Do NICE's HTA processes still lead to net improvements in NHS services?

James Wilson

November 2023. Final version published at http://doi.org/10.4324/9781003501268-6

NICE's health technology appraisal (HTA) processes were introduced to improve the overall running of the NHS by reducing postcode prescribing, bringing a greater degree of cost control, and taking specific decisions about which drugs to fund out of the hands of politicians. Since 2010, NHS budgets have not kept pace with rising healthcare needs, leading among other effects to record waiting times for routine surgery. Over the same period, NICE has increasingly focused on promoting innovation and has given less weight within its deliberations to the opportunity costs of recommending cost-ineffective interventions. Thus, as budgets have become tighter within the NHS as a whole, NICE has been increasingly willing to recommend health technologies that are of relatively low cost-effectiveness, so long as they are innovative.

NICE's greater emphasis on innovation has effects that are on average negative for population health. However, a good health and care system will have goals beyond the 1maximization of health, so a mean loss in population health would not automatically imply that the shift has been a net worsening for the NHS as a whole. Nonetheless, absent a plausible account of why it is ethically better to displace more health benefits than are created by prioritising innovation, it would be difficult to avoid the conclusion that this policy shift has contributed to a worsening of the care provided in the NHS. Unfortunately, NICE has not articulated such a rationale in detail. This chapter analyses NICE's approach to innovation, and to highly specialised technologies, arguing that even a sympathetic reconstruction of NICE's approach fails to uncover a convincing and consistent ethical framework. It is far from clear that NICE's current HTA processes lead to net improvements in NHS services.

1 Introduction

Section 233 of the Health and Social Care Act 2012 requires that in exercising its functions, NICE must have regard to three considerations. First, "the broad balance between the benefits and costs of the provision of health services or of social care",

second “the degree of need of persons for health services or social care in England”, and third, “the desirability of promoting innovation in the provision of health services or of social care in England”. The relationship between the benefits and costs of interventions delivered, and the degree of need for health and social care services have obvious—though contrasting—relationships to questions of fairness and equity in the delivery of services. I shall consider innovation shortly.

Focusing on the broad balance between benefits and costs requires above all, examining whether making a change such as introducing a new intervention, or revising guidance on management of a particular disease, would bring an overall improvement, or an overall worsening in the provision of health services or social care. The initial years of NICE coincided with an unprecedented rate of increase in the overall NHS budget, but since 2010 increases in health budgets have not kept pace with the rising costs of meeting health care needs. What counts as an overall improvement (or at least not worsening) needs to be understood against this background of rising scarcity within the health and social care system.

Waiting times, and a range of other service quality indicators, had been worsening significantly even before the Covid-19 pandemic. The diversion of healthcare resources away from non-Covid treatments exacerbated these problems in 2020–21, but performance against a range of indicators continued to decline once the pandemic subsided (Morris and Reed, 2022). Waiting lists for elective operations were larger by 2023 than at any point in the NHS’s history, which was combined with a crisis in adult social care (Schlepper and Dodsworth, 2023). The health and care system as a whole has failed to scale its ability to meet health and social care needs to match increases in levels of these needs—and in some cases the absolute level of system capacity has declined, for example in numbers of GPs (BMA, 2023).

Under these circumstances, considering opportunity costs — those benefits that are already being provided that would be displaced by the introduction of new interventions — is of even greater importance than at NICE’s inception. The simplest way to assess opportunity costs would be via a single measurement of value, such as a Quality Adjusted Life Year (QALY). On this simple view, both the QALYs to be gained, and those that would be displaced should be compared, and the intervention introduced

only if the expectation would be that the benefits created would be greater than those displaced.

However, NICE's evaluative second requirement, namely to consider the *need* of persons for health services or social care, provides reasons for thinking that NICE would be failing to deliver on its role as set down in statute, if it focused only on the maximisation of QALYs in considering opportunity costs. In particular, needs that are urgent and important, such as providing intensive care unit treatment can be costly to meet, as are drugs for very rare conditions. It is very plausible to think that QALYs would be maximised in ways that require diverting resources from such relatively cost-ineffective measures to ones that create additional health benefits more cheaply (for example, public health preventive measures). Nonetheless, the presumption that meeting health needs has an ethical importance that is separate from, and not reducible to maximising health benefits has a wide appeal: many worry that shifting resources from cost-ineffective treatment to more cost-effective prevention would *not* constitute an overall improvement in the health system, even though it would improve population health.[1]

