policy-makers should always be able to interrogate an algorithmic computation to understand why the recommendation came out as it did. In fact, policy-makers should review every decision if possible, particularly decisions with serious consequences, such as when patient lives are at stake.

10. Conclusion

I argued that although benefit maximisation is a consideration in healthcare allocation, we also need to take account of other considerations in order to allocate healthcare fairly. We need to avoid unfair discrimination; we also need to give priority to the worst-off and those who would otherwise be excluded.

In addition, we need to be aware that benefit maximisation doesn't mean maximising just *any* kind of benefit; some kinds of benefit are irrelevant in healthcare allocation. Furthermore, although I have endorsed quality-adjusted life years as a measure of benefit, I have rejected NICE's method of calculating them based on the public's preferences. Instead I argued that health matters for a reason, viz., the chance it gives its possessor of living a good life, and health outcomes should be evaluated in light of that.

I will now summarise my conclusions from each chapter, and then conclude the thesis with an outline of policy proposals.

10.1 Summary of Conclusions from Each Chapter

In more detail, in Chapter 2 I asked whether discriminating[428] by prognosis is wrong in the same way as racist discrimination. I started by offering prima facie reasons for thinking that the way an act is motivated plays a critical role in determining its moral status. I endorsed an account from Larry Alexander according to which the reason wrongful discrimination is wrong is that it is based on false or unwarranted judgments regarding the group who lose out by the discrimination. I concluded that discriminating by life expectancy and quality of life is not wrong in the same way as racism/sexism, as long as it is not motivated by unwarranted judgments regarding the patients who lose out by the discrimination.

I also argued that not just individual acts but also policies can be racist/sexist such that they are intrinsically wrong, to the extent that they are dependent

---

[428] As usual, I used the term "discrimination" in a morally neutral sense according to which certain acts of discrimination may be permissible.

upon their supporters' biases and unjustifiable judgments of lesser moral worth. If a policy was an expression of policy-makers negative attitudes towards people with a given trait, and the policy or law wouldn't have said what it does without those supporters' biases, then the policy or law is racist.

Another consideration which could be grounds for questioning a policy is whether the traits it discriminates by are HSD traits (traits which have been discriminated against in the past). It can be argued that groups with those traits should be protected from further disadvantage. Policies should be assessed in terms of this pro tanto consideration. However, generally, healthcare allocation policies which discriminate by prognosis will not involve discrimination by HSD traits.

In Chapter 3 I asked what benefit healthcare allocators should aim to produce. I concluded that healthcare allocators produce the right kind of benefit when they help patients get into more valuable health states. I endorsed Hausman's view that health states are valuable or not partly depending on the extent to which they constrain the possibilities of living well and pursuing valuable objectives.

I argued that in assessing whether a health state constrains the possibilities of living well, we must understand what living well consists of. We must understand wellbeing in a pluralist way which involves considerations to do with preference satisfaction, subjective contentment, and goods which anyone rational can see to be good. The way these different goods trade off in different circumstances will vary in unpredictable ways. So to know how good or bad a health state is, we must apply our judgment to determine how the different types of wellbeing balance against each other.

I argued that usually, the most authoritative view on how good or bad a health state is will be the view of people in that state, viz., patients. This is because patients are always in a better position to know what a health state is like, and usually they are also in a better position to evaluate it appropriately. But sometimes, we should be guided by the public's informed evaluation of health states, or perhaps even the judgments of experts; in particular, this is so where patient preferences are influenced by evaluative adaptation or where patient preferences are irrational or have the wrong kind of provenance.

I argued that as far as practically possible, we must evaluate health states individually for the way they constrain the possibilities of living well, since health conditions have sui generis impairments and because the impact of a combination of impairments is not deducible from the their impact individually. Contrary to Hausman, the impact of a health state on someone's wellbeing cannot be determined algorithmically from principles derived from a few cases. However, in practical terms, it is unlikely that it will be possible for healthcare policy-makers to evaluate every health state. Thus there is a balance to be struck between getting preferences regarding every single health state of interest and alternatively being guided by general principles.

In many systems, given constrained resources for doing health state valuations, it will be defensible to rely on an algorithmic valuation method of assessing health outcomes, such as QALYs. However, there should be two key differences in how they are calculated. First, QALYs are currently based on public preferences, whereas I argued that often, patient preferences would be better. Second, QALYs are based on health state preferences rather than a deliberative assessment of whether health states constrain the possibilities of living well.

In Chapter 4 I offered a new argument to show that policy-makers have reason to institute orphan drugs policies by relaxing their normal criteria in the case of rare diseases, when deciding whether to licence and fund treatments for them.

One argument that could help rare disease patients in some circumstances was a quite general argument, derived from Taurek and Lawlor. This argument gives us a pro tanto reason to show "equal concern and respect for each person"[429] by giving rare disease patients *some* share of healthcare resources, even if their treatments are expensive so that this results in fewer people benefitting from those resources. However, this is quite a general principle which does not give us reason to favour rare disease patients over common disease patients whose cost-effectiveness is the same.

The second argument was a new, more targeted defence of orphan drugs policies. I argued that if healthcare allocation policy is driven purely by benefit-maximising principles, unmitigated by orphan drugs policies, rare disease patients will stand very little chance of ever benefiting from future treatments, because of a problem to do with economies of scale. In effect, they will be permanently excluded from access to new treatments. I argued that orphan drugs policies can therefore be justified by appeal to a pro tanto reason to mitigate such a permanent exclusion of a group from these important benefits.

---

429 Taurek (1977), p. 303

However, on the pluralist view I defended, this does not answer the question of whether to institute orphan drugs policies in any given case, or which policies to institute. The considerations favouring orphan drug policies are only pro tanto considerations. Other, strong considerations can come into conflict with them: in particular, the "efficiency" consideration that we ought to maximise either the benefits we produce or the number of people we benefit. Nevertheless I conclude that orphan drugs policies will not always be arbitrary, as some of their opponents argue.

In Chapter 5 I investigated whether it was ever permissible to discriminate against disabled patients in healthcare resource allocation. I concluded that it is impermissible to discriminate against patients who have a disability caused by something other than the disease we are treating or its treatment, because, when we assess the benefits of a treatment, only the effects of the disease we are treating or its treatment are relevant. We must ignore symptoms of conditions unrelated to the condition we are treating, *unless* those causally unrelated costs and benefits involve dimensions of health impaired by the disease/treatment we are treating.

My case for this account depended on something like Kamm's "separate spheres" principle. I offered a new argument for such a principle. I based my argument on the mutual insurance rationale for publicly funded healthcare. Citizens can only be expected to support a mutual insurance fund to the extent that it delivers the kind of benefits that they might need, regardless of what they are doing with their life. So governments should prioritise services that deliver universally beneficial resources or capabilities.

However, I acknowledged that healthcare *also* fulfils a general obligation to help people in need; healthcare is over-determined in terms of the reasons we have for instituting and maintaining it. Thus when there are non-universal benefits in play, this rationale may pull against the mutual insurance rationale as conflicting pro tanto considerations.

I also noted two other caveats relating to disabled patients. First, we should consider the historical disadvantages suffered by disabled people and avoid reinforcing any sense that they are excluded. Second, in the case of trivial disabilities, egalitarian considerations of the kind characterised by Taurek may outweigh considerations to do with benefit maximisation. However, to summarise, if a serious disability is caused by the disease we are treating or its treatment, and such caveats do not outweigh considerations of benefit maximisation, we may take account of the disability in assessing the benefits of the treatment.

In Chapter 6 I asked whether it could be permissible to discriminate by age in healthcare allocation. I considered Daniels' argument that a system that treats different age groups differently may nevertheless be just as long as it treats people of the same age the same, and it is in people's lifelong interests. I endorsed the argument, but also argued that Daniels does not give a complete account of our obligations in healthcare. Daniels argues from our prudential reasons for preferring a system that favours the young, whereas we also have non-self-interested reasons for supporting publicly funded healthcare.

In line with this, in Chapter 7 I defended a fair innings principle, giving us reason to favour the young even if it is not in our interests; specifically, a reason based on prioritarian considerations. I clarified that this did not commit me to discriminating directly by the patient's current age, as has traditionally been assumed. I showed that, given the justifications that have been offered for age-based discrimination, the morally relevant consideration is not the patient's current age, but the age at which they are expected to die.

I noted an objection from John Harris against some ageist theories, to the effect that it is counterintuitive to discriminate between youngsters who are close in age. In response I distinguished two types of principle: "straight line" principles according to which the strength of your claim to life-extending treatments declines at a steady rate according to your expected age at death, and threshold or "stepped" principles according to which your entitlement declines faster as your expected age at death rises past a certain age. I adopted a threshold principle (a "fair innings" principle) according to which the fair innings in a given society is equal to that society's average life expectancy. This largely avoids the discrimination Harris complains about while permitting discriminating between patients above the threshold and patients below it. I defended this principle with an innovative set of cases in which society's average life expectancy is varied to show its importance for intuitions. I also offered a new set of egalitarian arguments to defend the same principle.

Finally, I endorsed a version of Alan Williams approach, which doesn’t only apply the fair innings principle to life years, but to QALYs. I noticed that Williams' proposal has some counterintuitive consequences, but I found that if Williams accepts some relatively painless constraints on his weighting factors, he can dodge the counterintuitive consequences. I concluded that if a patient's quality-adjusted life expectancy is below society's average, the benefits they can expect from treatment should be up-weighted by a small factor.

