On Discount Rates in the Cost-Benefit Analysis of Climate Change*

Edmund Tweedy Flanigan

2011

Introduction

There has been much discussion lately about the selection of an appropriate discount rate in the cost-benefit analysis of climate change. This is the rate at which the value of future stuff—commodities, harms, benefits—is weighed against the value of present stuff. To 'discount' in cost-benefit analysis is to select a positive discount rate, thereby devaluing future stuff. A number of factors determine what rate is appropriate, and theorizing about these factors has led eminent economists to select very different discount rates for the cost-benefit analysis of climate change. This seemingly small matter is of great practical importance. Nicholas Stern and William Nordhaus are two economists who make sharply divergent recommendations for action based on the results of their respective analyses; the divergence is due primarily to differences in the discount rate used.[1]

Philosophers appear to stand largely against the practice of discounting, which is also called 'discounting the future.' They oppose it on ethical grounds. Different philosophers oppose the practice in different ways. Most commonly, philosophers tend to oppose the incorporation of what is called a 'pure time discount rate' or a 'rate of pure time preference' into the overall discount rate. This is discounting future stuff *simply because* it comes in the future.

I stand with most philosophers in thinking the practice to be misguided, at least as it is commonly construed. Nevertheless, I also think that philosophers and economists are, as John Broome has said, in important respects talking at cross-purposes on the matter (see Broome 1992, 1994). Discounting is not quite so

*I am grateful to Patrick Connolly, Simon Feldman, and especially to John Broome and David Miller for comments on an earlier draft of this paper. A version of this paper made up Chapter 4 of my 2011 M.Phil. thesis (Department of Politics, University of Oxford).

[1]Stern uses a discount rate of 1.4%; Nordhaus recommends something in the range of 3–5%. Stern recommends strong action now to mitigate the effects of climate change, whereas Nordhaus recommends a more slowly ramped-up response. These differences in policy prescription are due to the differences in view on the discount rate. See Nordhaus (2007).

objectionable as writing by philosophers may make it seem. I think much of the disagreement arises from differing views about the purpose of cost-benefit analysis.

I will argue here for that claim. I begin by giving a preliminary account of what I take to be the purpose of cost-benefit analysis. This is, perhaps surprisingly, quite difficult to do. Cost-benefit analysis is ultimately a way of weighing harms and benefits,[2] but it's hard to specify exactly what counts as a harm and what counts as a benefit in a way that is both precise and consistent with economic practice. Since cost-benefit analysis is not my primary subject here, I give a loose characterization. I also introduce the idea of an *analysis-given reason*, which is the reason for action produced by a cost-benefit analysis. How forcefully that reason counts in favor of a plan of action depends on two things: (1) whether the analysis accurately evaluates and aggregates harms and benefits, and (2) whether the conception of *harm* and *benefit* adopted by the analysis captures enough of the goods and bads associated with the plan being considered (or maybe just enough of the *important* ones).

Discounting is done with respect to time. There are at least two very good reasons for discounting; these have to do with the way that economists use prices to approximate the values of harms and benefits. I describe these reasons. Next, I describe some other reasons that are commonly given for discounting. These reasons are ethical and political ones, and they have nothing to do with the price evaluation method that supports the first two discounting practices. They have to do with the force of the analysis-given reason. I suggest that these should be considered separately.

This sets up my discussion of 'pure time discounting.' I review an argument recently given for the practice by the economist Martin Weitzman. While the motivation for the argument seems to come from the method of economic cost-benefit analysis, its justification appears to be moral or political. I consider the merits of this justification.

This leads me to turn, finally, to one political reason for discounting that I think is especially important. It has not received enough attention. The reason is that nations, whose responsibility it is to enact climate policy, have special obligations to their citizens. These special obligations may extend to future citizens. As such, I'll consider some arguments for and against including future citizens among those people to whom nations owe special obligations.

One further point before beginning. There is quite a lot to say about discounting,

[2] I use the term 'harm' rather than 'cost' to reflect common philosophical practice. 'Benefit,' of course, is used by both philosophers and economists.

and I will not attempt to say anywhere near all of it. Some of what will go unsaid here is of great importance to the economic analysis of climate change. A few of these issues are (a) hyperbolic and gamma discounting,[3] (b) implausibly high implied savings rates of low discount rates,[4] and (c) the appropriate formalization of discounting formulas.[5] These matters are, in my view, better left to economists, who have the expertise necessary to really address them. I don't think my omission of them here will have any adverse effects on the arguments I offer.

1 Cost-Benefit Analysis

The conventional view of cost-benefit analysis is that it is a tool to assess the value of a possible plan of action (see Drèze and Stern 1987, Broome 2000).[6] When a person, or a group of persons, wants to know whether to undertake some plan, one way of beginning to answer that question is to ask how good or bad the plan is. That is, we can *evaluate* the plan. Oftentimes a plan will have various bad consequences associated with it; call these *harms*. If it is a plan we are seriously considering, it probably has some good consequences associated with it as well; call these *benefits*. Cost-benefit analysis is the adding up and weighing — the *aggregation* — of the harms and benefits of a plan of action as a method of evaluating it.[7]

Why do we want to assess the value of a plan of action? We do so because if the plan is a good one, in the sense that it will result in more benefits than harms, it may be that it is one that should be undertaken. If, on balance, a plan will result in a net harm, this fact may be said to count against the plan; and if, by contrast, the plan will result in a net benefit, this fact may be said to count in the plan's favor. In this sense, cost-benefit analysis aims at producing a practical reason. Call the reason generated by the cost-benefit analysis of some plan the *analysis-given reason* for or against that plan.

The analysis-given reason is a derivative reason. It is derivative of some of the

[3] See Cowen (2007, p. 12), Weitzman (1998), and Dasgupta and Maskin (2005).