In short, there should be no doubt that creating entitlements to novel interventions that are significantly less cost-effective than the mean cost-effectiveness of interventions within the health and social care system will tend to reduce, rather than increase, the overall health benefits the system generates. However, if QALY maximisation is an implausible goal for a health system, this loss of potential QALYs may yet be associated with an overall improvement rather than a worsening. This suggests that our conception of opportunity costs needs to shift in one of two ways. Either opportunity cost is just one of a number of ethically relevant concerns — such that, as in the case of funding intensive treatment units, it may be not just ethically permissible to pursue a policy that leads to more health being displaced than is created, but arguably ethically

[1] I argue in Chapter 7 of Wilson (2021), and in Wilson (2022), that while it is legitimate for a democratic deliberation process to give a strong preference for treatment over prevention, there is no necessity to do so. Considered in themselves, the ethical arguments advanced for preferring expensive treatments to cheaper prevention are far from compelling.

required to do so. Or alternatively, a broader vision of what counts as an opportunity cost could be adopted, such that we think of the opportunity costs of particular interventions in terms of the displacement of *value*, rather than health. On the broader view, opportunity costs map onto an account of whatever counts as an improvement for a health system, rather than presupposing that it is only the maximisation of health that matters. I suggest the latter is a more perspicuous framing of opportunity costs, as it foregrounds the crucial decisions in prioritisation, which is how to weigh the maximisation of health against other desiderata for a health system.

2 What counts as an improvement?

I use the concept of health system improvement to refer to the processes by which policymakers aim to deliver the best health system possible, given financial, broader resource, and other constraints (Wilson, 2023). I construe improvement in a deliberately broad and inclusive way, with the intention that it will be uncontentious (indeed truistic) that policymakers should aim to improve health systems. What will be contentious is the principles that should guide improvement, and what improvement will require in practice. Thus, while the concept of health system improvement serves to indicate neutrally a topic of enquiry, there will be a variety of competing conceptions of health system improvement.

Within health system improvement, we should distinguish between means-improvement and values-improvement. Means-improvement involves mapping the ways in which a system converts inputs into outputs, and examining whether these could be reconfigured to allow the system as a whole to better to achieve the values it aims to instantiate and promote. So, for example much of NICE's Clinical Guidance would count as means-improvement — aiming at better aligning treatment and clinical processes with the existing body of evidence. Values-improvement involves specifying and reconciling the values that a system should instantiate and promote. NICE's historical use of its Citizens' Council to provide advice and guidance in respect of particular ethical questions such as whether NICE should use the Rule of Rescue (NICE Citizens Council, 2004), would count as values-improvement.

Attempts at means-improvement always presuppose an account of values. Sometimes these values may be left unarticulated — perhaps because they are believed to be too obvious to be worth reflecting on — but the fact remains that values are presupposed in any attempt at means improvement. A set of changes to a health system will count a successful means-improvement only if it allows the system as a whole to move in the direction of its values, as articulated and reconciled through earlier processes of values-improvement. It also follows that what counts as successful means-improvement will also shift if the values as articulated in a process of values-improvement change.

I have described values-improvement as a process. Just as an attempted means-improvement may make a system worse, if for example, it reduces the ability to respond to patient need or makes the system more brittle, so attempts at values-improvement can also make a system worse. At a high level, a process of values-improvement fails if the revised set of values it leads to are worse than the values from which the revisions set out.

This raises an important question about how we determine what makes one set of values *better* than another, rather than just different? Obviously, this is a large question, but one thing that is paramount is the kind of story those leading the process are able to tell about *why* moving from the former set of values to a new one counts as an improvement. Things that can be appealed to are both features of the process, for example broad and deep inclusion of a wide range of relevant stakeholders within the process that leads to the value revisions, and substantive improvements in values, such as reduction in unfair exclusions, improved alignment with broader sets of values widely accepted elsewhere such as human rights standards, decreases in errors, or increase in capacity. Similarly, in criticising a set of values as a worsening, we could appeal to elements of process or substantive failures in the values. For example, Charlton, Lomas and Mitchell (2022) argue that the changes made to the way that social values are handled in NICE's Methods Manual in 2022 could be criticised on both procedural and substantive grounds, and if they are correct in this, the changes would amount to a values-worsening rather than a values-improvement.

Distinguishing means-improvement from values-improvement allows us to pose a set of questions about how NICE's role relates to the rest of the NHS. NICE has a job to do, and

in performing its role it must have regard to the three considerations with which this chapter began. However, from a broader perspective, the crucial question is whether if NICE performs its health technology assessments (HTA) in line with these three considerations, this will lead to a means-improvement in the health and social care system as whole.

