

# 6 A new defence of probability discounting<sup>1</sup>



*Kian Mintz-Woo*

## Introduction

Probability discounting (or probability weighting) in decision-making is multiplying the value of an outcome by one's subjective probability that the outcome will occur.<sup>2</sup> The broader import of defending probability discounting is to help justify cost–benefit analyses in contexts such as climate change. This chapter, however, addresses only decisions under risk. Decisions can be made under conditions of certainty, risk, and uncertainty.<sup>3</sup> Under certainty, the decision-maker knows the outcome of any of his or her decisions.<sup>4</sup> Under risk, they know the potential outcomes of any of their decisions, and can assign probabilities to any of those outcomes dependent on particular decisions. Finally, under uncertainty, at least some outcome for at least some decision cannot be assigned a probability. This defence of probability discounting applies only under risk. However, this is not unduly limiting. At a theoretical level, the possibility of probability discounting does not arise under the other conditions. At a practical level, evaluations of climate policy are almost never done under conditions of full certainty, and many are performed under risk. (For an overview of evaluations of climate policy under risk, please see Chapter 3 of this volume.) So addressing probability discounting under risk is widely applicable. The purpose of this chapter is thus to justify the use of probability discounting in contexts of risk – but also cost–benefit analyses more broadly – using a new argument.

While such probability discounting is sometimes taken to be objectionable, it is also sometimes taken to be uncontroversial. This chapter addresses both camps.

As a representative of the first sceptical camp, I consider an objection by Caney (2008, 2009) in the following section.<sup>5</sup> Caney has been particularly influential in the climate debate and argues that, if certain conditions obtain, it is morally impermissible to discount for probability. But his conditions – which are meant to indicate when probability discounting (and cost–benefit analysis) is impermissible – fail, since they overgeneralize. Since climate change is plausibly construed as putting many people's rights at risk, one question which arises is how to assess the risk of a violation compared to a violation which is certain to occur. Caney's answer is that, at least under the assumption that certain

conditions (“R1–R4”) obtain, there is no difference in how we ought to assess them: even putting human rights *at risk* is impermissible.

These conditions do not manage to do the work that Caney needs, and, more generally, this type of approach is unworkable; each individual action may increase risks of violation only in very small increments and it is only once these risks are aggregated that initially plausible conditions like Caney’s R1–R4 are satisfied. At least in theory, cost–benefit analyses with probability discounting are sensitive enough on the individual level to compare the small incremental risks and the marginal benefits of individual actions.

Among those in the second camp who take probability discounting to be uncontroversial, many refer to the decision-theory literature. However, probability discounting has not received a more catholic defense. Here I offer a defense of probability discounting which does not rely upon decision-theoretic axioms that those outside of decision-theory may not accept. My intention in my positive argument for probability discounting is to address those skeptical of probability discounting with a new normative defence. This defense can be taken as complementary to decision-theoretic defences. The complementarity, for those who are already sympathetic to the decision-theoretic defences of probability discounting, lies in this argument’s reinforcement of this conclusion.

This argument involves distinguishing between causal responsibility and moral responsibility, where moral responsibility for an effect can be cashed out in terms of blameworthiness and praiseworthiness and causal responsibility for an effect requires only that that effect is part of a causal chain issuing from one’s act. With this distinction in hand, moral responsibility can be seen as coming in degrees. Given that we can limit our deliberation and consideration to that which we are morally responsible for, and that our moral responsibility for outcomes is limited by our subjective probabilities, it follows that our subjective probabilities can ground probability discounting.

The concluding section includes some suggestions about what this means for long-term decision-making. The argument allows us to ground probability discounting in normative terms such as moral responsibility. Evaluations of climate change can be grounded, on this argument, in our moral responsibility which – on the assumption that our subjective probabilities decrease with respect to time – itself decreases with respect to time. This argument also shows that it is permissible for individuals to engage in (mini) cost–benefit analyses with probability discounting.

### **Caney’s objections to probability discounting**

In an argument against discounting for probability in the context of climate change, Caney (2009: 176) writes “A sound response to the current climate change, I suggest, would prescribe exactly the same course of action ... to mitigating climate change as would be appropriate if it were known that the malign effects would definitely occur.”<sup>6</sup> Essentially, the point is that, from the perspective of rights-based theories, “high” probabilities of rights violations are just as

impermissible to bring about as certainty of rights violations.<sup>7</sup> In effect, this means that cost–benefit analyses should not or need not be performed in these circumstances, and the actions which lead to high probabilities of violations should be avoided.