[4] See Arrow (1999) and Dasgupta (2007).

[5] By this I mean the appropriate way to incorporate discounting into the economic cost-benefit model. Economists tend to use an equation from Frank Ramsey's (1928) growth model. I concur with Beckerman and Hepburn (2007, pp. 194–195), who think parts of that formalization are "overworked" and "not rich enough to separate the key ethical elements relevant to climate change." For this reason, I have decided not to discuss discounting in terms of the Ramsey discounting equation, which some readers might expect to see given prominent treatment.

[6] One welfare economics textbook says right at its outset that the entire discipline "can be viewed as an investigation of methods of obtaining a *social ordering*"—a betterness ranking—"over alternative possible *states of the world*" (Boadway and Bruce 1984, p. 1; original italics).

[7] This may seem an unusual conception of *harm* and *benefit*. Many take harms and benefits to be differences in the wellbeing of particular persons. This conception is much broader than that, though it includes it. I discuss the narrower conception of harms and benefits *to people* in a note on page 12.

reasons linked to the harms and benefits of a plan. Each benefit counts in favor of a plan and is thus a reason for undertaking the plan; and conversely, each harm counts against the plan and is thus a reason not to undertake the plan. The weighing of these harms and benefits in the analysis is also an aggregation of reasons associated with each. But the analysis-given reason may not be derivative of *all* the reasons linked to the harms and benefits associated with a plan; indeed it is typically not.[8] How large the slice of reasons is that is captured by the analysis-given reason depends on how ambitious the analysis is and how good its methods are. This may seem opaque. What I mean should become clearer shortly.

Whether the analysis-given reason produced by a cost-benefit analysis counts strongly in favor of or against the plan depends mostly on two factors. The first is how accurate the evaluation of the harms and benefits of a plan is. This depends on whether the estimates made by the analysis of the value of each harm and benefit reflect the true values of those harms and benefits (Broome 2000, p. 954). Because of this, a cost-benefit analysis needs a theory of value. The second factor is how well the analysis-given reason captures the entire set of reasons for or against the plan. If it captures only a small slice of those reasons, or if the reasons it does capture are unimportant reasons, then the analysis-given reason may not do much to guide us practically. Because of this, a cost-benefit analysis needs to have an idea of the set of reasons it aims to capture.

This second factor is of special importance. It is important because what reasons are captured by a cost-benefit analysis is determined at least in part by the specification of *harm* and *benefit* that it employs; that is, by which of the bads and goods associated with a plan are aggregated by the analysis. This specification is constrained in two directions. It is constrained first by the ambitions of the analysis. The cost-benefit analysis of climate change is particularly ambitious. Judging at least by the analyses of the *Stern Review* (2007) and the *Garnaut Report* (2008), the cost-benefit analysis of climate change aims to provide a strong analysis-given reason for action. It thus needs to adopt a fairly wide specification of what counts as harms and benefits, so that it can account for harms to the environment, harms resulting from premature deaths, and so on.

The specification is constrained second by the demands and possibilities of its methods. Economic cost-benefit analysis uses the market price method of evaluation (see Broome 1994, p. 133). This method uses market prices—or estimated

[8]Beckerman and Hepburn (2007, p. 188) correctly point out that "the standard welfare economic approach has no room ... for ethical dimensions concerning the processes by which outcomes are reached." We should thus not expect a cost-benefit analysis to capture any procedural reasons for or against a plan, including those having to do with "rights, justice and freedoms." This limitation is very important, and is one example of a way in which the analysis-given reason fails to capture the full set of reasons for or against a plan.

market prices for stuff that is not traded—to value harms and benefits. Because the price of a benefit is a ranking of that benefit's value to people compared to other benefits, the price that people pay (or are willing to pay) to avoid a harm or receive a benefit does not give a value to that harm or benefit that is independent of the values of other harms and benefits. Cost-benefit analyses using the market price method of evaluation are thus committed to the full range of specifications of *harm* and *benefit* associated with market prices, and they may be committed to not weighing harms and benefits that cannot be priced.

For a cost-benefit analysis as ambitious as the cost-benefit analysis of climate change, this makes things very tricky. When weighing the harms and benefits associated with premature deaths and population change, for instance, the theory of value underpinning the analysis needs to be more explicitly worked out. This is because we need to know how bad it is if a life is cut short, and whether it is good or bad if certain people come into existence rather than others. In an ambitious analysis it is also more difficult to see how strong the analysis-given reason is for action, since it is more difficult to tell what portion of the reasons for or against the proposed plan of action the evaluated harms and benefits capture. For instance, it is hard to say whether the prices it assigns to certain benefits (such as the non-destruction of certain species) reflect attitudes about matters of justice. If so, we need to know whether those attitudes are correct. If not, matters of justice need to be considered exogenously when evaluating the strength of the analysis-given reason.

2 Value in the Future

Cost-benefit analysis weighs harms and benefits that come at different times. Most such analyses, as practiced by economists, adopt a rate of discount. As I have said, using a discount rate is a way of counting future harms and benefits for less than similar harms and benefits in the present.

This should seem surprising. Harms and benefits are differences in goodness. But whether something is good, or bad, seems independent of timing. A harm is no less a harm if it takes place now, at some time in the future, or at some time in the past. Call this the *time-insensitivity* of harms. Benefits, too, are insensitive to timing in this way. Similarly, we could say that harms and benefits are *subject-insensitive* and *location-insensitive*. They are subject-insensitive because a harm is just as bad no matter who endures it, just as a benefit is just as good. A harm is just as bad for me as it is for you; and as to the kind of harms that can be endured by both humans and other animals, those harms are just as bad no matter whether they are endured by a human or an animal, a person or a non-person.[9] The argument

[9]It may be thought that a harm is less bad when endured by a person who deserves to be harmed.

is the same for location-insensitivity. How bad or good a harm or benefit is does not depend on where it is endured or received. I take these features to be part of the meaning of *harm* and *benefit*.

