

DOES A DISCOUNT RATE MEASURE THE COSTS OF CLIMATE CHANGE?

CHRISTIAN TARSNEY*

Abstract: I argue that the use of a social discount rate to assess the consequences of climate policy is unhelpful and misleading. I consider two lines of justification for discounting: (i) ethical arguments for a ‘pure rate of time preference’ and (ii) economic arguments that take time as a proxy for economic growth and the diminishing marginal utility of consumption. In both cases I conclude that, given the long time horizons, distinctive uncertainties, and particular costs and benefits at stake in the climate context, discount rates are at best a poor proxy for the normative considerations they are meant to represent.

Keywords: climate change, discount rate, environmental ethics, Ramsey formula, time preference

1. CLIMATE CHANGE

The reality and the dangers of anthropogenic climate change are a matter of increasingly robust scientific consensus (Cook *et al.* 2013; IPCC 2014). It is widely agreed that global mean surface temperature (GMST) has begun to increase and that this increase will continue over the coming century. The most recent report of the Intergovernmental Panel on Climate Change (IPCC), for instance, estimates that by the year 2100, GMST will probably have increased by at least 1.5 °C relative to the baseline period 1850–1900, across a wide range of emissions scenarios, but that the increase could exceed 4 °C, especially under more pessimistic assumptions about the path of emissions (Stocker *et al.* 2013: 20).

Nevertheless, climate change remains a matter of enormous predictive uncertainty. On the climate science side, these uncertainties

* Department of Philosophy, Skinner Building, University of Maryland, College Park, MD 20742, USA. Email: ctarsney@umd.edu. URL: <http://www.philosophy.umd.edu/people/ctarsney>.

centre on the *equilibrium climate sensitivity* (the long-term change in GMST that results from a doubling of atmospheric greenhouse gases),¹ the net effect of feedbacks in the carbon cycle, and the effects of a given warming scenario on regional climates. On the social science side, there is of course great uncertainty with respect to public policy (what sorts of abatement measures will be adopted over the coming decades by the world's carbon-emitting economies), the speed at which low-carbon energy technologies will improve in quality and cost-competitiveness, and the effects that a given policy/technology package will have on actual emissions flows. But even if we could assess all the preceding factors with confidence, it would remain extremely difficult to predict the various human effects of a given path of change in global and regional climates. The compounded effect of these uncertainties is that, from our present standpoint, the long-term costs of failing to sufficiently mitigate climate change could range from relatively mild (Tol 2009, 2014) to utterly catastrophic (Weitzman 2009; Sherwood and Huber 2010).

Along with these open empirical questions, debates over climate policy raise important normative questions, in particular because of the very long timeframes involved in climate policy. Since the most appreciable effects of present emissions manifest a century or more in the future, how we ought to respond to climate change depends heavily on the nature and strength of our obligations to future generations.

A widely employed method for managing many of these uncertainties is the application of a *social discount rate* to future costs and benefits associated with policy responses to climate change. Representing the various effects of climate change in the common medium of discounted consumption simplifies difficult questions both about the economic and other human effects of climate change and about the normative significance of those effects. Discount rates therefore figure in nearly all prominent economic analyses of climate change (e.g. Nordhaus 2007; Stern 2008). But a number of philosophers and economists have expressed concerns about the use of discount rates for intergenerational questions in general (Parfit 1984; Cowen ms) and for climate policy analyses in particular (Broome 1994; Stern 2014). In the present paper, I aim to advance the debate over discounting in the context of climate policy by expanding on extant criticisms of the discount rate and by addressing some recent efforts both to justify (Heath 2016) and improve on (Posner and Weisbach 2010; Fleurbaey and Zuber 2013) the use of discount rates in the climate policy context. I will conclude that, although the discount rate is at least in part intended to approximate genuine normative considerations that must influence any evaluation of costs and benefits in

¹ The range given as 'likely' for this figure in the latest IPCC report is 1.5 °C to 4.5 °C (Stocker *et al.* 2013: 16).

the distant future, nevertheless in the particular context of climate policy the discount rate offers such a poor approximation of those concerns that it ought to be discarded and better methods found for representing the considerations that underlie it.

In the next section I introduce the standard formula for calculating the discount rate and describe two broad lines of justification for the practice of discounting, one ethical and the other economic, corresponding to the two terms of that formula. Section 3 examines the ethical justifications for discounting, focusing in particular on recent work by Joseph Heath that attempts to defend a 'pure rate of time preference' against standard philosophical criticisms. Section 4 turns to the economic case for discounting, based on the expectation of growth and the diminishing marginal utility of consumption. I argue that the assumptions underlying this economic rationale are inapplicable to the most important long-term costs of climate change. Finally, in Section 5 I sketch an alternative method for comparing costs and benefits in the climate policy context that aims to keep the nature of the costs and benefits as well as basic normative questions about the comparison of disparate values and the demandingness of our obligations to future generations more clearly in view.

2. THE SOCIAL DISCOUNT RATE

The practice of discounting is meant to express the fact that, in general, we attach less (dis)value to the prospect of certain gains or losses, the further those prospects lie in the future. Most commonly, the idea of temporal discounting will be expressed in the form of a single *discount rate* that gives a percentage by which our concern for a fixed (in some sense) gain or loss diminishes, year over year. If, as is typical, the discount rate is taken to compound at a fixed rate over an unlimited time horizon, even a comparatively low discount rate will instruct us to attach vanishingly little weight to the events of future centuries or millennia.

The discount rate is generally understood to represent the sum of two components, as described by the *Ramsey formula* (Ramsey 1928),

$$r = \delta + \eta g$$

where r is the overall discount rate, δ is the so-called 'rate of pure time preference', η expresses the 'elasticity of marginal utility of consumption' (the rate at which the marginal utility of additional units of consumption diminishes as absolute levels of consumption increase), and g is the expected rate of future economic growth.

Pure time preference is meant to express a discount rate on cardinal utility itself. It reflects the fact that we are often, as Derek Parfit puts it, 'biased towards the near' (Parfit 1984), preferring near-future gains to far-future gains *merely* because they happen sooner. Sometimes (e.g. in Stern

2008) δ is also used to incorporate certain kinds of risk, e.g. the risk of an exogenous global catastrophe that would destroy or so greatly damage human civilization as to render prospective gains and losses moot. I will discuss both these factors at greater length in §3.

The second term of the discount rate is meant to express not a discount on utility but rather the fact that, if future people are richer than present people, the same absolute gain or loss of consumption will matter less to them than it would to us. This reflects the familiar fact that most goods have diminishing marginal utility: one apple a day keeps the doctor away, a second might be a nice afternoon snack, but a third is simply monotonous. Because most goods have diminishing marginal utility, wealth in general (expressed in monetary terms) does as well: the marginal dollar makes a greater difference to the well-being of someone with less wealth or income than to someone with more.

These two terms, so conceived, capture some but not all of the justifications that have been offered for temporal discounting. Others will be explored below. But the discussion that follows will be structured around the two terms in the traditional discount formula. I will argue, for each term, that whatever genuine normative considerations it captures, it captures poorly at best in the context of climate policy. The implication is not that the discount rate for climate policies should be zero, in the sense that the factors behind the discount rate should be ignored. Rather, it is that these factors should be dealt with in a systematically different way, and that representation in terms of discount rates does not helpfully capture their relevance to debates over climate policy.

3. PURE TIME PREFERENCE

This section examines several arguments for a rate of pure time preference that have been offered in the recent literature. While each argument rests on potentially significant normative considerations, it will be seen that these considerations are each essentially orthogonal to the question of discounting.

Perhaps the most common argument for a discount rate in general and pure time preference in particular (e.g. in Nordhaus 2007) is that public policymakers in a democratic society are accountable to the body politic, whose members individually and collectively appear to exhibit pure time preference, insofar as pure time preference is an established feature of the psychology of human decision-making.

This argument is subject to several well-known and compelling objections. First, although it is true in general that individuals discount their own utility, this is often not behaviour they would reflectively endorse. Hume famously spoke of the 'natural infirmity' that led him to prefer the lesser but nearer good to the greater but further, as something

'I may very much regret, and . . . may endeavor, by all possible means, to free my self from' (Hume 2000 [1739]). This attitude toward personal time bias is extremely widespread: more often than not, we would rather be free of it, or at least reduce its effects. It is not so clear, then, that we wish for our agents in government to enact this preference on our behalf, or that they are obliged to do so.