I begin by addressing Caney’s particular conditions for cases where these actions simply should not be performed, before explaining what is wrong with this approach. The problem is that, if we look at practices at the social level using his approach, we lose subtle gradations in levels of risk for individual actions and the corresponding individual benefits of introducing small increments of risk. And, for Caney’s approach to work, we *do* have to look at practices at the social level, because his argument is about climate change and any given individual’s emissions in isolation do not materially increase the probability of rights-violations (let alone introduce *high* probabilities of rights violations).

Caney suggests four conditions meant to distinguish between cases in which probability discounting is warranted or unwarranted. He writes that the four conditions may not all be necessary, but if all are satisfied in some instance, then this justifies treating risky outcomes the same way as certain outcomes. But Caney’s conditions, which are meant to apply to the social practice of excessive carbon emitting, apply in an analogous manner to another case: the social practice of driving cars. The first clarification to make is that the relevant comparison is about the permissibility of pervasive social practices. The second clarification to make is that, obviously, driving cars is not in *all* respects analogous to excess emitting, but that the practice satisfies the four conditions Caney endorses, which is enough to show that the argument generalizes beyond what he intends. In other words, this is meant as a *reductio* – someone who endorses the conclusion that the practice of driving cars is impermissible need not be troubled by these remarks. However, there are few who would be willing to endorse this claim and Caney himself denies it (Caney 2009: 179–80).

First, here are Caney’s four rights conditions applied to emissions and climate change (2009: 177–9):

- R1 The changes to the climate involve both (a) a high probability of severe threats to large numbers of persons’ fundamental human rights and (b) a possibility of even more catastrophic threats to fundamental human rights.
- R2 Affluent members of the world can abstain from emitting high levels of greenhouse gases, and thereby exposing others to risk, without loss of their own human rights.
- R3 The risks of dangerous climate change will fall disproportionately on those whose human rights are already violated.
- R4 The benefits that arise when the affluent of the world emit high levels of greenhouse gases falls almost entirely to them, and not to those most at risk from climate change.

Now consider an everyday example of risky behavior: driving cars. This example of a social practice satisfies analogous conditions. This case certainly satisfies (R1): drivers kill many people daily, through either irresponsibility or bad luck. In the United States, the costs amount to tens of thousands of lives, orders of magnitude more of non-fatal injuries, and a total of hundreds of billions of dollars annually in economic losses and social harms. It is not the case that *any given* driver has a very high probability of killing other road users; however, it is the case that the *social practice* of driving leads to a high probability of severe threats to the fundamental human rights of many. Caney is right that driving satisfies (R2): “the emissions stemming from driving cars, taking plane flights, poorly insulated housing, and inefficient energy use [must be cut to] avoid dangerous climate change, but the loss involved cannot be said to compromise any human rights” (2009: 179). In defending (R3), Caney appeals to the income disparity between nations more at risk from climate change (mainly in the global south) and those less at risk (mainly in the global north). Since those most at risk from climate change are – independently of climate change – disproportionately poorer, they are more likely to suffer from human rights violations. Similarly, we can consider the risks from driving. On average, those most at risk from the practice of driving cars are those who do not themselves drive cars (the so called “vulnerable road users” such as pedestrians and cyclists) (e.g., Shinar 2012; Tiwari 2015).<sup>8</sup> Those who do not drive cars are *disproportionately* likely to be unable to afford cars so will be – independently of car driving – disproportionately poorer. Thus, by analogous reasoning, they too will be more likely to suffer from human rights violations. However, Caney is aware that driving cars might be a case that satisfies his conditions, and this is not a conclusion he endorses. So he addresses car driving in the discussion for (R4):

one might think that it is permissible for some to drive motor vehicles even though they pose a fatal threat to some because the risky activity is part of an equitable scheme which is generally beneficial, including to the risk-bearers. Even non-drivers might value the practice of driving – it means that their friends and family can visit, goods can get transported promptly from one end of the country to another, food gets delivered to shops still fresh, business and personal correspondence can arrive swiftly, and so on.

(Caney 2009: 179–80)

The purpose of (R4), Caney continues, is that it distinguishes between cases like climate change and driving. But it is doubtful that such a consideration is enough to distinguish between the two cases.

How can we measure the benefits from excess emissions or driving? We can consider possible worlds in which we do not benefit from the risky behaviors. We are not considering *ideal* possible worlds, but closer *non-ideal* worlds where people also often fail their duties, but scale down dramatically (or even completely) the risky behaviors under consideration. So, in a world where there is (almost) no driving, perishable foodstuffs, for example, would have to be

consumed sooner and would have to be sourced locally. In this world, there would be less capability to visit those far away. It would mean that serious injuries might not be treated in time (this is a friendly addition to Caney's argument). If we stopped (almost) entirely the practice of driving cars, Caney is right that some of these costs would be borne by risk-bearers (non-drivers).