There is thus a moral presumption against discounting. How can the practice be justified? I will now very briefly survey two of the ways in which it can be, since it is important to get clear on exactly what the justification is in these cases. These two cases represent extremely common practices in economic cost-benefit analyses.

2.1 Discounting Prices

Discounting can be justified when it is applied to certain harms and benefits within the market price method of evaluation. This method is, as John Broome and Joseph Stiglitz have said, a way of using "reasonable shortcuts" to approximate the aggregation of harms and benefits I described above (Broome 1994, Stiglitz 1983). The central feature of this method is that it uses market prices—a 'money metric'—to give a value ranking of harms and benefits. The prices, however, are not prices of benefits but rather prices of *commodities* that *count as* benefits. Because of certain facts about economies (namely, the growth of economies and the difference in the value of money at different times), prices at different times need to be adjusted if they are to form a reasonable basis for accurate comparisons of value. A commodity with some price X in the present may not have the same value as a commodity priced at X in the past or the future.

Consider a simple example, which I've taken from Broome (1994). Suppose that the price of a bottle of wine is \$5. With \$100, I could buy twenty bottles of wine. Alternatively, I could save that money and earn interest on it. Say the interest rate for saving is 10% and the price of wine a year from now rises to \$5.25. By saving the money and earning interest on it (bringing my holdings to \$110), I could buy twenty-one bottles a year from now. The interest rate of the wine is 5%, which is reflected in its price increase. Call the wine I can buy now *present wine* and the wine I can buy next year *future wine*. Present wine and future wine have both a present and future price. If I want to compare their value using the market price method of evaluation, I need to choose a time index. Presently, I can buy twenty bottles of present wine. How much future wine can I buy now? Twenty-one bottles. That is because future wine is worth five percent less in present prices. This example can be generalized, since most commodities behave this way.[10] The

This may be true, but it trades on a different sense of 'good' and 'bad' than the one I am using. I would say that while a harm is as *bad* no matter who endures it, *causing* it may be less *wrong* (or indeed not wrong at all) if it is deserved. This seems to me a very plausible distinction.

[10]Some commodities have interest rates that are zero or very near zero. The value of these commodities

present prices of future commodities are usually less than the present prices of present commodities. This represents a justified reason for discounting the value of future commodities.

Another justified reason for discounting within the market price evaluation method is because of the law of decreasing marginal value. Many commodities are less valuable to people 'at the margin'. In other words, the value to a person of an additional unit of some commodity depends on how many units of that commodity the person already has: the more of it had, the less valuable an additional unit. My tenth orange is less of a benefit than my second; a restauranteur's ten thousandth dollar less of a benefit than her first; a family's fourth home bathroom less of a benefit than its second; and so on. Additional stuff counts as less of a benefit to those who have more stuff to begin with. (The converse is true as well. A family losing one of its four home bathrooms counts as less of a harm than a family losing its only bathroom, etc.) Since economies grow, people in the future will be better off in the sense that they will have more commodities. Many benefits to them will thus be less valuable. In general, then, we can justifiably discount the value of stuff that will be received by people in the future as long as those people will have more stuff to begin with.

The thing that's important about these reasons for discounting is that they're methods of correcting the price of commodities at different times so as to reflect the fact that the harms and benefits brought about by them count for the same no matter when they're endured or received. This supports the time-insensitivity of harms and benefits. There can thus be reasons for discounting within the market price method of evaluation that are not reasons for discounting generally.[11] That is, there can be reasons for discounting price valuations that are not reasons for discounting harms and benefits.

---

should not be discounted. Many of these commodities, like Derek Parfit's example of a beautiful stretch of countryside (1984, p. 483), need to be accounted for in the cost-benefit analysis of climate change, which means that this kind of price adjustment discounting in that analysis needs to be done with special care.

[11]There are other reasons unrelated to time (and thus to discounting) for correcting price valuations that reflect different kinds of insensitivity of harms and benefits. Here is one that is very important to the cost-benefit analysis of climate change. One plausible way of valuing the harm represented by loss of life is to find out how much people are willing to pay to insure their lives. But people in rich countries are willing to pay much more than people in poor countries. This does not reflect a difference in the value of lives between rich and poor countries; it reflects a difference in the value of money. This difference of valuation needs to be corrected for precisely because the badness of loss of life does not depend on the value of money to a person whose life is lost.

2.2 Discounting Harms and Benefits

Neither of the immediately preceding justifications for discounting, then, justify discounting harms and benefits. They justify discounting prices. Are there other reasons for discounting harms and benefits? Some have been proposed. I'll discuss three of these now.

Risk and Uncertainty

The further into the future we look, the less likely we are to be correct about what benefits will or will not come about. There are lots of factors that may prevent the realization of a benefit. Storms might cause delays in construction that raise the cost of a construction project. Natural disasters may have similar but much more severe effects. In the cost-benefit analysis of climate change, we need to account for the possibility that humanity will not survive far into the future. In all of these cases, unforeseen events prevent the realization of a benefit, so in adding up expected future benefits, we might discount them due to this uncertainty. The same goes, *mutatis mutandis*, for harms.

I can think of just one justification for discounting expected future harms and benefits in this way. This justification requires us to assume that the likelihood of each (of every) future harm and benefit occurring decreases uniformly with time. If an expected benefit has an 85% chance of occurring at time t_1 and a 65% chance of occurring at time t_2, then they should not be valued equally; a good cost-benefit analysis should take account of this difference. Thus, the appropriate rate of discount for risk and uncertainty would be the rate according to time at which the expectancy of each future harm and benefit decreases.