In Chapter 8 I defended a rule of rescue obligation in healthcare allocation, this being an obligation to fund life-extending treatment for someone who faces imminent death, without applying our normal criteria in terms of cost-effectiveness or benefit maximisation. I rejected an account of our rule of rescue obligations in terms of identifiability. However, I endorsed an account in terms of the distribution of risk: an egalitarian principle that says that we should prioritise help for people who face a high risk of an early death. For example, this implies a preference for an intervention where a lot of people have a small risk over an intervention where a few people have a lot of risk, all else being equal. I found that this principle can be derived from another principle I defended in Chapter 7, a principle saying that we should prioritise help for people whose quality-adjusted life expectancy is below society's average.

However, I also found that this principle does not explain all of our rule of rescue intuitions. In particular, it does not explain why we prioritise life-extending treatment for those who are given very little notice of their death rather than those who are given more notice. We are very averse to letting people die suddenly. I argued that this is because we recognise the value of people having time to reconcile themselves to their death, sort out their affairs, say their goodbyes to friends and family. So the value of an extra month or year depends on the context: an extra month is worth a lot for someone who has just been diagnosed and who is only expected to live a month to live without treatment. But it is worth less for someone is expected to live two years without treatment.

In Chapter 9 I considered how we should we trade off all the conflicting considerations I'd canvassed up to that point. I considered whether we should rely on general trade-off algorithms, or instead deliberate over the best policy case by case. I considered arguments on each side from Williams and Toulmin and concluded that we should rely on algorithms where we have good evidence that they are reliable indicators of the best policy. I also argued that the process of agreeing an algorithm often produces better quality deliberation, so even the opponents of algorithms have some reason to support them.

10.2 Policy implications

What are the policy implications of all this? To an extent it would be premature to try and determine precise policies in advance of a process that has accountability for reasonableness,[430] such as a process in which stakeholders such as patients, public, clinicians and technology manufacturers have a chance to have an influence. However, some general policy implications can be outlined.

I start with the most general constraint that I discussed, namely the constraint debarring unfair discrimination. I argued that we may not to allocate healthcare based on false or unwarranted judgments regarding the group who lose out by the discrimination.

Relatedly, I argued that we must also consider any historical disadvantages suffered by the patient groups affected by our decisions and avoid reinforcing any sense that they are excluded. However, this does not preclude all discrimination by HSD traits, especially when such discrimination favours the historically disadvantaged groups. As an example, intuitively it seems morally permissible to bias diabetes screening to south Asians rather than whites, given that diabetes is 3.6 times more prevalent in that population, so such a screening policy might be more cost-effective.[431] As another example, it would seem permissible to take account of differences between the genders in their responsiveness to treatments for chronic depression.[432] Similarly it may be permissible to segment the population by gender, race or age when assessing treatments on the basis of cost-effectiveness differences.

However, policy-makers must also assess the combined effects of their policies. Suppose, for example, that south Asians consistently lost out in healthcare allocation decisions across a wide range of expensive cancer treatments, and this was because their treatments were just marginally less effective than treatments for whites and therefore also less cost-effective. This would give policy-makers reason to relax the cost-effectiveness threshold for south Asians as applied to cancer treatments in such circumstances. This is specifically because the case involves a historically disadvantaged group. Any

---

[430] For example, "the decision must rest on reasons that stakeholders can agree are relevant" (Daniels & Sabin, 2008)

431 Diabetes UK (2014)

[432] "Women were significantly more likely to show a favorable response to sertraline than to imipramine, and men were significantly more likely to show a favorable response to imipramine than to sertraline (Kornstein et al, 2000, p. 1445)

policy which consistently leads to that group losing out will have wider negative consequences than in the case of other patient groups (for example, in terms of the group feeling excluded). The same relaxation would not be justified if the patient group in question consisted of people with a certain gene, who didn't even know they had that gene.

With regard to rare diseases, there is a range of permissible ways of mitigating the exclusion of rare disease patients. One way is to relax the cost-effectiveness criteria. For example, a policy could say that if one disease is a hundred times rarer than another then, all else being equal, the cost-effectiveness threshold for the rarer disease is doubled. Such a policy would result in a smooth gradient of increasing weights, avoiding sharp cut-offs and thereby unacceptably arbitrary discrimination.

However, the most popular policies have been those that allow pharmaceutical companies a longer period of exclusivity before competitors are allowed to produce the same molecules.[433] Assuming this kind of policy is effective, as it has been in Europe and the US, then perhaps there is less reason for those who implement it to implement special provisions in their funding policies.

How should we take account of the quite general egalitarian considerations derived from Taurek and Lawlor? One way of balancing the need to maximise benefit with the need for equality is to run a lottery, but to confine it to treatments which are borderline in terms of cost-effectiveness. Treatments which are clearly cost-effective (by a wide margin) continue to be funded without further ado, and treatments which are clearly not cost-effective continue to be denied funding. But treatments whose cost-effectiveness is marginal are entered into a lottery (for example, this could apply to treatments costing between, say, £20k and £40k per QALY[434]). These lotteries are held at irregular intervals (say, each time a backlog of five borderline cost-effective treatments has accumulated).[435]

With regard to disability, I have already mentioned one important principle, which confines our attention to the effects of the disease we are treating or its treatment. In addition, we should consider the historical disadvantages

---

[433] For example, see FDA (2015) and Kanavos et al (2011)

[434] Officially the cost-effectiveness threshold is currently £20k - £30k per QALY.

[435] This avoids a difficulty with a system of lotteries at regular intervals, such as yearly. In such a system, patients would get an unfair advantage if their treatment is launched in a year when few other treatments get launched.

suffered by disabled people and avoid reinforcing any sense that they are excluded.

With regard to age, I argued that we should give extra weight to QALYs for patients whose quality-adjusted life expectancy, without treatment, is less than society's average.

Finally on rule of rescue considerations, in Chapter 8 on rule of rescue, I argued that the above principle regarding patients with a low quality-adjusted life expectancy would favour policies which result in risk being distributed widely rather than concentrated in a few people. In addition, I argued that patients who have had very little notice of their expected death should be prioritised for life-extending treatment.

I considered the detailed policy implications of this by considering how NICE's end of life premium would need to be amended to capture rule of rescue considerations. Patients who should qualify would be those who, without treatment, are expected to live less than a set period from diagnosis. The period should be decided on the basis of how long it take most people to sort out their affairs and come to terms with their death (this will not be a precise science, but perhaps three months or a year).

The same period should also constitute the maximum life expectancy with treatment to qualify for the premium. There should be no weighting for life expectancy improvements beyond the stipulated period. So if a treatment extends life beyond the stipulated period (say, beyond three months after diagnosis), the extra QALYs after the stipulated period should be counted as a benefit in the usual way, but the QALYs should not be up-weighted.

If a patient is expected to live less than, say, three more months, the weightings for additional life expectancy within that period should be high. Thereafter, the weightings should get lower.

The principle should also be understood as implying priority, not only for those clearly at risk, but also for those who are non-obviously at risk, where such individuals can be picked up with screening. In terms of preventive interventions for healthy patients, the priority should go to heart attack prevention rather than preventing slow growth cancers which get detected early.

Finally, how should we assess the benefit that is delivered by alternative policies? To the extent that we rely on generalisations about how the different features of the outcomes determine their value, I argued that we should rely on something like QALYs. On the view I defend, QALYs should be based on deliberative assessments of how much a health outcome constrains the possibilities of living well, usually done by patients in the state unless there are reasons to doubt the reliability of their preferences. However, we must confine our attention to the effects of the disease we are treating or its treatment and ignore symptoms of conditions unrelated to the condition we are treating.

Bibliography

Abelson, J., Bombard, Y., Gauvin, F.P., Simeonov, D. and Boesveld, S. 2013. Assessing the Impacts of Citizen Deliberations on the Health Technology Process. International Journal of Technology Assessment in Health Care, 29(03), Pp.282-289.

Adams, B. 2011. Cancer Drugs Should Not Receive Special Funding - MS Society. Pharmafile. (London, UK: Samedan Ltd). http://www.pharmafile.com/news/169703/cancer-drugs-should-not-receive-special-funding-ms-society. Accessed 3 November 2015.

Adams, B. 2015. Gilead's Sovaldi Wins NICE Approval but Faces Access Delay. Pmlive. http://www.pmlive.com/pharma_news/gileads_sovaldi_wins_nice_approval_but_faces_controversial_access_delay_662269. Accessed 2 June 2017

Adler, M. 2012. Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis. Oxford: Oxford University Press.

Adler, M.D. 2007. Well-Being, Inequality and Time: the Time-Slice Problem and Its Policy Implications.

Alcimed. 2005. Study on Orphan Drugs. Paris: Alcimed; 2005.

Alexander, L. 1992. What Makes Wrongful Discrimination Wrong? Biases, Preferences, Stereotypes and Proxies. University of Pennsylvania Law Review 141: 149-219.

American College of Emergency Physicians. 2017. EMTALA. https://www.acep.org/news-media-top-banner/emtala/. Accessed 2 June 2017.

Anand P. 2005. Capabilities and Health. J Med Ethics 2005:31:299-303.

Anderson, E. 1995. Value in Ethics and Economics. Harvard University Press.

Anscombe, G.E.M. 1957. Intention. Harvard University Press.

Appiah, A. 2008. Experiments in Ethics. Harvard University Press.

Arneson, R. 2013. Egalitarianism. The Stanford Encyclopedia of Philosophy (Summer 2013 Edition), Edward N. Zalta (Ed.), URL = http://plato.stanford.edu/archives/sum2013/entries/egalitarianism/. Accessed 1 August 2016.

