13

When does precision matter? Personalized medicine from the perspective of public health

James Wilson

Introduction

Much recent health policy in the UK, USA, and elsewhere presupposes a tech-led vision of the future of medicine—as exemplified by the US NIH Precision Medicine Initiative, and the UK government's AI Early Diagnosis Mission.[1] This new approach aims to use large-scale datasets and rich information about individuals (particularly genomic or multi-omics datasets) to detect abnormalities that would have caused clinically relevant cases of disease, and enable early intervention. The underlying model of medicine is thus anticipatory rather than reactive: early intervention will ensure that a disease never eventuates, or that its effects are less severe. A lot of the excitement has coalesced around the model of P4 medicine: medicine which is predictive, preventive, personalized, and participatory (Hood and Friend, 2011; Vogt et al., 2016).

P4's focus on prevention is used to position it as superior when compared to previous models of clinical care.[2] However, the idea of prioritizing prevention is not new. Public health in particular has long focused on preventing disease and illness through interventions such as childhood vaccination or health and safety legislation. Within public health, primary, secondary, and

[1] For further on these policies, see Collins and Varmus (2015) and Her Majesty's Government (2019).

[2] See, for example, the following indicative passage: 'Surgery and other aspects of traditional medicine will be informed by the stratification of disease and of people where relevant. Focusing on the causes rather than the symptoms of disease will enable intervention to occur much earlier in the disease process, in many cases preventing disease from occurring in the first place.' (Flores et al., 2013, p. 567).

James Wilson, *When does precision matter? Personalized medicine from the perspective of public health*
In: *Can precision medicine be personal; Can personalized medicine be precise?*. Edited by: Yechiel Michael Barilan, Margherita Brusa, and Aaron Ciechanover, Oxford University Press. © Oxford University Press 2022.
DOI: 10.1093/oso/9780198863465.003.0013

tertiary prevention are standardly distinguished. Primary prevention aims to reduce the likelihood of disease or injury before it occurs (for example, anti-smoking or road safety campaigns). Secondary prevention aims to diagnose disease early in order to allow interventions that will minimize its effects (for example, cancer screening; training of employees to spot the signs of workplace stress). Tertiary prevention aims to minimize the effects of disease or injury that is already severe enough to have made a noticeable difference to the patient's life (e.g. rehabilitation programmes after a stroke; antiretroviral drugs after an HIV diagnosis).

Thus, what is distinctive in P4 medicine is not the idea of prevention per se, but highly targeted secondary prevention through the use of rich and large-scale datasets. Looking at P4 medicine within a broader account of prevention raises a question about whether the approach to prevention embodied in P4 medicine should be preferred to the well-established approaches to prevention embodied in classic public health. There is something contentious about P4 medicine's framing of secondary prevention as *early intervention*, where the contrastive 'later' intervention would be tertiary prevention. Secondary prevention counts as 'early' relative to tertiary prevention, but *late* relative to primary prevention.

If all that mattered was intervening as early as possible, then the 'early intervention' rationale for shifting from conventional medicine to P4 medicine would, by the same token, be an argument for shifting from P4 medicine to public health interventions that target primary prevention. What this shows is that if P4 medicine amounts to a superior model of medicine, it cannot be solely because it advocates acting early within a disease trajectory. It must be because it aims to do so in an appropriately *targeted* way.

The value of more precise targeting of interventions is usually taken for granted, rather than argued for explicitly; and it is assumed that the more tightly targeted an intervention, the better. The main argument of this chapter is that the value of precision targeting is more equivocal and contextual than is often thought. Precision targeting of an intervention is important only where there would be bad effects if the intervention misses its intended target—and this is far more likely to occur with some kinds of interventions than with others. Even where biomedical interventions are such that precision in their application would be valuable, the model of early intervention at an individual level on the basis of predictive modelling may struggle to give sufficiently accurate information. It is likely to lead to

overtreatment, and to be less cost-effective and less equitable than public health approaches.

The overall conclusion is that ensuring that interventions have the optimal balance of benefit over harm is a much more important goal than precision, and precision medicine will tend to be a less effective way of pursuing the more important goal than public health-led approaches. In short, contrary to what is implied by policies such as the US NIH Precision Medicine Initiative, and the UK government's AI Early Diagnosis Mission, precision medicine could easily amount to a step backwards.