But of course there is no single rate at which the expectancy of future harms and benefits decreases. Different events, and therefore different expected harms and benefits, are more and less likely to occur in the future, and the rate at which that likelihood changes with time is similarly differential.

Consider: If a benefit occasions every 'heads' in a series of coin tosses extending into the future, the likelihood of that benefit occurring remains 50% as far into the future as the series of coin tosses does. This is a very different kind of uncertainty from the kind that accounts for harms due to certain natural disasters that are, let's say, 50% likely to occur every four years. To discount both of these kinds of uncertainty according to a single rate would be to commit a serious error.

Moreover, economists have a mechanism for accounting for risk and uncertainty that is a much finer implement than a discount rate. It is called expected utility

theory[12], and it weights individual future harms and benefits according to their likelihood of occurring. It works, more or less, like this. Instead of aggregating future harms and benefits, which there is an epistemic barrier to doing accurately, in practice what we do is to aggregate *expected* harms and benefits. If Plan A has a 90% chance of bringing about some benefit and Plan B has a 95% chance of bringing about the same benefit, then Plan B has a greater expected benefit. That fact ought to count in favor of Plan B in a cost-benefit analysis comparing the two. Expected utility theory is the right tool for this job; a discount rate is not.[13]

Future People Will Be Better Off

I mentioned just above that one good reason for discounting the future within the market price method of evaluation is that if future people will have more stuff, additional stuff that comes to these people will count as less of a benefit. This is due to the law of decreasing marginal value. But there is another argument that says harms and benefits themselves count for less to people who are better off. This argument is similar but distinct, and worth considering on its own.

It is widely held that it is either usually or always more important to benefit people who are worse off. Conversely, it is widely held that it is either usually or always more important to avoid harming people who are worse off. John Rawls held something like this view; prioritarians certainly hold this view; and most egalitarians presumably hold this view, if only derivatively. We might think that, if we hold this view, and if future people will be better off, we have reason to discount future harms and benefits.

This is a good argument, but some care is warranted here. There is an ambiguity when we speak of harms being 'worse' or benefits being 'better' under certain circumstances or for certain people. Take a harm, such as the infliction of some amount of pain, and consider it as befalling two people: a poor child and a rich adult. In one sense, we could say that the harm to the poor child is worse than the harm to the poor adult, either because the child is a child (and harms to children are worse), because she is poor (and harms to the poor are worse), or both. But each endures the same amount of pain, so in another sense the harm is just as bad for each. When I say that harms and benefits are subject-insensitive, I mean that

[12]I don't want to spend time talking about 'utility' here. For a discussion of that term as it's used in the market price method (and indeed in welfare economics generally), see Broome (1991). For an analog in ethics, see *expectabilism* in Parfit's *On What Matters* (2011, p. 106).

[13]Savvy readers will know that the risk of human extinction is sometimes factored into the discount rate of the cost-benefit analysis of climate change rather than accounted for using expected utility theory. This practice, which I don't understand, is adopted in the *Stern Review*. Stern recognizes that this is disputable; see Stern (2007, p. 663 fn8). See also Broome (1992, p. 102) for support on this point.

in this latter way. Another way of formulating this would be to say that a harm or a benefit *qua* harm or benefit is just as bad or good no matter who endures it.

To avoid this ambiguity, I will say that a harm is never *worse* to one person or another, but that it can be *less important*. The same goes for benefits. According to this way of speaking, if we say we should discount harms or benefits because future people will be better off, what we must mean is that a harm or benefit that comes to someone who is better off is less important; that is, that it counts for less as a reason for or against adopting some plan of action.

This may well be a good reason for discounting harms and benefits. If it is, it is because (a) the ethical view about the importance of benefitting people who are worse off is the right one, and because (b) enough future people will indeed be better off, and in a way that bears a good correlation to time. The first of these two justifications has to do with ethics; the second concerns a matter of empirical fact.

It should be clear that this reason for discounting is very different from the reasons discussed in the preceding section. Those were reasons for discounting prices; they operated within the price method of evaluation as a way of making that method's evaluations accurate. This reason represents a view about how to weigh certain kinds of harms and benefits as they are aggregated. It is a way of giving greater force to the analysis-given reason by weighing harms and benefits in a way that better captures the reasons associated with them. It is an explicitly ethical view for discounting, and it should be considered on its merits as an ethical view. It should be considered separately from price discounting reasons.

Self-Interest and Special Ties

There is a third argument we might take to justify discounting future harms and benefits that I think is especially important. It is the argument from special ties. Special ties are obligations that we have to ourselves or to certain other people because of a relationship that exists between us. We have special ties to ourselves because we each have a special interest in our own wellbeing. We have special ties to our families, friends, and fellow citizens because we have a special interest in their wellbeing. These interests are different from the interests I have in everyone's doing well or being free from suffering; those interests are not special.

The argument claims that since we have these special ties, and since these special ties drop off in the future, we should discount future harms and benefits.[14] As

[14]John Broome has pointed out to me that this argument could be right about the nature of our special ties but wrong about discounting. Suppose benefits to grandchildren count for half as much as

above, this is not a claim about harms being any less bad or benefits being any less good to future people. Rather, it's the claim that harming or benefitting future people is less important. Since it's less important, the reasons generated by future harms and benefits count for less as part of the analysis-given reason.

The same argument could be given in political terms: Nations will be the actors undertaking climate policy. A nation has special ties to its own citizens. So nations should discount harms that come to non-citizens, including to future people.