Second, as Hume points out in the same passage, an especially regrettable feature of our individual time preference is that it is *not* simply exponential, as a fixed discount rate would imply:

In reflecting on any action, which I am to perform a twelve-month hence, I always resolve to prefer the greater good, whether at that time it will be more contiguous or remote . . . But on my nearer approach, those circumstances, which I at first over-look'd, begin to appear, and have an influence on my conduct and affections. A new inclination to the present good springs up, and makes it difficult for me to adhere inflexibly to my first purpose and resolution. (Hume 2000 [1739]: 382)

Deriving guidelines for policy directly from revealed time preference would commit us to just such changes of course: prefer the greater good to the lesser when both are in the distant future, then suddenly abandon our resolution as the prospect of the lesser good looms large. But it is at least *prima facie* implausible that public policy should mirror this feature of our individual preferences.²

Revealed rates of pure time preference may also vary substantially between individuals, with age, and across different choices and contexts.³ Even if a rate of pure time preference could be read off the economic

² An anonymous referee points out that time inconsistency can arise from certain attractive axioms sets for social welfare relations over infinite utility streams (that is, assignments of utility to an infinite sequence of generations). The problem of defining such a social welfare relation is extraordinarily difficult. For my part, I am sympathetic to approaches along the lines of Lauwers and Vallentyne (2004), whose 'Full Weak Catching-Up' principle preserves time-consistency (among other desiderata) at the expense of completeness and strong anonymity. Nevertheless, these questions are still far from settled and it would doubtless be a mistake to treat any desideratum, including time consistency, as sacrosanct in infinite contexts until a fully satisfactory theory is arrived at. In any case, the empirical unlikelihood of an infinite number of future human generations provides some reason to doubt that policy choices should be powerfully influenced by considerations that arise only in such contexts.

³ For instance, in their review of the empirical literature on discount rates, Frederick *et al.* (2002) conclude that '[w]hile the [discounted-utility] model assumes that people are characterized by a single discount rate, this literature reveals spectacular variation across (and even within) studies. The failure of this research to converge toward any agreed-upon average discount rate stems partly from differences in elicitation procedures. But it also stems from the faulty assumption that the varied considerations that are relevant in intertemporal choices apply equally to different choices and thus that they can all be sensibly represented by a single discount rate'. (Frederick *et al.* 2002: 352).

behaviour of individuals, it is not clear that this is the same rate they would apply to questions of public policy and intergenerational justice. The popularity of temporal neutrality (i.e. no rate of pure time preference) in academic discussions of discounting in public policy seems to be evidence of this, unless it is found that the advocates of this view are all exceptionally prudent and foresightful in their private lives.

Additionally, it is far from obvious whether the present electorate has the right to simply impose its rate of pure time preference on future persons (as well as those present persons not yet of voting age). As Broome (1994) has argued, future generations seem to be inappropriately disenfranchised when matters of public policy that so profoundly affect their interests are simply referred to the preferences of a majority of present persons.

But there is a more decisive objection to the argument from majority will. Even if it is true that legislators should be entirely subservient to their constituents in setting a rate of pure time preference, this is merely a shell game from the standpoint of substantial debates over public policy. It is no more interesting than the observation that the executive branch of government ought not unilaterally impose a carbon tax without legislative authorization. To the extent that one accepts a 'delegate' view of representation on which political representatives are bound by the preferences or interests of their constituents rather than by their own judgements of the good (McCrone and Kuklinski 1979), questions of public policy are questions about what the people should choose and not what their legislators should choose. But this has no substantive effect on the reasons for or against a particular policy choice – it only redirects those reasons so that they apply to a different set of agents.

The argument from democracy might be reframed so that it concerns not the preferences of a policymaker's constituency but rather their *interests* and the legitimate claims those interests create. One modest aim of climate policy is *efficiency*, specifically the elimination of deadweight loss that is achieved by forcing polluters to internalize their negative externalities. Where deadweight loss exists, it is in principle possible to realize Pareto improvements over the status quo, making at least some people better off while making no one any worse off. But climate policies that aim for impartial utility maximization will be more ambitious, and are very unlikely to represent Pareto improvements relative to the status quo. For one thing, given the diminishing marginal utility of consumption and the scale of global inequality, optimistic policies will likely impose as much of the cost of climate change mitigation as possible on the global rich, and as little as possible on the global poor, leaving the rich at least slightly worse off than they are in the status quo. For another thing, the enormous returns to future generations from mitigating the risk of catastrophic climate change may justify levels of emissions reduction that

would leave the present generation, even the present poor, significantly poorer on balance. If one thinks that policymakers ought to serve the interests of their (present) constituents, however, it may seem perverse to enact policies that make those constituents worse off on balance, for the benefit of non-constituents.⁴

It seems to me, however, that the strongest objections to the original ('majoritarian') argument still apply to this ('fiduciary') argument. First, it is implausible to hold that policymakers ought never redistribute away from their constituents, even when those constituents approve of and consent to such redistribution. For instance, humanitarian aid like famine relief need not meet the litmus test of national interest, providing indirect benefits to the sender country that outweigh its costs, if the electorate of the sender country is simply willing to shoulder some cost for the sake of the humanitarian benefit in the receiver country. So again, even if one holds that policymakers should never redistribute away from their constituents *without their consent*, the question simply becomes whether the constituency in question ought to consent to the potentially redistributive policy.⁵

Second, along the lines of Broome's 'disenfranchisement' argument described above, it does not seem that the present generation has the right to simply maximize its own interests in the context of climate policy. Climate change mitigation concerns not (or at least not primarily) the redistribution of benefits to future persons but rather the prevention of serious harm in the form of droughts, heatwaves, ecological destruction and the like. It therefore implicates *negative* rights claims, claims that we ordinarily take to trump a policymaker's pursuit of the interests of her constituency. While one nation may not be obligated to provide famine relief to another, it is certainly obligated not to *cause* a famine in the other nation for the sake of its own economic interests (or at least, so we intuit). Thus, while aggressive climate policies may constitute redistribution relative to a status quo in which negative rights of future people are being violated, they are not redistributive relative to a baseline in which those rights are respected – they are redistributive, one might say, only in the sense in which compelling a thief to return stolen property or a tortfeasor to pay damages is redistributive.

Finally, as I will argue at greater length below, worries about the acceptability of imposing large costs on a constituency for the sake of non-constituents do not justify a *social discount rate*, but if anything simply

⁴ I thank an anonymous referee for pressing this objection.

⁵ Of course, political libertarians hold that the state ought never redistribute from any of its citizens without his or her *personal* consent, and hence will deny that the consent of the majority is sufficient to legitimate redistribution away from the constituency as a whole. But the next two arguments, I think, should be convincing to libertarians as well.

justify a flat discount on the interests of all non-constituents, regardless of when those interests are realized.

A second argument for pure time preference presents greater difficulties. Joseph Heath (2016) defends a rate of pure time preference for climate policies on the grounds that we may otherwise find ourselves called upon to make extreme and unreasonable sacrifices for the sake of future persons. The argument runs as follows: We can always, at the margin, trade present consumption for savings and investment. Since investment (in capital, research, etc.) produces compounding returns, the undiscounted benefits of delaying consumption in favour of investment will always outweigh the immediate value of consumption, over long enough time horizons. If we aim to maximize undiscounted utility over an unlimited time horizon, then the best strategy would be to restrict ourselves to a Spartan existence while working feverishly to build the infrastructure and discover the technologies that will allow enormous numbers of our descendants to enjoy enormously high levels of well-being.⁶

I will not attempt to examine the empirical plausibility of this story, but it is plausible enough at least that time-neutral maximizing consequentialist theories like classical utilitarianism enjoin on us substantial sacrifices for the sake of future persons (see for instance Beckstead (2013) for a more friendly exposition of this view). The argument notably parallels demandingness objections to maximizing consequentialist theories that arise in the context of present-day concerns like global poverty: if we treat the interests of others as placing moral demands on us that do not diminish in force with (spatial or temporal) distance, we quickly discover that there is a great deal of good that we are called upon to do, but that to do all the good we can will impose on us more personal sacrifice than we can be entirely sanguine about.