Arneson, R. J. 2006. What Is Wrongful Discrimination? San Diego Law Review 43: 775-808

Arneson, R.J. 1989. Equality and Equal Opportunity for Welfare. Philosophical Studies, 56(1), Pp.77-93.

Arras, J. 2016. Theory and Bioethics. The Stanford Encyclopedia of Philosophy (Winter 2016 Edition), Edward N. Zalta (Ed.), URL = <https://plato.stanford.edu/archives/win2016/entries/theory-bioethics/>. Accessed 1 March 2018

Attema, A.E., Edelaar-Peeters, Y., Versteegh, M.M. and Stolk, E.A. 2013. Time Trade-Off: One Methodology, Different Methods. The European Journal of Health Economics, 14(1), Pp.53-64.

Austin, D and Russell E. M. 2003. Is There Ageism in Oncology? Scottish Medical Journal Vol 48, Issue 1, Pp. 17 - 20

Baltussen, R. and Niessen, L. 2006. Priority Setting of Health Interventions: The Need for Multi-Criteria Decision Analysis. Cost Effectiveness and Resource Allocation, 4(1), P.14.

Barham, L. 2016. End-of-Life Treatments: What Will NICE Accept? https://pharmaphorum.com/articles/end-of-life-treatments-what-will-nice-accept/. Accessed 15 December 2017

BBC1. 1986. The Heart of The Matter. 1986, Oct.

Beauchamp T. L., Childress J. F. 2009. Principles of Biomedical Ethics. 6th Ed. New York: Oxford University Press.

Bell, J. M., & Mendus, S. 1988. Philosophy and Medical Welfare. Royal Institute of Philosophy. Lecture Series, (23) Bennett, R. 2017. Are Universities Bastions of Anti-Brexit Brainwashing? https://www.thetimes.co.uk/article/are-universities-bastions-of-anti-brexit-brainwashing-kn2f3ffl8. Accessed 8 January 2018

Bentham, J. 1789. An Introduction to The Principles of Morals. London: Athlone.

Bognar, G. 2008. Age-Weighting. Economics and Philosophy, 24(02), 167-189.

Bognar, G. 2010. Does Cost Effectiveness Analysis Unfairly Discriminate Against People with Disabilities?. Journal of Applied Philosophy, 27(4), Pp.394-408.

Bognar, G. 2015. Fair Innings. Bioethics, 29(4), Pp.251-261.

Bozeman, B., & Sarewitz, D. 2011. Public Value Mapping and Science Policy Evaluation. Minerva, 49(1), 1-23.

Brauer, S. 2009. Age Rationing and Prudential Lifespan Account in Norman Daniels' Just Health. Journal of Medical Ethics, 35(1), 27-31.

Brazier, J., Akehurst, R., Brennan, A., Dolan, P., Claxton, K., Mccabe, C., Sculpher, M. and Tsuchyia, A. 2005. Should Patients Have A Greater Role in Valuing Health States?. Applied Health Economics and Health Policy, 4(4), Pp.201-208.

Brock, D. 2004. Ethical Issues in the Use of Cost Effectiveness Analysis for the Prioritization of Health Resources. In Handbook of Bioethics (Pp. 353-380). Springer Netherlands.

Brock, D.W. 2002. Priority to the Worse Off in Health-Care Resource Prioritization. Medicine and Social Justice, Pp.362-372.

Brockmann, H. 2002. Why Is Less Money Spent on Health Care for the Elderly Than for the Rest of the Population? Health Care Rationing in German Hospitals. Social Science & Medicine, 55(4), 593-608.

Broome, J. 1988. Good, Fairness and QALYs. In Bell & Mendus (1988), 57-73.

Broome, J. 1992. Deontology and Economics. Economics and Philosophy, 8(02), 269-282.

Broome, J. 1997. Is Incommensurability Vagueness?. In Chang (1997)

Buchanan, A. E. 1984. The Right to A Decent Minimum of Health Care. Philosophy & Public Affairs, 55-78.

Buckwalter, W, and Stich, S. 'Gender and Philosophical Intuition.' Experimental Philosophy, Vol. 2. Eds. Joshua Knobe, and Shaun Nichols. Oxford: Oxford University Press, 2013. 307–346.

Burns, P.B., Rohrich, R.J. and Chung, K.C. 2011. The Levels of Evidence and Their Role In Evidence-Based Medicine. Plastic and Reconstructive Surgery, 128(1), P.305.

Business Insider. 2014. The Most (and Least) Athletic Countries. http://www.businessinsider.com/the-most-and-least-athletic-countries-2014-2?ir=t. Accessed 5 March 2018

Busschbach, J.J.V., Hessing, D.J., De Charro, F.T. 1993. The Utility of Health at Different Stages In Life: A Quantitative Approach. Social Science and Medicine 37 (2), 153±158

Callahan, D. 1995. Setting Limits: Medical Goals in an Aging Society with a Response to My Critics. Georgetown University Press.

Callahan, D. 2012. Must We Ration Health Care for the Elderly? The Journal of Law, Medicine & Ethics, 40(1), 10-16.

Carl, N. 2017. Why Do Academics Lean Left? London: Adam Smith Institute.

Carlsson, P., Hoffman, M., Levin, L.-A., Sandman, L., & Wiss, J. 2012. Prioritization and Financing of Drugs for Treatment of Patients with Rare Diseases. (In Swedish: Prioritering Och Finansiering Av Lakemedel for Behandling Av Patienter Med Sallsynta Sjukdomar.) Appendix 4 In the Official Governmental Inquiry (SOU) 2012:75. Price, Access, and Service-Continued Developments In the Drug and Pharmaceutical Markets. (In Swedish: Pris, Tillga˚ng Och Service-Fortsatt Utveckling Av La¨Kemedels- Och Apoteksmarknaden). http://www.regeringen.se/content/1/c6/20/25/65/36dbeb36.pdf. Accessed 1 September 2016.

Carter, D., Gordon, J. and Watt, A.M. 2016. Competing Principles for Allocating Health Care Resources. Journal of Medicine and Philosophy, P.Jhw017.

Chang, R. (Ed.). 1997. Incommensurability, Incomparability, and Practical Reasoning. Harvard University Press.

Chang, R. 1997. Incommensurability, Incomparability, and Practical Reasoning: An Introduction. In Chang (1997)

Christiano, T. 2015. Democracy. The Stanford Encyclopedia of Philosophy (Spring 2015 Edition), Edward N. Zalta (Ed.), URL = <https://plato.stanford.edu/archives/spr2015/entries/democracy/>. Accessed 30 June 2017

Churchill, L. R. 2005. Age-Rationing In Health Care: Flawed Policy, Personal Virtue. Health Care Analysis, 13(2), 137-146.

Citizens Advice. 2017. Indirect Discrimination. https://www.citizensadvice.org.uk/law-and-courts/discrimination/what-are-the-different-types-of-discrimination/indirect-discrimination. Accessed 26 February 2017.

Citizens Advice. 2018. Equality Act 2010 - Discrimination and Your Rights. https://www.citizensadvice.org.uk/law-and-courts/discrimination/about-discrimination/equality-act-2010-discrimination-and-your-rights/. Accessed 4 March 2018.

Claxton, K. , & Culyer, A. J. 2007. Rights, Responsibilities and NICE: A Rejoinder to Harris. Journal of Medical Ethics, 33(8), 462-464.

Claxton, K. , & Culyer, A. J. 2008. Not A NICE Fallacy: A Reply to Dr Quigley. Journal of Medical Ethics, 34(8), 598-601.

Clayton, M. , & Williams, A. (Eds.). 2000. The Ideal of Equality. Basingstoke, Palgrave Macmillan: 81-125.

Cleemput, I. and Neyt, M. 2015. Which Quality of Life Measures Fit Your Relative Effectiveness Assessment?. International Journal of Technology Assessment In Health Care, 31(3), Pp.147-153.

Cohen, G. A. 1989, On the Currency of Egalitarian Justice, Ethics 99, Pp. 906–944.

Collins, P.Y., Patel, V., Joestl, S.S., March, D., Insel, T.R., Daar, A.S., Bordin, I.A., Costello, E.J., Durkin, M., Fairburn, C. and Glass, R.I. 2011. Grand Challenges In Global Mental Health. Nature, 475(7354), P.27.

Conaghan J. 1999. Feminism and Labour Law: Contesting the Terrain. In: Morris & O'Donnell T (1999)

Cookson, R. 2013. Can the NICE "End-of-Life Premium" Be Given A Coherent Ethical Justification?. Journal of Health Politics, Policy and Law, 38(6), Pp.1129-1148.

Cookson, R., Drummond, M. and Weatherly, H. 2009. Explicit Incorporation of Equity Considerations Into Economic Evaluation of Public Health Interventions. Health Economics, Policy and Law, 4(02), Pp.231-245.

Cookson, R., Mccabe, C., & Tsuchiya, A. 2008. Public Healthcare Resource Allocation and the Rule of Rescue. Journal of Medical Ethics, 34(7), 540-544.

Cropper, M.L., Aydede, S.K., Portney, P.R. 1994. Preferences for Life Saving Programs: How the Public Discounts Time and Age. Journal of Risk and Uncertainty 8, 243±265.

Dakin, H., Devlin, N., Feng, Y., Rice, N., O'Neill, P. and Parkin, D. 2015. The Influence of Cost-Effectiveness and Other Factors on Nice Decisions. Health Economics, 24(10), Pp.1256-1271.