One thing to notice about this argument is that whether it affects the force of the analysis-given reason depends on just who is considering the reason. That is, it makes cost-benefit analysis *agent-relative*. Imagine a cost-benefit analyst who gives extra weight to benefits that will come to her own family. If it is the cost-benefit analyst who is considering undertaking the plan of action being evaluated, this way of weighing will indeed capture a fuller set of reasons for or against the plan. For her, the analysis-given reason will have greater force than it would have had otherwise. But if it is someone else who is considering the plan of action, the cost-benefit analyst will have made an error. For the analysis will have failed to accurately represent the importance of the harms and benefits being aggregated, as it will have given undue extra weight to certain benefits.

Discounting for special ties, then, can be tricky, since it may be unclear which agent-relative reasons should be incorporated and what their strength is against other reasons. One plausible view would be that since the business of cost-benefit analysis is a 'pure' evaluation, it should take account of no agent-relative reasons to weigh some harms and benefits differently. Special ties are agent-relative, so on this view it would be improper to discount (or otherwise give special weight to) harms and benefits that will be received by those to whom we have special ties. Another plausible view, however, says that the agent for whom the analysis-given reason is meant to count should be considered as part of the analysis, precisely by giving extra weight to harms and benefits that will be received by those to whom the agent has special ties. These are not compatible views, so they must be chosen between.

As I said, I think this argument is especially important, and I will say more about it soon. Now, however, I just want to reiterate a point: this reason, like the reason from the argument that future people will be better off, represents a view about how to weigh future harms and benefits as they are aggregated. It is a view that is, depending on its formulation, ethical or political, and it should be considered on its merits as such a view. It too should be considered separately from price

benefits to children. If a person has twice as many grandchildren as children, the aggregate importance of benefits is constant across these two generations, thus failing to support a rate of discount. I do not wish to discuss issues relating to population size here, but this is a point well-worth noting.

discounting reasons.

There are other, similar arguments for discounting that I will not continue to go through now.[15] I have gone through these three to bring out my central claim in the first part of this paper. It is that there are reasons for discounting within the market price method of evaluation and reasons for discounting outside it. The reasons within the method have to do with the accuracy of the evaluation given the method of economic cost-benefit analysis. They are excellent reasons, and correspond in precise ways with time. The reasons outside the method are very different. When they do justify discounting, they justify it in ways that recommend different rates of discount. The grounds for this kind of justification are ethical and political, and are ways of altering the force of the analysis-given reason. While an ambitious cost-benefit analysis should take account of these reasons, since it aims to produce an analysis-given reason with great force, they need to be considered separately and on their own merits. Grouping both kinds of reason under the heading 'discounting' invites confusion.

3 Pure Time Discounting

We are now in a position to assess an argument given very recently by Martin Weitzman (2007). He argues for what is called a rate of pure time preference. Such a rate is sometimes incorporated into the discount rate. I said earlier that pure time discounting devalues future harms and benefits just because they come in the future. Though this is how the practice is commonly described, it should be clear by now that that cannot be exactly right. Let's get clearer on what it means. Consider two possible interpretations of the practice.[16]

1. Future harms and benefits are discounted because they count for less. They count for less just because they come in the future.

2. Future harms and benefits are discounted because they count for less. They count for less for some reason(s) which correlate(s) with time.

I think the only plausible interpretation is the second. (1) contradicts the time-insensitivity of harms and benefits. It offers a poor justification, since timing as

[15] One that merits mention even though I do not want to address it here is that we may not in fact benefit or harm certain future people at all; namely, those particular people whose very existence is owed to the policy options being evaluated by a cost-benefit analysis. This is what Derek Parfit (1984) calls the non-identity problem. There are different views on the extent to which this is a problem for cost-benefit analysis, though I cannot think of a reason why it would justify a rate of discount.

[16] The practice does indeed warrant 'interpretation.' See Stern (2007, p. 664) and Frederick et al. (2002).

such does not affect how much a harm or benefit counts for. So we should expect a reason why future harms and benefits should count for less. Many philosophers have pointed this out. While philosophers who point this out are not wrong to do so, some take the unwarranted further step of rejecting pure time discounting on that basis. It's unwarranted because there's a better interpretation of the practice. We should address that interpretation.[17] With that in mind, let's turn to Weitzman.

Weitzman's focus is the discount rate selected by Nicholas Stern in the *Stern Review*. Stern discounts mostly just for the reasons that correct prices in the market price method; that is, he discounts for growth. The rate of discount he uses is significantly lower than observed market interest rates. Since market interest rates reflect revealed preferences about the value of future stuff, Weitzman objects to the discrepancy. It is inappropriate, Weitzman writes, to select a low discount rate

> ...by relying mostly on a priori philosopher-king ethical judgments about the immorality of treating future generations differently from the current generation—instead of trying to back out what possibly more representative members of society ... might be revealing from their behavior is *their* implicit rate of pure time preference. An enormously important part of the "discipline" of economics is supposed to be that economists understand the difference between their own personal preferences for apples over oranges and the preferences of others for apples over oranges. Inferring society's revealed preference value ... is not an easy task in any event (here for purposes of long-term discounting, no less), but at least a good-faith effort at such an inference might have gone some way towards convincing the public that the economists doing the studies are not drawing on conclusions primarily from imposing their own value judgments on the rest of the world. (2007, p. 712)

There's a lot going on in this argument, and I won't try to address all of it.[18] The main claim seems to me to be that rather than make ethical or political judgments about the discount rate, we should try to discover the rate at which people actually discount the value of future stuff in their day-to-day lives. This could mean trying to infer a discount rate from the way that individuals value harms and benefits that will come in the future in their own lives. This claim is sometimes argued for. But in the cost-benefit analysis of climate change, we are thinking about

[17] Beckerman and Hepburn (2007) seem to have this in mind when they suggest we should consider agent-relative ethical theories when selecting a rate of pure time preference.