Precision

The more precisely targeted an intervention is, the more information those implementing it need to have about the targets at whom it is directed. Targeting interventions thus requires: (i) data collection, (ii) modelling or stratification and (iii) selecting or prioritizing an intervention based on risk profile. Depending on how precise the targeting is, more or less data, and a more or less sophisticated stratification model will be required. At the limit where an intervention is entirely untargeted, then little or no data collection will be required for targeting purposes. The degree of high dimensional variation in the phenotypic data collected (stage 1) will often be orders of magnitude greater than the variation in the kinds of interventions that can be applied (stage 3). This is one of the reasons why, after some initial enthusiasm for the terminology of 'personalized medicine', many shifted to talking of precision or stratified medicine, as it became clear that it is unlikely—at least in the short to medium term—that clinicians will be able to deliver on the promise of fully individualized therapies (Juengst et al., 2016).

For illustrative purposes, I shall use two formulations to contrast with precision targeting of interventions. First, *broadly targeted*, to suggest that an intervention is targeted on particular groups but not that tightly—for example where a diagnostic assay is run on a sample from a bacterial infection, and a suitable antibiotic chosen. Second, *untargeted* to suggest that an intervention is made available with minimal targeting at individuals—for example, a childhood vaccination campaign, in which it is recommended that all children receive the vaccine except the small number for whom this is biomedically contraindicated. Whereas the three categories of precisely targeted, broadly targeted, and untargeted will be useful for the purposes of

exposition, it is important to be aware that there is a continuum from completely untargeted interventions to fully individualized interventions.[3]

Why think that precisely targeted interventions should be preferred to broadly targeted, or untargeted interventions? The value of greater precision is usually taken for granted, rather than argued for explicitly by proponents of P4 medicine. Where we do see arguments, they centre around *arguments from non-maleficence*, namely that untargeted interventions are harmful—or at least sub-optimal in the amount of benefit they provide—precisely because they are untargeted, and *arguments from waste*, namely that untargeted interventions are wasteful, just because they are applied to people who will not benefit from them.

Precision medicine is argued to be less harmful than untargeted approaches, because it will reduce 'spillover' effects on those individuals who would not have been selected to receive the intervention if it were very precisely targeted. Given the importance of the duty of non-maleficence, this might be taken to be a strong argument in favour of the importance of precision targeting of interventions. However, as the next sections explore, things are not so straightforward. First, it is wrong in theory to think that effective and ethically justifiable preventive activity requires precise targeting. Second, as currently constituted, P4 medicine is likely to generate models and interventions that are more precise than they are accurate—leading to overdiagnosis and overtreatment, and this itself raises questions about non-maleficence.

Are precision interventions less likely to be harmful?

P4 aims for a medicine that is predictive, preventive, personalized, and participatory. This might seem to suggest that there is an inherent connection between being personalized and being effectively preventative, i.e. that effective prevention is individually targeted. However, there is no intrinsic connection between prevention and personalization. While it is easy to think

[3] The distinction between precisely targeted, broadly targeted, and untargeted interventions describes the way that individual level data is used to tailor interventions, rather than the scale or richness of the datasets used to inform the intervention. It may be useful to make use of large and rich datasets even in broadly targeted or untargeted interventions. For example, post-marketing surveillance for pharmaceuticals may make use of rich data derived from a large number of individual patients. Insofar as this data is used to decide whether a drug should be withdrawn on safety grounds (or its indication narrowed), the resulting intervention would nonetheless be a broadly targeted or untargeted one.

of examples of health interventions that would be very harmful if they were rolled out in an untargeted fashion—such as chemotherapy, or brain surgery—there are also a range of health-related interventions where targeting is unimportant because the intervention would be either beneficial or at worst have no effect if it were made the subject of a blanket recommendation.

Alongside the rise of 'big data' early intervention approaches, recent years have also seen an increased emphasis on social prescribing for long-term conditions. That is, clinicians referring patients for health-system-recommended non-medical activities and support services, which might include exercise classes, weight management advice, visiting museums, joining a choir, or volunteering in a community gardening project (Drinkwater et al., 2019). Social prescribing embodies a very different approach to early intervention than is indicated by P4 medicine. Many of the interventions that fall under the category of social prescribing are non-specific. Clinicians could fruitfully suggest them to patients with a wide variety of conditions. Many of these are activities that individuals might want to do anyway regardless of whether they were ill—and so there seems to be little risk to patients in using such interventions in a broadly targeted or even an untargeted way. Some pharmaceutical interventions, such as statins, can also be prescribed safely in a broadly targeted way.

