

# Discounting for Epistemic Reasons

Kian Minz Woo

University of Oxford, UK

## Abstract

Utility discount rates in intergenerational economic modelling have been viewed as problematic, both for descriptive and normative reasons. However, positive discount rates can be defended normatively; in particular, it is rational for future generational utility to be discounted to take uncertainty about radical societal re- construction into account. Social discount rates are defended against objections from Parfit (1986) and Broome (2005, 2012).

**Keywords:** Climate change, intergenerational justice, discount rate, rational decision-making, applied ethics *Word Count:* 7000.

# 1 Introduction

A moral commonplace in everyday life is that one is not morally responsible for *all* the consequences of one's actions. If cashed out in terms of reactive attitudes, we might say that one is not blameworthy for indefinitely many effects one's actions lead to.<sup>1</sup> It is true that one is responsible for the immediate and predictable consequences, but as one's knowledge of the effects lessen, it is less and less reasonable to hold one responsible. This is a vague intuition, but it is widely shared. When I call in sick to work, I am responsible for the probable and expected consequences—my colleagues having a larger workload, my boss having to reschedule a meeting, etc. The unexpected effects of these effects I am less responsible for: that my harder working colleague has to miss a date; that my boss ends up having to stay overtime as a result. The effects of these effects I am even less responsible for. And so, when I am deciding whether to call in sick, I do not have to include these further knock-on effects in my calculation, both because (a) it would be irrational and impossible for me to consider them all; and because (b) I have no or diminished responsibility for these effects.

What is the import of this observation? I argue that it can be used to defend an important aspect of intergenerational justice—the utility discount rate<sup>2</sup> (or just “discount rate”)— $\delta$ , which indicates the level of discounting for future generational utility.<sup>3</sup> The discount rate is an important aspect in modelling the value of policies on future generations, and is a consequential topic for climate policy. The level of discounting may influence the urgency of actions both to mitigate and to adapt to climate change (cf. Dasgupta 2012; Dietz et al. 2007; Godard 2009; Wahba and Hope 2006). With high rates of discounting, future generations' harms have less weight in decision-making, whereas low rates of discounting raise the valuation of future utility. However, discounting has been objected to on grounds both descriptive (Frederick et al. 2002; Nordhaus 2007) and normative (Broome 2005; Cowen and Parfit 1992; Parfit 1986; Ramsey 1928). Roughly, the descriptive objections are rooted in market and psychological data, whereas the normative arise from principles and intuitions about moral value (Arrow et al. 1996).

In this essay, I argue that, on normative grounds, a (non-zero) discount rate is defensible for the purposes of decision-making, which is how it is used in economics. In economics, discounting does not imply that some items are objectively worth less: it is used to evaluate or weigh some objects more highly. For instance, when *offered* two goods of  $\pounds x$ —one now and one in a year—it is not true that one is worth more than the other.<sup>4</sup> They are by hypothesis worth the same, i.e.  $\pounds x$ . But we discount to show that they are not equally valuable *to me* in my temporal position, since I can use the intervening time to take advantage of possession of the good and the productive nature of capital. Discounting does not change the objective value of either good, it is used to guide decision-making.

Similarly, I argue that discount rates are temporally and agentially relative. In particular, an agent's limited epistemic access is a primary salient factor. The appropriate discount

---

<sup>1</sup> Moral responsibility is thus not the same as responsibility simpliciter. Moral responsibility is the extent of praiseworthiness or blameworthiness. Responsibility may correspond roughly to being in the causal chain of an *effect*. So it may be that something *issued* from my actions (i.e. I am *responsible* for it), but it is unreasonable to praise me for this *effect* (i.e. I am not *morally responsible* for it). For instance, this might happen if the *effect* was unintended or unexpected.

<sup>2</sup> This is sometimes called the “social rate of time preference” or “pure time preference” but these locutions suggest a single agent with preferences over times. I avoid such phrases since my argument paints a very different picture:  $\delta$  acts as a heuristic for evaluating policies and is not a function of *preferences* at all.

<sup>3</sup> I do not address the complex issue of how generations are meant to be delineated; the arguments in this paper do not depend on any particular specification. A helpful discussion of the range of possibilities can be found in (Gardiner 2011, esp. Ch. 5).

<sup>4</sup> Abstracting away from phenomena like inflation and deflation, &ct.

depends on the actions or policies being evaluated as well as the epistemic reasons at play; it is not a measure of objective value. Ultimately, this analysis will not justify high discount rates, but it will justify limited *positive* discount rates.