[18] For some considerations that call into question the value of revealed preferences in this context, see Beckerman and Hepburn (2007, pp. 202–205).

discounting across generations; that is, we are discounting across lives, not within them. Alternatively, then, the claim might be that we should try and discover at what rate people actually discount across generations. This latter claim looks to be what Weitzman has in mind.

Why would we care what 'representative members of society' reveal to be their implicit rate of pure time preference with regard to intergenerational projects? If we were considering an *intra*generational project, one justification might seem appealing. It is that it is important to respect the attitudes that people have toward their future wellbeing, since they have the right to count future harms and benefits for less, even at an overall harm to themselves. So it only matters whether people prefer benefits to come sooner and harms to come later, not whether people *should* prefer them that way. We must respect the autonomy of individuals.

This argument could be extended to cover *inter*generational projects. If we can discover the rate at which people value harms and benefits that will be received by their children and grandchildren, we should respect those valuations even if they appear to be morally unjustified.

But to argue from respect for people's autonomy is to misunderstand the purpose of cost-benefit analysis. As I have said, the conventional view is that it is a tool for evaluating plans of action. It produces an analysis-given reason for or against the plan that reflects this evaluation. People should be free to respond to, disregard, or otherwise react to the analysis-given reason, but the analysis itself need not take into account *how* they are likely to react.[19] Indeed, to do so would be to view cost-benefit analysis not as an evaluative tool but as a decision-making procedure—a preference aggregator. But it is not one. Rather, an agent should take the analysis-given reason into account as part of its decision-making procedure. If this seems unclear, consider a parallel. Suppose a government asks for a report on the security dangers posed by a rogue nation. Key members of the government are hawkish, so they are likely to give more weight to reasons for military action. It is obvious that the report should not take into account the political leanings of the members of the government. That is because the report has been asked to give an evaluation of a threat. After receiving the report, the government should *then* be free to weigh the advice of the report as part of its policy decision, or to disregard it. To incorporate the political leanings of the members of government into the evaluation would be to turn the evaluation into a decision-making procedure, which it is not.

[19]This is different from saying that we should or should not take agent-relative reasons into account in cost-benefit analysis, which I have suggested is an open question. The argument from autonomy says that people should be free to respond or not respond to the analysis-given reason. This seems right to me, but does not bear on the cost-benefit analysis itself.

Besides, the intergenerational variant of the argument is problematic on its own terms. It seems doubtful at best that it would be morally permissible for people to consistently benefit themselves at the expense of their children or grandchildren, living or future. This is because while people typically have the right to make poor decisions that affect themselves, they are not permitted to do so when those decisions adversely affect others. Autonomy is no grounds for including the rate at which present people value harms and benefits to future generations as part of a cost-benefit analysis. We have no right of autonomy to bring harms to future people. If we should value harms and benefits that come to future people for less, it is not because we have a right of autonomy to do so.

4 Future Generations and National Ties

So should the cost-benefit analysis, as an evaluation, give harms and benefits equal weight no matter what generation will be affected? To think so would be to endorse the

> *intergenerational neutrality thesis*, which says that we ought to remain neutral between harms and benefits to present and future generations.

I'm not sure that the intergenerational neutrality thesis is true. I can think of several considerations that suggest it might not be and some others that may suggest it is. These considerations have to do with discounting for what I called 'special ties' above. Discounting for this reason is to adopt the view that cost-benefit analysis should account for agent-relative differences in the value of harms and benefits to different people.

One reason to doubt the intergenerational neutrality thesis has to do with a special feature of what we could call the

> *individual neutrality thesis*, which says that we ought to remain neutral between harms and benefits that will come at different times in our own lives.

Consider first that many people think that we have special ties to ourselves. These people think that each person has "one supremely rational ultimate aim": that her life go, on balance, as well as possible (Parfit 1984, p. 4)—or something similar.[20] These people think that this interest generates normative reasons for action; they think that each should act in such a way as to ensure that her life go as well as possible.

Does this view contradict or support the individual neutrality thesis? It supports it. It demands that I assess possible actions from the point of view of my entire

[20] Parfit, of course, argues against this view.

life, rather than from the present point of view or through what we could call a 'myopic lens' (giving extra weight to the nearer future). Since, on this view, my one supremely rational ultimate aim is that *my entire life* go as well as possible, any preference for one time in my life over another can only be derivative; that is, it can only matter insofar as preferring my wellbeing at some time is instrumental to the overall wellbeing of my life. In this sense, individual self-interest view requires that we support the individual neutrality thesis. It requires that we remain rationally neutral between the wellbeing of our present and future selves.

Now consider that there is no plausible direct parallel of the individual self-interest view to support the *intergenerational* neutrality thesis. Such a parallel would recommend, from the point of view of *humanity* (*i.e.* all generations), that we *as humanity* act with the supremely rational aim that things go, for humanity itself, as well as possible. The humanity self-interest view would entail that we should remain rationally neutral between the welfare of generations; that is, it would endorse the intergenerational neutrality thesis. But the humanity self-interest view is untenable. Humanity cannot act *as humanity*, nor is it clear that humanity has robust interests. I know of no one who endorses the humanity self-interest view, and I know of no reason why anyone should.

Nevertheless, there may be a different special ties or 'self-interest' view that supports the intergenerational neutrality thesis. It is a self-interest view that I believe many people endorse. It is the

> *national self-interest view*, which says that nations have one supremely rational interest: that things go, for that nation, as well as possible.

Does this view support the intergenerational neutrality thesis? The answer is that it depends. It depends on the nature of the special ties that nations have. It is possible that the nature of a nation's special ties depends on the correct view about the identity of nations through time. On one view of this kind, nations have special ties only to the people who make it up at a time; that is, only to present people. This first view is the view that nations do not have an identity that persists through time. On another such view, nations have a special obligation to these people, but that obligation extends through time to future generations. This second view is the view that nations do have identity through time.