Knowledge dissemination and infectious disease control provide other cases of untargeted interventions that are widely thought to be good practice. Knowledge such as healthy eating advice is a non-rival good; if such information is provided freely to everyone, then many will be able to benefit without anyone losing anything. It is a feature rather than a bug of non-rival goods that they provide benefits to many beyond those at whom they were originally targeted (Wilson, 2012).

Goods such as herd immunity are network goods: the more who are vaccinated, the more beneficial it is for everyone. The greater the proportion who are vaccinated, the greater the overall benefit (and especially to those who would otherwise be at high risk). So in general it is a good rather than bad thing if a lot of people who would themselves be at low risk are vaccinated.[4] Careful targeting is ethically required only where activity undertaken for preventive purposes would create disproportionate risks of harm if undertaken

[4] Obviously, vaccine campaigns require at least a minimal form of targeting, to ensure those for whom they would be medically inappropriate, such as the immunocompromised, are not vaccinated.

in an indiscriminate way. Where an intervention's spillover effects are beneficial rather than harmful, then this is a good reason *not* to target it.

What implications do these thoughts have for thinking about precision medicine? One objection would be that there may be relatively few cases in clinical medicine where untargeted interventions are likely to be neutral or beneficial. Within clinical medicine it would be generally thought to count as a breach of nonmaleficence to provide an intervention which had no effect, where an effective intervention is available.[5] There is something to this worry—but addressing it requires us to think more deeply about where and how in disease trajectories it is best to intervene, and how health systems should prioritize between different modalities of prevention.

Are untargeted interventions wasteful?

Untargeted interventions by their nature are undiscriminating about whom they affect, and so it is easy to assume that precisely targeted interventions will be more efficient. However, this does not follow. What counts as a wasteful (or an efficient) use of a resource should be defined relative to the overall goals of the health system. If the most efficient way of pursuing the goals of the health system involves applying an untargeted intervention, then such an intervention is by definition not wasteful.

It is true that there is a sense in which the intervention is 'wasted' if it is applied to many who do not benefit from it. However, this sense of waste would be relevant for policy decision-making only if the main goal of the health system were to avoid applying interventions unless they could be shown to be biomedically beneficial for each individual to whom they are applied. Assuming without argument that this is a good goal for health systems would be question-begging.

Health system goals such as maximizing population health, and reducing health inequalities are often better served by untargeted interventions than by targeted interventions. (For a classic explanation of why, see Rose, 1981.) The World Health Organization-recommended standard of care for various neglected tropical diseases (NTDs), including schistosomiasis, is mass drug

[5] Hence the use of placebos in clinical trials has often been thought controversial, and it has been argued that there is an ethical requirement on doctors not to enrol patients in clinical trials unless they think that the trial intervention is equally as good as the standard intervention. (See Freedman (1987) for this view, and Miller and Brody (2003) for a critique.)

administration (MDA). MDA is explicitly untargeted, and has been defined as 'delivering safe and inexpensive essential medicines based on the principles of preventive chemotherapy, where populations or sub-populations are offered treatment without individual diagnosis' (Webster et al., 2014). MDA programmes can be *much* more cost-effective than testing individuals and providing treatment only where an individual tests positive. Indeed, Webster et al. (2014) argue that an integrated MDA programme targeting several different infections can be delivered for less than US$0.50 per person per year. Far from being wasteful, MDA is much more efficient than a targeted approach would be.

It is thus much too hasty to make the blanket assumption that precise interventions provide a better use of resources than untargeted ones. Whether precisely targeted interventions use available resources more efficiently than untargeted ones depends on what the goals of a health system are, alongside contextual features such as the cost of case-finding versus the cost of the intervention itself, and the cost per intervention of the untargeted as opposed to the targeted intervention.