However, a world where there were dramatically reduced emissions would similarly have costs which would be borne by risk-bearers (non-affluent countries). Roughly one-third of global emissions stem from transportation. To cut this down significantly would greatly change the world. Assuming that all of the global emissions are scaled down proportionately, this would restrict many aspects of global transportation. For instance, Bangladesh, which is at severe risk from climate impacts, would have limited access to world markets for clothing; the workers in Hong Kong would have limited access to global markets for electronic goods; countries with disparities in resources would be less able to trade. There would also be less transportation that allows for foreign aid and medical support: mosquito nets would be less distributable to sub-Saharan Africa; fewer tetanus and measles vaccinations would be sent to Iraq; in the case of sudden disasters, it would be less feasible to react internationally with aid; and it may not have been possible to eradicate polio. These are just some of the consequences of reducing transportation emissions; reductions in other sectors might well have other impacts for the risk-bearers. In short, like driving, emissions from affluent countries are part of a scheme which generates benefits, including to the risk-bearers.

Obviously, it is only a fraction of the gains associated with excessive emissions that are distributed in these laudable ways, but we are considering a world where all of the activities associated with significant affluent emissions are ratcheted down.<sup>9</sup> My claim is that, just as in (R4), there are substantial gains to non-drivers and society at large from transport as well. Or enough gains that it is plausible that the analogue of (R4) fails in the case of excess emissions.

Caney might respond that these are not the "excessive" emissions he was considering, so they are not part of the relevant social practice. However, this response will not work for three reasons. First, such a response would raise a worry about ad hoc emendations to the theory. If we define "excess" emissions as only the emissions that we have moral or social objections to, then not only would this be less possible to apply, it would invite worries about being merely dialectically motivated.

Second, his R1–R4 are not sensitive enough to pick out the excess emissions generated which are not necessary to maintain the human rights of those emitting, and it is R1–R4 that are under consideration. For instance, they do not distinguish between emissions by the affluent who are travelling with NGOs to distribute bed nets, medicine, trade or aid and those who are merely travelling for holidays.

Finally, and most importantly, he *cannot* add a condition that says we are including only those cases where the emissions are for a sufficiently good cause or have positive outcomes which outweigh their increase to the threats to

human rights (via climate change), since that is precisely the kind of cost–benefit analysis condition that he is trying to avoid or supplant with the conditions R1–R4. The purpose of a cost–benefit analysis is precisely to be sensitive to the cumulative effect of smaller (or larger) benefits and costs. On this reading of Caney’s argument, his intention is to suggest conditions which supplant cost–benefit analyses by being grounded in human rights. If he appeals to the aggregation of more mundane considerations, then he is engaged in cost–benefit analysis by another name.

So either (R4) is not a plausible condition to determine whether a cost–benefit analysis applies, or, following Caney, it implies that on rights-based theories, drivers are rights-violators to precisely the same extent as if we were all *certain* that we would kill when driving, i.e., it is not permissible to drive cars. Most people would view this as a reduction of Caney’s conditions.<sup>10</sup>

This third point leads to a general worry: that, while such a rights-based view does not greatly differ from a probability discounter at the level of social practice – since the non-discounter does take into account that the practice must in aggregate have a “high” probability of threats to fundamental human rights – it does not give any weight to small individual benefits. Thus, this view does not allow one to compare (i) practices where each instance of the practice has an infinitesimal marginal addition to the probability of fundamental violations to (ii) practices where each instance of the practice has a more substantial marginal addition to the probability of fundamental violations. If the practices are widespread enough, the non-discounter will not be able to distinguish between these practices because, in aggregate, both will involve high probabilities of threats to large numbers of persons’ fundamental human rights. Furthermore, if we are focussed on rights-violations, we will not be able to compare practices where the *benefits* are minor or major in each instance of the practice, so long as the benefits are never required for upholding fundamental human rights. The type of view that Caney endorses does not allow for fine-grained evaluation, unlike cost–benefit analyses with probability discounting. This coarseness matters, since small costs and benefits can in aggregate be very important, even if they do not directly relate to human rights – for instance, the social benefits of driving and extra emissions.

So far, we have considered why Caney’s argument against probability discounting is unsuccessful. But there is more: a positive moral argument that discounting for probability is permissible, which we turn to in the next section.

### **A positive argument for probability discounting**

A moral commonplace in everyday life is that one is not fully 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 fully blameworthy (or praiseworthy) for indefinitely many effects of one’s actions. It is true that one is morally 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 fully morally responsible.<sup>11</sup>

This commonplace about moral responsibility can ground a defence of probability discounting in rational decision-making. The conclusion to this section will consider how this argument might apply both at the individual level and the social level, with an eye towards social applications to evaluations of climate change policy. Responsibility at the social level, though not argued here, is not unlike responsibility at the individual level. There are some interesting differences, mainly related to greater limitations on individual decision-making, which we discuss in the concluding section.