Dancy, J. 1993. Moral Reasons. Oxford: Blackwell.

Daniels N, Sabin JE. 2008. Setting Limits Fairly: Learning to Share Resources for Health. 2nd Ed. New York: Oxford University Press

Daniels N, Sabin JE. Setting Limits Fairly: Learning to Share Resources for Health. 2nd Ed. New York: Oxford University Press, 2008.

Daniels, N. 1979. 'Wide Reflective Equilibrium and Theory Acceptance In Ethics', Journal of Philosophy, 76(5): 256–82; Reprinted In Daniels, N. 1996, Justice and Justification: Reflective Equilibrium In Theory and Practice, Cambridge: Cambridge University Press, Pp. 21–46.

Daniels, N. 1985. Just Health Care. Cambridge University Press.

Daniels, N. 1988. Am I My Parents' Keeper?: An Essay on Justice Between the Young and the Old (Pp. 40-65). New York: Oxford University Press.

Daniels, N. 1989. Justice and Transfers Between Generations. Workers Versus Pensioners: Intergenerational Justice In An Ageing World, 57-79.

Daniels, N. 1996, Justice and Justification: Reflective Equilibrium In Theory and Practice, New York: Cambridge University Press.

Daniels, N. 2007. Just Health: Meeting Health Needs Fairly, New York: Cambridge University Press.

Daniels, N. 2008. Justice Between Adjacent Generations: Further Thoughts.Journal of Political Philosophy, 16(4), 475-494.

Daniels, N. 2013. Global Aging and the Allocation of Health Care Across the Life Span. The American Journal of Bioethics, 13(8), 1-2.

Daniels, N. and J. Sabin. 1997. Limits to Health Care: Fair Procedures, Democratic Deliberation, and the Legitimacy Problem for Insurers. Philosophy and Public Affairs, 26, 303–50.

Daniels, N. and J. Sabin. 2002. Setting Limits Fairly: Can We Learn to Share Medical Resources? New York: Oxford University Press.

Daniels, N. and J. Sabin. 2008. Setting Limits Fairly: Learning to Share Resources for Health. 2nd Edition. New York: Oxford University Press.

Daniels, N. and Sabin, J.E. 2008. Accountability for Reasonableness: An Update. BMJ: British Medical Journal (Online), 337.

Daniels, N. 1979. Wide Reflective Equilibrium and Theory Acceptance In Ethics. The Journal of Philosophy, 76(5), Pp.256-282.

Daniels, N. 2008. Justice Between Adjacent Generations: Further Thoughts. Journal of Political Philosophy, 16(4), Pp.475-494.

Darwall, S.L. 2002. Welfare and Rational Care. Princeton: Princeton University Press.

Darwall, S.L. 1977. Two Kinds of Respect. Ethics, 88(1), Pp.36-49.

De Wit, G.A., Busschbach, J.J. and De Charro, F.T. 2000. Sensitivity and Perspective In the Valuation of Health Status: Whose Values Count?. Health Economics, 9(2), Pp.109-126.

Department of Communities, Queensland Government. 2007. See Beyond the Disability: Queenslanders with A Disability Share Their Ideas for A More Inclusive Community. Brisbane: Queensland Government. https://www.communities.qld.gov.au/resources/disability/community-involvement/documents/see-beyond-disability.pdf. Accessed 31 May 2016.

Devlin, N. and Sussex, J. 2011. Incorporating Multiple Criteria In HTA. Methods and Processes. London: Office of Health Economics.

Diabetes UK. 2014. Diabetes: Facts and Stats. https://www.diabetes.org.uk/documents/about%20us/statistics/diabetes-key-stats-guidelines-april2014.pdf. Accessed 18 August 2016

Diener E, Lucas RE, Scollon CN. 2006. Beyond the Hedonic Treadmill: Revising the Adaptation Theory of Well-Being. Am Psycho/ 2006;61:305-14.

Dignity In Dying. 2008. A Fear of Dying Alone. London: Dignity In Dying

Dixon, J., & Welch, H. G. 1991. Priority Setting: Lessons From Oregon. The Lancet, 337(8746), 891-894.

Dolan P. 1997. Modelling Valuations for Euroqol Health States. Med Care 1997:35:1

Dolan, P. and Kahneman, D. 2008. Interpretations of Utility and Their Implications for the Valuation of Health. The Economic Journal, 118(525), Pp.215-234.

Dolan, P. and Metcalfe, R. 2012. Valuing Health A Brief Report on Subjective Well-Being Versus Preferences. Medical Decision Making, 32(4), Pp.578-582.

Dolan, P., Shaw, R., Tsuchiya, A., & Williams, A. 2005. QALY Maximisation and People's Preferences: A Methodological Review of the Literature. Health Economics, 14(2), 197-208.

Donagan, A. 1977, Informed Consent In Therapy and Experimentation, Journal of Medicine and Philosophy 2: 307–329.

Drummond, M. F., Wilson, D. A., Kanavos, P., Ubel, P., & Rovira, J. 2007. Assessing the Economic Challenges Posed By Orphan Drugs. International Journal of Technology Assessment In Health Care, 23(01), 36-42.

Dworkin, R, 2000, Sovereign Virtue: Equality In Theory and Practice, Cambridge: Harvard University Press.

Edmonds, D. 2006. Caste Wars: A Philosophy of Discrimination. London: Routledge.

Ellison, N., Loprinzi, C.L., Kugler, J., Hatfield, A.K., Miser, A., Sloan, J.A., Wender, D.B., Rowland, K.M., Molina, R., Cascino, T.L. and Vukov, A.M. 1997. Phase III Placebo-Controlled Trial of Capsaicin Cream In the Management of Surgical Neuropathic Pain In Cancer Patients. Journal of Clinical Oncology, 15(8), Pp.2974-2980.

Elster, J. 1985. Sour Grapes. Cambridge: Cambridge University Press.

Epatko, L. 2010. Chile’s Mine Rescue: Costs and Benefits. http://www.pbs.org/newshour/rundown/cost-of-mine-rescue/. Accessed 2 June 2017

Epstein, J.B. and Marcoe, J.H. 1994. Topical Application of Capsaicin for Treatment of Oral Neuropathic Pain and Trigeminal Neuralgia. Oral Surgery, Oral Medicine, Oral Pathology and Oral Radiology, 77(2), Pp.135-140.

Equality and Human Rights Commission. 2016. Equality Act FAQs. https://www.equalityhumanrights.com/en/equality-act/equality-act-faqs. Accessed 18 August 2016

European Commission. 2013. Rare Diseases – What Are They? http://ec.europa.eu/health/rare_diseases/policy/index_en.htm. Accessed 31 October 2015

European Commission. 2016. Orphan Medicinal Products. http://ec.europa.eu/health/human-use/orphan-medicines/index_en.htm. Accessed 17 October 2016.

FDA. 2015. Patents and Exclusivity. FDA/CDER SBIA Chronicles. May 2015. http://www.fda.gov/downloads/drugs/developmentapprovalprocess/smallbusinessassistance/ucm447307.pdf. Accessed 17 October 2016

Feinman S. 1978. The Blind As "Ordinary People". J Vis Impair B/Ind 1978;72:231-8.

Fishman, J. C., & Skrepnek, G. H. 2012. Pharmacoeconomic Analyses of Treatments for Rare Disease. Pharmaceuticals Policy and Law, 14, 51-67.

Fleck, L. M. 2010. Just Caring: In Defense of Limited Age-Based Healthcare Rationing. Cambridge Quarterly of Healthcare Ethics, 19(01), 27-37.

Fletcher, J. 1974, The Ethics of Genetic Control: Ending Reproductive Roulette, New York: Doubleday and Company.

Francis, L. P. 2002. Age Rationing Under Condition of Injustice. Medicine and Social Justice. Essays on the Distribution of Health Care. In: Rhodes R, Battin MP, Silvers A, Eds. Oxford: Oxford University Press, 270-7.

Frankfurt, H.G. 1988. Freedom of the Will and the Concept of A Person. In What is A Person? (Pp. 127-144). Humana Press.

Garau, M. and Devlin, N.J. 2017. Using MCDA As A Decision Aid In Health Technology Appraisal for Coverage Decisions: Opportunities, Challenges and Unresolved Questions. In Marsh (2017).

Gauthier D. 1982. Unequal Need: A Problem of Equity In Access to Health Care. President’s Commission for the Study of Ethical Problems In Medicine and Biomedical and Behavior Research Securing Access to Health Care, V. 2, Pp. 179-206. Government Printing Office, Washington,-DC, 1982

Gini, C. 1912. Variabilità E Mutabilità. Reprinted In Pizetti, E.; Salvemini, T., Eds. 1955. Memorie Di Metodologica Statistica. Rome: Libreria Eredi Virgilio Veschi.

Gladwell, M. 2007. Blink: The Power of Thinking Without Thinking. Back Bay Books.

Glasgow, J. 2009. Racism As Disrespect. Ethics 120: 64-93.

Gold, M.R., Siegel, J.E., Russell, L.B. and Weinstein, M. 1996. Cost-Effectiveness In Health and Medicine: Report of the Panel on Cost-Effectiveness In Health and Medicine. New York: Oxford Univ Press

Goodin, R. 1988. Reasons for Welfare: Economic, Sociological, and Political-But Ultimately Moral, In J. Donald Moon (Ed.) Responsibility, Rights, and Welfare, Boulder, Co.: Westview Press, Pp. 19–54.