NICE's statutory duty to promote innovation in the provision of health services has been interpreted by it in a way that would license paying a greater incremental cost for interventions that are held to be 'innovative' than would be proportionate to the health benefits they provide. One central ethical challenge for this interpretation, which NICE has failed to discharge, is explaining how and why giving innovation an independent weighting provides an improvement either to NICE's processes, or to the broader health and social care system. One important question is whether the rationale is that promoting innovation in this way is a good means to improve the health and social care system over the medium or long term (even if it leads to short-term losses), or whether the rationale lies in benefits outside the direct sphere of the health and social care system (for example, by supporting the life sciences industry in the UK). Whatever its rationale, installing innovation as a third priority within NICE has contributed to the NHS as a whole in England becoming *less* able to meet a variety of patient needs, and less capable of delivering good value for money than it was previously. That the rise in the importance of innovation has led to a reduction in the focus given to the other two goals is not seriously in doubt (Charlton and Rid, 2019). In other words, the increased focus on innovation as an end in itself has led to a deterioration in fairness and equity, as well as lost health throughout the system. And so, it would be plausible to think that the focus on innovation has led to an overall improvement only if a perspective is adopted that goes beyond that is best for meeting the health needs of patients.

3 Health Technology Assessments

NICE's aims at a high level for its health technology assessments (HTA) follow the three aims laid down in statute that were cited at the beginning of this chapter. The HTA process takes for granted that resources are limited, and not all the care that would provide health benefits will be affordable. Prioritisation is thus required, and the

process of HTA allows recommendations to be determined about whether a technology should be made available within the NHS. The chosen mechanism for doing this is to assess the incremental cost-effectiveness ratio (ICER) of interventions, with the intention that this, in association with a set of baselines and modifiers (which are different between a general process, and one for very rare diseases) can be used to determine whether to recommend an intervention.

HTA can be done piecemeal, or systematically. If it is done systematically, it allows for a ranking of all interventions, and thus a value threshold for a particular budget. Culyer (2016) provides a helpful simple model for this:

1. Work out how cost-effective each intervention is in £ per QALY.
2. Order all the interventions in order of cost-effectiveness. (Culyer asks us to visualise this as ordering books from left-to-right on a bookshelf in order of height; where height would represent cost-effectiveness).
3. In funding interventions, start by funding the most cost-effective, and keep moving to the right of the shelf until the money runs out.
4. If the budget limit changes, the cost-effectiveness threshold for which interventions can be afforded also changes — rising as the budget increases, and falling as the budget decreases.

Culyer's model presupposes that it is *health gain* as measured via QALYs that a health system should be aiming to maximise. As we have already discussed, both NICE's duties under statute (and broader concerns from ethical theory) push towards a more complex account of what counts as an improvement, and what counts as a worsening within the health and social care system as a whole.

Culyer's general point still stands if we relax the assumption that a systematic approach to HTA should aim at the maximisation of QALYs. When looked at more broadly, a systematic approach to HTA should allow for an ordinal ranking of technologies such that, by use of this ordinal ranking, it is possible to order technologies from better to worse in relationship to the criteria by which 'better' is defined for the particular health and care system. If such an ordering is performed, and technologies are funded in order of precedence according to this ranking function, then allocating the budget in this way will be guaranteed to generate at least as much value as any alternative method (at least

in so far as the value generated by each intervention is determined atomically rather than relationally or holistically[2]).

Whatever account of value is selected for a systematic approach to HTA, the same model needs to be applied *both* to new interventions *and* to any existing interventions that would be displaced by the new interventions. For example, if it is appropriate to allow severity to count as a modifier for the ICERs for new interventions assessed (thus allowing recommendations of interventions with higher ICERs where the condition ameliorated is severe), then any such modifier would also need to be applied in decisions about which interventions to allow to be displaced, even if the decision is only to allow implicit rationing by queuing (Wilson, 2023). Unless decisions about what to defund are also driven by the same process, then particularly where recommendations are made to fund interventions that are well above the mean cost of creating one QALY in the health system, such decisions are likely to displace more value than they create.

HTA is piecemeal, rather than systematic, if some but not all interventions are assessed. Where HTA is piecemeal, it does not allow for a complete ordinal ranking. As a result, on a piecemeal approach, it may not be immediately obvious whether, when new interventions are introduced as a result of the HTA process, and others displaced, this will in fact lead to an improvement rather than a worsening given the system's own stated goals.

The health and social care system of which NICE is part adopts a piecemeal rather than a systematic model of HTA. Some but not all novel interventions are assessed, and there are many interventions that have been in long use that have not received an HTA. Even those interventions that are assessed are not assessed in a way that would allow for a clear ordinal ranking between them: the HTA process delivers a result of either a recommended, optimised (recommended for a smaller group), only in research, or not

---

[2] I argue elsewhere that there are many circumstances in which the value generated by one technology or intervention is determined relationally or holistically rather than atomically. See Wilson (2023) for more on this point. However, I leave this argument on one side for the purposes of this chapter.

recommended — and so it is better to see it as supporting a satisficing rather than a maximising approach to value (Rumbold *et al.*, 2017).

