

## Discounting future health

Hilary Greaves, University of Oxford

### 1. Introduction

Suppose that society has to choose between receiving some benefit now, and receiving the same benefit in ten years' time. Is the earlier benefit to be preferred, either simply on the ground that it is earlier, or for some other reason that is reliably correlated with the benefit's timing?

To answer in the affirmative is to *discount* future goods. When the goods in question are run-of-the-mill "monetary" goods – rice, beans, concrete, double glazing, silk – discounting is ubiquitous, and rightly so. As is well-recognised, the issue is important for practical purposes: the choice of discount rate can make a big difference to cost-effectiveness and cost-benefit analyses, when some costs and/or benefits are located significantly into the future. Government guidelines typically recommend discount rates, for such monetary goods, in the region of 3-7% for developed countries, and sometimes as high as 10-15% for developing countries.

It is in principle quite another question whether future *health* should be discounted: whether future QALYs<sup>1</sup>, for example, should be discounted at all and, if they should, whether the discount rate for QALYs should be numerically the same as the discount rate for monetary goods. Here, too, the choice of discount rate can make an important difference. If future health is discounted at a rate of (say) 5% per annum, then health benefits that occur 10 years into the future are valued only 60% as much as otherwise equivalent health benefits occurring now. Health benefits 20 years in the future count for even less, viz. 37% of equivalent benefits today. This makes an important difference, in the first instance, to comparative cost-effectiveness judgments between interventions whose main benefits are *immediate* morbidity reductions on the one hand, versus those whose benefits include more morbidity reductions relatively far into the future (for example, preventive medicine programmes and in cures for longer-term conditions). The evaluation of premature mortality (particularly of children) is also very sensitive to the choice of discount rate, since the usual practice is to discount 'years of life lost' via premature death according to the time at which each of the lost years *would have been lived*.<sup>2</sup> For example, if years of life lost are discounted at 5%, then a premature death that occurs today and that results in a loss of 60 perfectly healthy life-years over the next 60 years is counted as a loss of only 19 'discounted QALYs'.

Current and recent practice of discounting in healthcare analysis is varied. The UK's National Institute of Health and Clinical Excellence formerly used a discount rate of 1.5% on future health; this was revised upward to 3.5% in 2004. Often, the uncertainty as to the correct figure is dealt with by performing a sensitivity analysis to check how sensitive one's more important conclusions are to the choice of discount rate. In this spirit, the 1996 Global Burden of Disease (GBD) study computed the implications of discount rates of 0%, 3% and 10% per annum; the 2012 Copenhagen Consensus similarly used 0% and 3%; Spanish guidelines (Lopez-Bastida et al, 2010) recommend a reference rate of 3%, but also recommend investigating 0% and 5%. In the more recent 2010 edition of the

---

<sup>1</sup> Or DALYs; none of the considerations in the present chapter distinguish between these two scales.

<sup>2</sup> See, however, section 2.2 below, for a reason for doubting the appropriateness of this usual practice.

Global Burden of Disease report, however, ethical arguments drove a decision simply to apply zero discounting to future health (Murray et al, 2013).

The question of discounting future health has proved quite divisive. A common and entirely understandable intuitive reaction is that discounting future health is extremely odd: why think that changes in future health and suffering matter less, simply because they are in the future? On the other hand, if we are right to discount so-called 'monetary' goods, one would need to make a special case for health to be an exception; and several arguments purport to show that it is borderline-paradoxical to discount monetary goods but not health. This chapter sets out the issues, and attempts to untangle the controversies. Its conclusion will be that (with some qualifications) it is indeed appropriate to discount future health, but there is no reason to think that the discount rate on health must be identical to that for monetary goods.<sup>3</sup>

## 2. Some reasons to discount future goods

### 2.1 The argument from diminishing marginal returns

Let us return, for a minute, to the case of common-or-garden 'monetary' or 'consumption' goods. (Throughout, I use the term 'monetary goods' rather than 'money'. An additional reason for discounting *money* is a positive rate of inflation, so that the same amount of money buys less of a typical 'monetary good' at a later time. But inflation is irrelevant to the discounting (or not) of monetary *goods*.) The contribution of most such goods to well-being exhibits diminishing marginal returns. That is to say, the increase in well-being that is generated by receiving an extra unit of the good is less, the more of the good in question one has to start with. Clearly, for example, the increase in well-being that results from receiving an extra sack of rice is far greater if one is on the brink of starvation than if one's family is already well provided for.