Gosain, S. 2015. Consideration for Rare Diseases In the Context of Drug Reimbursement Processes (Doctoral Dissertation, Université d'Ottawa/University of Ottawa).

Gould, C. 1988, Rethinking Democracy: Freedom and Social Cooperation In Politics, Economics and Society, New York: Cambridge University Press.

Gross, N. and Simmons, S. 2007, October. The Social and Political Views of American Professors. In Working Paper Presented at A Harvard University Symposium on Professors and Their Politics.

Grove, W. M., & Meehl, P. E. 1996. Comparative Efficiency of Informal (Subjective, Impressionistic) and Formal (Mechanical, Algorithmic) Prediction Procedures: The Clinical–Statistical Controversy. Psychology, Public Policy, and Law, 2(2), 293-323.

Habermas, J. 1990. Discourse Ethics: Notes on A Program of Philosophical Justification. In the Communicative Ethics Controversy, Ed. S. Benhabib and F. Dallmayr, 60–110. Cambridge, Mass: MIT Press.

Habermas, J. 1990. Moral Consciousness and Communicative Action. Cambridge, MA: MIT Press.

Hadorn, D.C. 1991. Setting Health Care Priorities In Oregon: Cost-Effectiveness Meets the Rule of Rescue. Jama, 265(17), 2218-2225.

Hadorn, D.C. 1991. The Role of Public Values In Setting Health Care Priorities. Social Science & Medicine, 32(7), Pp.773-781.

Haidt, J. 2001. The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment. Psychological Review, 108(4), P.814.

Haidt, J., Koller, S., and Dias, M. Affect, Culture, and Morality, or is It Wrong to Eat Your Dog? Journal of Personality and Social Psychology 65 (1993): 613–628.

Hájek, A. 2012. Interpretations of Probability. The Stanford Encyclopedia of Philosophy (Winter 2012 Edition), Edward N. Zalta (Ed.), URL = <https://plato.stanford.edu/archives/win2012/entries/probability-interpret/>. Accessed 1 October 2016.

Halldenius, L. 2005. Dissecting Discrimination. Cambridge Quarterly of Healthcare Ethics 14: 455-463

Hampshire, S. (Ed.). 1978. Public and Private Morality. Cambridge University Press.

Harris, J. 1985. The Value of Life. Psychology Press.

Harris, J. 1987. QALYfying the Value of Life. Journal of Medical Ethics, 13(3), 117-123

Harris, J. 2005. It’s Not NICE to Discriminate. Journal of Medical Ethics, 31(7), 373-375.

Harris, J. 2009. Deciding Between Patients. In Kuhse & Singer (2009), Pp. 333-350

Harris, J. 1985. The Value of Life: An Introduction to Medical Ethics. Routledge.

Harris, J. 1987. QALYfying the Value of Life. Journal of Medical Ethics, 13(3), Pp.117-123.

Harris, J. 1994. Does Justice Require That We Be Ageist?. Bioethics, 8(1), Pp.74-83.

Harsanyi, J. C. 1953. Cardinal Utility In Welfare Economics and In the Theory of Risk-Taking. The Journal of Political Economy, 61(5), 434.

Haugen, A.S., Søfteland, E., Eide, G.E., Sevdalis, N., Vincent, C.A., Nortvedt, M.W. and Harthug, S. 2013. Impact of the World Health Organization's Surgical Safety Checklist on Safety Culture In the Operating Theatre: A Controlled Intervention Study. British Journal of Anaesthesia, 110(5), Pp.807-815.

Hausman, D. 2000. Cost-Value Analysis In Health Care: Making Sense Out of QALYs, Erik Nord. Cambridge University Press, 1999, 157 + Xxiii Pages. Economics and Philosophy, Issue 02, October 2000, Pp 333 - 378

Hausman, D. M. 2007. What's Wrong with Health Inequalities? Journal of Political Philosophy, 15(1), 46-66.

Hausman, D.M. and Mcpherson, M.S. 2009. Preference Satisfaction and Welfare Economics. Economics & Philosophy, 25(1), Pp.1-25.

Hausman, D.M. 2002. The Limits to Empirical Ethics. Summary Measures of Population Health: Concepts, Ethics, Measurement and Applications. Geneva, Switzerland: World Health Organization, Pp.641-6.

Hausman, D.M. 2006. Valuing Health. Philosophy & Public Affairs, 34(3), Pp.246-274.

Hausman, D.M. 2010. Valuing Health: A New Proposal. Health Economics, 19(3), Pp.280-296.

Hellman, D. 2008. When is Discrimination Wrong? Cambridge, Harvard University Press

Heyes, C. 2016. Identity Politics. The Stanford Encyclopedia of Philosophy (Summer 2016 Edition), Edward N. Zalta (Ed.), URL = <http://plato.stanford.edu/archives/sum2016/entries/identity-politics/>. Accessed 31 May 2016.

Holm, S. 1995. Not Just Autonomy--the Principles of American Biomedical Ethics. Journal of Medical Ethics, 21(6), 332-338.

Hope, T. 2001. Rationing and Life-Saving Treatments: Should Identifiable Patients Have Higher Priority?. Journal of Medical Ethics, 27(3), 179-185.

Hope, T. 2004. Medical Ethics: A Very Short Introduction. Oxford University Press.

Hope, T., Østerdal, L.P., & Hasman, A. 2010. An Inquiry Into the Principles of Needs-Based Allocation of Health Care. Bioethics, 24(9), 470-480.

Hopkins, R. 2007. What is Wrong with Moral Testimony?. Philosophy and Phenomenological Research, 74(3), 611-634.

Hughes, D. A., Tunnage, B., & Yeo, S. T. 2005. Drugs for Exceptionally Rare Diseases: Do They Deserve Special Status for Funding?. Qjm, 98(11), 829-836.

Hughes, J. and Walker, T. 2009. The Rule of Rescue In Clinical Practice. Clinical Ethics, 4(1), Pp.50-54.

Hume, D. 1739/1978. A Treatise of Human Nature. In L. A. Selby-Bigge (Ed.), Revised with Notes By P. H. Nidditch. Oxford: Clarendon Press.

Hunt, R. W. 1993. A Critique of Using Age to Ration Health Care. Journal of Medical Ethics, 19(1), 19-27.

Inman, T. 1861. Foundation for A New Theory & Practice of Medicine. John Churchill.

Insinga RP, Fryback DG. 2003. Understanding Differences Between Self-Ratings and Population Ratings for Health In the EuroQOL. Qual Life Res 2003 Sep; 12 (6): 611-9

Intellectual Property Watch. 2015. US Congressional Study Finds Excessive Profit-Seeking In USD84K Hepatitis Drug Sovaldi. Intellectual Property Watch. https://www.ip-watch.org/2015/12/01/us-congressional-study-finds-excessive-profit-seeking-in-usd84k-hepatitis-drug-sovaldi/. Accessed 2 June 2017

Jackson, J. C. 2006. Distributive Justice In Healthcare. In Jackson, J. C. 2006. Ethics In Medicine. Pp.102-115,213. Cambridge: Polity.

Jecker, N. S. 2013. Justice Between Age Groups: An Objection to the Prudential Lifespan Approach. The American Journal of Bioethics, 13(8), 3-15.

Jecker, N.S. 2013. The Problem with Rescue Medicine. Journal of Medicine and Philosophy, 38(1), 64-81.

Jenni, K. & Loewenstein, G. 'Explaining the 'Identifiable Victim Effect', Journal of Risk and Uncertainty 14 (1997): 235–257.

Johnson, S. 2016. Nurse Wins Fight for Cancer Drug After Call From Nicola Sturgeon's Office to Health Board. The Telegraph. Telegraph.Co.uk/news/2016/09/15/nurse-wins-fight-for-cancer-drug-after-call-from-nicola-sturgeon/. Accessed 2 June 2017

Johri, M., Damschroder, L. J., Zikmund-Fisher, B. J., & Ubel, P. A. 2005. The Importance of Age In Allocating Health Care Resources: Does Intervention-Type Matter?. Health Economics, 14(7), 669-678.

Jonsen A.R. 1986. Bentham In A Box: Technology Assessment and Health Care Allocation. Law Med Health Care 1986;14:172-4

Jonsen, A.R. and S. Toulmin, 1998, The Abuse of Casuistry, Berkeley: University of California Press.

Jonsen, A.R. 1995, Casuistry: An Alternative or Complement to Principles?. Kennedy Institute of Ethics Journal, 5: 237–51.

Jull, A., Parag, V., Walker, N. and Rodgers, A. 2010. Responsiveness of Generic and Disease‐Specific Health‐Related Quality of Life Instruments to Venous Ulcer Healing. Wound Repair and Regeneration, 18(1), Pp.26-30.

Juth, N. 2014. For the Sake of Justice: Should We Prioritize Rare Diseases?. Health Care Analysis, 1-20.

Kagan, S. 1989. The Limits of Morality, (Oxford: Clarendon Press)

Kahan, D.M. and Braman, D. 2006. Cultural Cognition and Public Policy. Yale Law & Policy Review, 24(1), Pp.149-172.

Kahan, D.M., Slovic, P., Braman, D. and Castil, J. 2005. Fear of Democracy: A Cultural Evaluation of Sunstein of Risk.

Kahneman, D., Knetsch, J.L. and Thaler, R.H. 1991. Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias. The Journal of Economic Perspectives, 5(1), Pp.193-206.