It is precisely this parallel, however, that should make us skeptical of the demandingness argument for temporal discounting. There is no reason to suggest adopting a 'spatial discount rate' in order to limit the demands made on us by the distant poor. Indeed, doing so would have undesirable and counterintuitive effects: while most of us do not wish to reduce ourselves, individually or nationally, to the 'point of marginal utility' (as Singer (1972) famously enjoins) in order to relieve global poverty, we might still wish to make some effort in that direction, and in choosing our efforts we are utterly indifferent to physical distance *per se*. An American choosing where to direct her modest charitable efforts should prefer to prevent two cases of malaria in Ethiopia rather than one

⁶ The worry that only a rate of pure time preference allows us to avoid unreasonably high savings rates goes back at least as far as Ramsey (1928). For a more recent articulation of the argument see Arrow (1999: 13–16).

case in Belize, relative distances notwithstanding. Analogously, as Parfit argues (Parfit 1984: 485), if we could choose at the same present cost (and with perfect certainty) either to prevent a minor catastrophe in 100 years or to prevent a much greater catastrophe in 200 years, there is no independent motivation to be found either in theory or in intuition for thinking that we should choose the former. Even if the rate of pure time preference is not held constant into the more distant future (and we accept the predictable reversals-of-course that this would entail), as long as the rate remains above some positive threshold to any future time (rather than asymptotically approaching zero), we can construct cases in which the difference in scale of the catastrophes is arbitrarily large, and yet our discount rate commits us to preventing the smaller, nearer-though-still-distant catastrophe. But it seems an obvious mistake to, say, prevent one loss of limb in a hundred years at the predictable cost of a thousand lives a million (or any other number of) years hence.

The spatial and temporal discounting cases are different, it should be acknowledged, from an egalitarian standpoint: Demandingness worries about present-day transfers across space typically concern transfers to the less well off, where the transfer is a 'leaky bucket' (a net loss of consumption) but nevertheless realizes a gain in aggregate utility due to the diminishing marginal utility of consumption. In contrast, worries about the demandingness of impartial utilitarianism as concerns future generations have an egalitarian tinge: on the assumption of continuing economic growth, future generations (perhaps, in the more distant future, even their poorest members) will be better off than the members of the current generation who are being asked to forgo consumption for their sake. But, it seems to me, this difference does not undermine the analogy drawn in the last paragraph: Worries about *demandingness* create no reason to prioritize the nearer future over the more distant future, just as they create no reason to prioritize the needs of nearer strangers over those of more distant strangers. The difference between the spatial and temporal cases is relevant insofar as it gives rise to an argument for discounting from inequality aversion, based on the expectation that the more distant future will be richer than the nearer future. This sort of argument, however, does not support a *pure* rate of time preference (a positive δ), but rather concerns the second term of the Ramsey formula (ηg). (I will address this sort of justification for discounting in the next section.)

It is better, then, to approach the demands of the distant future just as we approach the demands of the present: we compare undiscounted costs and benefits, acknowledge that a straightforward utilitarian reckoning favours our making significant sacrifices, and then ask in light of the relevant costs and benefits how much sacrifice we are *morally required* to make. However much sacrifice we feel required (or are prepared) to

make, we will try to do as much good as we can by that amount of sacrifice, distances notwithstanding. It is better to decide on a satisficing threshold at which we have met our obligations, rather than distorting the underlying utilities to create the appearance that we are perfect utility maximizers.

There is a pragmatic argument as well against letting demandingness concerns drive the choice of discount rate. Any prescription that fully incorporates these concerns is, by definition, no longer over-demanding and hence not supererogatory. It represents the non-negotiable, no-matter-what moral minimum we owe to others. But there is significant risk of this fact being lost in transmission, when economists and other academics make recommendations to policymakers. Since these policy recommendations are, at best, likely to be significantly weakened in practice, one ought to recommend the moral optimum rather than the moral minimum if it is important that the moral minimum be met.

A third argument for pure time preference, however, pushes the assumptions of the demandingness argument a step further and suggests a more principled objection to pure utility maximization: namely, that it might result in an 'indefinitely postponed splurge', a policy by which consumption is continually forgone in favour of greater investment, producing growth that *never* pays off as consumption or utility.⁷

The in-principle possibility of such a trap is illustrated by the following case (a variant of the 'Trumped' paradox discussed in Arntzenius *et al.* 2004): A clerical error is made at the time of your death, as a result of which you find yourself standing before the decidedly unpearly gates of Hell. You are met by Satan, who apologizes profusely, explaining that you should have been sent to Limbo, there to pass eternity in a state of mild contentment. He tells you that he has arranged to compensate for the inconvenience with a pass good for one day in Heaven, before sending you on to your intended destination. You gratefully accept his offer and make for the exit. But before you can leave, Satan points out to you that there is a loophole in the celestial day-pass system: for every day wrongfully spent in Hell, the value of your pass will double. That is, if you can endure the sufferings of the damned for a full day, you'll be entitled to two days in Heaven (a day in Heaven being exactly as good as a day in Hell is bad); if you stay for two days, you'll be entitled to four days in Heaven, and so on. Of course, you recognize the trap, but as a rational utility maximizer there is no way out: by iteratively eliminating suboptimal strategies, you end up condemned to eternal torment, running up the value of an investment vehicle you'll never cash in, and thereby forgoing eternal contentment.

⁷ This 'paradox of the indefinitely postponed splurge' is due to Koopmans (1967).

Heath (2016) suggests that this sort of conundrum, in only slightly less torturous form, might confront us when considering real problems of intergenerational allocation. A strategy that maximizes present growth at the expense of present consumption, for the sake of later consumption, will always offer expected gains over any strategy of non-minimal present consumption, but will be in turn dominated by another strategy that pushes consumption still further into the future, and so on *ad infinitum* (Heath 2016: 14–16). Heath offers no very concrete picture of what such an investment strategy might look like over the long run, but in the limit it could mean investing all our resources in building more and faster spaceships so that we can colonize greater portions of the universe, only to use their resources for further breakneck expansion.⁸ Of course, given our present understanding, even this process could not continue literally forever, and so the indefinitely postponed splurge would turn into merely a very long-postponed splurge (postponed until, say, the oncoming heat death of the universe negates the net gains of further expansion and tells us that the time has come to start consuming). So to think that we are really subject to this paradox requires that we accept more outré hypotheses concerning the physical possibilities for growth. But perhaps the mere logical possibility of such a conundrum requires some constraint on our attitudes toward future utility.

Like the broader demandingness worry, though, the problem raised by the paradox of the indefinitely postponed splurge is not a problem about time *per se*, and temporal discounting is not a principled solution. Similar puzzles arise when considering unbounded payoffs at a time, e.g. in the St Petersburg Lottery.⁹ More crucially, though, discounting is at best an unreliable solution to these worries, and depending on it may commit us to intuitively unacceptable results. Suppose that, to escape from Hell, you adopt a discount rate of 51% per day, and thus manage to leave Satan's offer on the table with a clear conscience. This strategy is derailed if Satan ups the ante by letting the value of your pass *triple* for each day you remain in Hell. No discount rate you can adopt will be immune from such an out-pacing manoeuvre. This suggests, at the very least, that discount rates do not get at anything like the *explanation* for our wish to avoid indefinitely postponed splurges or excessive moral demands more generally.

⁸ Or rather, *nearly* all our resources beyond some minimum required to sustain human life as well as the basic infrastructure and institutions that are requisite to pursuing any more ambitious project.

⁹ Granted, it is plausible to think that there is some upper bound on the utility that one can actually *experience* at a time. But if there are also relatively strong physical constraints on aggregate future utility imposed e.g. by the inevitability of heat death, it is not clear that this amounts to a difference in principle.

But moreover, the sorts of infinite-time-horizon scenarios that give rise to indefinitely postponed splurges place the would-be discounter between a rock and a hard place. When you turn down Satan's offer by invoking your 51% discount rate, all he has to do is offer a different deal: one *immediate* day in Heaven, followed by an eternity in Hell. At your chosen discount rate, this represents an improvement over the eternity in Limbo (assuming that the utility of a day in Heaven is at least 100 times greater than the utility of a day in Limbo). Since raising your discount rate will only make this new offer more appealing, and lowering it will lead you back into the original trap, reliance on discounts leaves you hemmed in: Satan can place a pair of offers on the table such that, whatever discount rate you adopt, you wind up spending an eternity in Hell. Temporal discounting, then, is not the way to escape the problems of infinite time horizons and indefinitely postponed splurges.¹⁰

Similarly, in the public policy context, if we face unlimited time horizons and opportunities both for indefinitely compounding investment and for present gratification with open-ended harmful consequences, no choice of discount rate will allow us to escape both of these bad options. Since discounting already produces counterintuitive results in cases like Parfit's, we are better off simply dealing with demandingness and unbounded utility problems on their own respective terms.