This observation does not *essentially* have anything to do with the timing of benefits *per se*, but a strong link is formed by the robust assumption that a society's material well-being increases over time. Future improvements in material conditions, in that case, correspond to smaller improvements in well-being than do objectively equal improvements to present material conditions, and for that reason the future improvements count as lesser benefits, in the sense of 'benefit' that is relevant to cost-benefit analysis (i.e., well-being increase).

This point is entirely uncontroversial as an argument for discounting *monetary* goods. Such goods should be discounted by at least an amount that reflects the changing conversion rate of marginal increases in consumption of monetary goods into marginal increases in well-being. That rate is in turn jointly determined by (i) the shape of the function that relates bundles of goods to well-being levels (that is, the extent to which the phenomenon of 'diminishing marginal returns' is real) and (ii) the rate at which the abundance of material goods is expected to grow over time.

---

<sup>3</sup> Some of the literature in health economics frames the issue of whether *benefits* should be discounted in addition to *costs*, and if so, whether the rates for benefits and costs should be equal. This is misleading, however: while we are often discussing health benefits that can be bought via monetary costs, the benefit/cost distinction does not necessarily line up with the health/money distinction, and the latter is the relevant one. Indeed, whether a given change (in either health or consumption of monetary goods) counts as a 'cost' or a 'benefit' in the first place depends on which state of affairs is regarded as the status quo, so that treating costs and benefits differently *per se* would be all but nonsensical.

No such argument, however, is applicable to the case of health. It may well be – it hopefully is – the case that the amount of health also increases over time. But this does not give rise to any phenomenon of diminishing marginal returns of health to well-being *when units of health are measured in QALYs*, because the QALY scale has already more-or-less been constructed to be linear in well-being.<sup>4</sup> There is therefore no reason *based in considerations of diminishing marginal returns* to discount future QALYs.<sup>5</sup>

## 2.2 The argument from uncertainty

The future is uncertain; generally, things get increasingly uncertain as we go further into the future. According to standard decision theory, in the presence of uncertainty one should directly seek to maximise *expected value*: a weighted sum of the values that might result from one's actions, with the weightings given by the probabilities of the respective outcomes.

The uncertainty of the future can ground a nonzero discount rate: the future *intended* benefits of one's actions might be discounted, and discounted more the further in the future they lie, on the ground that it becomes increasingly uncertain whether an intended increasingly far-future benefit would actually materialise. In the Stern Review on climate change, for example, Stern's discount rate includes an effective discount rate of 0.1% per annum on well-being. This is not because Stern thinks that far-future well-being is intrinsically less important than near-future or present well-being; it is rather because he estimates that there is a probability of 0.1% per annum of humanity going extinct for exogenous reasons, in which case any expenditures incurred now to mitigate or adapt to climate change would fail to deliver their intended well-being benefits.

As for climate change, so also for health: various types of uncertainty might ground a positive effective discount rate. The extinction of humanity is still one relevant possibility, but in the case of health there are also many others.

Most obviously, the project in question might just fail to deliver the intended benefit. If, for example, the success of one's intervention depends on a certain hospital administering follow-up treatment in 20 years' time, then one must take into account the probability that the hospital will by then have

---

<sup>4</sup> More precisely: well-being is approximately linear in QALYs *if and when* the health utility estimates that underlie QALYs are derived from subjects' responses to evaluative survey questions concerning time trade-offs, person trade-offs and/or standard gambles. This was formerly the universal practice. In the 2010 GBD report, the GBD project moves away from that notion of a health utility, instead simply asking respondents to compare health conditions in terms of "how healthy" the person or population in question is (Salomon et al, 2012). The motivation for this move stems from an "aspiration to quantify health loss rather than welfare loss". In general, not much of use can be said about the relationship between QALYs and welfare on this second method of quantifying degrees of health, except insofar as its results turn out approximately to coincide with those of the first method; in any case, though, there is still no general reason to expect diminishing as opposed to increasing returns here.