While the approach taken makes an ordinal ranking of interventions impossible, it is possible to estimate the less statistically demanding figure of the mean cost of creating one QALY within the health system, which Claxton *et al.* (2015) calculated at a little under £13,000. If this is anything like correct, then it would obviously be the case that NICE's standard HTA process, which tends to wave through any interventions that deliver QALYs at under £20,000 per QALY, and also funds interventions at up to £50,000 per QALY where there are relevant factors such as severity, frequently allow new interventions that would be predicted to displace more health benefit than they create.

For the reasons already noted, even if such interventions tend to displace more health benefit than they create, this is not yet to say that they amount to overall worsenings rather than improvements of the health and care system. Nonetheless, given both the lack of systematic and rigorous processes for deprioritisation, and the extent of implicit rationing and rising waiting lists, there is no strong positive reason to think that the present net result of NICE HTA recommendations is an overall improvement to the health and social care system. As we shall shortly examine, Highly Specialised Technologies (HST) raise this problem even more strongly, given that they are funded at a baseline of £100,000 per incremental QALY and may reach up to £300,000 per incremental QALY.

4 The value of innovation

One way in which one might attempt to explain how NICE HTA recommendations *do* lead to improvements within the NHS even if they end up displacing more health benefits than they create is via the value of innovation. As we noted, there is a statutory duty on NICE to have regard to "the desirability of promoting innovation in the provision of health services or of social care". The fact that the duty is there implies that policymakers thought that there is *something* that would be improved if NICE had reference to this duty in conducting its affairs.

Nonetheless it is not obvious that it is conducive to overall improvement within the health and care system for HTA processes to place a premium on innovation, and thus treat the fact that an intervention is deemed 'innovative' to justify spending significantly more per incremental QALY on it than would be the case if it were not deemed to be 'innovative'.

As Charlton and Rid (2019) argue, while the Health and Social Care Act 2012 led to a significant increase in mentions of innovation within NICE's HTA processes, neither the intention of NICE's policy on innovation nor what counts as innovative, has ever been spelled out in detail. Charlton and Rid suggest that the most coherent reconstruction of NICE's policy is that an intervention is innovative only if the following three conditions apply: (a) the intervention is novel, (b) the treatment gives rise to significant or substantial health benefits (a "step-change"), and (c) that the health benefit is not adequately captured in the ICER.

Charlton and Rid's account provides a plausible reconstruction of how the policy is intended to work, but as their own analysis underlines, if they are correct about the requirements of the policy, it would appear that the policy has often been misinterpreted by both NICE committees and applicants. On their analysis, there are significant inconsistencies in the way in which each of the three requirements has been interpreted by different committees, for example, the requirement that there should be a substantial benefit was sometimes interpreted in a way that almost any intervention would meet the bar and at other times in a stricter way (Charlton and Rid, 2019, p. 223). Whether there are benefits that are not fully captured in the QALY is something that HTA committees already need to consider in all cases; and so it is unclear what this requirement adds in theory, and as Charlton and Rid examine, in practice it has led to significant inconsistencies with some committees taking improving earning capacity of patients to be relevant while other committees ruling this out (Charlton and Rid, 2019).

A more fundamental point is that, even if the policy as interpreted by Charlton and Rid were applied consistently, it would remain unclear why 'innovation' thus circumscribed should be rewarded. It is worth noting that the conception of innovation implicit within NICE's processes is inconsistent with that used within patent law. No invention is patentable unless it meets a test of novelty. NICE does not appraise generic drugs

through its HTA processes, and the vast majority of interventions that NICE examines in HTA are pharmaceuticals that are under patent, so *all* these interventions have by definition already met a threshold of novelty. Thus, 'innovative' as used by NICE picks out a subset of those interventions that for the purposes of intellectual property law have been judged to be novel.

Patents on pharmaceuticals are not without ethical challenges (Wilson, 2012a), but there are good reasons for wanting to create incentive structures to encourage innovation through intellectual property law, and thereby make drug discovery attractive for businesses to invest in. However, given that such incentives for research and development are provided by intellectual property law, it is not immediately obvious why an additional policy of incentivisation would be required, which applies only once a new product has been brought to market.

Insofar as the presence of 'innovation' as defined by NICE allows higher prices to be paid than would otherwise be permitted, and the budget is limited, the greater the innovation premium or incentive, the greater the health losses will be elsewhere in the system. Even if these incentives are successful at bringing into use some innovative treatments that would otherwise have been excluded on grounds of poor value for money, and thereby benefit patients who can access the treatments, there is an obvious opportunity cost to other patients. Patients will not, on average, be benefited by an innovation premium, even if some are.