On the first view, nations do not exist above and apart from those people who currently count among its members. Though the membership of a nation may change through time, referring to some nation at a time is just the same as referring to a list of people—namely, those people who are members of the nation. The only special obligations that nations have are to a particular set of people who all exist at a time. The special obligation the nation has is to the interests or wellbeing of

these people. On this view, any regard that a nation has for future people must be simply derivative of an interest that present people have in the wellbeing of future people (*e.g.* their children and grandchildren). That is, of course, if present people indeed have an actual interest in the wellbeing of these future people at all. The special obligation of nations to act as the caretakers of the interests of future people on this view is thus a very shallow one. This view does not endorse the intergenerational neutrality thesis. It endorses something weaker, such as the claim that we ought to give as much weight to the wellbeing of future generations as is reflected in the interests of present people.

I am not inclined to endorse the first view, but I do not have a decisive argument to give against it. The second view seems to me the right view about the nature of nations' identities through time. On this view, a nation is a collection of people, and since nations persist through time (and indeed across generations), the collection of people that make up a nation must include people at all of those times. Thus, the constituency of a nation includes present and future people.

This view of the nature of nations means that if nations have special ties, they are to people who exist at different times. The national self-interest view demands that a nation remain neutral between its present and future interests. I take it that this means a nation must remain neutral between the interests or wellbeing of the people who constitute it in the present and the future. It thus endorses a particular variant of the intergenerational neutrality thesis. It says that a nation should remain neutral between its own present and future generations.

This view raises a further question about the identity of nations. Certainly, the character of nations changes through time. This has something to do with the nature of the people who are its citizens and also something to do with the institutions that govern it. We could hold different views about the extent to which this matters for the identity of a nation through time, and further views about how this affects the special ties that a nation has to its citizens. One plausible view is that while a nation may have identity through time despite a changing character, at time t the only special ties it has are to citizens of the nation when its character is sufficiently similar to its character at t, or that the strength of the special ties decreases at temporal distance from t. This view seems supported by historical cases. While France may have existed for a very long time, it seems odd to say that the people fighting in the Gallic Wars had any special ties to present-day citizens of France. Another plausible view is that while some continuity of character is a necessary condition for continued identity of a nation through time, only identity is required to generate special ties. A nation on this view has special ties to all of its past, present, and future citizens, and thus to past, present, and future generations. This view claims that the Gauls indeed had special ties to the

present-day French. It might explain the apparent oddity of that claim by saying that since the Gauls would not have known that their history would extend so far into the future, much less anything about the interests or wellbeing of the future citizens of their nation, they did not have reasons to take these future people into account.[21]

While I am inclined to endorse the latter view, I do not really know which of these views is right. There are interesting parallels to questions about personal identity, but I suspect that our intuitions about national identity are much less robust than our intuitions about personal identity. It may therefore be very difficult to argue for a position about national identity from these parallels. Indeed, the question of special ties of nationhood may come apart significantly from the question of the identity of nations through time; it may thus be a mistake to argue from the latter to the former. Which of these views we endorse determines which future generations we think nations have special ties to. If we discount for special ties, holding one view rather than another makes a serious difference.

There may be other self-interest views that support or do not support some variety of the intergenerational neutrality thesis. I won't search for them here. Because nations are the agents that are able to take the action required to mitigate the effects of climate change, the national self-interest view is the one worth reckoning with. But as I have argued, there are a number of plausible ways to interpret the national self-interest view. If we hold a view that does not give nations special ties to generations in the further future, that may justify a higher discount rate.

Whether *any* special ties justify discounting also depends, as I noted, on whether cost-benefit analysis should account for agent-relative valuations of harms and benefits. Whether agent-relative reasons are truly normative is a rather large and difficult question for moral philosophy in general. Thomas Nagel asks,

> For the purposes of ethics, should we identify with the detached, impersonal will that chooses total outcomes, and act on reasons that are determined accordingly? Or is this a denial of what we are really doing and an avoidance of the full range of reasons that apply to creatures like us? (Nagel 1986, p. 185)

[21] This raises another interesting question. If cost-benefit analyses are to take account of agent-relative reasons, as the argument for discounting for special ties suggests, then the cost-benefit analysis of climate change—generating an analysis-given reason for a nation—should consider the possibility that the nation will not exist far into the future. This is distinct from considering the possibility that humanity itself will not exist far into the future, which, as I've noted, the *Stern Review* does.

This is, as he says, a true philosophical dilemma. So is welfare economics, and particularly cost-benefit analysis, the right place to take a stand on it? Perhaps not, but then again perhaps the choice is forced. If we include agent-relative valuations in cost-benefit analysis and are right to do so, then we will increase the force of the analysis-given reason. But if we are wrong to do so, the analysis-given reason will be weakened since it will have weighed some harms and benefits wrongly. If we choose not to include agent-relative valuations in cost-benefit analysis, the risks are the same but in the opposite direction. I don't know which view should be adopted. Whether cost-benefit analysis should account for agent-relative valuations seems to me an open question.

Summary and Conclusion

Harms and benefits are just as bad whether they occur now or in the future, whether they occur to one person or another. This is because harms and benefits are differences in goodness. That a plan will occasion a harm counts against the plan; that it will occasion a benefit counts in its favor. In deciding what to do, we need to add up and weigh harms and benefits. This is what is meant when we say that the purpose of cost-benefit analysis is to *evaluate* a plan of action. This evaluation produces what I have called an *analysis-given reason*.

I have argued, however, that there are a number of grounds on which we might count future harms and benefits for less than present ones. Some of these are not very controversial: economic growth, or, that future people will be better off, gives reason to discount the future value of certain goods. These grounds for discounting have to do with the way that economists use prices to value commodities that count as harms or benefits. Specifically, they are ways of correcting prices so that they accurately reflect value.