In §2, I lay the moral groundwork for my argument. I argue that our status as agents with bounded epistemology limits our moral responsibility. In §3, I discuss the discount rate, its relation to temporal discounting, and several attempts which have been made to defend it. In §4, I argue that, given epistemic limitations, it is both normatively and rationally defensible to discount intergenerationally due to our uncertainty with respect to future radical societal reconstruction. I then apply this as a defence of  $\delta \neq 0$ . In §5, I consider objections from Parfit and Broome. Parfit's objection is that it is immoral to discount the utility of other individuals simply in virtue of the fact that they are temporally distant. Broome's objection is that discounting introduces objectionable time-relativity into judgments of goodness. I also briefly suggest a method for determining utility discount rates in accordance with the limited epistemic status that we have.

## 2 Examining the Intuition

In order to determine what is motivating this intuition, it is necessary to examine moral responsibility. Moral responsibility is a complex concept, and is clearly a function of several factors. My intention is not to provide an analysis, but merely to defend one particular claim about it, viz. epistemic uncertainty limits responsibility, so epistemic uncertainty should be factored into moral deliberation.

To motivate my claim—which I suggest is generally held and intuitive—I want to go out of my way to consider an everyday case, which is less likely to generate confused or distorted intuitions. If it works in every day cases like the following, I submit that the intuition is more robust than an intuition generated by an esoteric or science-fiction example.

*Sick*: I am considering whether to call in sick to work, and eventually decide to. This leads to an increased workload for my colleague. Furthermore, I miss a meeting that was scheduled with my boss Margaret, where I was supposed to give her some report. However, unbeknownst to me, as Margaret needed this report in order to ship a product urgently, she is forced to stay overtime. But, due to this overtime, she is unable to have her scheduled date with her partner Rosa. Rosa gets angry with Margaret, and they have a fight.<sup>5</sup>

In *Sick*, I clearly have moral responsibility for some of the consequences of failing to show up for work: I am morally responsible for the workload of my colleague (and blameworthy with respect to this *effect*). I am morally responsible for missing my meeting and failing to hand in the report, both of which were foreseeable consequences of not showing up for work. It is less intuitive that I am morally responsible for Margaret's overtime stress, and even less intuitive that I am morally responsible for Rosa's anger with Margaret, even though these effects are caused by my calling in sick.

Why might this be? I suggest that there is a natural criterion which *affects* my responsibility with reference to *different effects* of calling in sick: how foreseeable those *effects* are for me. It is easy for me to recognize that my colleague will have to do my work, and that I won't be attending the meeting I planned with my boss. However, as I stipulated that I am not aware that the report I am scheduled to give to Margaret is so important (I do not think that this importance is *impossible*, but I think it extremely unlikely). It's even less

---

<sup>5</sup> I wish to emphasize that (perhaps apart from me) no one need have acted wrongly in this instance. It might have been that Margaret could not have told me how important the report was because the product details were supposed to be kept secret from me. And while it was a breach of arrangement to miss the date, Margaret was clearly reasonable in staying in to work if the product launch and months of preparation were in jeopardy. As well, it's reasonable for Rosa to be disappointed and to get angry, even if Margaret had good work-related reasons to miss the date.

foreseeable that my calling in sick would lead to Rosa being angry with Margaret, under the assumption that I do not know that Rosa and Margaret have a date (or even that Rosa exists).

We can test this criterion by holding most details of the case constant, except that, in a modified case, I am aware of the import of my report for this product. In this case, I am now clearly morally responsible for Margaret's overtime stress, but still not morally responsible for Rosa's response. Finally, we can consider the case where I am deciding whether to call in sick and I know all of these consequences that will result from my calling in sick. I know that Margaret needs the report to ship the (urgent) product and that she has a date with Rosa for the evening. Then I am morally responsible for Rosa's disappointment and for Margaret's overtime stress. If so, the key *difference* is how predictable these outcomes are for me.

We can codify these intuitions into a principle about moral responsibility:

**Principle 2.1** Moral responsibility for the effects of one's actions is diminished by low subjective epistemic probability in the likelihood of those effects obtaining.

This principle does not imply that I have *no* moral responsibility for those outcomes for which I have low certainty; it implies that I have *less*. But as the likelihood of certain outcomes falls lower and lower, at some point I have—practically speaking—no responsibility even if—theoretically speaking—I may have infinitesimal levels of moral responsibility. In what follows, I sometimes use “not responsible” as a shorthand for “practically speaking not responsible” for simplicity.

Note also that this principle does not require that one explicitly performs a prediction when deliberating; the issue is the probability that the agent would assign to an outcome as a function of their current information state. As finite agents of the type we are, we do not have access to enough causal information to determine precisely what our actions will lead to. But we must take something to guide our actions and what we have are subjective probabilities.<sup>6</sup>

This leads to my second claim, about moral deliberation. I think that consideration of *Sick* leads us to another principle of deliberation. Suppose I am planning whether to call in sick for the day (and that I think that the consequences of my action are morally relevant to the decision). I could try to enumerate *every* contingency, assign them probabilities (while trying to avoid continuum problems), and then calculate expected value. But this is going to make living pretty *difficult*—by which I mean impossible.