So the challenge in providing an ethical rationale for the innovation premium is to explain either why initial appearances are deceptive, and the innovation premium is in fact on average beneficial for patients, or why despite not being on average beneficial for patients the innovation premium nonetheless would be expected to bring an overall improvement.

One obvious defence of the innovation premium is that to argue that it is only over the short-term that it appears to be on average harmful to patients. It could be argued to be a loss-leader: we pay now in order to make it possible for innovation to continue in the future and also because the innovations that are expensive now will become cheaper as they go off patent and become available as generics. On this line of thinking, paying more now is an investment in cheaper and more effective drugs in the future.

This is the general economic justification given for IP law, but it works much better at a global level than it does as an argument about how England should spend its health budget. England is a relatively small percentage of the global market for pharmaceuticals (around 3%); and so it is clearly not the case that pharmaceutical innovation would be slowed significantly, were NICE *not* to incorporate an innovation premium in its HTA processes. If the NHS relied heavily on other health systems subsiding innovation, there might potentially be an argument to be made about the ethics of free-riding (Wilson, 2012b), though it is notable that such an argument has not been advanced as the rationale for NICE's focus on innovation. As a result, in so far as an innovation premium benefits patients in England, it may be best to think of it not as providing incentives for innovation per se, but rather incentives for manufacturers to license drugs in the UK and to engage with the NICE HTA process. The risk that the innovation premium attempts to forestall may not be that drug discovery will not take place at all, but that manufacturers (particularly in a post-Brexit world) may not take the UK market to be worth engaging with.

It is far from clear that adopting an innovation premium is, in fact, on average beneficial for users of the NHS in the short or medium term. Considered from the perspective of benefiting patients, it would be perverse to have an innovation premium unless there is reason to think that the interventions funded under this premium will in the medium term tend to become available at a price and value point that will not displace more value than they create. If the prices of interventions that were initially funded as 'innovative' remain high over the medium term, then adopting such interventions makes it harder for a health system to achieve its goals.

While arguments for thinking that an innovation premium will lead to improved average health are unconvincing, the loss of health gains that could be obtained more cheaply through adopting other priorities does not yet entail that the innovation premium makes the health and care system as a whole worse. There may be features of specific patient groups or specific healthcare needs that justify very high prices, even though this has opportunity costs for others. For example, it could be argued that innovative technologies deliver especially valuable health benefits. Or it could be argued that the benefits brought by innovative technologies are either not properly captured by NICE's standard HTA processes, or difficult to measure.

These kinds of arguments are most plausible in the case of interventions necessary to treat serious but very rare conditions, such as those funded under NICE's Highly Specialised Technologies (HST) programme. However, as the next section explores, the factors that make it ethically plausible to think that a health system should be willing to pay significantly more per QALY in such cases are not closely related to the level of innovation that an intervention shows.

5 The Highly Specialised Technologies Programme

As NICE makes clear in the Methods Manual, the HST programme is "designed to be used in exceptional circumstances", namely very rare diseases that by their nature have small numbers of patients, and in addition have "limited or no treatment options", and "challenges for research and difficulties with collecting evidence, because of the uniqueness of the disease" (NICE, 2022, sec. 7.1.3). The Manual explains that the HST programme is required by fairness and equity: it aims to "secure fairer and more equitable treatment access for very small populations with very rare diseases", and to "recognise that an approach that maximises health gain for the NHS may not always be acceptable: it could deliver results that are not equitable" (NICE, 2022, sec. 7.1.4).

NICE's standard HTA process now (post 2022) is formalised in a way that the baseline ICER is £20,000 per QALY, which with the severity weighting can go out to an ICER of £50,000. The HST process has a baseline ICER of £100,000 per QALY, with the ability to go up to £300,000 per QALY. The Manual is frank that each time a HST is recommended, "the NHS must commit to allocate resources that would have otherwise been used on activities that would be expected to generate greater health benefits" (NICE, 2022, sec. 7.1.5). It talks of the need to "strike a balance between the desirability of supporting access to treatments for very rare diseases against the inevitable reduction in overall health gain across the NHS that this will cause" (NICE, 2022, sec. 7.1.5).

One means by which HST aims to strike the appropriate balance is to draw the eligibility criteria in a deliberately fairly narrow way: the processes "intentionally do not seek to capture every case when there are challenges in generating an evidence base or when there is a small population with a rare disease" (NICE, 2022, sec. 7.1.5). Another is that HST processes are used only when a disease is *very* rare: orphan drug designations are

available from the MHRA and the European Medicines Agency when the prevalence of a disease is no more than 5 in 10,000 people, but the NICE HST process only applies to diseases that have a prevalence that is twenty-five times smaller, at less than 1 in 50,000 people.