These points have not always been well recognized in the precision medicine literature, where it has too often been presupposed that precision interventions will be more cost-effective than broadly targeted interventions—even where precision medicine involves bespoke pharmaceutical interventions being applied to the targeted individuals.[6] In reality, bespoke pharmaceutical interventions are unlikely to be anywhere near as cost-effective as loosely targeted pharmaceutical interventions using existing pharmaceuticals of much wider indication (Cossu et al., 2018). This is because of an obvious corollary of stratification and precision targeting of pharmaceuticals: the size of the target market will be significantly reduced—a shift that has sometimes been described as one from blockbusters to niche busters (Dolgin, 2010). Given that the size of target markets will be very significantly reduced, while the costs of drug development will not, the stratification and precision targeting of patented treatments would be expected to significantly raise drug prices (Gronde et al., 2017).[7]

[6] For example, Flores et al. (2013, p. 567) write 'Systems medicine will make disease care radically more cost effective by facilitating the stratification of both people and disease into distinct subgroups. Genomic analysis stratifies people into subgroups with different reactions to drugs, different disease risks and other clinically relevant factors . . . . These stratifications are providing increasingly more accurate diagnoses and cost-effective interventions based on the underlying causes of disease.'

[7] Where health systems take cost-effectiveness analysis seriously—as in the English National Health Service (NHS)—it is clear that introducing new drug therapies frequently leads to a net reduction of health gains within the health system. In a widely cited study, Claxton et al. (2015)

Secondary prevention, overtreatment, and overall benefit

Precision implies fine calibration and reproducibility, but it need not imply accuracy (Tal, 2017, section 8.3). A model that is the basis for interventions can be precise to an arbitrarily high degree, but yet be deeply unhelpful, if the model fails to provide a reliable basis for making interventions that will improve patients' lives. As Paul Kilanthi argued, there is good reason to think that it may be 'irresponsible . . . to be more precise than you can be accurate' (Kalanithi, 2014).

An ideal scenario would be to base interventions on measurements that are both precise and accurate, but this combination is difficult to achieve in the domain of clinical medicine. Whenever an intervention is targeted, there needs to be a test that matches individuals to interventions (or determines that no intervention is indicated). Tuning a test to minimize the chance of a type II error (a false negative), will tend to increase the chance of a type I error (a false positive), and vice versa. Thus, increasing the sensitivity of a test will typically also increase the chances of detecting cases that would not have led to clinically relevant disease—whether because they are false positives (misdiagnosis) or cases in which what is detected is genuinely incipient disease, but where this would not have gone on to cause a clinical problem (overdiagnosis).

Overdiagnosis can occur for a variety of reasons including that the disease would self-resolve, that it would have remained confined to a particular area (e.g. ductal carcinoma *in situ*), or that it is sufficiently slow progressing that the person would have died from another cause before it becomes problematic (Brodersen et al., 2014). Precision medicine, in so far as it aims to support targeted early intervention, is likely to make the problem of overdiagnosis much worse:

> monitoring many features of the body with highly sensitive technologies is bound to detect many abnormalities but without the ability to tell which,

calculated that the most likely value for the cost it takes to create one quality-adjusted life year (QALY) within the NHS is £12,936. However, when the National Institute for Health and Care Excellence (NICE) appraises new drugs to see if they should be recommended for use in the NHS, it rarely rejects pharmaceuticals that are priced at under £30,000 per QALY. The rather uncomfortable implication is that most of the drugs approved by NICE would be expected to displace larger health benefits than they create. Precision pharmaceuticals would, for the reasons given, be expected to be even more likely to lead to net health losses within the system than new non-precision interventions.

> if any, will become clinically manifest. As a result, more people may be labelled with more harmless conditions. (Vogt et al., 2019, p. 2)

In many cases the underlying processes of disease progression are stochastic. For example, it may be indeterminate whether a given very small cancer would have remained insignificant or if it would have become virulent if left untreated (Wu, 2020). It is thus difficult, if not impossible, even to determine whether overdiagnosis has occurred at an individual level. The extent of overdiagnosis can only really be explained and quantified at a population level (Hofmann, 2018).

P4 medicine, in other words, has a profile of risks that is much more similar to traditional secondary prevention policies such as breast cancer screening than its advocates sometimes realize. In both cases, the intervention targets individuals who do not consider themselves unwell. In both cases, despite the best of intentions, the policy may end up intervening in ways that do not make individuals' lives better, and crucially, it may not be straightforward for individuals to determine whether it is in their interests to accept the intervention. Due to this profile of risks, policies such as breast cancer screening have become much more controversial in recent years (Marmot et al., 2013).