If we think that ought implies can, then it cannot be that we *ought* to include every contingency in our deliberation over actions. This follows if it is impossible for beings like us to do such calculations. I want to say something stronger, which is that it would be irrational for us to try to do such calculations, given that we are beings with limited cognitive capacities.<sup>7</sup> But this leads to the need to remove some of the complexity of our deliberation process.

Here is a suggestion: it is rational to limit deliberation to that which one is morally responsible for.<sup>8</sup>

---

<sup>6</sup> Lenman (2000) discusses the problem of indefinite (and radically unforeseeable) consequences. He takes it to be fatal to the consequentialist. One can view Principle 2.1 as an olive branch to the consequentialist in light of Lenman's objections. If the consequentialist accepts a principle of this sort, then they are able to address Lenman's worry about consequentialist inaction in the face of indefinite and indeterminate consequences.

<sup>7</sup> If we were beings with unlimited cognitive capacities, or with *sufficient* capacities to execute such calculations, then I would not think it would be irrational. In such scenarios, it might even be required.

<sup>8</sup> It may be that in a *sufficiently* complex situation, all the morally salient *effects* are overwhelming for deliberation purposes. Then further restrictions may be necessary. Furthermore, I hold that individual and group decision-making should be parallel in this respect since they share the relevant characteristics: both are finite, bounded decision-makers that can in appropriate circumstances be morally responsible.

**Principle 2.2** When morally deliberating (either as individuals or groups), it is rational to limit consideration to the effects which one is morally responsible for.

Why is this? I think this principle gives us the right result in *Sick*. Since I am aware that I am inconveniencing my colleagues and that I will miss the meeting with my boss Margaret, I am morally responsible for these *effects*. By this principle, I should weigh and consider these when I am choosing whether to call in sick. In contrast, I am not aware (or almost unaware) that this action will lead to Margaret's agenda filling up, or that it will lead to Rosa's anger with Margaret. So it is reasonable for me to ignore these possible *effects* when deciding to call in sick.

One might *offer* an objection: maybe we should not limit our consideration to the *effects* which we are responsible for—i.e. that we give *sufficient* subjective probability to—because our subjective probabilities may be mistaken.<sup>9</sup> This objection is built on the idea that what really matters are *objective* probabilities, not subjective. For example, consider a doctor who thinks that a medicine lacks harmful side-*effects*, but is not completely sure. In such a case, they assign low probability to this harmful outcome, but we are readily likely to assign blame to them if there are side-*effects* despite this low probability.

There are at least two styles of response that could be *offered*; I am agnostic between them. The first response to this objection is that one is epistemically responsible for assigning the right subjective probabilities, but this is distinct from moral responsibility from acting upon these probabilities. So we can epistemically criticize the agent for having the wrong probabilities, but—as the first principle states—the moral responsibility is a function of the subjective probabilities the agent assigns.

Another response would be to distinguish between *who* is responsible and the *extent* of their responsibility. We are interested in the latter, but this objection might conflate the two. When someone has a low credence in a bad outcome from their action, we still want to hold *them* responsible, but this is distinct from holding them fully or largely *morally* responsible. It may be that their epistemic status mitigates their moral culpability.

Regardless, it very well may be that one may be obligated to act to improve one's epistemic position. There are certainly situations where one should act to improve one's epistemic position. But failing to check medical details is not *always* blameworthy; we do not always need to be improving our epistemic position. In particular, we are responsible for failing to do so when we have a high subjective epistemic probability that doing so would *affect* our levels of moral responsibility. We do not blame doctors for failing to check things that they are (almost) certain of, and this is what my principle predicts.

### 3 Introducing $\delta$

In this section, I begin by defining and discussing discounting. I then discuss a couple normative methods of determining the value from the literature in order to contrast my own suggestion.

I take discounting to be a very thin concept: a discount rate (e.g.  $\delta$ ) is a time-indexed mathematical factor which is used to weight temporally disparate consequences. The justification for a certain value of  $\delta$  is not *itself* a discount rate. So to say that I take a positive discount rate does not tell you anything about *why* or *what justifies* that discount rate. There are also *different* kinds of discount rates, meaning discount rates for *different* streams, but primarily discount rates are used for consumption (measured in currency units) or utility (often measured in utils or QALYs).

---

<sup>9</sup> I thank Alison Hills for suggesting this objection.

So I might have a consumptive discount rate that I justify using psychological data about impatience, or I might have a consumptive discount rate that I justify by appeal to empirical market prices (i.e. productive nature of capital), or I might have a utility discount rate that I justify using a politically democratic rationale, or I might have a utility discount rate that I justify by appealing to uncertainty.