The HST programme thus clearly raises ethical questions, which are alluded to with the references to equity and the discussion of opportunity costs in the Manual, but these ethical questions are not faced head-on. This is odd, given that similar questions *were* taken up by the NICE Citizens' Council when it examined ultra orphan drugs in 2004 and the rule of rescue in 2006, and as the next section examines, the Citizens Council's conclusions are not easy to reconcile with NICE's current approach.

The costs of drug development provide obvious reasons why drugs for orphan diseases tend to have higher ICERs than drugs for common diseases, and why drugs for ultra orphan diseases tend to have even higher ICERs. But it is nonetheless surprising from an ethical perspective to adopt the approach that NICE has by segmenting HTA into two separate processes with wildly different baseline ICER requirements, especially as the eligibility requirements for HST only apply when the prevalence of a disease is 25 times lower than the threshold for an orphan drug designation. One implication of NICE's approach is that there are many diseases that are rare enough that ICERs for novel interventions are likely to be significantly higher than for more common diseases, but which will not be nearly rare enough to receive an HST designation. Such cases of disease that are rare, but not rare enough for HST, include cystic fibrosis and Fragile X syndrome.

Clarke, Ellis and Brownrigg (2021) found that orphan drugs assessed via NICE's standard Single Technology Assessment (STA) seemed to suffer disadvantages both relative to non-orphan drugs, and relative to those assessed via HST. As compared to non-orphan drugs, the assessment of orphan drugs took significantly longer within the STA process; and while the chance of receiving a positive recommendation was comparable between orphan and non-orphan drugs, the orphan drugs were meeting an unmet healthcare need. Relative to HST, chances of receiving a positive recommendation were lower and decisions took longer.

In short, NICE's current processes end up disadvantaging those with rare, but not very rare, diseases. And so, even if it is true that it would be unfair not to make special provision for very rare diseases, the way that this has been done risks creating other unfairnesses for groups of patients who have rare, but not very rare, diseases. We will now consider the question of the ethical status of rarity per se: does the rarity of a disease sometimes provide good reasons to adopt a different approach to HTA, in particular by shifting baseline ICERs?

5.1 Rarity and severity

If the rarity of disease is separated from all other features, and considered on its own, it is far from obvious that it matters ethically. McCabe (2005) asks us to imagine a case where all features are held the same between two diseases, except that one is ten times rarer and ten times less cost-effective to treat than the other: those with the rarer condition "have the same personal characteristics, the same prognosis without treatment, and the same capacity to benefit from the treatments" (McCabe, 2005, p. 1018). If we are right to think that cost-effectiveness matters in the usual case, then other things being equal, we should prefer a health gain that is in all respects similar but ten times cheaper to one that is ten times as expensive. Indeed, if all else is equal, we should rank the one much more highly than the other, as the opportunity cost of treating one individual with the more expensive condition would be that ten individuals with the same prognosis and capacity to benefit would go untreated. Assuming the budget is limited, it might well be the case that the first should be funded for many patients, while the second will lose out and not be funded at all.

Thus, the ethical challenge posed by special provisions for very rare diseases is to explain why it is ethically better if a greater number of individuals are denied treatment in order to allow treatment for a smaller number with rarer conditions, even though the prognosis and capacity to benefit of individuals in both groups is the same. When put in such stark terms, it becomes clear that absent such a justification, paying a much higher ICER for very rare conditions is *unethical*.

When the NICE Citizens Council examined ultra orphan drugs in 2004, more than half thought that much higher ICERs should be paid for rare conditions, but those who did favour special arrangements tended to be sceptical that it was rarity itself that made the

difference: “people stressed that it was because these rare diseases are often so severe that it is justified. On its own, the rarity of a disease, would not give it special status” (NICE Citizens Council, 2004, p. 7).

NICE’s standard HTA processes and HST both provide approaches (albeit different approaches) to responding to severity. The HST programme responds to the value of severity, in so far as the HST pathway is defined in such a way that technologies can be routed to it only if the condition that will be treated is severe. While there will be non-severe conditions that could be treated for very rare diseases, such interventions are out of scope for HST. NICE’s standard HTA processes also include severity as a value: in the 2022 Methods Manual, severity is a value that justifies a modifier of up to 1.7x on the baseline cost-effectiveness ratios. The presence of severity as a value in both NICE’s standard HTA and HST processes works to undermine the claim the HST programme is justified *because* rare diseases are severe.

If we step back a bit from NICE’s actual processes, it is clear that the severity of a condition is orthogonal to its rarity. A condition can be common but very severe, as well as very rare but mild. There is no reason to think that the suffering caused by very rare conditions is per se *more severe* than for more common conditions, and so it is not obvious how giving due weight to the ethical importance of severity would justify different ICERs for rare conditions. There is thus no reason to think that severity provides an additional reason for giving extra weight to HST, over and above the ethical weight that would be due to severity in other cases. In so far as severity matters by itself, it is an argument to fund treatments for very severe conditions at a premium regardless of their rarity, rather than a consideration that would apply only to very rare conditions. So, even if severity matters intrinsically, this does not by itself provide an ethical justification for special processes for very rare diseases.