In the case of cancer screening, the number needed to screen in order to save one life is in the thousands, so 'saving' one identifiable individual simultaneously means that there are thousands of individuals who are not benefited and who may be made worse off by the screening as a result of misdiagnosis or overdiagnosis. Those whose early stage cancer is detected by screening and are successfully treated may feel that the screening has saved their lives, but no one is really in a position to say this, as in each case the treatment could have been overtreatment. So, on closer analysis, even the common assumption that screening *does* benefit identified individuals is contestable (Wu, 2020). The same worry would seem to apply to early intervention as promoted by P4 medicine.

Overtreatment occurs where an intervention would be expected to make the patient's life worse rather than better.[8] Whether a particular intervention counts as overtreatment would depend in part on the individual's values and preferences. For example, if an individual understands the relevant facts

[8] Overdiagnosis and overtreatment are closely aligned: to the extent that a disease is detected that would not have gone on to harm the patient, it is highly likely that an intervention that is applied to prevent the progress of that disease may make the patient's life overall worse—for example if it causes pain or other side-effects, or even if it is merely inconvenient.

and risk factors, but prefers to live with a particular health risk than to have the intervention, then giving this patient the intervention regardless would amount to overtreatment, even if other patients with a similar condition would prefer to receive the intervention.

Overtreatment is a potential problem for all kinds of medicine—but is a particular challenge in cases such as chronic disease, where interventions cannot provide a cure, but rather a way of helping to manage symptoms, prevent exacerbations, and improve quality of life. Some kinds of intervention will be more susceptible to giving rise to overtreatment than others. The profile of P4 medicine, like that of screening—in which interventions are applied to hundreds or thousands in order to prevent a single case, set up these conditions perfectly. Where an intervention provides a benefit that would also be valued by an individual quite separately from whether it would be necessary for improving their health, then it is much less likely that it will amount to overtreatment. Thus, interventions such as social prescribing are much less likely to lead to overtreatment than P4 medicine's early interventions.

Prioritizing public health, P4 medicine, and social prescribing

Much of the proposed benefit of P4 medicine comes from secondary prevention, and much of the benefit of public health comes from primary prevention. This raises an important question of how health systems should prioritize: which of the different potential preventive approaches should be given most weight?

I address these questions at much greater length in Wilson (2021), but we can sketch the following points here. Precision matters only where there is a likelihood of harm from an intervention being applied to those to whom it is not intended to be applied, so precision should not be a fundamental goal of health systems. However, goals such as cost-effectiveness, improving equity, and avoiding iatrogenic harm are all of core importance in health policy.[9] Given these goals it is plausible to think that some governments may

[9] I have argued elsewhere that primary prevention work through public health activity needs to be designed in a way that is ethically justifiable *given that* it is untargeted and that spillovers will occur. For example, any iatrogenic harms need to be justifiable *ex ante* to individuals who are harmed—perhaps explained by the benefits to health equity improvement in population health, or the overall prospects for representative individuals (Wilson, 2016). A similar principle will apply to precision medicine too: good early intervention requires finding an appropriate

currently be *over-valuing* P4 medicine relative to the benefits that other preventive modalities can bring.

A number of researchers have objected to this line of argument by claiming that it should not be a matter of either public health *or* P4 medicine, but *both*. On this view, supplanting public health is not the intention of P4 medicine, and so P4 medicine practitioners can agree that primary prevention through population-level public health measures is a superior bedrock for a health system strategy to secondary prevention through P4 medicine, without threatening their own position.

In response, I agree that primary and secondary prevention are compatible and complementary in ideal scenarios. For example, it is possible to support both tobacco control for primary prevention of cancer, as well as secondary prevention of cancer through personalized risk profiling and early intervention. However, primary and secondary prevention will often be in competition in practice due to limited resources—for example, additional resources for genomic sequencing and early intervention may come at the cost of reduced resources for smoking cessation.

Wherever there is population-level primary prevention activity that could be undertaken but is not as a result of cost (as will always be the case), and P4 medicine *is* funded to a significant degree, it is difficult to deny that the health system has implicitly prioritized between the two. Thus, I am unrepentant about the need to assess and prioritize P4 medicine against other preventive interventions including primary prevention, which compete for the same resources.[10]

Primary prevention through public health measures provides a more ethically defensible bedrock for a prevention strategy than P4 medicine for various reasons. It is simply easier to do more good more cost-effectively via primary prevention. Public health primary prevention is usually aimed at significant and common health problems, focusing on the social determinants of health, or reducing pollution, or operating childhood vaccination programmes. As such interventions operate to reduce or remove generic health risks, they do not usually need to be precisely targeted, as they will act to secure public goods and network goods, and to control generic risk

balance of the potential benefits of providing a diagnosis and getting a patient who benefits into treatment, against the potential harms of doing so.