The point is that there is the value of the variable  $\delta$  in the cost-benefit analysis and then there is the defence of that value, and that conflating them is a problem. Unfortunately, this problem occurs in both philosophy and in economics. For instance, Damon et al. (2013, p. 47) write that “Discounting builds on the simple fact that money earns interest” which neglects alternative reasons one might discount consumption. In response to the arguments of Dasgupta and Heal (1979), Broome (1992, p. 102) simply asserts that “Uncertainty is to be accounted for separately. It cannot *affect* the value of [the utility discount rate]  $r$ ”. Both of these writers take discounting to be intrinsically tied to some justifications and not capable of a variety of defences.

With this in mind, it is necessary to consider what value of  $\delta$  to take, but also what the justification for that value is. The value is important, but the justification is of more philosophical import. The primary concern with this essay is with the justification for a discount rate for utility, but first it is worth surveying some previous arguments about justifications in order to contrast my view.

Why let  $\delta > 0$ ? Here are two significant normative arguments that have been *offered* in the literature: we might think that future generations are less valuable or less valuable to *us* (Beckerman and Hepburn 2007), or that future generations might not exist (Stern 2007). Let us consider these considerations in reverse order.<sup>10</sup>

The second consideration is that we should discount because this implies odds about the survival of humanity. Stern (2007) famously argued that this is the *only* legitimate reason to have a discount rate, generating a comparatively small value for  $\delta$  (i.e.  $\delta = 0.1$ ).<sup>11</sup> There is uncertainty about the existence of future human beings, and this uncertainty is time-dependent. If future generations will not exist, then certainly their utility should not be counted.

In fact, I endorse a similar reason to have a positive value for  $\delta$ , although my argument differs importantly from Stern’s. Stern takes us to be decision-making under risk, whereas I take us to be decision-making under uncertainty over the long term. This uncertainty has roots including the possibility of existential risk (i.e. population collapse), but I take the sources of uncertainty to be considerably larger. Part of the impetus for my argument comes from the variety of pertinent outcomes: extinction is just the beginning. The variety of outcomes makes the type of risk analysis that Stern endorses to be far more complex than he allows for; I suggest that our limited epistemic status prevents us from being able to explicitly assess these outcomes. This will be my response to Broome and Parfit in §5.

The other defence of a positive  $\delta$  is due to Beckerman and Hepburn (2007), who argue that one possibility is that we should only care (or care more) for those with whom we have privileged relationships.<sup>12</sup> Traditional economics has a tacit utilitarian and impartial underpinning, but this need not be. If we consider a sentiment-based ethics (à la Adam Smith or David Hume), then it is natural to discount because we justifiably care for those with whom we have closer bonds.

---

<sup>10</sup> I am not including the account of Rawls (2001) since it bypasses arguments for  $\delta$ , arguing instead about values of the social discount rate  $\delta + \eta g$ .

<sup>11</sup> This echoes an argument by Dasgupta and Heal (1979), which held that external existential risk of population collapse can be appealed to in defence of a discount rate.

<sup>12</sup> Dasgupta et al. (1999) survey the *different* economic models that take this approach.

It is certainly true that many accept that it is permissible (or even obligatory) to treat one's family members in a preferential manner, even if that means slightly more appropriate candidates receive less goods as a result. But this should not be applied to the case at hand. For one thing, such preferential treatment is allowed—if it is so allowed—only in the *private* sphere.<sup>13</sup> So the fact that we will know (and care about) the coming generation or two, but will not know succeeding generations is not of moral import from a public perspective. This important disanalogy between the two cases shows that treating close generations closer as a matter of public policy is ethically problematic.

#### 4 Defending $\delta$ by Appeal to Normative Responsibility

So can a positive value for  $\delta$  be morally defended? I suggest it can. If we accept Principles 2.1 and 2.2, then it is rational to discount when considering those actions for which we have uncertainty. Intergenerational *effects* have decreasing levels of certainty, and increasing levels of uncertainty. In the context of climate change, we should discount future utility since the policies we select may have *different* or trivial *effects* on future utility in the face of radical societal reconstruction, leaving us with uncertainty about how those *effects* will obtain. Since this uncertainty limits our moral responsibility for those *effects*, it is rational for us to discount to the extent that we are uncertain. That uncertainty increases with respect to time. These remarks justify a positive discount rate for future well-being.

Uncertainty with respect to climate change is often mentioned in concert with climate denialism. Many denialist doctrines rely upon discrepancies between various modelling forecasts for climate change impacts (often while highlighting those forecasts which are least troubling). While it is true that there are significant uncertainties about the probabilities we should assign to *different* scenarios, that does not mean we should deny that climate change will have significant and almost certainly negative *effects*. What it does mean is that we should be aware that we have epistemic attitudes *about* our credences: sometimes we are highly certain of the value of assigned probabilities and sometimes we are almost fishing in the dark. In the near term, there is more convergence, but in the long term, we should accept our epistemic limitations and try to report them more explicitly (cf. Sluijs 2012, reprinted in this volume).