Kamm, F.M. 1994. 'to Whom?' Hastings Center Report 24, 4: 29–33.

Kamm, F.M. 1998. Précis of Morality, Mortality, Vol. 1: Death and Whom to Save From It. Philosophy and Phenomenological Research, 58(4), Pp.939-945.

Kamm, F.M. 2000. Does Distance Matter Morally to the Duty to Rescue. Law and Philosophy, 19(6), 655-681.

Kamm, F.M. 2001. Morality, Mortality Volume I: Death, and Whom to Save From It. Oxford University Press.

Kamm, F.M. 2001. Morality, Mortality Volume II: Rights, Duties, and Status. Oxford University Press.

Kamm, F.M. 2008. Responses to Commentators on Intricate Ethics. Utilitas, 20(1), Pp.111-142.

Kamm, F.M. 2009. Aggregation, Allocating Scarce Resources, and the Disabled. Social Philosophy and Policy, 26(01), Pp.148-197.

Kanavos, P., Vandoros, S., Irwin, R., Nicod, E. and Casson, M., 2011. Differences in costs of and access to pharmaceutical products in the EU. http://eprints.lse.ac.uk/41732/. Accessed: 1 February 2018

Kant, I. 1785/1996. Groundwork of the Metaphysics of Morals In Immanuel Kant: Practical Philosophy, Trans. Mary Gregor (Cambridge, 1996), P. 429, Italics Omitted. I Follow Kerstein (2009) In Referring to Preussische Akademie Edition (Vol. IV) Pagination

Kappel, K. and Sandøe, P. 1992. QALYs, Age and Fairness. Bioethics, 6(4), Pp.297-316.

Kappel, K. and Sandøe, P. 1994. Saving the Young Before the Old-A Reply to John Harris. Bioethics, 8(1), Pp.84-92.

Karikios, D., Schofield, D., Salkeld, G., Mann, K.P., Trotman, J. and Stockler, M.R. 2014. Rising Cost of Anticancer Drugs In Australia. Internal Medicine Journal, 44(5), Pp.458-463.

Kerstein, S. 2009. Treating Others Merely As Means. Utilitas, 21(2), Pp.163-180.

Kerstein, S.J. 2015. Dignity, Disability, and Lifespan. Journal of Applied Philosophy.

Khanka, S.S. 2009. Motivational Orientation of Assamese Entrepreneurs In the SME Sector. The Journal of Entrepreneurship, 18(2), Pp.209-218.

Klarman, M.J. 2006. From Jim Crow to Civil Rights: The Supreme Court and the Struggle for Racial Equality. Oxford University Press.

Kornstein, S.G., Schatzberg, A.F., Thase, M.E., Yonkers, K.A., Mccullough, J.P., Keitner, G.I., Gelenberg, A.J., Davis, S.M., Harrison, W.M. and Keller, M.B. 2000. Gender Differences In Treatment Response to Sertraline Versus Imipramine In Chronic Depression. American Journal of Psychiatry, 157(9), Pp.1445-1452.

Kuhse, H. , & Singer, P. (Eds.). 2009. A Companion to Bioethics. John Wiley & Sons.

Kuhse, H., & Singer, P. 1988. Age and the Allocation of Medical Resources.Journal of Medicine and Philosophy, 13(1), 101-116.

Ladd, P. 2003. Understanding Deaf Culture: In Search of Deafhood. Multilingual Matters.

Lance, M., & Little, M. 2008. From Particularism to Defeasibility In Ethics. In Potrc Et Al, 2007, 53-74.

Largent, E.A. and Pearson, S.D. 2012. Which Orphans Will Find A Home? the Rule of Rescue In Resource Allocation for Rare Diseases. Hastings Center Report, 42(1), Pp.27-34.

Lawlor, R. 2006. Taurek, Numbers and Probabilities. Ethical Theory and Moral Practice, 9(2), Pp.149-166.

Lazenby, H. 2011. Is Age Special? Justice, Complete Lives and the Prudential Lifespan Account. Journal of Applied Philosophy, 28(4), 327-340.

Lee, J. 2015. Gilead's 2014 Profit Margin Nears 50%, Fueled By Hep C Drugs. Modern Healthcare. http://www.modernhealthcare.com/article/20150203/news/302039949. Accessed 2 June 2017

Levine, M. P. , & Pataki, T. (Eds.). 2004. Racism In Mind. Cornell University Press.

Lewis, P. A., & Charny, M. 1989. Which of Two Individuals Do You Treat When Only Their Ages Are Different and You Can't Treat Both?. Journal of Medical Ethics, 15(1), 28-34.

Linley, W. G., & Hughes, D. A. 2013. Societal Views on NICE, Cancer Drugs Fund and Value-Based Pricing Criteria for Prioritising Medicines: A Cross-Sectional Survey of 4118 Adults In Great Britain. Health Economics, 22(8), 948-964.

Lippert-Rasmussen, K. (2006a).The Badness of Discrimination. Ethical Theory and Moral Practice 9: 167-185

Lippert-Rasmussen, K. 2014. Born Free and Equal?: A Philosophical Inquiry Into the Nature of Discrimination. Oxford University Press.

Lockwood, M. 1988. Quality of Life and Resource Allocation. In Bell & Mendus (1988), 33-55.

Machery, E. 2004. 'Semantics, Cross-Cultural Style.' Cognition 92 (2004): B1–B12.

Machery, E. 2012. 'Expertise and Intuitions About Reference.' Theoria 27 (2012): 37–54.

Magnifico, L. 2016. What's the Outlook and Life Expectancy for Hepatitis C? http://www.healthline.com/health/hepatitis-c/hepatitis-prognosis-and-life-expectancy. Accessed 2 June 2017

Marsh, K., Goetghebeur, M., Thokala, P. and Baltussen, R. 2017. Multi-Criteria Decision Analysis to Support Healthcare Decisions. Cham: Springer International Publishing

Mccabe, C., Claxton, K., & Tsuchiya, A. 2005. Orphan Drugs and the NHS: Should We Value Rarity?. BMJ: British Medical Journal, 331(7523), 1016.

Mccabe, C., Tsuchiya, A., Claxton, K., & Raftery, J. 2006. Drugs for Exceptionally Rare Diseases: A Commentary on Hughes Et Al.

Mccabe, C., Tsuchiya, A., Claxton, K., & Raftery, J. 2006. Orphan Drugs Revisited. Qjm, 99(5), 341-345.

Mccabe, C., Tsuchiya, A., Claxton, K., Raftery, J. 2007. Assessing the Economic Challenges Posed By Orphan Drugs: A Comment on Drummond Et Al. International Journal of Technology Assessment In Health Care, 23(3), 397-404.

Mcguire, A., Drummond, M., Martin, M. and Justo, N. 2015. End of Life or End of the Road? Are Rising Cancer Costs Sustainable? Is It Time to Consider Alternative Incentive and Funding Schemes?. Expert Review of Pharmacoeconomics & Outcomes Research, 15(4), Pp.599-605.

Mckerlie, D. 2001a. Justice Between the Young and the Old. Philosophy & Public Affairs, 30(2), 152-177.

Mckerlie, D. 2001b. Dimensions of Equality. Utilitas, 13(3), 263-288.

Mckerlie, D. 1992. Equality Between Age-Groups. Philosophy & Public Affairs, 275-295.

Mckerlie, D. 2006. Egalitarianism and the Difference Between Interpersonal and Intrapersonal Judgments.

Mckie, J. , Kuhse, H. , Richardson, J. , & Singer, P. 1996. Another Peep Behind the Veil. Journal of Medical Ethics, 22(4), 216-221.

Mckie, J. , Kuhse, H. , Richardson, J. , & Singer, P. 1996. Double Jeopardy, The Equal Value of Lives and the Veil of Ignorance: A Rejoinder to Harris. Journal of Medical Ethics, 22(4), 204-208.

Mckie, J., & Richardson, J. 2003. The Rule of Rescue. Social Science & Medicine, 56(12), 2407-2419.

Mcmillan, J. , Sheehan, M. , Austin, D. , & Howell, J. 2006. Ethics and Opportunity Costs: Have NICE Grasped the Ethics of Priority Setting? Journal of Medical Ethics, 32(3), 127-128.

Mcnaughton, D. 1988. Moral Vision: An Introduction to Ethics.

Menzel P, Dolan P, Richardson J, Olsen JA. The Role of Adaptation to Disability and Disease In Health State Valuation: A Preliminary Normative Analysis. Soc Sci Med. 2002;55:2149–58.

Mill, J.S. 1861. Considerations on Representative Government. The Philosophy of John Stuart Mill, New York: Modern Library, Pp.401-406

Moerman, D.E. 2002. Meaning, Medicine, and the Placebo Effect. Cambridge: Cambridge University Press.

Moreau, S. 2010. What is Discrimination? Philosophy & Public Affairs 38(2): 143-179.

Morris A, O'Donnell T, Eds. 1999. Feminist Perspectives on Employment Law. London: Cavendish Publishing

Murray, C.J. and Lopez, A.D. 1996. The Incremental Effect of Age-Weighting on YLLs, YLDs, and DALYs: A Response. Bulletin of the World Health Organization, 74(4), P.445.

Murray, C.J.L. 1994. Quantifying the Burden of Disease: The Technical Basis for Disability-Adjusted Life Years. Bulletin of World Health Organization 72, 429-445.

Nado, J. 2014. Philosophical Expertise. Philosophy Compass, 9(9), Pp.631-641.