One potential escape from these problems is to claim that the interests of future persons (not yet in existence) or merely potential persons (whose

¹⁰ Landesman (1995) discusses this sort of problem in the context of charitable investment, and suggests that the appropriate response is to adopt a satisficing threshold whenever you find yourself faced with an option set containing no best option. If demandingness concerns force us to think in terms of satisficing anyway, then this move becomes all the more natural. See also Moller (2006: 244–245), who argues that, although there are powerful reasons to save for the sake of the distant future, we are unlikely to find ourselves in the conditions (of scope for fruitful philanthropic investment over an unlimited time horizon) that might force naïve utility maximizers into an indefinitely postponed splurge.

It is worth noting that other, non-temporal forms of discounting can avoid the paradox of the indefinitely postponed splurge without falling prey to any reciprocal temptation toward an imprudent immediate splurge. For instance, Zuber and Asheim (2012) defend a *rank*-discounted utilitarian criterion for comparing intergenerational utility streams that amounts to a prioritarian weighting of the utilities of generations enjoying different levels of wealth. Nothing I say in this paper will constitute an argument against prioritarian inequality-aversion or, therefore, against rank-discounting of the interests of groups enjoying different levels of consumption/well-being. In the next section, however, we will find reasons to doubt that a social discount rate for climate policy can reasonably approximate these prioritarian concerns, even given an expectation of future economic growth, *inter alia* because much of the harm from climate change is likely to be borne by the poorest members of future generations whose interests should not be discounted based on the wealth of their generation as a whole. I thank an anonymous referee for encouraging me to address this point.

eventual existence or nonexistence will be determined by our own present choices) have greatly diminished weight or no weight at all from the standpoints of morality and public policy. While examination of such views is beyond the scope of this paper, suffice it to note that indifference to future persons is not equivalent to a social discount rate, and that over the timescales of centuries implicated by climate policy, a discount rate will be an exceptionally poor way of approximating our unconcern for persons-not-yet-present. Over short timescales, a discount rate might very roughly approximate the rate at which present members of the population die off and are replaced. But attaching lower weight to the interests of future persons than to those of present persons cannot explain why we should *continue* to apply the same discount rate beyond, say, the 90 or 100 year timeframe within which nearly all present persons will have been replaced. In doing so, we will end up giving the interests of people 200 years hence much less weight (or the interests of people 100 years hence much more) than we intend. And in any case, if we are decided to draw a distinction between present and not-yet-present persons, the rate at which present persons are expected to die off over the next century is not so difficult a thing to represent in our models that we need to resort to an approximation as crude as a social discount rate.¹¹

Heath proposes a final argument for a pure rate of time preference involving, as he puts it, the way in which abstract moral principles are 'institutionalized'. The argument, in brief, is this: We may believe, as a matter of moral principle, that spatial distance is morally irrelevant (e.g. that it is no less important to relieve distant suffering than to relieve nearby suffering), yet nevertheless approve of institutional arrangements (like the division of the world into states or the division of states into provinces and municipalities) that distribute responsibilities geographically, giving public institutions greater responsibility for the needs of those within their prescribed territories, and perhaps more responsibilities for nearby neighbours (the next town over) than for those on the other side of the world. Possible reasons for favouring this kind

¹¹ Heath writes that he 'will be ignoring the distinction between arguments for discounting that discount for time because they regard the passage of time as morally significant, and arguments that discount for time because time is a good proxy for some other phenomenon (such as people passing in and out of existence) that is regarded as morally significant', claiming that such a distinction amounts to 'unproductive hair-splitting' (Heath 2016: 11n). The distinction he dismisses strikes me as important even in principle, but this may be just a measure of one's taste for a certain variety of hair-splitting. In any case, the key point is that the discount rate can only be justified as a proxy for some other morally significant phenomenon if it is in fact a *good* proxy for that phenomenon. The general theme of the objections raised in this paper to both ethical and economic justifications for discounting is that, given the timeframes and particular costs and benefits at stake in the climate context, the discount rate is a poor proxy for the phenomena that it might be used to stand proxy for.

of scheme include the greater costs of rendering aid over a distance and the greater knowledge that institutions and their agents may be expected to have of local and neighbouring populations compared with distant populations (Jackson 1991: 473-5). Likewise, then, even if we judge that *temporal* distance is irrelevant as a matter of pure moral principle, we might still endorse a social discount rate as a feature of institutional cost-benefit analysis, since logistical and informational constraints make it easiest for our public institutions to do good for the present generation, harder to do good for the next generation not yet born, and harder still to do good for the further distant future, suggesting a division of labour under which each generation should be mainly responsible for itself and the generation or two after it.

There are at least two objections to this argument in the context of climate policy. First, to reiterate a point made earlier in this section, what may be plausible with respect to the provision of benefits is much less plausible with respect to the infliction of harm. A state's responsibility to provide disaster relief, say, may be strongest with respect to its own citizens and diminish with distance beyond its borders, but its responsibility not to *cause* disasters cannot plausibly be said to diminish with distance. Climate change mitigation concerns responsibilities of the latter sort: The hurricanes, droughts, floods, heat waves, and various downstream consequences thereof that climate change may give rise to in the next century are not simply unlucky calamities that we might generously undertake to assist our great-great-grandchildren in mitigating. They will be, at least in part, calamities that we have inflicted on them through our actions.

Second, the considerations that count in favour of an institutional localization of responsibilities – namely, uncertainty about distant consequences and the relative cost of doing good for more distant times and places – are already included in economic analyses of climate policy (e.g. Nordhaus 2007; Stern 2007). Uncertainties about the effects of climate change and the efficacy of proposed mitigation policies, and the present economic costs that those policies create in order to generate a given increment of future benefit, are both explicit features of cost-benefit analysis and of any economic model by which a proposed climate intervention might be judged. To treat these considerations as justification for a social discount rate, then, amounts to double-counting. Moreover, it is unclear that these considerations have much force anyway in the climate context: while there is uncertainty about the *magnitude* of the harms that future generations would be spared by timely reductions of greenhouse gas emissions, there is very little uncertainty that some harm would be averted. And indeed, in making climate policy, it is both cheaper and more certain to aim at averting harms in the more distant future (say, 150 years from now) than in the nearer future (say, 50 years from now): Even

extremely aggressive and costly climate policies would probably have little effect within the 50-year timeframe, while even fairly modest climate policies are likely to avert substantial harm in the 150-year timeframe. So the considerations that might normally favor an institutional policy of greater responsibility for the nearer future are, if anything, inverted in the context of climate change to favour greater attention to the more distant future. As with demandingness, then, a rate of pure time preference is not a good proxy for considerations of institutional competence and uncertainty.

The arguments for δ as pure time preference, then, are unsuccessful. But what of the thought that δ can be seen instead as a measure of exogenous global catastrophic risks (Stern 2008)? As long as the risks in question are truly exogenous to our policy choices, and roughly constant with time, I have no great quarrel with someone who wishes to treat them as a discount rate. But very few risks fit this profile: the most common example of exogenous risk, for instance, namely the threat posed by large Earth-crossing asteroids or comets, is subject to mitigation in varying degrees depending on rates of economic and technological progress and hence is dependent to some extent on our policy choices. The best candidates for risk factors that are truly exogenous over the relevant timeframes (e.g. nearby supernovae or gamma-ray bursts) are all extraordinarily unlikely on an annual basis, and so justify at most a negligible discount rate.

4. GROWTH AND CONSUMPTION

The second term of the Ramsey formula, ηg , measures the diminishing marginal utility of consumption over time as a result of expected economic growth. Because the rate of economic growth measures the average consumption of future people relative to present people, it also provides a rough approximation of the utility gains to be had by shifting potential future consumption to the present.