But I will make a stronger claim, which helps to emphasize the implications in a climate policy context. Even if we were deciding under scientific *certainty* about the *effects* that climate change would lead to, we would still have significant sources of uncertainty which are often overlooked. I argue that this would hold even if we knew *precisely* the increased probability of sea level rises, extreme weather events, biodiversity loss, and other consequences. In short, if we had certainty about our risk.

That is because the shape of society that these *effects* impact may *differ* considerably from our own today; society might be radically reconstructed. One might think that no matter what shape society takes, we can predict that the consequences of climate change policy will have strongly negative consequences on future peoples regardless of how their society looks.

But this is false. As Stern (2007) notes, it impacts our calculations if there are no future human beings extant at some point. From that point, we would not adduce any value from human beings. So this is the first radical societal reconstruction: extinction (or a similarly catastrophic population outcome).

---

<sup>13</sup> Although it may be permissible for governments to favour their own citizenry, this is because the governments are usually meant to represent their citizens. If they were meant to represent more than one nation, they should also treat the larger group impartially. Since climate change is an issue where policies are determined internationally, we should not be preferential. This comment is not conclusive, but addressing this issue in full would take me beyond the scope of this essay.

One might think that it is so unlikely as to be negligible that humanity goes extinct. A naïve argument might make an inductive appeal to our lack of extinction thus far, but just as my existence up until this moment gives me no reason to generalize that I will always exist, neither does our species' survival amidst the background of species extinction.

But we can adduce positive reasons to think that our risk of extinction is significant. As Bostrom (2013, p. 15) notes, “Estimates of 10–20 per cent total existential risk in this century are fairly typical among those who have examined the issue”. These are obviously dependent on assumptions and subjective assessments, but we have reasons to suppose that this is a non-trivial risk. Rees (2003) concurs that, in light of our increasing ability to harness technological, biological, and chemical science, the risk of catastrophic use is growing. Ultimately, Rees thinks that the odds are no better than even that society will survive this century. Finally, Quammen (2013) argues that zoonoses—microbes which break the species barrier—are increasingly dangerous as humanity disrupts and expands into new habitats, thus coming into contact with new zoonotic agents while being increasingly globally interconnected. These types of risks are relatively new, making our past survival less predictive.

It might seem that if we do not go extinct, then we *do* know how climate change *effects* will *affect* future humans. But this is not true, either. When considering long-time scales, and assuming a surviving *homo sapiens*, it becomes plausible that humanity relocates partially (or significantly) to other planets. This is our second family of radical societal reconstructions, since climate change policies will be *ineffectual* if humanity is not located on the planet Earth.<sup>14</sup>

Finally, it might appear that if we survive and stay on Earth that we would be certain of the *effects* of climate change policies. But once again, there are scenarios of radical societal reconstruction. This scenario is that we will be able to generate self-replicating machines, or that we will be able to augment ourselves in significant ways, uploading our “software” into supercomputers. This could lead to intentional evolution, and possible transhuman states (Kurzweil 2005). If this occurred, the environmental conditions of Earth might have very little *effect* on the beings we would become since our current biological needs could be lessened or eliminated. Future technological advances could change us in ways we are unable to imagine at present (Ord et al. 2010). Such possibilities are highly speculative, but they are instances of radical societal reconstruction which would become more probable as technological ability increases (almost by definition).

This list is not meant to be exhaustive, nor is it meant to present likely scenarios. It is meant to demonstrate that—even if we removed all scientific uncertainty—there are important sources of uncertainty about how our policies will *affect* others. This is because we do not know for certain what future generations will look like, and that uncertainty grows with respect to time. Furthermore, radical societal reconstruction may not exhaust the outcomes that a *complete* climate change accounting of expected utility would have to include.

If radical societal reconstruction occurs, then we would be *indifferent* between climate policies at the point which it occurred. These scenarios have a non-negligible chance of occurring, but quantifying the probabilities is hard to impossible. We lack the epistemic status to be able to give informed probabilities of these outcomes. Since we are more and more uncertain of the consequences of our climate policies further into the future, by Principle 2.1, we are less morally responsible for those *effects*. As Principle 2.2 states, it is rational to decide on policies while ignoring the *effects* for which we have less responsibility. This justifies

---

<sup>14</sup> Here, I am operating under the assumption that climate change *effects* are relegated to this planet.



counting those policies less with respect to time, to the extent that we are uncertain about their impact; hence, a positive  $\delta$ .

Does this argument support a particular value of  $\delta$ ? No, it is just supposed to justify a non-zero  $\delta$  on normative grounds. The argument indexes  $\delta$  to our current epistemic state with regard to outcomes like radical societal reconstruction. Given that we think that we are quite certain about our *effects* in the short to medium term (roughly, less than fifty years), it supports a relatively small value of  $\delta$ . Note, however, that this argument suggests that  $\delta$  should be strictly larger than the value of  $\delta$  chosen in (Stern 2007), since Stern considers only a subset of the scenarios in radical societal reconstruction. A consequence of my claim is that, if our knowledge about the future increases, our valuation of future generations would rise or fall accordingly; for instance, because we might have a better idea of whether there *will be* future generations.<sup>15</sup> The indexing means that it is possible that  $\delta$  fluctuates over time; for instance we should increase  $\delta$  if some rogue individual manages to control weapons of mass destruction or engineer a deadly fast-spreading virus.