5.2 Bad luck and nonabandonment

While the fact that an ailment is rare does not per se make it more agonising or more deadly than one that is common, nonetheless there is a non-accidental relationship between the rarity of a disease and the low cost-effectiveness of the drugs needed to treat it. Even if, other things being equal, cheap health benefits should be ranked more highly than less cost-effective health benefits, and even if other things being equal, the

fact that a condition is rare is not intrinsically ethically significant, we might still think that the combination of the two creates an ethical claim: if there is a non-accidental correlation between treatments that are available for rare conditions and expensive treatments.

There may be cases where the intersection of two factors has unfair implications, even if by themselves neither have this implication separately. This is sometimes explained via the idea of nonabandonment — the idea that everyone deserves or has a strong claim to some kind of effective treatment for their condition, even if it would be very expensive to provide it (Gericke, Busse and Riesberg, 2005). It might be argued that giving a significant weight to nonabandonment is a plausible way of interpreting the idea of treating citizens as equals in this context.

Luck egalitarians claim that it is unfair if someone is worse off than others through no fault of their own. It looks to be a matter of bad luck that any particular individual (a) develops a fatal or debilitating disease that they had no ability to affect, and it is a further piece of bad luck that (b) in virtue of its rarity, their disease is highly likely to be much more expensive to treat than a common disease. This might provide an argument for thinking that something like HST should be put into operation in order to counteract this bad luck. However, the idea of luck is itself somewhat nebulous and subject to different interpretations (Lippert-Rasmussen, 2023). While developing a fatal or debilitating disease that one had no ability to affect is often given as a paradigm case of bad luck, it is not immediately clear whether the badness of luck should be graded such that *very* bad luck creates a stronger claim to aid than ordinarily bad luck (and if so, by how much).

In addition, what counts as bad luck in the context of HTA decisions about which particular interventions to include within the scope of public funding is not easily separable from the policy decisions that are taken. As was discussed earlier, the NICE HST programme employs has a cut-off of 1 in 50,000. Given this decision, it would then seem to be a matter of bad lack for those individuals with a condition that is rare enough that is too expensive to be funded under the normal HTA rules, but not rare enough to benefit from the HST process. So while it might be the case that some special

provision for rare diseases would follow from luck egalitarianism, it is not obvious that NICE's specific approach would be the best way of doing so.

6 Conclusion

NICE began with a remit to reduce postcode prescribing, bring a greater degree of cost control, and take specific decisions about which drugs to fund out of the hands of politicians. In its early years, NICE was able to stand up for the importance of cost-effectiveness analysis and consideration of opportunity costs in making decisions about when interventions should be recommended for use in the NHS, even while doing so would have been tricky for politicians. Both NICE and the broader NHS context have changed significantly since 2010. NHS budgets have failed to keep pace with rising healthcare needs, but during the same time period, NICE has de-emphasised the opportunity costs of recommending cost-ineffective interventions. The combined effect is that NICE recommendations are increasingly likely to lead to the displacement of more health benefits than they create, for example, leading to cuts to other services or degradation of service quality.

From an ethical perspective, the important question is whether there are reasons to think that the decreased salience given to opportunity costs has brought net ethical benefits despite being on average negative for population health. This chapter analysed NICE's approach to innovation, and to highly specialised technologies, arguing that even a sympathetic reconstruction of NICE's approach fails to uncover a convincing ethical justification for the policies currently adopted. While it is plausible to think that sometimes the benefits of a health technology are not well captured by the QALY, and that a health system should be willing to pay more per QALY for interventions that treat severe conditions, neither consideration is specific to innovative technologies. As Victoria Charlton examines in Chapter [X], the increased emphasis on innovation is most plausibly explained by a perceived need on NICE's part to align with the political priorities of the government's industrial strategy.

HST provides the most radical departure from a maximising approach, and here also there is much that is puzzling, or ethically questionable about NICE's policy — from the eligibility requirements to the conception of fairness it draws upon. In both cases, what

is striking is that while the processes as set out in the Methods Manual and realised in committee review have become ever more well articulated and rigorous, the ethical framework within which these processes unfold has become sketchier. NICE itself seems to have retreated from the project of debating and publicly justifying the core values at the heart of its approach, just as those values have become more ethically questionable. It is much less clear than it was 25 years ago that NICE's health technology assessments lead to an overall improvement rather than worsening of the service that the NHS as a whole can provide.