But there are other reasons for discounting that are more controversial. They aim to increase the force of the analysis-given reason by changing the way in which certain harms and benefits that come in the future are weighed against harms and benefits in the present. One such reason came from special ties, or self-interest views. These views may give reason to think that a harm or benefit received by a future person may count for less than the same harm or benefit to a present person. Whether this is the case depends on the particular self-interest view adopted.

The cost-benefit analysis of climate change recommends action by nations, which could recommend adopting a discount rate based on the national self-interest view. This view gives nations putatively normative reasons for acting in the interests of their own citizens. I argued for an interpretation of the national-self interest view that supports a version of the intergenerational neutrality thesis. That version

says that a nation should remain neutral between its own present and future generations. However, I did not have a decisive argument against a rival interpretation of the national self-interest view that does not support any version of the intergenerational neutrality thesis, nor against an interpretation that supports a weak version of the intergenerational neutrality thesis. Whether nations ought to support the thesis thus depends on the right view about the nature of the special ties nations have to their members; namely, whether they have a strong obligation to their future members. The issues here are very complex and do not appear to have straightforward answers. That makes discounting for special ties related to nationhood a very difficult task.

We might think that it is immaterial whether any self-interest view supports the intergenerational neutrality thesis since we could think that self-interest views do not give normative reasons for action; we could reject agent-relativity as a source of normative reasons. The intergenerational neutrality thesis is certainly true from an impartial perspective—Sidgwick's 'point of view of the Universe' or Nagel's 'view from nowhere'—since an impartial perspective does not admit of self-interested reasons of any kind. I am sympathetic to this view, but I know of many people who support a national self-interest view. Indeed, nations themselves certainly support it. As I have said, I regard this as a serious dilemma to which there is as yet no clear answer.

I end therefore with a pragmatic question. Cost-benefit analyses produce analysis-given reasons, which are reasons for action. But more than that, they are usually tasked with producing reasons *for someone*. In the case of the cost-benefit analysis of climate change, that someone is nations. If the analysis is to be *true*, we may well think it should adopt an impartial perspective and endorse the intergenerational neutrality thesis. This is to reject discounting for special ties; this is to deny that agent-relative reasons are normative. This is the view I tend toward. On this view, the discount rate should be very low. But normativity may admit of agent-relative reasons, and anyway nations certainly take account of their own agent-relative reasons for action. So if we want the analysis to be *accepted*, we may well think that it should adopt some more partial perspective, such as the national self-interest view. On this view, the discount rate may or may not be very low. Whether it is or not depends on which is the right view of national self-interest.

Since it is true on any accounting that climate change will have some very bad effects, perhaps, as the saying goes, we should not let the perfect stand in the way of the good. Our efforts may be best spent by arguing for the interpretation of the national self-interest view that supports the intergenerational neutrality thesis, rather than by arguing against self-interested views in cost-benefit analysis

altogether.

References

Kenneth Arrow. Discounting, morality, and gaming. In Paul R. Portney and John P. Weyant, editors, *Discounting and Intergenerational Equity*, chapter 2, pages 13–22. Resources for the Future, Washington, 1999.

Wilfred Beckerman and Cameron Hepburn. Ethics of the discount rate in the stern review on the economics of climate change. *World Economics*, 8(1), 2007.

Robin Boadway and Neil Bruce. *Welfare Economics*. Blackwell, Oxford, 1984.

John Broome. Utility. *Economics and Philosophy*, 7:1–12, 1991.

John Broome. *Counting the Cost of Global Warming*. White Horse Press, 1992.

John Broome. Discounting the future. *Philosophy and Public Affairs*, 23(2):128–156, 1994.

John Broome. Cost-benefit analysis and population. *The Journal of Legal Studies*, 29(2):953–970, June 2000.

Tyler Cowen. Caring about the distant future: Why it matters and what it means. *The University of Chicago Law Review*, 74(1):5–40, 2007.

Partha Dasgupta. Commentary: The stern review's economics of climate change. *National Institute Economic Review*, 199(4):4–7, 2007.

Partha Dasgupta and Eric Maskin. Uncertainty and hyperbolic discounting. *The American Economic Review*, 95(4):1290–1299, 2005.

Jean Drèze and Nicholas Stern. The theory of cost-benefit analysis. In Alan J. Auerback and Martin Feldstein, editors, *Handbook of Public Economics*, volume 2, pages 909–989. North-Holland, 1987.

S. Frederick, G. O. Lowenstein, and T. Donoghue. Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40:351–401, 2002.

Ross Garnaut. *The Garnaut Climate Change Review: Final Report*. Cambridge UP, Cambridge, 2008.

Thomas Nagel. *The View From Nowhere*. Oxford UP, Oxford, 1986.

William Nordhaus. A review of the "stern review on the economics of climate change". *Journal of Economic Literature*, 45(3):686–702, September 2007.

Derek Parfit. *Reasons and Persons*. Oxford UP, Oxford, 1984.

Derek Parfit. *On What Matters*. Oxford UP, Oxford, 2011.

Frank Ramsey. A mathematical theory of saving. *Economic Journal*, 38:543–559, 1928.

Nicholas Stern. *The Economics of Climate Change: The Stern Review*. Cambridge UP, Cambridge, January 2007.

Joseph Stiglitz. The rate of discount for benefit-cost analysis and theory of the second best. *SSRN eLibrary*, 1983.

Martin L. Weitzman. Why the far-distant future should be discounted at its lowest possible rate. *Journal of Environmental Economics and Management*, 36:201–208, 1998.

Martin L. Weitzman. A review of "the stern review on the economics of climate change". *Journal of Economic Literature*, 45(3):703–724, September 2007.