<sup>5</sup> Prioritarians would dissent here, since prioritarians believe that well-being itself has diminishing marginal moral value: that is, that a given marginal unit of *well-being* matters less when it accrues to a better-off person than when it accrues to a worse-off person. Provided that future people are better off overall (i.e. not only in terms of health), this could in principle provide a diminishing-marginal-returns rationale for discounting future health. My own view is that prioritarianism is either mistaken or will not in fact generate that verdict (depending on how it is made precise), but this is not the place to defend that claim. For discussion, see e.g. Broome (1991), McCarthy (2008), Otsuka and Voorhoeve (2009), Parfit (1997), Parfit (2012), Porter (2012), Rabinowicz (2002), Greaves (2014), Greaves (2016a).

been destroyed, or the program otherwise abandoned. In this case, if one is working with a naïve model that (before the application of any discounting) ignores this possibility, then the consequences of *going ahead with the intervention* may be otherwise than the naïve (pre-discounting) model assumes.

More subtly, the program can fail to make the intended *difference* to future health states – the difference being, of course, the relevant thing for the purposes of cost-effectiveness analysis – if the naïve model inaccurately predicts what the future health state *would have been* in the *absence* of the program. Naïve models generally assume that future cases of each disease are treated just as they would be today, so that, in particular, the QALY loss associated with each such case is as it would be today. But, since the state of medical technology improves over time, this may well not be the case. Of course, one cannot now *know* which treatments will be available in the future; but the *expected* QALY benefit of a future-directed intervention is generally less than the naïve estimate. An uncertainty-based discount rate, applied to the health benefits estimated by the naïve model, is therefore appropriate to correct for the fact that, in expectation, that model overestimates the benefits.

There is, however, one important respect in which the uncertainty-based rationale for discounting suggests a practice that deviates from the current standard. The issue concerns the relationship between discounting future morbidity-related QALYs (YLDs), versus future life-years lost via premature death (YLLs). When a nonzero discount rate is used, the normal practice, as we noted above, is to discount each future life-year lost not based on the year in which the death occurs, but rather, for each future life-year lost, based on the year in which that life-year would have been lived. But the majority of the uncertainty-based case for discounting provides no rationale for discounting beyond the time at which the death occurs. Granted, the nonzero probability per unit time that the whole of humanity goes extinct does provide such a rationale (since, if the individual in question would have been wiped out by this catastrophe anyway, then the amount she loses via premature death is reduced). But we can safely assume (cryogenics aside) that the probability of future technology providing a ‘cure for deadness’ is zero, so there are no issues of present mortality-reducing interventions crowding out any such cures. Insofar as the discount rate is based on uncertainty, therefore, YLLs should be discounted only to the time of death.

### 2.3 Discounting future well-being

Unlike Stern, some authors *do* take the view that future well-being should be discounted, not only because of considerations of uncertainty, but as a fundamental matter: the decision-maker should regard future well-being improvements as less important than present ones, and far-future well-being improvements as less important than near-future ones. This *positive pure discounting*, in the realm of health or otherwise, tends to be unpopular among moral philosophers and more foundationally-minded economists, but is embraced in particular by many more ‘applied’ theorists.<sup>6</sup>

---

<sup>6</sup> For a more detailed survey of the arguments on both sides, together with onward citations, see Greaves (2016b).

The central argument for positive pure discounting involves the idea of *permissible partiality*.<sup>7</sup> The basic idea here is that a zero pure discount rate would require excessive sacrifice on the part of present persons, for the sake of future persons. In the monetary context, for example, Arrow and others have claimed that thanks to the productivity of capital, a zero pure discount rate would enjoin the present generation to save (for instance) two thirds of their income for the enrichment of future generations: future generations who, we predict, will already be richer than us even without such extreme saving on our part (Arrow, 1999). In the context of health, the disease eradication/healthcare research and Keeler-Cretin paradoxes (cf. section 4 below) are often taken to show that a failure to employ pure discounting in the context of health would lead to an implausible degree of bias towards future health interventions at the expense of treating present health-related suffering. In response to these thought-experiments, one very natural reaction is to hold that just as one is (arguably) morally permitted, or even required, to favour the interests of family and friends over the interests of distant strangers, so a present-day decision-maker is morally permitted (or even required) to favour the interests of present persons or present health conditions over the improvement of future health. Further, one might think, the inclusion of a positive rate of pure discounting in one's algorithm for cost-benefit analysis is the appropriate way to encode this partiality towards the present.

In reply to this, two comments are in order.