This idea is sometimes expressed by the metaphor of a 'leaky bucket' for wealth transfers: if we could simply transfer consumption from rich to poor, all else being equal, how much of that consumption could be lost in the transfer while still realizing a net gain in utility or well-being? Plausibly, net gains could be realized even with very high leakage rates: the utility a multibillionaire derives from the marginal million dollars of added consumption might be comparable to the utility of just a few dollars of added consumption, to someone below the global poverty line. Likewise, a single loaf of bread or article of simple clothing might be worth more to the average 15th-century European than many such loaves or articles would be to the average contemporary European. The many-fold increases in mean and median consumption levels brought about by

economic growth mean that, in developed countries, even the poor of today are rich by the standards of the moderately distant past, and likewise the poor of the coming centuries may be rich by our own standards.

Despite this close analogy between synchronic and diachronic wealth disparities, those who argue for the application of 'leaky bucket' reasoning to intergenerational utility comparisons and discounting are seldom eager to apply the same logic to present-day transfers. A fairly standard η of 2, for instance, if applied uniformly to all questions of public policy, would dictate an extraordinarily high rate of transfer from the global rich to the global poor, even at great costs to total consumption (Stern 2008: 53–54).

But this is by the by. *Modulo* demandingness intuitions, one might accept this implication and remain consistent. Or one might argue that the cases are distinct in other ways, for instance that enormous public transfers to the global poor would violate negative constraints or the distinctive responsibilities of rich governments to their own citizens.

The real reason, I will argue, why this second term of the Ramsey formula should not be used to fix a discount rate for the evaluation of climate policy is that, in the climate context, it either fails drastically to measure the relevant utilities, or must be offset and gerrymandered so extensively to produce a plausible measure that it loses any simplifying value.

No one will deny, of course, that the Ramsey formula in its pure form relies on idealized assumptions about economic growth, consumption and utility. But it is worth seeing just how many such idealizations are involved. For ηg to correctly measure the effects on aggregate welfare of policies that trade present for future consumption or vice versa, the following conditions (*inter alia*) would have to hold: (i) a constant, globally uniform rate of economic growth; (ii) total consumption distributed equally across all persons at a given time; (iii) identical functions from consumption to utility across individuals at a time; (iv) identical functions from consumption to utility across times; (v) smoothly exponential diminution of marginal utility with increases in consumption (measured by η); (vi) stability in the relative prices of consumables (which would require, presumably, equal rates of growth for all existing goods, no fundamentally new goods, and no disappearance of goods); and (vii) only marginal changes in growth and consumption (that is, changes small enough to be subject to marginal analysis) as a result of policy choices, as well as no relevant changes in the inputs to well-being apart from consumption of those economic goods included in measures of economic growth (meaning *inter alia* no changes in population).¹²

¹² For now, I assert without argument that these idealizations are contained in the second term of the discount formula. This claim will be explicated as I examine the individual idealizations below.

The fact that none of these assumptions come even close to obtaining, for any significant length of time, at any time and place in human history, should at least motivate some *prima facie* suspicion that the outputs of the discount formula might be misleading. Since reality diverges from many of these idealizations more significantly over longer periods of time, that suspicion is strongest when the discount rate is used to consider shifts in consumption across decades or centuries. Of course, any economic model simple enough to be useful must involve some degree of idealization, and in many contexts even wide mismatches between theoretical assumptions and reality are compatible with theoretical *outputs* that provide a reasonably close and predictively useful approximation of the relevant realities. But, I will argue, there is good reason to believe that in the climate policy context, the unrealism of the discounting model is not mere benign inelegance but results in significant misrepresentations of the likely costs of climate change.

4.1. The growth rate

Rates of economic growth, of course, are not globally uniform but vary dramatically between and within regions. But even if we focus solely on the global aggregate rate of economic growth, there are grounds for substantial uncertainty and for predicting that this rate will not remain uniform over long timeframes. Expectations of long-term global growth vary significantly, even ignoring extreme possibilities like the collapse of civilization or a positive technological singularity. Some economists (e.g. Cowen 2011) expect growth to remain relatively anaemic over the coming decades, while others (e.g. Baily *et al.* 2013) expect it to be far more robust. Depending on the length of recent history from which one chooses to extrapolate, historical growth rates could presage either an economic and technological plateau (extrapolating from the past decade) or continued acceleration to new and unprecedented future growth rates (extrapolating from the past several centuries). Even if these uncertainties could be smoothed out into a single expected growth rate for some interval of time, there is no obvious reason why we should expect the same rate of growth over, say, 20-year and 200-year time horizons.¹³

¹³ A more fundamental objection to any cost-benefit analysis that relies on estimates of long-term growth rates is that comparisons of GDP and other aggregate economic measures across long enough timeframes are simply meaningless. These measures are useful for year-to-year comparisons where the major changes in economic output are quantitative rather than qualitative—a larger or smaller basket of the same goods. But as the basket of goods changes and fundamentally new goods are introduced, these comparisons lose validity. Thus, talk of economic growth rates on timescales of centuries may not be tracking anything at all, let alone anything that has to do with human well-being.

Of course, climate change itself may significantly impact global economic growth and could conceivably result in long-term economic contraction. This speaks to the fact that climate policy choices do not represent marginal changes around a baseline, a point to which we will return shortly.¹⁴

In any case, assuming no rate of pure time preference, the discount rate is simply a linear function of the expected rate of future economic growth (with η as its coefficient), and given that even very small changes in the discount rate can make an enormous difference to how we evaluate large costs to be borne decades or centuries hence, the enormous uncertainty involved in predicting economic growth rates on a scale of decades or centuries, coupled with the unlikelihood that these rates are truly exogenous from our climate policy choices, are already good reasons to be wary of discounting in the climate context.

4.2. Inequality

Global consumption is not now equally distributed, nor is it likely to be in the foreseeable future. Thus, even if the future is richer than the present, this does not guarantee that the future beneficiaries of climate change mitigation will enjoy higher consumption levels than the present people who bear the costs of those policies. Indeed, an important argument for aggressive climate change mitigation is that the costs will be borne largely by the present rich, without great loss of utility, and will stave off significant dangers to the future poor.¹⁵

In two recent papers, Posner and Weisbach (2010) and Fleurbaey and Zuber (2013) attempt to incorporate growth rate variability and inequality into the discount framework. Posner and Weisbach argue that, given uncertainty about economic growth, long-term discount rates aggregated over those uncertainties will converge with the discount rate for the most pessimistic scenario alone. Fleurbaey and Zuber extend this reasoning to considerations of inequality, reaching the conclusion that over the long term discount rates will converge to whatever rate is appropriate for the least-well-off future beneficiaries of the policy, under the worst possible scenario. Thus, long-range climate change mitigation policies ought to be treated much like transfers to the present-day global poor (since, in

¹⁴ It is also worth pointing out that we should not make the mistake of taking growth in aggregate consumption as the relevant variable, in any case, but rather growth in *per capita* consumption (since more people consuming at the same rate does not decrease the marginal utility of an additional unit of consumption).

¹⁵ I take no stance on whether this argument is correct. One might certainly worry that the sort of policy needed to effectively reduce greenhouse gas flows would slow economic growth substantially in the developing world and harm the present or future global poor.

the worst-case scenario, the future least well off will remain at about this consumption level).

These claims are no doubt correct as far as they go, but they do not rescue the discount rate from its defects. Rather, they illustrate its essential futility in this context as an analytical tool. Once the assumption of a constant, positive rate of economic growth has been abandoned, time is no longer a proxy for a relevant variable (namely, consumption levels) in calculations of expected utility, and attempts to weight the value of transfers in terms of time will be theoretically and practically misleading.

Suppose that I can transfer a unit of consumption either to A, who will receive it in 50 years, or to B, who will receive it in 100 years. Suppose that both A and B will be fairly poor, subsisting at only half my present consumption level, and that there are no other relevant differences between them. And suppose $\eta = 1$, so that without leakage the transferred unit of consumption will have twice the value to its recipient as it would have had to me. We can then calculate 'discount rates' for the transfers as follows: for person A, $r = 'g' = .5^{1/50} \approx .986$; and for person B, $r = 'g' = .5^{1/100} \approx .993$ (implying annual discount rates of -1.4% and -0.7% respectively). When we *apply* these discount rates, however, we simply cancel the exponent, and reach the unsurprising conclusion that the transfers have identical net utility, and that (absent pure time preference) the choice between them is a matter of indifference. Given that, in Fleurbaey and Zuber's analysis, the 'discount rate' for transfers of consumption to the least-well-off future persons under the least optimistic consumption scenario turn out to dominate the overall discount rate, and the consumption level of those potential-least-well-off does not change over time in expectation, the discount rate will simply come out in the wash in the manner just described.