When I call in sick to work, I am morally responsible for the foreseeable consequences: my colleagues having a larger workload, my boss having to reschedule her meeting, etc. The unexpected effects of these effects I am *less* morally responsible for: my harder working colleague has to miss dinner with her wife; my boss ends up with no space in her agenda. 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 fully include these further effects in my calculation, because I have diminished (or no) moral responsibility for these effects.

Moral responsibility is a complex concept. These arguments are not meant to characterize it, let alone define it. They are based on a particular claim about moral responsibility, which is consistent with many different substantive accounts: that one can have less than full moral responsibility – understood as levels of, *inter alia*, blame-worthiness and praiseworthiness – for some outcomes. Whatever conditions one takes to be necessary and/or sufficient for moral responsibility, this argument is intended to demonstrate an additional layer: the level of moral responsibility is *lessened* by certain kinds of ignorance. So, although the focus on subjective probability may be surprising in a discussion of moral responsibility, it is best thought of as an emendation. The argument is also meant to apply widely, not just to actions with significant moral import, such as those involving pain and death, but also everyday actions, such as listening to loud music in public.<sup>12</sup> The argument, schematically, is very simple:

- 1 Sometimes, we do not have full moral responsibility for the effects of our actions.
- 2 The most plausible explanation for (1) is that moral responsibility is limited by the subjective probabilities of particular effects.
- 3.: Moral responsibility is limited by the subjective probabilities of particular effects. (1,2, abduction)
- 4 It is rational to exclude that which you are not morally responsible for from deliberation.
- 5.: It is rational to exclude from deliberation outcomes to the extent that they are limited by subjective probabilities of particular effects. (3,4, deduction)

Premise (1) captures the idea that moral responsibility for effects or consequences of an action are not always full or complete. There is a separate sense of responsibility which is always full, which is sometimes called *causal responsib-*

ility, for which *any* effect of one's act is fully one's responsibility. It is in this latter causal sense that responsibility is fully transitive or iterative. For instance, if I scribble down my grocery list and some biologist happens to read the grocery list and has a flash of inspiration and develops a new strain of guava, I (or my writing of the list) might be *causally responsible*, although it would be odd to say that I am *morally responsible*. Common sense has it that there is a sense of responsibility which is not always full, unlike this causal sense. According to this common sense, one can diminish one's responsibility for the effects of an action by showing that one did not know that those effects would follow. In other words, if the effects are surprising or unforeseen, then this can have some exculpatory force. It is this moral sense of responsibility that this argument relies upon, and I use this sense from here on except where explicitly noted.

To motivate Premise (2), i.e., that it is subjective probability that is changing the level of moral responsibility, we will consider several cases. Here is the first:

**Chemistry – Certainty:** Tamara, a high school student, has a lab experiment to conduct. She is following the instructions she has written down on the chemicals to mix: X and Y. However, she knows that X and Y explode when combined, and when she mixes them, they do indeed explode, leading to significant property damage.<sup>13</sup>

Tamara intuitively is morally responsible (and blameworthy) for this property damage, and also causally responsible. We can change the case so that she does not know about this consequence (i.e., she assigns the probability of an explosion less than 1, maybe even close to 0):

**Chemistry – Low/High Risk:** Tamara does not know that X and Y explode when combined, but when she mixes them, they do explode, leading to significant property damage. There are two subcases. In the first, Tamara has a low subjective probability that X and Y explode (e.g., she copied her instructions for mixing them out of her textbook or another generally reliable source). In the second, Tamara has a high subjective probability that X and Y explode, although she also does not know (e.g., she knows that sometimes she accidentally writes down chemical Y instead of chemical Z, or that she knows she was not paying full attention when she was taking notes).

In the Low Risk subjective probability case, where Tamara thinks it very unlikely that she has made a mistake, Tamara has less responsibility. In contrast, in the High Risk subjective probability case, where Tamara has reasons to believe that her instructions might be in error, she has greater responsibility.

Since the only differences between Low and High Risk cases is the subjective probability Tamara would assign to the explosion, that is a good explanation for the assigned level of responsibility.

An objector might suggest a different explanation: in particular, that the relevant probability is not subjective probability, but *objective* probability. Such an objector may remind us that individuals sometimes have very inaccurate probability assignments. As Smart (1973: 40) writes,

The ordinary man is frequently irrational in his moral thinking. And if he can be irrational about morals why cannot he be irrational about probabilities? The fact that the ordinary man thinks that he can weigh up probabilities in making prudential decisions does not mean that there is really any sense in what he is doing.