Nagel, T. 1978. Ruthlessness In Public Life. In Hampshire (1978)

National Institute of Health. 2010. Quarterly Newsletter of the NIH Rare Diseases Clinical Research Network, December 2010.

NHS. 2016. Body Dysmorphic Disorder (BDD). http://www.nhs.uk/conditions/body-dysmorphia/pages/introduction.aspx. Accessed 25 June 2016.

NHS. 2016. Body Dysmorphic Disorder (BDD). NHS. http://www.nhs.uk/conditions/body-dysmorphia/pages/introduction.aspx. Accessed 13 January 2017.

NICE. 2009. Appraising Life-Extending, End of Life Treatments. London: NICE

NICE. 2013. Guide to the Methods of Technology Appraisal 2013. London: National Institute for Health and Care Excellence

NICE. 2014. Value Based Assessment of Health Technologies (Consultation Paper). London: NICE

NICE. 2017. Measuring Effectiveness and Cost Effectiveness: The QALY. London: NICE. https://web.archive.org/web/20090809124835/http://www.nice.org.uk/newsroom/features/measuringeffectivenessandcosteffectivenesstheqaly.jsp. Accessed 21 February 2018

NICE. 2018. History of NICE. https://www.nice.org.uk/about/who-we-are/history-of-nice. Accessed 1 March 2018

Nichols, S, Stich, S., and Weinberg, J. 'Metaskepticism: Meditations In Ethno-Epistemology.' the Skeptics. Ed. Steven Luper. Burlington, VT: Ashgate, 2003. 227–248.

Nickerson, R.S. 1998. Confirmation Bias: A Ubiquitous Phenomenon In Many Guises. Review of General Psychology, 2(2), P.175.

Nietzsche, F. 1886/1966. Beyond Good and Evil. Trans. W. Kaufmann. New York: Random House.

NOP World. 2005. NOP World Culture Score Index Examines Global Media Habits ... Uncovers Who's Tuning In, Logging on and Hitting the Books. https://www.prnewswire.com/news-releases/nop-world-culture-scoretm-index-examines-global-media-habits-uncovers-whos-tuning-in-logging-on-and-hitting-the-books-54693752.html. Accessed 5 March 2018

Norcross, A. 2008. Off Her Trolley? Frances Kamm and the Metaphysics of Morality. Utilitas, 20(1), Pp.65-80.

Nord, E. 1999. Cost-Value Analysis In Health Care: Making Sense Out of QALYs. Cambridge University Press.

Nord, E. 2005. Concerns for the Worse Off: Fair Innings Versus Severity. Social Science & Medicine, 60(2), Pp.257-263.

Nord, E., Richardson, J., Street, A., Kuhse, H., Singer, P. 1996. The Signi®Cance of Age and Duration of Effect In Social Evaluation of Health Care. Health Care Analysis 4, 103±111.

Nussbaum, M. 2011. Creating Capabilities: The Human Development Approach. Cambridge, MA: Harvard University Press. Pp. 33–34

O'Brian, P. 1990. Master and Commander. London: WW Norton & Company.

Office for National Statistics. 2016. Births By Parents' Characteristics In England and Wales: 2015. https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/livebirths/bulletins/birthsbyparentscharacteristicsinenglandandwales/2015. Accessed 22 May 2017

Okin, SM. 2003. Poverty, Well-Being, and Gender: What Counts, Who’s Heard? Philosophy and Public Affairs 31 (3): 280-316

Orphanet. 2016. What is An Orphan Drug? http://www.orpha.net/consor/cgi-bin/education_aboutorphandrugs.php?lng=en. Accessed 17 October 2016

Orr, S.W., & Wolff, J. 2014. Reconciling Cost-Effectiveness with the Rule of Rescue: The Institutional Division of Moral Labour. Theory and Decision: An International Journal for Multidisciplinary Advances In Decision Sciences.

Pager, D. , & Shepherd, H. 2008. The Sociology of Discrimination: Racial Discrimination In Employment, Housing, Credit, and Consumer Markets. Annual Review of Sociology, 34, 181.

Palmer, E. 2015. Gilead Considered $115K Price for Sovaldi, Senate Investigation Says. Fiercepharma. http://www.fiercepharma.com/regulatory/gilead-considered-115k-price-for-sovaldi-senate-investigation-says. Accessed 2 June 2017

Parfit, D. 1984. Reasons and Persons. Oxford: Clarendon Press.

Parfit, D. 1997. Equality and Priority. Ratio, 10(3), 202-221.

Parfit, D. 1978. Innumerate Ethics. Philosophy & Public Affairs, Pp.285-301.

Parfit, D. 1997. Equality and Priority. Ratio, 10(3), Pp.202-221.

Peeters, Y., Vliet Vlieland, T.P. and Stiggelbout, A.M. 2012. Focusing Illusion, Adaptation and EQ-5D Health State Descriptions: The Difference Between Patients and Public. Health Expectations, 15(4), Pp.367-378.

Pemberton, M. 2011. Finger on the Pulse. The Telegraph, 18 Apr 2011. http://www.telegraph.co.uk/lifestyle/wellbeing/maxpemberton/8450877/finger-on-the-pulse-max-pemberton.html. Accessed 3 November 2015.

Phillips, L.D. 2017. Best Practice for MCDA In Healthcare. In Marsh (2017).

Popper, KR. 1957, The Propensity Interpretation of the Calculus of Probability and the Quantum Theory. In S. Körner (Ed.), The Colston Papers, 9: 65–70.

Quigley, M. 2007. A NICE Fallacy. Journal of Medical Ethics, 33(8), 465-466.

Rabin, R., Oemar, M., Oppe, M., Janssen, B. and Herdman, M. 2011. EQ-5D-3L User Guide: Basic Information on How to Use the EQ-5D-3L Instrument. Rotterdam: Euroqol Group, 22.

Radcliffe Richards, J. 2000. Practical Reason and Moral Certainty-the Case of Discrimination. In Ullmann-Margalit (2000), P. 151-163.

Rare Genomics Institute. 2015. Dr. Emil Kakkis, CEO of Ultragenyx – Rare Disease Drug Developer. http://raregenomics.org/emil-kakkis-ceo-ultragenyx-rare-disease-drug-developer. Accessed 31 October 2015.

Rawlins, M. , & Dillon, A. 2005. NICE Discrimination. Journal of Medical Ethics, 31(12), 683-684.

Rawls, J. 1971. A Theory of Justice, 2nd Edition 1999, Cambridge, MA: Harvard University Press.

Reed V. Reed, 404 U. S. 71, 92 S. Ct. 251, 30 L. Ed. 2d 225 (1971).

Rémuzat, C., Mzoughi, O., Rodrigues, J., Korchagina, D., & Toumi, M. 2014. HTA Assessment Comparison of Orphan Drugs In France and Germany. Value In Health, 3(17), A236-A237.

Richardson, J. 2002. The Poverty of Ethical Analyses In Economics and the Unwarranted Disregard of Evidence. Summary Measures of Population Health: Concepts, Ethics, Measurement and Applications, Geneva: World Health Organization, Pp.627-640.

Risse, M. and Zeckhauser, R. 2004. Racial Profiling. Philosophy & Public Affairs, 32(2), Pp.131-170.

Ross, W. D. (1930/2002). The Right and the Good. Oxford University Press.

Rulli, T., & Millum, J. 2014. Rescuing the Duty to Rescue. Journal of Medical Ethics, Medethics-2013.

Samuelson, P.A. 1954. The Pure Theory of Public Expenditure. Review of Economics and Statistics. 36 (4): 387–389

Sartre, J. P., & Elkaïm-Sartre, A. 2007. Existentialism is A Humanism. Yale University Press.

Sayers, G.M. and Nesbitt, T. 1998. Ageism In the NHS and the Human Rights Act 1998: An Ethical and Legal Enquiry. European Journal of Health Law, Volume 9, Issue 1]

Scanlon, T. 2003. Value, Desire, and the Quality of Life, In the Difficulty of Tolerance: Essays In Political Philosophy (Cambridge: Cambridge University Press), Pp. 175-76

Scanlon, T. M. 1998. What We Owe to Each Other, Cambridge, MA: Harvard University Press.

Scanlon, T. M. 2008. Moral Dimensions. Harvard University Press.

Schicktanz, S., Schweda, M. and Wynne, B. 2012. The Ethics of 'Public Understanding of Ethics' - Why and How Bioethics Expertise Should Include Public and Patients' Voices. Medicine, Health Care and Philosophy, 15(2), Pp.129-139.

Schlander, M., & Beck, M. 2009. Expensive Drugs for Rare Disorders: to Treat or Not to Treat? the Case of Enzyme Replacement Therapy for Mucopolysaccharidosis VI. Current Medical Research and Opinion, 25(5), 1285-1293.

Schlander, M., Garattini, S., Holm, S., Kolominsky-Rabas, P., Nord, E., Persson, U., Postma, M., Richardson, J., Simoens, S., Solà Morales, O.D. and Tolley, K. 2014. Incremental Cost Per Quality-Adjusted Life Year Gained? The Need for Alternative Methods to Evaluate Medical Interventions for Ultra-Rare Disorders. Journal of Comparative Effectiveness Research, 3(4), Pp.399-422.

Schmidt, F. L. , & Hunter, J. E. 1998. The Validity and Utility of Selection Methods In Personnel Psychology: Practical and Theoretical Implications of 85 Years of Research Findings. Psychological Bulletin, 124(2), 262.

Schokkaert, E. 2015. How to Introduce More (Or Better) Ethical Arguments In HTA?. International Journal of Technology Assessment In Health Care, 31(03), Pp.111-112.