Of course, there are optimistic growth scenarios in which the position of the least well off (or at least, of the potential beneficiaries of climate change mitigation) is substantially improved relative to their present counterparts. The probability distribution over possible consumption profiles thus changes with time. First, though, the results cited above suggest that these changes are largely irrelevant: the worst-case consumption scenario tends to dominate the discount rate, and the character of this worst-case scenario does not change substantially over time.¹⁶ And second, while summing over scenarios characterized

¹⁶ The worst-case scenario for the worst off cannot get any worse with time, because hundreds of millions of people already live at or very near the consumption 'floor': it is difficult to imagine a unit of consumption having greater marginal utility among some future population than it does among the poorest of the present-day global poor. And it is surely possible (though, one hopes, far from certain) that a significant portion of the world population will remain at this level of bare subsistence for a significant time to come.

by different constant growth rates may be mathematically tractable (assuming we can come up with meaningful probability estimates), we still face the problem that on the most plausible scenarios, expected growth rates may *not* remain constant.

4.3. Utility functions

The third of the idealizations listed above (identical functions from consumption to utility between individuals at a time) is difficult to avoid, and not a distinctive defect of discount framework. But the fourth (identical utility functions across times) and fifth (smoothly diminishing marginal utility describable by a constant η) deserve attention. There is substantial debate among economists, psychologists and philosophers as to the empirical connection between consumption and utility. Much depends on the choice between hedonic and preference-based notions of utility. It has been famously argued (Easterlin 1995) that above a certain level, uniform growth within an economy has no measurable effect on measured happiness or well-being, which could be taken to suggest that consumption gains above a certain level should simply count for nothing.¹⁷

But it is also not unreasonable to imagine that, over the timeframes implicated by climate change, the relationship between consumption and utility might undergo substantial changes. Psychological or pharmaceutical interventions might substantially increase overall levels of happiness/satisfaction (Bostrom 2009), either rendering overall consumption levels irrelevant (such that η is extremely large or infinite above a certain consumption level) or allowing us to derive more enjoyment from our consumption (thereby decreasing η). In considering states of the world 100 or 200 years from now, it is unreasonable to ignore such considerations and to assume the sort of simple relationship between consumption and utility required by the Ramsey formula.

4.4. Substitution and inelastic goods

The Ramsey formula treats the marginal utility of consumption as diminishing smoothly, at the same rate for all categories of consumables, as total wealth increases through economic growth. But this assumption can fail, *inter alia*, when some significant category of goods is relatively supply-inelastic (meaning that large increases in price beget only small increases in supply) and weakly substitutable (meaning that there are no substitutes that can readily assuage demand for goods of the category in question). With respect to such inelastic and weakly substitutable goods,

¹⁷ This phenomenon, the so-called Easterlin paradox, remains a matter of significant empirical dispute (see e.g. Hagerty and Veenhoven 2003).

even as total wealth rises the marginal utility of consumption may remain quite high. If the inputs to the good become more scarce over time, shifting the supply curve to the left, then this effect is exacerbated and, because there is less of the good to consume, the marginal utility of its consumption may *increase* even as total wealth increases.

Many of the environmental goods endangered by climate change – e.g. natural beauty, species diversity, and the various secondary goods derived from these – fit the above pattern. They are supply-inelastic: The price consumers are willing to pay to vacation on a tropical island may increase substantially with economic growth, but the supply of tropical islands cannot be expected to increase commensurately.¹⁸ And they are only weakly substitutable: The values we derive from pristine wilderness and healthy ecosystems generally cannot easily be substituted, e.g. by other forms of entertainment or recreation.

Further complicating the picture is the difficulty of pricing many environmental goods. Some features of the natural environment are, at least in part, public goods (non-rivalrous and non-excludable), meaning that market prices will tend to understate their value. The value we derive from the beauty of the Arctic wilderness, for instance, comes partly in the form of ordinary economic goods like cruise holidays and nature documentaries. But it is very plausible that we also derive value simply from our *awareness* that places like this exist. Because this awareness has the characteristics of a public good, there is no market price by which to measure it. For different reasons, the value of species diversity and robust ecosystems is difficult to measure simply because we don't yet know with any degree of confidence what goods (e.g. medical discoveries) may be derivable from the existing stock of species or, on the other hand, what negative downstream consequences may ensue from the loss of potential keystone species or the collapse of terrestrial and marine ecosystems. The existence of certain species also seems to have a public-good-like dimension: If tigers go extinct in the wild, for instance, many would judge that we have suffered a loss not captured in the bottom lines of the tourism or nature film industries.

The point for our purposes is that neither the easily measured dimensions of environmental goods, such as tourism, recreation and entertainment, nor their difficult-to-measure values as public goods, sources of discovery, and keystones of the wider ecosystem are likely to exhibit diminishing marginal value with increasing wealth as do other goods and services. This claim seems to be borne out by history: In the past century, our collective willingness to pay to preserve wilderness and

¹⁸ More precisely, the supply of beachfront resorts and even beaches may respond significantly to rising demand, but it is capped by the relatively (though not completely) inelastic supply of tropical island coastline.

prevent species loss has increased dramatically, from almost nothing to billions of dollars a year in the developed world, despite the increasing abundance and affordability of other forms of entertainment. Though it is a difficult question to answer empirically, it seems plausible that the marginal value (and not just the marginal price) we attach to, for instance, preserving a given region of wilderness or endangered species has at least remained constant, if not increased, over the last century of economic growth.

It remains a matter of great uncertainty to what extent climate change may contribute to the destruction of natural ecosystems and the loss of species, and to what extent ecosystems will adapt effectively to rising temperatures and other environmental changes. Nevertheless, extant research suggests that the negative effects on both species and ecosystems will be substantial (e.g. Thomas *et al.* 2004; Thuiller *et al.* 2005; Parmesan 2006; Carpenter *et al.* 2008). While stocks of environmental goods are not simply doomed to inevitable decline, and in many cases may be substantially replenished by the investments that growing wealth enables (e.g. in reforestation efforts), there is nevertheless reason to think that one of the significant consequences of climate change will be losses in environmental goods that cannot be easily offset and will not diminish in marginal significance as a result of GDP growth.

4.5. Non-marginal changes

All of the above are reasons why, in the context of climate policy, the discount rate does not serve its intended purpose: namely, to compute the impact on aggregate utility of marginal gains and losses in economic goods. But even if it did serve this purpose, it would still be largely irrelevant to debates over climate policy, because the changes brought about by any significant climate policy (in particular, the sort of global emissions control scheme needed to effectively curb greenhouse gas flows) are decidedly non-marginal. Comparison of these policies involves choice between qualitatively different global scenarios, differently shaped growth paths, and potentially different population sizes, none of which are accommodated by marginal analysis.

To begin with, climate change poses a variety of direct and indirect risks to human life: natural disasters, drought and famine, the spread of tropical disease further from the equator, increased conflict over scarce natural resources, and deaths from increasingly frequent and intense heat waves. Loss of life does not represent a loss of consumption spread evenly across a wealthy population: rather, the greatest cost in terms of consumption is borne by a single individual whose future consumption is reduced to zero. Furthermore, like the environmental goods discussed in the last subsection, years of healthy life are a highly non-substitutable

good, both for the person living the life in question and for their friends and loved ones. As individuals, we are willing to forgo a great deal of highly valued consumption or substitute inferior alternatives (spinach instead of french fries, exercise instead of television) in order to avoid an early demise, and as with environmental goods our willingness to make these tradeoffs does not decrease but increases as we grow wealthier. Likewise, to others, the death of a loved one represents a difficult-to-quantify but obviously substantial disutility the size of which does not diminish with increasing wealth. Therefore, with respect to another of the largest potential costs associated with climate change, GDP growth and the diminishing marginal utility of consumption do not justify a discount rate since they do not reduce the significance of those costs.¹⁹

This is just to say that lost lives, and changes in population generally, are not subject to marginal analysis because the effect on well-being does not consist, even to a rough approximation, of marginal decrements to individual consumption. To take the point a step further, the ultimate non-marginal effect of climate policy is a change in the risk of an eventual globally catastrophic outcome, up to and including human extinction. In a recent paper, Martin Weitzman (2009) offers the rough estimate that, accounting for carbon cycle feedbacks, there is at least a 5% chance of long-term GMST increases from climate change in excess of 11 °C, and at least a 1% chance of increases in excess of 22 °C. Even an 11 °C increase in global mean surface temperature could render much of Earth's surface uninhabitable to human beings (Sherwood and Huber 2010), making human extinction a serious possibility. Smaller temperature changes, while not directly threatening human survival, might still increase the likelihood of other scenarios for extinction (e.g. from disease or armed conflict).