Firstly, it is unclear whether or not it is ethically appropriate for the decision-makers in question to exhibit partiality towards the present. Permissions or requirements for partiality are very plausible in the domain of *private* morality: one generally expects private individuals to favour their family and friends over strangers, for example, when giving time and material assistance. However, the considerations that make the permissibility of partiality plausible in the private context – concerning, for example, the limits of human motivation, the psychological *inability* of individuals to be consistently impartial between friends/family and strangers in their private decisions, and the damaging effects that would ensue to personal relationships if they tried – do not apply in the context of public decision-making algorithms. Very plausibly, present-day public decision-makers have just as much responsibility towards those unborn future persons who will be affected by their decisions as they do towards present persons. Against this, however, one might argue on grounds of some sort of contributory principle that it is appropriate to favour contributors to the present health budget when allocating that same budget; and of course the as-yet-unborn have not contributed to the present budget.

Secondly, if one does decide to discount for partiality, one should distinguish between partiality towards present *persons*, and partiality towards present *health*. What the 'permissible partiality' argument most directly seems to suggest is partiality towards *persons* who are already alive today, but not any preference for improving the health of those persons today versus improving the health of the same persons in (say) 40 years' time. To discount future QALYs therefore risks being an imperfect match to the moral principle that the advocate of permissible partiality is trying to

---

<sup>7</sup> A different argument appeals to the empirical claim that people do in fact prefer otherwise-similar health benefits to occur sooner rather than later (within their own lifetimes) as a purely intrinsic matter, even when uncertainty indirect effects (cf. sections 2.2 and 2.4) have been accounted for. Even if this empirical claim is true, however, it is doubtful that it is appropriate for public decision-making bodies to defer to this particular kind of personal preference, since the preferences in question appear to be irrational.

implement. In fact, since most preventive interventions concern the future health of already existing persons, it is doubtful that any discounting on grounds of partiality is warranted at all in practice, even if one does think that cost-effectiveness and cost-benefit analysis should encode a degree of partiality towards present *persons*.<sup>8</sup>

## 2.5 The instrumental benefits of earlier health improvements

Health is important for two reasons. First, it is important *intrinsically*: one's health status at any given time  $t$  is an important component of well-being at that same time  $t$ . But second, it is also important *instrumentally*: that is, it may in turn cause additional benefits for oneself or others. When those additional benefits are themselves more valuable if they occur earlier, this can generate an additional reason for discounting the explicitly modelled health effects, beyond those we have discussed above.

Most notably, one's productivity is affected by one's health, sometimes dramatically so. At least at the level of society, general well-being is positively correlated with the general level of productivity. If some members of society are sufficiently disabled by health impairments that they are unable to engage in any productive work at all, for example, then either those same persons will suffer from severe financial distress, or they will be supported financially by the remainder of society, so that (either way) their impairment reduces the average level of consumption in the society. At a less extreme level, generally speaking, we all work more effectively when we are in good health.

In terms of this taxonomy, we have so far been focussing entirely on the intrinsic value of health. A full cost-benefit analysis, however, should attempt to take into account all the benefits of a proposed intervention, instrumental as well as intrinsic.<sup>9</sup> But economic productivity is indeed, in general, more valuable the earlier it occurs. Just as a unit of physical capital (say, a hammer) is more valuable at an earlier time because it can be put to work to generate a greater amount of capital at a later time, so equally a unit of *human* capital (a given amount of working time in a good state of health) is more valuable at an earlier time because its outputs can, in general, similarly be put to work. This is simplest to see via the observation that the outputs of human labour often *are* additional units of physical capital (one's labour might contribute, for example, to a greater number of hammers being made, or to a new factory being built). But the effect is equally real for other types of labour. For instance, once some occupations exhibit this effect in a relatively direct way, it

---

<sup>8</sup> It is well-recognised (in the literature on discounting in general) that applying a single, *time-independent* discount rate tends to lead to assigning an implausibly low degree of importance to well-being in the *far* future (e.g. Weitzman, 1998). A more plausible model, if one is to encode a principle of partiality in one's cost-benefit analysis, would employ a discount rate that declines as one moves further into the future, so that the near future is assigned a lower moral weight than the present, but the far future is assigned a moral weight that is equal to, or only slightly below, that of the near future. (See, e.g., Gollier, 2013, chapters 7 and 8.) Those who are unconvinced by my tentative rejection of the 'permissible partiality' argument may therefore want to consider using a declining, rather than a constant, discount rate. Such non-constant discount rates lead to time-inconsistency; when different times correspond to different interests, as here, however, such time-inconsistency is no more puzzling (or worrying) than the "inconsistencies" between the preferences of distinct persons.