It is manifest that individual subjective probabilities may diverge significantly from the actual ones. This is true in these Chemistry cases; in fact, the objective probabilities that *X* and *Y* will explode in *each* of these cases is 1 and that objective probability diverges from Tamara's subjective probabilities in each Risk case. However, if you believe that having more or less credence in this outcome (as in High Risk/Low Risk) affects her level of moral responsibility, then it is subjective probability and *not* objective probability that is relevant for responsibility.

An objector could press a separate case. One could say that in Low/High Risk it is relevant whether Tamara has tried to *improve* her subjective probabilities. She is fully morally responsible, one might claim, in all the cases because she should do all that she can to improve her probabilities, say by rechecking about the explosive potential of *X* and *Y* with another authoritative source.<sup>14</sup> In this manner, the "excess" moral responsibility attaches to her epistemic practices.

There are at least three responses to this objection:

- First, one is *epistemically* responsible for assigning the right subjective probabilities, but this is distinct at the point of decision from any *moral* responsibility for acting upon these probabilities. So we can epistemically criticize the individual for having generated the wrong probabilities up to the point of decisions, but the moral responsibility is still a function of the subjective probabilities the individual had assigned at the point of decision. This might also act as a response to someone who endorses *evidential probability*, a theory according to which the correct probability is the probability that draws on the evidence that is available to the individual (e.g., Zimmerman 2008, 2014).<sup>15</sup> On this theory, it is sensible to say that the individual is being epistemically irresponsible (or irrational) if their evidential probabilities diverge from their subjective probabilities. But their moral responsibility is still a function of their subjective probabilities at the point of action.
- Second, it is easy to conflate moral responsibility for some action  $\phi$ -ing and moral responsibility for a different action  $\psi$ -ing that would have changed one's information *about*  $\phi$ -ing. These may be conflated, for instance, because, in both cases, the responsibility attaches to the same individual, and  $\psi$ -ing affects whether or not the individual  $\phi$ s. Because of this potential

conflation, it is important to distinguish between *which party* is morally responsible and *their level* of moral responsibility. For the purposes of this argument, we are primarily concerned with the latter. This account can also handle actions such as updating one's information – one is responsible to the extent that one would have a subjective probability that an action which updates one's information (e.g.,  $\psi$ -ing) would have the effect in question.<sup>16</sup> To illustrate, suppose that a doctor prescribes some medication believing there are no relevant contraindications. The doctor turns out to be wrong. An objector might say that, even though the doctor had (for instance) a low subjective probability that the medication had contraindications (say, because they had checked the contraindications recently), we are likely to hold him responsible, i.e., blame him. However, it is not important here to determine *whom* we are blaming (that is not in doubt), but *why* and *to what extent* we are blaming. Presumably, the reason we blame the doctor is not for prescribing a medicine they thought was safe. Indeed, *that* act is praiseworthy. We blame the doctor for failing to check the indications. But the subjective probabilities matter again: if there were good (subjective) reasons to think that there were no contraindications, then we should blame the doctor less for failing to check, even if those reasons were ultimately misleading. If we assigned full blameworthiness for failing to check, then we would also have to assign full blameworthiness for failing to *recheck* each time the medicine is prescribed, or perhaps even more often. But this seems implausible; at some point, it no longer makes sense to spend time ascertaining contraindications and to spend time on patients who need it. So the blameworthiness must be diminishing as well. Again, this is what one should accept if one judges Tamara less responsible in Low Risk than in High Risk.

- But these first two responses may be too concessive. A third response would be to try to undermine the intuition that one is morally responsible for poor epistemic practices. For instance, one could argue that such an intuition is problematic, since it is asymmetrical. Epistemic practices which lead to mistaken beliefs in cases with good outcomes do not seem blameworthy in the same manner that practices in cases with bad outcomes seem blameworthy.

The intuition we need to undermine, however, is that, even if Tamara was in Low Risk, she is still morally responsible (blameworthy) for her epistemic state of being in Low Risk. Let us consider a contrasting case:

**Donation (Low/High Risk):** Meili has several different codes for all the accounts or payments she makes. She intends to load a substantial amount of money onto her prepaid coffee card (e.g., her expected annual coffee budget), but she enters the wrong code and instead sends it to a charity which does good works.<sup>17</sup> She is aware that it is possible that she used the wrong code. There are two subcases. In the first, Meili has a low subjective probability that she used the wrong code (e.g., she recently used her codes

and knows that her recent attempts were successful). In the second, Meili has a high subjective probability that she used the wrong code, although she also does not know she did (e.g., she knows that regularly she forgets which of these codes is which).