While it remains relatively unlikely that the effects of climate change will reach this level of severity, extinction risks on the order of 1–5%

¹⁹ To put the point slightly differently, a loss of the same quantity of consumption will generally produce a much greater loss of utility when it comes about through a decrease in lifespan rather than a reduction in consumption per year over a fixed lifespan. Compare three possibilities: (i) \$50 000 of consumption per year, generating 100 utiles of well-being per year, over a lifespan of 100 years; (ii) \$40 000 of consumption per year, generating 90 utiles of well-being per year, over a lifespan of 100 years; (iii) \$50 000 of consumption per year, generating 100 utiles of well-being per year, but over a lifespan of only 80 years. Possibilities (ii) and (iii) represent the same loss of consumption relative to (i), but (iii) represents a greater loss of utility than (ii), because the diminishing marginal utility of consumption blunts the impact of the lost consumption in (ii) (a 20% reduction in consumption generating only a 10% reduction in lifetime utility) but not in (iii) (a 20% reduction in consumption through lost lifespan generating a 20% reduction in lifetime utility). Thus, the diminishing marginal utility of consumption might provide a justification for discounting losses of the sort illustrated by possibility (ii), but cannot justifying discounting losses of the sort illustrated by possibility (iii).

(and indeed, well below this level) ought to be taken very seriously. As Bostrom (2013) points out, absent pure time preference (or other reasons to severely discount/ignore the interests of potential future people as such), even very small changes in the risk of extinction simply overwhelm more ordinary consequentialist considerations. He estimates (p. 18) that a future version of humanity might be able to generate up to the equivalent of 10^{54} subjective life-years of experience, at levels of utility substantially higher than our own (although, here again, much depends on one's notion of utility, and Bostrom's is tendentiously hedonic). Even a vanishingly small chance of this outcome makes even vanishingly small changes to the likelihood of near-term extinction many orders of magnitude more important, from an aggregative consequentialist standpoint, than nearly any other goal imaginable.²⁰ A significant rate of pure time preference might allow us to ignore these considerations, but discounting for the diminishing marginal utility of consumption is simply misleading – if applying a discount rate changed things enough to make any practical difference, we would know at once that the discount rate was too high.²¹

These considerations are salient because they drive home the point that the standard use of discount rates cannot be made compatible with standard aggregative consequentialism. To the extent that economists have anything like a classical utilitarian ethical understanding in the

²⁰ As Bostrom puts it: 'Even if we give this allegedly lower bound on the cumulative output potential of a technologically mature civilisation a mere 1 per cent chance of being correct, we find that the expected value of reducing existential risk by a mere *one billionth of one billionth of one percentage point* is worth 100 billion times as much as a billion human lives' (Bostrom 2013: 19).

²¹ I can see three ways of resisting Bostrom's argument inside a broadly consequentialist framework. First, one can go the empirical route by claiming that the likelihood of extinction from runaway climate change is so low, the likelihood of exogenous extinction so high, or our ability to change these likelihoods so minimal that the potential magnitude of the stakes is nullified. None of these claims seem plausible, given the relevant magnitudes. As long as epistemic (decision-relevant) probabilities of these claims are not unreasonably extreme, the expected utility of reducing existential threats from climate change will remain overwhelmingly large.

Second, one can appeal to demandingness considerations. If one simply refuses to believe that ethics could demand anything from us as significant as a couple of percentage points of forgone economic growth to reduce extinction risk, then this argument is unanswerable. But I have already argued that demandingness concerns do not justify a discount rate.

Finally, one might deny that the interests of the merely potential future people whose existence extinction would preclude carry any weight at all. Consideration of these so-called 'person-affecting' views lies beyond the scope of this paper, but it is enough to reiterate that this view, even more clearly, will not support a general discount rate. The effects of our activities on future people, or the effects of preventing some future person from being born (including lost consumption), will simply be ignored. This indifference does not increase with time, and it has no effect on how we treat whatever future consequences we do care about.

background, the use for example of market rates of return as off-the-shelf discount rates simply ignores the most salient feature of the climate change problem. If the proponents of economic models that apply discount rates to the costs of climate change are *not* classical utilitarians, then it is worth asking for an articulation and defence of the alternative normative framework that motivates the discount approach. As the previous section has argued, supplementing utilitarianism with a rate of pure time preference will not do the trick.²²

5. AN ALTERNATIVE TO DISCOUNTING

Any scenario in which the major effect of climate change is a diminished supply of the generic, substitutable consumables subject to Ramsey-style discounting is one in which climate change turned out to be extraordinarily and unexpectedly benign. In this future, and on anything like the assumptions that make the Ramsey formula applicable, these losses will not count for much—the world will still contain plenty of stuff. We are worried about climate change precisely because of the effects that do not fit this profile, that are non-marginal and do not diminish in significance as global consumption increases: disease, displacement, loss of basic goods among the still-impooverished, degradation or destruction of irreplaceable environmental goods, and perhaps an eventual climate catastrophe that results in mass death and delays or curtails human progress.

Of course, some of these effects may be offset directly or indirectly by economic growth, or rather by technologies that result from economic growth. The possibility of technological silver bullets (e.g. geoengineering techniques) that mitigate or erase individual effects of climate change cannot be dismissed out of hand. But there is simply no reason to think that the discount rate gives a good approximation of these tradeoff effects. Market rates of return, for instance, are not a reliable measure of the speed at which these technologies are being developed or the size of their effects.

²² I have not considered here a different rationale for discounting, offered by Posner and Weisbach (2010), which views the discount rate as a measure of the default opportunity cost of forgone investment at market rates of return. This rationale for discounting differs fundamentally from any of the arguments so far considered. But I am content to reiterate the response given by Fleurbaey and Zuber (2013), namely, that the use of discounts to measure opportunity cost is simply superfluous: we can compare the returns of possible projects (including the null project of non-interference with markets) without such a discount rate as long as we remember to directly account for the cost of taxes or regulations in terms of forgone investment. So by declining to discount at the market rate we are not putting ourselves at any extra risk of irrationally selecting suboptimal endeavours.

How, then, should we go about evaluating the costs and benefits of climate policies, given our limited foreknowledge? My tentative suggestion is as follows:

- (1) The potential costs of climate change should remain disaggregated among a few qualitatively distinct categories. These include in particular: loss of 'ordinary' economic consumption (which, as I have argued, is a relatively minor consideration compared with those that follow); loss of environmental goods that are resistant to substitution or replenishment; disutilities like disease and displacement that will fall most heavily on the least well off; loss of life; and risk of global catastrophe that might permanently curtail the potential of human civilization. Some of these cost categories may be subject to reasonable aggregation for certain purposes, but there are irreducibly normative questions about how to compare them, and for the last two items in particular, quantification in economic terms adds little to our understanding of their significance.
- (2) Similarly, in assessing the costs of climate change *mitigation*, costs to the global rich and global poor should be kept mostly separate. From an impartial standpoint, the costs of a global carbon tax to present First Worlders seem fairly trivial compared with the possible costs of delayed economic growth in China, India and Africa. From the standpoint of First World governments, which may be thought to have special responsibilities to their citizens, the relative weighting of these costs raises another set of irreducibly normative questions.
- (3) Finally, the problem of demandingness should be given separate and explicit consideration. When evaluating the astronomical utilities associated with existential risk mitigation, in particular, we may well decide that the future can only ask so much of us. But this judgement should not be hidden in the details of an economic model, nor does it have anything to do with time *per se* (as Parfit's 'catastrophe' case illustrates).