<sup>9</sup> The 'separate spheres' argument urges against doing this, but that resistance is arguably misguided: see e.g. Persad (2016, this volume), Broome (MS).

will follow that many other occupations exhibit the same effect derivatively: insofar as various services (the provision of easily accessible food and coffee, haircuts, bicycle fixing, train servicing) facilitate the productivity of others, those services are more instrumentally valuable earlier whenever the work they facilitate is itself more valuable earlier.

There should, therefore, be an additional component of the discount rate on future health, corresponding to the generally greater instrumental value of earlier but otherwise similar health improvements.

### **3 Discounting future health versus explicit modelling**

Section 2 concluded that a discount rate should be applied to future health for two reasons. First, to correct for the possibility that the beneficial future effects of a proposed health intervention, in the relevant sense of 'effect', are smaller than the model's naïve calculation suggests. Second, to take into account the fact that earlier health benefits are generally more *instrumentally* valuable than otherwise-equal later ones.

In both cases, the use of a discount rate for future health is not the only way to model the effects in question. Alternatively, one could (i) *explicitly* represent the various possible actual effects of one's intervention and their probabilities, and perform a corresponding explicit expected-value calculation, eschewing the use of the 'naïve calculation' as an estimate of effects on future health, and (ii) *explicitly* model the instrumental as well as the intrinsic benefits of the health improvements one is considering. If one did both of those things, there would be no reasons of the kinds we have accepted for discounting future health. (Of course, if one also thought, contrary to the conclusions drawn in section 2, that a nonzero *pure* discount rate was appropriate, for example to represent the decision-maker's partiality towards the present and nearer-future, then that component of the discount rate would still remain.)

From a purist's point of view, such explicit modelling is arguably more satisfactory: it more completely represents the structure of the complex reality that is actually driving the calculation, and (relatedly) would encourage decision-makers to carefully consider how the factors in question are likely to play out in the specific context of the intervention they are analysing. The result, *if* the explicit modelling were carried out consistently and correctly, would generally be more accurate than results obtained via the discount-rate methodology suggested here.

Notwithstanding these points, however, three factors count in favour of modelling at least a significant proportion of the effects in question via a discount rate for future health, rather than more explicitly, in practice.

The first concerns calculational complexity. Many of the effects in question (for example, concerning the rate of advance of medical technology and the future state of healthcare budgets) will not vary across cases. For these effects, it is more efficient to have a single discussion of what degree of correction needs to be carried out to account for them, and thereby to settle on a general-purpose discount rate, rather than repeating the same discussion countless times, once within each cost-benefit analysis.

Secondly, and relatedly: since there is inevitably a range of reasonable estimates for the effects in question, modelling on a case-by-case basis would be likely to lead to these effects being treated

inconsistently across cases. This would introduce unhelpful noise into our attempts to compare interventions *against one another*; in contrast, modelling via a discount rate encourages a uniform treatment where uniformity is appropriate.

The third factor is more pragmatic: in practice, cost-benefit analysts almost invariably *do not* explicitly model the factors we have been discussing, nor (given their constraints) does it seem likely that they will start to. In that case, the choice we actually face is between taking account of the factors in question via a discount rate on future health, and not taking account of them at all. Of those two alternatives, the former is clearly superior.

All this said, the arguments of the present chapter do suggest that for the purpose of designing appropriate guidelines for discount rates, greater attention should be paid to differences that are relevant from the point of view of those particular arguments. For example, we have already suggested that if discounting is rationalised in the way suggested in this chapter, a modification to the currently-standard way of discounting YLLs (as opposed to YLDs) is indicated. In taking the discount-rate approach, we must guard against the danger that the discount rate is mistaken for a simple representation of pure time preference, so that the relevance of these factors is forgotten. In particular, sources of uncertainty and/or instrumental effects that are significantly idiosyncratic to the case in question will not be well dealt with via a general-purpose discount rate, and will still need to be modelled explicitly.

#### **4 Consistency between the discount rate for health and the discount rate for monetary goods**

We have concluded that the discount rate for health (given the normal modelling techniques for health interventions) should be positive; but we have found no argument for the claim that the discount rate for health should be equal to that for monetary goods. Some have claimed, however, that it is 'paradoxical' or 'inconsistent' for these two discount rates to differ. The aim of this section is to lay these inconsistency worries to rest.