As noted before, this is *different* from previous defences, such as those of (Dasgupta and Heal 1979) or (Stern 2007). That is because they claim that discounting is reasonable since we are decision-making under risk with the risk of existential threats. My claim is that we are decision-making under uncertainty and that the range of pertinent outcomes is far wider than death. With this range of outcomes, I suggest that it is rational to count those outcomes less in our reasoning for which we have limited epistemic status. This is because we have reduced moral responsibility and accounting for that reduced responsibility justifies a discount rate.

Is it reasonable to model these vague epistemic uncertainties with an annual percentage discount rate? Parfit (1986) phrases the objection in this manner: if the point of  $\delta$  is to reflect uncertainty, why not just introduce a value that *is* uncertainty? Uncertainty and an annual discount rate may sometimes, even often, coincide, but there is no necessary connection.

But this makes the mistake about discount rates that I mentioned before: the value of a discount rate (positive) and its justification (uncertainty) are distinguishable. There is not some overriding justification for a value of  $\delta$ . Since I defend  $\delta$  with a justification of uncertainty, the question is whether my justification is reasonable, not whether discounting can be for uncertainty. It is worth noting that the discounting does not depend on whether the uncertainty comes from sources that generate positive or negative utility, compared with our envisioned outcome. Insofar as scenarios reduce our certainty, they reduce our moral responsibility, regardless of the valence of changes.

On the more technical question of whether a discount rate is the appropriate tool for modelling what I want to express—growing uncertainty with respect to time, I suggest that it is an excellent fit. Of course,  $\delta$  is at best an approximation. Uncertainty does not necessarily generate smooth curves. But there are certain properties uncertainty will have relative to time: it will have a positive first derivative, since uncertainty increases with (future) temporal distance.<sup>16</sup> Second, it is likely to have an asymptote, because it seems implausible to assign absolute uncertainty about even the distant future. We cannot have more than 100% uncertainty, and it would be arbitrary to imply that there is a particular date beyond which we know nothing. Beyond this arbitrariness, it appears empirically false. Society knows *some* things about even the distant future, e.g. functions of nomological facts like that gravity won't repel and thermodynamics will hold. With these criteria, an exponential or hyperbolic function is a natural fit, since they have asymptotes and are increasing.

---

<sup>15</sup> Under the assumption that Stern generates the right probabilities for the scenarios he considers.

<sup>16</sup> At least for beings like us who tend to experience time in a linear fashion and can more easily predict the near term than the far term.

## 5 Objections

Many objections can, and have been, raised against arguments of the preceding type. In this section, I address three.

The first objection is raised by Parfit (1986). He claims that it is morally abhorrent to discount future people (or their utility) *simply* in virtue of the fact that they are temporally distant. Would it be morally permissible to discount people who are *spatially* distant? Modulo the details of what occurs in a life, it is worth as much as any other no matter when or where it occurs.

Of course this last point is right. But the conclusion (that we should not have a social discount rate) does not follow from the premise (that future people's utility are worth the same as present people). Discounting is a technique used for modelling and for assessing policy or decision choices. It is not, nor is it meant to be, an objective picture of the moral status of people. Discounting for uncertainty is compatible with the claim that people in *different* time periods have the same objective value. I am defending the claim that, despite (possibly) equal objective value, evaluating with temporal *differences* is defensible. The discounting I defend is for the sake of decision-making; it is not a moral judgment.

This is directly analogous with the economic case. As I wrote in § 1, when given the choice between two equally priced goods at *different* times, we do not discount to indicate that the goods are *different*. They have equal prices, and that is *unaffected* by when they are given. However, we discount to indicate that we should evaluate those good *differentially*.

This leads to a second objection. Broome (2012); Cowen and Parfit (1992) argue that the modelling of uncertainty should be in terms of *different* scenarios which have *different* probabilities attached to them, i.e. using traditional expected utility. Once we take the *different* options, we can evaluate options by the weighted sum of the outcomes that may result. In private correspondence, Broome reiterates this point, writing that if we are not discounting purely for time, that should be made clearer.

The reasons that I favour using  $\delta > 0$  are mentioned in §4. It may be a simplification to use exponential curves, but the alternative—expected utility—demands more than our epistemic or cognitive conditions can bear when dealing with an issue like climate change.

This requires some defence. Let us begin with the observation that in order to include outcomes in a formula of expected utility, it is necessary to be able to enumerate those outcomes. This first step is already extremely hard: most of the outcomes that I considered under the rubric of radical societal reconstruction have been overlooked, and I have no confidence that I am aware of all of the outcomes which would change our calculations for climate policies. This is already under the implausible assumption that we have removed any uncertainty or confusion with *climate* impacts which, as we know, are subject to tipping points and perturbations that are already unforeseeable.