[10] Beyond the practical incompatibility, there may also be disagreements between the literatures of public health and precision medicine—for example, about the extent to which genomics is relevant for disease causation and risk reduction (Chowkwanyun et al., 2018).

factors. Moreover, whole-population policies tend to be more equitable in their effects than ones that rely heavily on individual choice.

P4 medicine's interventions are much more precisely targeted than other interventions, but this does not suffice to give them a more optimal balance of benefit and harm. As we have seen, precision in detecting early signs of disease is compatible with (and may lead to) harm through overdiagnosis and overtreatment (Vogt et al., 2019). The underlying model of massive data collection in order to intervene in individual lives through secondary prevention is unlikely to be as cost-effective, or as beneficial for the worst off, as a model that emphasizes primary prevention through public health.

The discussion of social prescribing showed that there are some broadly targeted interventions, which provide individuals with access to a range of activities that are useful for improving wellbeing and resilience. These interventions are unlikely to make lives worse and are well placed to provide support for some of those with long-term conditions. So such interventions may score highly on desiderata of equity, and of avoiding iatrogenic harm, even though their current lack of a solid evidence base may make them weaker in terms of cost-effectiveness, and not an effective substitute for either large-scale public health, or for targeted clinical interventions.

In other words, P4 medicine may compete unsuccessfully with a range of other prevention modalities including public health and social prescribing, in so far as health systems prioritize cost-effectiveness, equity, and avoidance of iatrogenic harm. This conclusion should not be understood to be a claim of necessity, but rather a challenge for whether, and, if so, how P4 medicine could rethink itself to place more emphasis on these factors. For example, the cost-effectiveness of precision targeting could be greatly improved by focusing on how best to recombine or calibrate existing pharmaceuticals or other interventions that would be used widely in any case, or to select interventions that should not be performed. In such cases, so long as case-finding is reasonably cheap, it is plausible that precision approaches could be cost-saving relative to some less targeted approach.[11] Perhaps also, there will be some way in the medium term of much better distinguishing the signal from the noise when monitoring the vitals of millions for incipient risk factors.

[11] Thus, Nelson et al. (2019) provide a useful example of how a high-dimensional machine-learning model that predicts which patients are most likely fail to attend their appointments could save costs by prioritizing patients to receive telephone appointment reminders.

Conclusion

There are strong reasons to think that the availability of massively larger linked datasets should create a step change in the way that medicine is conceived, but it is too hasty to suppose that P4 medicine in its current incarnation provides the best future for medicine. Risk factors can be cost-effectively reduced for large numbers of individuals without individual targeting; and the fact that an intervention is untargeted and improves prospects for all can be an important part of what makes it equitable. If policy-makers are interested in creating a cost-effective and equitable health system that makes the best use of limited budgets, then P4 medicine as it is currently conceived looks like a questionable means of achieving this. Moreover, given the risks of overdiagnosis inherent in current iterations of P4 medicine, interventions such as social prescribing that aim to improve wellbeing through generalized interventions that are unlikely to cause harm may sometimes be preferable.

Much of this chapter has sought to deflate some of the over-optimistic estimations of what precision medicine can offer, but the aim of the critique has not been to provide in-principle objections to precision medicine. Rather, the aim has been to guide governments towards a more sober estimation of where the relative priorities should lie within the money they invest in prevention. Making perspicuous the ways in which primary, secondary, and tertiary prevention will in practice compete for funding should focus attention on the value for money each can offer. It is a dereliction of government's duty to divert funding to the new, high-tech solutions that will benefit relatively few, at the expense of getting the basics of primary prevention right.[12]

[12] If the genomics-led focus of P4 medicine is to be justified, this would need to be done in the light of the fact that it is likely for the foreseeable future to be significantly less cost-effective than some less targeted interventions. Perhaps what Norman Daniels (2007) has sometimes described as the 'specialness' of healthcare justifies spending much more per QALY on precisely targeted healthcare interventions than on public health interventions, for instance. For what is probably the best current attempt to make this case, see Badano (2016). I criticize such approaches in Wilson (2009).