The most prominent argument for equality between the health and monetary discount rates is the so-called 'Keeler-Cretin paradox' (Keeler and Cretin 1983).<sup>10</sup> To set up the 'paradox', suppose that we have a fixed amount of money,  $m$ , and are trying to decide between two ways of spending this money. One option is to spend the money on immediate treatment of a given disease  $F$ , in which case we spend the money  $m$  now; the program then yields some fixed number  $h$  of QALYs, also with immediate effect. Alternatively, we could invest the sum  $m$  for some number of years, with a rate of return that is equal to the discount rate on money, planning to use the resulting increased sum to treat a larger number of cases of  $F$  in the future. Suppose now that we discount health at a lower rate than money. Then (the argument runs) the decision-maker will always prefer a later over an earlier implementation of the  $F$ -treatment program, since, as the program is moved further and further into the future, the present value of the required cost is constant, but the present value of the resulting health benefits increases. But this seems paradoxical: surely indefinite postponement

---

<sup>10</sup> The 'disease eradication and healthcare research paradoxes' are also often discussed in the literature on discounting health (see e.g. Murray and Acharya (1997)). Since, however, these latter 'paradoxes' claim only to show that there is something paradoxical about a *zero* discount rate for health (rather than, as in the Keeler-Cretin case, claiming to establish the stronger conclusion that the discount rate for health must be equal to that for monetary goods), their conclusion is compatible with the ones we have reached above by an independent route.



of a health program cannot be the optimum choice? In contrast, if the discount rate on health is equal to the rate of return on monetary investments, then the decision-maker will be indifferent between any two possible timings of the program under consideration.

The setup of the Keeler-Cretin paradox, however, requires assuming that there is some disease F with both of the following features: (i) at least prior to Keeler-Cretin arguments, it is considered cost-effective to fund treatment of F today, and (ii) unless the decision-maker chooses to invest the money in question for the purpose of treating future cases of F, those future cases of F will not be treated. But assuming that the healthcare budget at each time is set and then prioritised correctly, this combination of features is inconsistent with the assumption that the discount rate on future money is positive. The latter assumption means that the future is materially richer, in which case the future healthcare system should be willing to fund treatments with a higher cost per QALY than those at the "cost-effectiveness threshold" today. It follows that if treatment of F is cost-effective today, the same treatment will be even further inside the cost-effectiveness threshold in the future, and so *will be funded regardless of the decision of today's decision-maker*. The actual effect of investing money to increase the future healthcare budget, therefore, is *not* that some future cases of F get treated that otherwise would not have been treated; it is instead that some *other*, less cost-effective, future interventions get funded. The ratio of the cost-effectiveness of these new treatments to the present cost-effectiveness of marginal present-day health interventions will indeed (in idealised cases) be equal to the discount rate for money, but this has little to do with the discount rate for health.

It is also worth noting that, although arguments that it is somehow "inconsistent" to discount future health and monetary goods at different rates are common in the literature on health economics, from a fundamental point of view any such claim would anyway be extremely puzzling. This is because talk of a single discount rate for "monetary goods" is in any case only an approximation. More fundamentally, in this context "monetary goods" is a shorthand for a long list of tradeable goods (rice, beans, concrete, double glazing, silk...). We define the discount factor (between time 0 and time t) for any *individual* such good as the ratio of the amount of welfare that is generated by an additional marginal unit of that good at t versus at time 0; the rate at which the discount factor for the given good declines with time defines that good's 'own rate of discount'. But given these definitions, it is obvious that different goods will have different own rates of discount. If rice is scarce at time 0 but abundant at time t, for instance, while beans increase in abundance only gradually, then the own rate of discount for rice is likely to be substantially higher than that for beans. The single figure that we call the discount rate for 'money' is some unholy average of these own rates of discount for tradeable goods; precisely *which* average is determined by (e.g.) which goods are included in the bundle that defines the consumer price index, a matter about which there is significant arbitrariness. One should anyway be highly suspicious, therefore, of any claim that *as a matter of fundamental principle* the discount rate for health has to equal that for 'money'.