For the objector whose intuition is that Tamara has moral responsibility for her epistemic practices which led to her being in Low Risk, an analogous claim about Meili in Low Risk is unconvincing. It sounds absurd to claim that Meili is *morally* blameworthy for failing to check the codes in order to switch her payment *away from charity*. This undermines the intuition in Tamara's Low Risk case, since the intuition does not generalize.

Since it is beyond the scope of this chapter, these are only brief and tentative remarks about how to explain the intuition. One is a (reverse) halo effect. The halo effect is the bias to associate extra positive properties to individuals with positive properties. Reverse halo means that, since we already blame Tamara for causing the property damage, we may be prone to judge her negatively in other ways as well. In particular, we might judge her epistemic practices as morally problematic, even though epistemic responsibility and moral responsibility should be kept separate. A more interesting and speculative possibility is that this asymmetry is connected to the Knobe (2006) effect, whereby attributions of praiseworthiness/blameworthiness (and intentionality) are asymmetrically a function of the moral desirability of side-effects. It is taken as given that this asymmetry is, at least *prima facie*, problematic. Here, the asymmetry may come from a similar set-up; instead of them being side-effects of a stated goal, the outcomes are accidental outcomes which occur instead of their goals. In short, we blame an individual for not avoiding a negative outcome, whether that outcome was a side-effect of a separate goal or unexpectedly occurred instead of the goal. In the former case, this blame means imputing intentionality to the individual, whereas in the latter case, since blame for the action may be limited by their expectations, the "excess" blame is shifted to subjective probabilities. That this shift is asymmetrical can be seen by considering Donation, and that the shift is asymmetrical suggests that the shift is illegitimate.

Now, we can move on to Premise (4), which tells us that, in deliberation, it is rational to exclude that which you are not responsible for. This premise is probably straightforwardly or conceptually true, but some justification is necessary. Just as it is irrational to spend unlimited time and energy trying to attain certainty about what consequences will follow from an action (recall that the consequences might be extremely spatially or temporally distant), it is irrational to include potential consequences for which one is not responsible. This is most obvious when considering consequences which one cannot affect. Such consequences are not one's moral responsibility and so, given we are finite beings, it is irrational for us to consider them. If we also accept that this is the case with consequences which one could affect, but which are not one's moral responsibility, then this is sufficient for (4). In other words, the reason to adopt Premise (4) is that it unifies or systematizes the considerations that may be permissibly

excluded from decision-making. I think it is a sensible limitation on decision-making, as it helps to prevent it from becoming implausibly onerous on finite beings like ourselves.

With (3) and (4), it is easy to see how rational deliberation will limit outcomes to the extent of the subjective probabilities, giving us (5). Finally, since probability discounting is a way of representing this exclusion mathematically, it can be used to exclude outcomes from deliberation. Mathematically, the way to exclude the *outcomes* to the extent given by subjective probability is just to weight (i.e., multiply) the outcomes by their subjective probability. This argument demonstrates that, when decision-making under risk, probability discounting is defensible when taking responsibility into consideration.

### Additions and extensions

There are interesting decision-theoretic arguments and axiomatic derivations of the claim that it is rational to weight outcomes by the subjective probabilities that those outcomes will occur. Many of these hold that rationality implies that decision-makers will try to maximize the satisfaction of their preferences, assuming that their preferences are sufficiently coherently structured. These are targeted towards idealized decision-makers with rational preferences. The arguments presented here are targeted in a slightly different manner; they are aimed at interlocutors who find these decision-theoretic conclusions and axioms less plausible. They are more plausible since they are based on the claim that it is rational for decision-makers to *focus* their decision-making by *limiting* things in their decision-making.

Now, in practice, it is important to hypothesize that such cost–benefit analyses are a good model for actual decision-making for individuals, with two considerations or qualifications:

- Individuals have very coarse-grained probability increments, so it is probably descriptively more accurate that individuals work with categories such as “unlikely” or “almost definitely not” or “almost certainly” and that these can only roughly be mapped onto a range of probabilities in percentage terms (e.g., 10–25 percent or 0–10 percent or 90–100 percent, respectively). Responsibility on this argument would then be just as coarse-grained as the probability categories a given individual is working with.
- Individuals tend to round probabilities up and down, such that small probabilities get rounded down to 0 percent (impossibility) and high probabilities get rounded up to 100 percent (certainty) (cf. Kahneman 2011). This matters greatly for simplifying decision-making, since most of the distant consequences or effects of an action are very difficult to predict (or have very small effective subjective probabilities). For individual decision-making, this means that many potential consequences are given negligible probabilities and, due to this psychological quirk, tend to round those to 0 percent, weighting those outcomes by 0 and discounting them from the decision-making entirely.<sup>18</sup>

It seems that both of these considerations make the decision-making account given here descriptively more accurate. But adopting these considerations is more normatively defensible or rational for individuals than for social groups, since a proper cost–benefit analysis involving all of the imaginable or foreseeable potential outcomes – even those with very small probabilities – would be so cognitively taxing for individuals as to be unworkable.