Schöne-Seifert, B. 2009. The 'Rule of Rescue'in Medical Priority Setting: Ethical Plausibilities and Implausibilities. Rationality Markets Morals.

Schulz, E., Cokely, E.T. and Feltz, A. 2011. Persistent Bias In Expert Judgments About Free Will and Moral Responsibility: A Test of the Expertise Defense. Consciousness and Cognition, 20(4), Pp.1722-1731.

Schwitzgebel, E, and Cushman, F. 2012. 'Expertise In Moral Reasoning? Order Effects on Moral Judgment In Professional Philosophers and Non-Philosophers.' Mind & Language 27 (2012): 135–153.

Schwitzgebel, E. and Cushman, F. 2015. Philosophers' Biased Judgments Persist Despite Training, Expertise and Reflection. Cognition, 141, Pp.127-137.

Segall, S. 2012. What's So Bad About Discrimination? Utilitas, 24(01), 82-100.

Sen, A.K. 1992. Inequality Re-Examined, Oxford: Clarendon Press.

Sen, A.K. 1987. On Ethics and Economics. Oxford: Oxford University Press.

Shah, K.K., Tsuchiya, A. and Wailoo, A.J. 2014. Valuing Health at the End of Life: An Empirical Study of Public Preferences. The European Journal of Health Economics, 15(4), Pp.389-399.

Shaw, A. B. 1994. In Defence of Ageism. Journal of Medical Ethics, 20(3), 188-194.

Sheehan, M. 2007. Resources and the Rule of Rescue. Journal of Applied Philosophy, 24(4), 352-366.

Silver RL. 1982. Coping with An Undesirable Life Event: A Study of Early Reactions to Physical Disability (Unpublished Doctoral Dissertation). Evanston, IL: Northwestern University, 1982.

Simpson, J. 1997. Touching the Void. Random House.

Sinclair, S. 2012. How to Avoid Unfair Discrimination Against Disabled Patients In Healthcare Resource Allocation. Journal of Medical Ethics, 38(3), Pp.158-162.

Singer, P. 1972. Famine, Affluence, and Morality. Philosophy & Public Affairs, 229-243.

Singer, P. 2005. Ethics and Intuitions, The Journal of Ethics 9: 331-352.

Singer, P. 1999. Practical Ethics, 2nd Edition, Cambridge: Cambridge University Press.

Singer, P. 2005. Ethics and Intuitions. The Journal of Ethics, 9(3-4), Pp.331-352.

Smith, R. 2014. Most Terminally Ill Patients Not Told They Are Dying, Says Damning Report. https://www.telegraph.co.uk/news/health/news/10829690/most-terminally-ill-patients-not-told-they-are-dying-says-damning-report.html. Accessed 1 March 2018

Snider, B. 1993. A Global View. Gallaudet Today, Summer 1993: 28–31.http://saveourdeafschools.org/paddy_ladd_1993.pdf. Accessed 31 May 2016.

Soto, C. 2012. The Veil of Ignorance and Health Resource Allocation. Journal of Medicine and Philosophy, 37(4), 387-404.

Splash, C.L. 2007. Deliberative Monetary Valuation (DMV) In Theory (No. 2007-01). CSIRO Sustainable Ecosystems.

Stein, M. S. 2002. The Distribution of Life-Saving Medical Resources: Equality, Life Expectancy, and Choice Behind the Veil. Social Philosophy and Policy,19(02), 212-245.

Steinpreis, R., Anders, K., and Ritzke, D. 1999. The Impact of Gender on the Review of the Curricula Vitae of Job Applicants and Tenure Candidates: A National Empirical Study. Sex Roles, 41: 7/8, 509-528.

Stewart, F. 2001. Book Review Women and Human Development: The Capabilities Approach, By Martha Nussbaum Journal of International Development 13 (8): 1191-1192

Taurek, J.M. 1977. Should the Numbers Count?. Philosophy & Public Affairs, Pp.293-316.

Temkin, L. 2003. Equality, Priority, or What? Economics and Philosophy 19: 61–88.

The Scottish Government. 2015. Fund for New Medicines Doubles. http://news.scotland.gov.uk/news/fund-for-new-medicines-doubles-18eb.aspx Accessed 31 October 2015

Thomson, J. J. 1976. Killing, Letting Die, and the Trolley Problem. The Monist, 59: 204–17.

Times Higher Education. 2015. Almost Half of Sector Back Labour, THE Election Poll Suggests. London: Times Higher Education.

Times Higher Education. 2016. EU Referendum: 9 Out of 10 University Staff Back Remain. London: Times Higher Education.

Tobia, K., Buckwalter, W. and Stich, S. 2013. Moral Intuitions: Are Philosophers Experts?. Philosophical Psychology, 26(5), Pp.629-638.

Tolstoy, L. (1877/2014). Anna Karenina. Yale University Press.

Tordrup, D., Tzouma, V., & Kanavos, P. 2014. Orphan Drug Considerations In Health Technology Assessment In Eight European Countries. Rare Diseases and Orphan Drugs: An International Journal of Public Health, 1(3), 86-97.

Toulmin, S. 1981. The Tyranny of Principles. Hastings Center Report, 11(6), Pp.31-39.

Triggle, N. 2005. Hewitt 'Has Left NHS Toothless'. BBC. news.bbc.co.uk/1/hi/health/4422662.stm. Accessed 2 june 2017

Tsuchiya, A. 1996. The Value of Health at Different Ages. Journal of Health Care and Society 6 (3), 123±136 In Japanese.

U.S. Food and Drug Administration. 2013. Orphan Drug Act: Relevant Excerpts. Last Updated August 2013. http://www.fda.gov/forindustry/developingproductsforrarediseasesconditions/howtoapplyfororphanproductdesignation/ucm364750.htm. Accessed 3 November 2015

Ubel, P.A., Hirth, R.A., Chernew, M.E., & Fendrick, A.M. 2003. What is the Price of Life and Why Doesn't It Increase at the Rate of Inflation?. Archives of Internal Medicine, 163(14), 1637-1641.

Ullmann-Margalit, E. 2000. Reasoning Practically. Oxford: Oxford University Press.

Unger, P. K. 1996. Living High and Letting Die: Our Illusion of Innocence (P. 133). Oxford: Oxford University Press.

University of Arizona. 2017. How a Wrong Theory Brought Good Results. University of Arizona. http://ircamera.as.arizona.edu/natsci102/natsci102/images/extsteadystate.htm. Accessed 7 January 2018

Van Der Sluijs, J.P., Kaiser, M., Beder, S., Hosle, V., Kemelmajer De Carlucci, A. and Kinzig, A. 2005. The Precautionary Principle.Paris: United Nations Educational, Scientific and Cultural Organization

Van Reenen, M. and Janssen, B. 2015. EQ-5D-5L User Guide. Basic Information on How to Use the EQ-5D-5L Instrument. Version 2.1, April 2015.

Vanhuysse, P. 2014. Intergenerational Justice and Public Policy In Europe.European Social Observatory (OSE) Paper Series, Opinion Paper, (16).

Väyrynen, P. 2009. A Theory of Hedged Moral Principles. Oxford Studies In Metaethics 4 (2009), 91-132.

Veatch, R. M. 1993. How Age Should Matter: Justice As the Basis for Limiting Care to the Elderly. Facing Limits: Ethics and Health Care for the Elderly. Boulder: Westview, 211-29.

Voorhoeve, A., & Fleurbaey, M. 2012. Egalitarianism and the Separateness of Persons. Utilitas, 24(03), 381-398.

Wagland, R. 2012b. Social Injustice: Distributive Egalitarianism, The Complete Life View, and Age Discrimination. eprints.soton.ac.uk. Accessed 1 October 2016.

Wagland, R. 2012a. A Fair Innings or A Complete Life: Another Attempt at An Egalitarian Justification of Ageism.

Wahlberg, A. and Rose, N. 2015. The Governmentalization of Living: Calculating Global Health. Economy and Society, 44(1), Pp.60-90.

Weimer, D.L. and Vining, A.R. 2015. Policy Analysis: Concepts and Practice. Fourth Edition. Pearson: Prentice Hall. P. 72

Weinberg, J., Nichols, S. and Stich, S. 'Normativity and Epistemic Intuitions.' Philosophical Topics 29 (2001): 429–460.

Wells, T. 2018. Sen's Capability Approach. Internet Encyclopedia of Philosophy. http://www.iep.utm.edu/sen-cap/. Accessed 18 February 2018

Williams, A.H. 1997. Intergenerational Equity: An Exploration of the 'Fair Innings' Argument. Health Economics, 6(2), Pp.117-132.

Williams, A.H. and Cookson, R.A. 2006. Equity–Efficiency Trade-Offs In Health Technology Assessment. International Journal of Technology Assessment In Health Care, 22(01), Pp.1-9.

Williams, B. 1981. Persons, Character and Morality. In Williams, B. 1981. Moral Luck. Cambridge: Cambridge University Press

Wirtz, V., Cribb, A., & Barber, N. 2003. Understanding the Role of The Hidden Curriculum In Resource Allocation - the Case of the UK NHS. Health Care Analysis, 11(4), 295-300.

World Health Organization, 2001. The World Health Report 2001: Mental Health: New Understanding, New Hope. World Health Organization.

Zamzow, J.L. and Nichols, S. 2009. Variations In Ethical Intuitions. Philosophical Issues, 19(1), Pp.368-388.