## 5 Conclusions

While discounting of future money is (rightly) ubiquitous, there is significantly more controversy over whether or not it is appropriate to discount future health, and, if it is, whether the discount rate on future health should be equal to that on money. I have argued that: (1) 'pure' discounting should be rejected, at least in the health context (if not already in the monetary context). (2) Since QALYs

(more or less by definition) do not make diminishing marginal contributions to well-being, the second major argument for discounting monetary goods (based on the assumption that society will be richer in the future) is equally inapplicable to the case of health. (3) There are, however, two good reasons to discount future health, relating respectively to the representation of uncertainty and the instrumental benefits of health interventions. In both cases, it would be *possible* instead to conduct cost-benefit analysis using explicit models of all relevant forms of uncertainty as to the actual health benefits of the program and all relevant instrumental effects; and there would be no remaining reason for a non-zero discount rate for future health within that approach to modelling. I have no deep quarrel with theorists who prefer this more explicit approach – the most important thing is that the effects of uncertainty and instrumental benefits be included in the model *somehow*. But fully explicit modelling seems unlikely to happen in practice, and there are some reasons for thinking that it would in any case be an inferior approach. (4) There is no reason, and cannot be any fundamental reason, to think that the discount rate for health must be equal to that for ‘monetary goods’. (5) Insofar as the rationale for discounting future years of life lost (YLLs) lost is based on uncertainty, those YLLs should be discounted only to the time of death, not to the time at which each life-year would have been lived.

I have not attempted to settle, here, the question (crucial for practical purposes) of precisely what the numerical value of the discount rate for future health should be. To settle this, further work is required to clarify what sort of figures the above considerations of uncertainty and of instrumental benefits would imply in practice: as inputs to the calculation of an appropriate discount rate we would need to quantify, for instance, the expected rate of advance of medical technology, and the typical ratio between the intrinsic and instrumental benefits of health interventions (which will often be specific to the country or region in question). We can be confident that the resulting discount rate would be positive, but it could easily be either lower or higher than that for monetary goods.

## References

- Arrow, K. J. (1999). ‘Discounting, morality, and gaming.’ In P. R. Portney & J. P. Weyant (eds.), *Discounting and intergenerational equity*. Washington: RFF Press.
- Broome (1991). *Weighing Goods*. Oxford: Blackwell.
- Broome (MS). ‘Measuring the burden of disease.’
- Gollier, 2013 *Pricing the planet's future: The economics of discounting in an uncertain world*. Princeton University Press.
- Greaves (2014). ‘Antiprioritarianism.’ *Utilitas* 27(1), 1-42.
- Greaves (2016a). ‘A reconsideration of the Harsanyi-Sen-Weymark debate on utilitarianism’. Forthcoming in *Utilitas*.
- Greaves (2016b). Discounting for climate change: A survey. Forthcoming in *Economics and Philosophy*.
- Keeler, E. B., & Cretin, S. (1983). Discounting of life-saving and other nonmonetary effects. *Management Science*, 29(3), 300-306.

- López-Bastida, J., Oliva, J., Antoñanzas, F., García-Altés, A., Gisbert, R., Mar, J., & Puig-Junoy, J. (2010). 'Spanish recommendations on economic evaluation of health technologies.' *The European Journal of Health Economics*, 11(5), 513-520.
- McCarthy (2008). 'Utilitarianism and Prioritarianism II.' *Economics and Philosophy* 24, 1-33.
- Murray, C. J., & Acharya, A. K. (1997). 'Understanding DALYs.' *Journal of Health Economics*, 16(6), 703-730.
- Murray, C. J., Ezzati, M., Flaxman, A. D., Lim, S., Lozano, R., Michaud, C., ... & Wikler, D. (2013). 'GBD 2010: design, definitions, and metrics.' *The Lancet*, 380(9859), 2063-2066.
- Otsuka, M. and A. Voorhoeve (2009). 'Why it matters that some are worse off than others: An argument against the priority view.' *Philosophy and public affairs* 37, 171-99.
- Parfit (1997). 'Equality and priority.' *Ratio* 10, 202-21.
- Parfit (2012). 'Another defence of the priority view.' *Utilitas* 24, 399-440.
- Persad (2016). [Chapter in this volume]
- Porter (2012). 'In defence of the priority view.' *Utilitas* 24, 349-64.
- Rabinowicz (2002). 'Prioritarianism for prospects.' *Utilitas* 14(1), 2-21.
- Salomon, J. A., Vos, T., Hogan, D. R., Gagnon, M., Naghavi, M., Mokdad, A., ... & Farje, M. R. (2013). 'Common values in assessing health outcomes from disease and injury: disability weights measurement study for the Global Burden of Disease Study 2010.' *The Lancet*, 380(9859), 2129-2143.
- Weitzman, M. L. (1998). 'Why the far-distant future should be discounted at its lowest possible rate.' *Journal of environmental economics and management*, 36(3), 201-208.