Presenting the costs and benefits of climate policy in this disaggregated form need not imply that the various categories of cost and benefit are incommensurable, nor need it deny the existence of an objective, complete, and transitive ordering of policy options in terms of all-things-considered desirability. The question I have been pursuing in this paper is simply whether the attempt to represent all costs and benefits in the medium of discounted consumption improves our understanding of the relevant comparisons and guides us toward the correct ordering of policy options.

Economic tools can and should be brought to bear on questions of public policy, even where there are seemingly 'non-economic' goods like human life at stake and even where there are irreducibly normative questions as well as empirical questions to be answered. But we must have some confidence that economic models are giving us a reasonable approximation of the true costs and benefits. As I have argued, there is simply no reason to think that representing considerations like the diminishing marginal utility of consumption, aversion to inequality, and concerns about over-demandingness by means of a social discount rate can meet this requirement in the context of climate policy. The framework suggested above, which essentially puts to the side the kind of costs captured by the discounting approach and attempts to keep qualitative distinctions between cost categories as well as key normative questions clearly in view, seems to me to promise a closer approximation of the salient realities. The dangers of climate change are multifarious and unpredictable, and this makes the use of the simple damage functions and discount rates a tempting simplification. But if these methods do not approximate the reality, they make no useful contribution to the important decisions that these dangers force on us.

ACKNOWLEDGEMENTS

My thanks to Mathias Frisch, Dan Moller, Juli Tarsney, Zachary Taylor and two anonymous referees for their helpful comments on earlier drafts of this paper.

REFERENCES

- Arntzenius, F., A. Elga and J. Hawthorne. 2004. Bayesianism, infinite decisions, and binding. *Mind* 113: 251–283.
- Arrow, K. J. 1999. In *Discounting and Intergenerational Equity*, ed. P. R. Portney and J. P. Weyant, 13–21. Baltimore, MD: Johns Hopkins University Press.
- Baily, M. N., J. Manyika and S. Gupta. 2013. US productivity growth: an optimistic perspective. *International Productivity Monitor* 23(3): 3–12.
- Beckstead, N. 2013. *On the Overwhelming Importance of Shaping the Far Future*. PhD thesis. Rutgers University Graduate School, New Brunswick.
- Bostrom, N. 2009. The future of humanity. In *New Waves in Philosophy of Technology*, ed. J. K. B. Olsen, E. Selinger, and S. Riis. London: Palgrave Macmillan.
- Bostrom, N. 2013. Existential risk prevention as global priority. *Global Policy* 4: 15–31.
- Broome, J. 1994. Discounting the future. *Philosophy & Public Affairs* 23: 128–156.
- Carpenter, K. E., M. Abrar, G. Aeby, R. B. Aronson, S. Banks, A. Bruckner, A. Chiriboga, J. Cortés, J. C. Delbeek, L. DeVantier, et al. 2008. One-third of reef-building corals face elevated extinction risk from climate change and local impacts. *Science* 321: 560–563.
- Cook, J., D. Nuccitelli, S. A. Green, M. Richardson, B. Winkler, R. Painting, R. Way, P. Jacobs and A. Skuce. 2013. Quantifying the consensus on anthropogenic global warming in the scientific literature. *Environmental Research Letters* 8(2): 024024.
- Cowen, T. 2011. *The Great Stagnation: How America Ate All the Low-Hanging Fruit of Modern History, Got Sick, and Will (Eventually) Feel Better*. London: Penguin.

- Cowen, T. Unpublished. What is the correct intergenerational discount rate? URL: <<https://www.gmu.edu/centers/publicchoice/faculty%20pages/Tyler/DISCOUNT.pdf>>.
- Easterlin, R. A. 1995. Will raising the incomes of all increase the happiness of all? *Journal of Economic Behavior & Organization* 17: 35–47.
- Fleurbaey, M. and S. Zuber. 2013. Climate policies deserve a negative discount rate. *Chicago Journal of International Law* 13: 565–685.
- Frederick, S., G. Loewenstein and T. O'Donoghue. 2002. Time discounting and time preference: a critical review. *Journal of Economic Literature* 40: 351–401.
- Hagerty, M. R. and R. Veenhoven. 2003. Wealth and happiness revisited: growing national income *does* go with greater happiness. *Social Indicators Research* 64: 1–27.
- Heath, J. 2016. Climate ethics: justifying a positive social time preference. *Journal of Moral Philosophy* (forthcoming).
- Hume, D. 2000 [1739]. *A Treatise of Human Nature*. Oxford: Oxford University Press.
- IPCC. 2014. *Climate Change 2014: Synthesis Report*. Contribution of Working Groups I, II, and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Geneva: IPCC.
- Jackson, F. 1991. Decision-theoretic consequentialism and the nearest and dearest objection. *Ethics* 101: 461–482.
- Koopmans, T. C. 1967. Objectives, constraints, and outcomes in optimal growth models. *Econometrica* 35: 1–15.
- Landesman, C. 1995. When to terminate a charitable trust? *Analysis* 55: 12–13.
- Lauwers, L. and P. Vallentyne. 2004. Infinite utilitarianism: more is always better. *Economics and Philosophy* 20: 307–330.
- McCrone, D. J. and J. H. Kuklinski. 1979. The delegate theory of representation. *American Journal of Political Science* 23: 278–300.
- Moller, D. 2006. Should we let people starve – for now? *Analysis* 66: 240–247.
- Nordhaus, W. 2007. *The challenge of global warming: economic models and environmental policy*. URL: <http://aida.econ.yale.edu/~nordhaus/homepage/dice_mss_091107_public.pdf>.
- Parfit, D. 1984. *Reasons and Persons*. Oxford: Oxford University Press.
- Parmesan, C. 2006. Ecological and evolutionary responses to recent climate change. *Annual Review of Ecology, Evolution, and Systematics* 37: 637–669.
- Posner, E. A. and D. Weisbach. 2010. *Climate Change Justice*. Princeton, NJ: Princeton University Press.
- Ramsey, F.P. 1928. A mathematical theory of saving. *Economic Journal* 38: 543–559.
- Sherwood, S. C. and M. Huber. 2010. An adaptability limit to climate change due to heat stress. *Proceedings of the National Academy of Sciences* 107: 9552–9555.
- Singer, P. 1972. Famine, affluence, and morality. *Philosophy & Public Affairs* 1: 229–243.
- Stern, N. 2007. *The Economics of Climate Change: The Stern Review*. Cambridge: Cambridge University Press.
- Stern, N. 2008. The economics of climate change. *American Economic Review* 98(2): 1–37.
- Stern, N. 2014. Ethics, equity and the economics of climate change paper 1: science and philosophy. *Economics & Philosophy* 30: 397–444.
- Stocker, T. F., Q. Dahe and G.-K. Plattner. 2013. *Climate Change 2013: The Physical Science Basis. Working Group I Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Summary for Policymakers*. Geneva: IPCC.
- Thomas, C.D., A. Cameron, R. E. Green, M. Bakkenes, L. J. Beaumont, Y. C. Collingham, B. F. Erasmus, M. F. De Siqueira, A. Grainger, L. Hannah *et al.* 2004. Extinction risk from climate change. *Nature* 427: 145–148.
- Thuiller, W., S. Lavorel, M. B. Araújo, M. T. Sykes and I. C. Prentice. 2005. Climate change threats to plant diversity in Europe. *Proceedings of the National Academy of Sciences USA* 102: 8245–8250.

- Tol, R. S. 2009. The economic effects of climate change. *Journal of Economic Perspectives* 23(2): 29–51.
- Tol, R. S. 2014. Correction and update: the economic effects of climate change. *Journal of Economic Perspectives* 28(2): 221–225.
- Weitzman, M. L. 2009. Some basic economics of extreme climate change. In *Changing Climate, Changing Economy*, ed. J.-P. Touffut, 106–134. Cheltenham: Edward Elgar Publishing.
- Zuber, S. and G. B. Asheim. 2012. Justifying social discounting: the rank-discounted utilitarian approach. *Journal of Economic Theory* 147: 1572–1601.

BIOGRAPHICAL INFORMATION

Christian Tarsney is a PhD candidate in philosophy at the University of Maryland, College Park. His dissertation concerns problems of rational decision-making under moral uncertainty. He also has research interests in the philosophy of time, decision theory, and normative/applied ethics.