References

BMA (2023) *Pressures in general practice data analysis*. Available at: https://www.bma.org.uk/advice-and-support/nhs-delivery-and-workforce/pressures/pressures-in-general-practice-data-analysis (Accessed: 10 May 2023).

Charlton, V., Lomas, J. and Mitchell, P. (2022) 'NICE's new methods: putting innovation first, but at what cost?', *BMJ*, 379, p. e071974. Available at: https://doi.org/10.1136/bmj-2022-071974.

Charlton, V. and Rid, A. (2019) 'Innovation as a value in healthcare priority-setting: the UK experience', *Social Justice Research*, 32(2), pp. 208–238. Available at: https://doi.org/10.1007/s11211-019-00333-9.

Clarke, S., Ellis, M. and Brownrigg, J. (2021) 'The impact of rarity in NICE's health technology appraisals', *Orphanet Journal of Rare Diseases*, 16(1), p. 218. Available at: https://doi.org/10.1186/s13023-021-01845-x.

Claxton, K. *et al.* (2015) 'Methods for the estimation of the National Institute for Health and Care Excellence cost-effectiveness threshold.', *Health technology assessment*, 19(14), p. 1.

Culyer, A.J. (2016) 'Cost-effectiveness thresholds in health care: a bookshelf guide to their meaning and use', *Health Economics, Policy and Law*, 11(4), pp. 415–432.

Gericke, C.A., Busse, R. and Riesberg, A. (2005) 'Ethical issues in funding orphan drug research and development', *J Med Ethics*, 31(3), pp. 164–168. Available at: https://doi.org/10.1136/jme.2003.007138.

Lippert-Rasmussen, K. (2023) 'Justice and Bad Luck', in E.N. Zalta and U. Nodelman (eds) *The Stanford Encyclopedia of Philosophy*. Spring 2023. Metaphysics Research Lab, Stanford University. Available at: https://plato.stanford.edu/archives/spr2023/entries/justice-bad-luck/ (Accessed: 8 May 2023).

McCabe, C. (2005) 'Orphan drugs and the NHS: should we value rarity?', *BMJ*, 331(7523), pp. 1016–1019. Available at: https://doi.org/10.1136/bmj.331.7523.1016.

Morris, J. and Reed, S. (2022) *How much is Covid-19 to blame for growing NHS waiting times?*, *Nuffield Trust*. Available at: https://www.nuffieldtrust.org.uk/resource/how-much-is-covid-19-to-blame-for-growing-nhs-waiting-times (Accessed: 22 April 2023).

NICE (2022) *NICE health technology evaluation topic selection: the manual*. NICE. Available at: https://www.nice.org.uk/process/pmg37/chapter/about-this-guide (Accessed: 12 April 2023).

NICE Citizens Council (2004) *Ultra Orphan Drugs*. London: National Institute for Health and Care Excellence (NICE) (NICE Citizens Council Reports). Available at: http://www.ncbi.nlm.nih.gov/books/NBK401721/ (Accessed: 1 May 2023).

Rumbold, B. *et al.* (2017) 'Public Reasoning and Health-Care Priority Setting: The Case of NICE', *Kennedy Institute of Ethics Journal*, 27(1), pp. 107–134. Available at: https://doi.org/10.1353/ken.2017.0005.

Schlepper, L. and Dodsworth, E. (2023) *The decline of publicly funded social care for older adults*, *Nuffield Trust and Health Foundation*. Available at: https://www.nuffieldtrust.org.uk/resource/the-decline-of-publicly-funded-social-care-for-older-adults (Accessed: 22 April 2023).

Wilson, J. (2012a) 'On the value of the intellectual commons', in A. Lever (ed.) *New frontiers in the philosophy of intellectual property*. Cambridge, UK: Cambridge University Press.

Wilson, J. (2012b) 'Paying for patented drugs is hard to justify: An argument about time discounting and medical need', *Journal of Applied Philosophy*, 29(3), pp. 186–199. Available at: https://doi.org/10.1111/j.1468-5930.2012.00567.x.

Wilson, J. (2021) *Philosophy for Public Health and Public Policy: Beyond the Neglectful State*. Oxford, New York: Oxford University Press.

Wilson, J. (2022) 'When does precision matter? Personalized medicine from the perspective of public health', in Y.M. Barilan, M. Brusa, and A. Ciechanover (eds) *Can precision medicine be personal; Can personalized medicine be precise?* Oxford University Press, pp. 173–186. Available at: https://doi.org/10.1093/oso/9780198863465.003.0013.

Wilson, J. (2023) 'What makes a health system good? From cost-effectiveness analysis to ethical optimisation in health systems', *Medicine, Health Care and Philosophy* 26, pp. 351–365.