But the problem has another order of complexity to add: once we have enumerated these outcomes, it is necessary to provide them with credences or probability values (or ranges). Here, I wish to emphasize once again that we are creatures with very limited epistemic status: we have *difficulty* determining what will happen in our world in the coming months, let alone years. When it comes to centuries, any credences *offered* for outcomes like the possibility of a technology Singularity should be ridiculed. It is tempting to bring to mind how likely the shape of our society today would have looked to those a century ago or even fifty years ago. Just because we might be able to generate some credences (e.g. by finding the mean or mode of some “expert” opinions) should not give us any credence *in* those credences. As Dasgupta (2008, p. 164) warns, “We shouldn’t believe any model that explicitly models risk when the time horizon extends 100–200 years into the future. We simply don’t know what the probabilities are”.

But this isn't the end of the complexity for the expected utility model. Let us say that we are able to identify the types of scenarios that would impact our accounting of climate policies *and* are able to give some credences for those outcomes. Then we have to repeat this process for each *different* temporal unit: the calculation that presumes we go to Mars in 2100 is a *very different* calculation from one which presumes we go to Mars in 2150 which is also *different* from going to Mars in 2125, and so forth. This process should be repeated for all of the relevant outcomes (with appropriately varying credences).

In short, it is implausible that such outcome enumerations and probability-ascriptions are tractable, either epistemically or cognitively. Assigning an exponentially increasing discount rate is far more mathematically tractable.

That's the negative work. But the more positive (and fundamental) response is that we *are and should be* discounting for time. This is because, the further into the future we model, the less we know. The less we know, the less we can be held responsible. This follows from the decision-making principles from §2.

A closely related objection is that I am not making calculations more tractable, because the discount rate is just a summary parameter of all of the probabilities of various radical societal reconstructions. The final formulation *looks* simpler, but that is because all the work is done behind the scenes to generate a value of  $\delta$ .<sup>17</sup>

I was not intending to have  $\delta$  be a summary of these probabilities, as I doubt that we have access to such probabilities at long timescales. But the obvious question this raises is how to determine  $\delta$  in light of our epistemic limitations. Here is a suggestion. I do think that on shorter timescales (e.g. less than a century), we can assign probabilities of societal reconstruction. If so, the way to set a value for  $\delta$  is by giving explicit probabilities for various societal reconstruction outcomes at shorter timescales, adding these probabilities at *different* time points, and then fitting an exponential discount curve through these points. Since discount rates are such blunt tools, once one fits the curve to data points close to the present, there is no need to find data points beyond a century out, which is good because I deny that explicit risk values are reliable so far in the future.

The final objection also comes from Broome (2005, 2012). His objection to temporal discounting is that it introduces temporal relativity, and that such relativity is absurd. Since discounting is only with respect to future persons, he claims that a model that includes discounts can (correctly) value a future person less than an individual who lived after that person would (correctly) value that person. Broome's example is people who die in war. Individuals prior to Caesar's wars (call them Early Individuals) would judge the utility lost in the World Wars as less bad than Caesar's wars since the discount rate would lead to heavily discounting the harms from the World Wars due to their large temporal distance. But to us (Later Individuals), this seems absurd. Surely the World Wars are many times more harmful than Caesar's wars. So the relativism of discounting is problematic.

But the *difficulties* compound, Broome continues. The Early Individuals could be *aware* that the World Wars would be so much larger and that they (were they in our shoes) would view the Early Individuals' judgment as absurd. So the relativism is genuine; there is no factual information that *differs*—just temporal placement.

But how do we explain our Later Individual judgments about the World Wars' (much) greater disutility when compared to Caesar's wars? Well, recall that we discount *for epistemic reasons*. Being the type of creatures that we are, we do not have uncertainty about the past. So it is irrational for us to discount the utility of events we are certain of. If we were beings that knew only the future, and did not remember any of the past, this would be reversed. Thus, a being with supreme knowledge—i.e. one for whom nothing was

---

<sup>17</sup> Thanks to Teru Thomas for suggesting this.

uncertain—any discounting would be irrational, since everything would be an *effect* of their action and they would know all of their *effects*.

When Broome claims there is no factual *difference* between the Early Individuals and the Later Individuals, he is begging the question. For the Later Individuals, the past *is* a fact, whereas for Early Individuals the future is only conjecture, which they can assign probabilities to as best they can. So the relativism is not surprising or problematic at all: it is united in the epistemic capabilities of the *different* agents. The Later Individuals count past lives equally, since these past lives are *unaffected* by the policies of Later Individuals; the Early Individuals discount in line with their uncertainty over how and whether their actions will *affect* those yet to come. Just like the simple case of deciding whether to call in sick, if we forget our bounded epistemic status, we are bound to make irrational choices.