However, when the stakes are high, it is best for individuals to try to mitigate both of these caveats and take the cognitive effort required to do so. Furthermore, at the social and political level, decision-makers should always try to avoid these effects. Social and political decision-making do not, at least in theory, have as many limitations as individuals do in day-to-day decision-making and so should include, with as fine probability assignments as possible, all the outcomes that can be predicted. There may well be practical limitations, but simplifications should be tolerated in fewer cases than for individual decision-making. This is the case with evaluations of projects in the context of climate change.

This argument implies that, if our subjective probabilities of the effects of our climate-altering actions decrease with respect to time, we have decreasing moral responsibility for effects which do occur. As for what the subjective probabilities are, and who the relevant “we” are, reports issued by the Intergovernmental Panel on Climate Change (IPCC) should be seen as a legitimate first approximation of the state of published science so anyone who has reasonable access to the results of these reports should be included in this “we” and the subjective probabilities should reflect the IPCC reports. In fact, due to the level of agreement required for information to be included in the reports, they are likely to be conservative in their assessments.

A final point to make is that the probability discounting endorsed here does not conflict with or double-count probability discounting for axiomatic reasons (for instance). It offers an alternative route to the claim that one ought to discount by subjective probabilities, a route which is meant to be normatively and intuitively acceptable. By buttressing the conclusion in a context divorced from the decision-theoretic axioms, it is intended to make this conclusion more palatable to a wider audience. In this way, we can justifiably use cost–benefit analyses even in contexts which concern costs and benefits over the very long-term, as with climate change.

## Notes

- 1 I would like to thank Joanna Burch-Brown, Marc Davidson, Eike Düvel, Rachel Fraser, Clare Heyward, Alison Hills, Lukas Meyer, Daniel Petz, Thomas Pölzer, Matthew Rendall, Dominic Roser, Harald Stezler, and Teru Thomas for helpful comments and suggestions; audiences at the University of Graz, the University of Oxford, and the European Consortium for Political Research’s Joint Session organized by Matthew Rendall and Dominic Roser; and the very helpful editors of this volume. My subjective probabilities are such that responsibility for remaining errors lie with me. Funding from the Austrian Science Fund (FWF) under research grant W 1256-G15 (Doctoral Programme Climate Change – Uncertainties, Thresholds and Coping Strategies) is gratefully acknowledged.

- 2 It is not meant to be confused with either *utility discounting* or *consumption discounting*, such as discussed in (Broome 1992, 1994; Dasgupta 2008; Parfit 1984; Stern 2007), among many others. In particular, it is not the same as the special category of consumption discounting that Parfit (1984) calls *probabilistic discounting*. The latter picks out consumption discount rates which are justified by reference to probabilities. The type of discounting under consideration could potentially be used to justify probabilistic discounting, but it is not directly about discount rates at all. The kind of probability in question is also sometimes called *ex ante* probability.
- 3 See Ellsberg (1961) for a seminal paper on these distinctions.
- 4 I use the terms “consequence” and “effect” synonymously, but “outcome” indicates the entire set of consequences (or effects) that result from an action. For expositional simplicity, I assume that the probability an individual assigns to an outcome is the same as their credence in that outcome.
- 5 Caney (2009: 196) agrees that probability discounting only applies under risk, and mentions that this is a limitation, since we do not always have solid grounds for subjective probability assignments, especially in the context of climate change. However, Caney’s (2009) argument which I engage with here presumes that we have probabilities of rights violations, so the argument in question does presume decision-making under risk.
- 6 Caney writes this in terms of money spent, but this means either he is suggesting no *consumption* discounting, which he explicitly denies, or, more plausibly, it is simply a *means* to the morally relevant consequences.
- 7 Caney is not explicit about what probabilities constitute “high,” but for the purposes of the critique here, I do not focus on the arbitrariness of setting any particular special threshold, although I believe that any such specification would invite new objections.
- 8 Of course, when drivers are not driving cars, they may be at risk just as much as non-drivers. However, *over a lifetime*, or any other given span of time, non-drivers will spend more time than drivers as vulnerable road users and at disproportionate risk.
- 9 There would be good consequences in these worlds, too. For instance, world was might become technically infeasible. However, in line with (R4), we are only considering the actual benefits of risky behaviors, i.e., the costs in these possible worlds which are avoided in the actual world.
- 10 I should point out that I am not making this argument simply because I intend to justify probability discounting; I actually think, to be consistent, that rights-based theorists *should* be this concerned about driving cars, along with many other technologies which are potentially rights-violating. Although it may well not be feasible to go back to a society without *any* rights-endangering transportation technologies, I think rights-based theories lead us to the conclusion that such a society would be in that respect morally superior.
- 11 Of course, I am not intending to imply that knowing is always extensionally equivalent with assigning a probability of 1, nor that not knowing is in general extensionally equivalent with assigning a probability of less than 1, but using the term in this manner is, I believe, perspicuous in the current context.
- 12 If one takes the concept of moral responsibility to apply only in a narrow range of morally weighty circumstances, please treat my wider usage as stipulative.
- 13 Many of the examples in the literature involve death (e.g., Harman 2015; Parfit 2011). I think this is problematic, since we have laws against killing, even when it is accidental or unintended (“manslaughter”), or when it is intended but unsuccessful (“attempted murder”). So one could think that another is morally responsible for some unsuccessful or unintentional killing, when one is instead conflating responsibility for the other’s killing for these other legally punishable offences. Another issue is that moral responsibility for murder might be so great that it is difficult to tell the difference between an intuition of complete moral responsibility for murder and (say)