## References

- Arrow, Kenneth, W.R., Cline, Mäper, K.G., Squitieri, R., and Stiglitz, J.E. 1996. "Intertemporal equity discounting, and economic *efficiency*." In J.P. Bruce, H. Lee, and E.F. Haites (eds), *Climate Change 1995—Economic and Social Dimensions of Climate Change*, 125–144. Cambridge: Cambridge University Press.
- Beckerman, Wilfred and Hepburn, Cameron. 2007. "Ethics of the Discount Rate in the Stern Review on the Economics of Climate Change." *World Economics* 8:187–210.
- Bostrom, Nick. 2013. "Existential Risk Prevention as Global Priority." *Glob Policy* 4:15–31. doi:10.1111/1758-5899.12002.
- Broome, John. 1992. *Counting the Cost of Global Warming*. Cambridge: White Horse Press.
- . 2005. "Should We Value Population?" *The Journal of Political Philosophy* 13:399–413.
- . 2012. *Climate Matters: Ethics in a Warming World*. New York: Norton.
- Cowen, Tyler and Parfit, Derek. 1992. "Against the Social Discount Rate." In Peter Laslett and James S. Fishkin (eds), *Justice Between Age Groups and Generations*, 144–161. New Haven: Yale University Press.
- Damon, Maria, Mohlin, Kristina, and Sterner, Thomas. 2013. "Putting a Price on the Future of Our Children and Grandchildren." In Michael A. Livermore and Richard L. Reversz (eds), *The Globalization of Cost-Benefit Analysis in Environmental Policy*, chapter 4, 47–65. Oxford: Oxford University Press.
- Dasgupta, Partha. 2008. "Discounting climate change." *J Risk Uncertain* 37:141–169.
- . 2012. "Time and the Generations." In Robert W. Hahn and Alistair Ulph (eds), *Climate Change and Common Sense: Essays in Honour of Thomas Schelling*, chapter 7, 101–130. Oxford: Oxford University Press.
- Dasgupta, Partha and Heal, Geoffrey. 1979. *Cambridge Economic Handbooks: Economic Theory and Exhaustible Resources*. Cambridge: Cambridge University Press.
- Dasgupta, Partha, Mäler, Karl-Göran, and Barrett, Scott. 1999. "Intergenerational Equity, Social Discount Rates and Global Warming." In John Weyant and Paul Portney (eds), *Discounting and Intergenerational Equity*. Washington: Resources for the Future.
- Dietz, Somon, Hope, Chris, and Patmore, Nicola. 2007. "Some economics of 'dangerous' climate change: Reflections on the Stern Review." *Global Environmental Change* 17:311–25.
- Frederick, Shane, Loewenstein, George, and O'Donoghue, Ted. 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature* 40:351–401.
- Gardiner, Stephen. 2011. *A Perfect Moral Storm: The Ethical Tragedy of Climate Change*. Oxford: Oxford University Press.
- Godard, Olivier. 2009. "Time discounting and long-run issues: the controversy raised by the Stern Review of the economics of climate change." *OPEC Energy Review* 33:1–22.
- Kurzweil, Raymond. 2005. *The Singularity is Near*. Viking.
- Lenman, James. 2000. "Consequentialism and Cluelessness." *Philosophy & Public Affairs* 29:342–370.
- Nordhaus, William. 2007. "Review: A Review of the "Stern Review on the Economics of Climate Change"." *Journal of Economic Literature* 45:686–702.
- Ord, Tony, Hillerbrand, Rafaela, and Sandberg, Anders. 2010. "Probing the improbable: Methodological challenges for risks with low probabilities and high stakes." *Journal of Risk Research* 13:191–205. doi:[10.1080/13669870903126267](https://doi.org/10.1080/13669870903126267).

- Parfit, Derek. 1986. *Reasons and Persons*. Oxford: Oxford University Press.
- Quammen, David. 2013. *Spillover: Animal Infections and the Next Pandemic*. New York: W. W. Norton.
- Ramsey, Frank P. 1928. "A mathematical theory of saving." *The Economic Journal* 38:543–559.
- Rawls, John. 2001. *Justice as Fairness*. Cambridge, Mass: Harvard University Press.
- Rees, Martin. 2003. *Our Final Hour*. New York: Basic Books.
- Sluijs, Jeroen P Van Der. 2012. "Uncertainty and Dissent in Climate Risk Assessment: A Post-Normal Perspective." *Nat Cult* 7:174–195. doi:[10.3167/nc.2012.070204](https://doi.org/10.3167/nc.2012.070204).
- Stern, Nicholas. 2007. *The Economics of Climate Change*. Cambridge and New York: Cambridge University Press.
- Wahba, Mohammed and Hope, Chris. 2006. "The marginal impact of carbon dioxide under two scenarios of future emissions." *Energy Policy* 34:3305–16.