- quarter responsibility for murder: both might be so much greater than any day-to-day level of moral responsibility as to be intuitively indistinguishable in the assignment of blame. For these reasons, none of my examples involve death.
- 14 Sepielli (2009), for instance, distinguishes between *narrow-scope* and *wide-scope* norms, depending on whether the norm applies to an individual given the set of subjective probabilities they have or whether the norm is satisfied by the individual also doing some action which leads them to revise their subjective probabilities. Harman (2015) makes a similar distinction (between “blameworthiness for behavior” and “blameworthiness for causing that behavior”). My first response is addressed to narrow-scopers or those who want to retain Harman’s distinction; my second is aimed at undermining these distinctions. One advantage of losing this distinction – making actions aimed at epistemic updating just extra choice possibilities – is that we have a more unified concept of moral responsibility (or blameworthiness) without it. However, both Sepielli’s and Harman’s discussions are about blameworthiness in the presence of moral uncertainty; mine is about non-moral uncertainty, which I believe to be more relevant in the context of climate change since most plausible moral systems converge decisively on urgent action (e.g., Stern 2014, 2015).
  - 15 Of course, there is the further worry for evidential probability advocates about how to specify which evidence is “available to the individual,” but that is well known (e.g., Timpe 2009).
  - 16 On my account, one always has *some* moral responsibility for not gathering more information, but that responsibility diminishes the less helpful one thinks information gathering would be.
  - 17 Let me stipulate that such charities exist, and, in climate contexts, an existence proof for such a good charity is one that helps indigenous people to protect rainforest, CoolEarth, [www.coolearth.org](http://www.coolearth.org).
  - 18 On the assumption that the subjective probabilities will decrease with respect to time, perhaps consistently, these two considerations allow us to answer the worries of Lenman (2000), about one’s inability to determine all of the consequences of an action (see also Burch-Brown 2014). If one cannot know all of the consequences of one’s actions, Lenman argues, then one cannot ever know that one acts rightly. My argument shows that, even for a consequentialist, it is rational to act even if one does not know rightly since it is rational to probability discount according to one’s subjective probabilities. For very difficult to foresee future events, rounding down to 0 percent – even if they may occur – is rationally defensible for finite beings like ourselves.

## References

- Broome, J. (1992), *Counting the Cost of Global Warming*, Cambridge: White Horse Press.
- Broome, J. (1994), “Discounting the future,” *Philosophy & Public Affairs*, 23(2): 128–56.
- Burch-Brown, J. (2014), “Clues for Consequentialists,” *Utilitas*, 26(01): 105–19.
- Caney, S. (2008), “Human rights, climate change, and discounting,” *Environmental Politics*, 17(4): 536–55.
- Caney, S. (2009), “Climate change and the future: discounting for time, wealth, and risk,” *Journal of Social Philosophy*, 40(2): 163–86.
- Dasgupta, P. (2008), “Discounting climate change,” *Journal of Risk Uncertainty*, 37: 141–69.
- Ellsberg, D. (1961), “Risk, ambiguity, and the savage axioms,” *The Quarterly Journal of Economics*, 75(4): 643–69.
- Harman, E. (2015), “The irrelevance of moral uncertainty,” in R. Shafer-Landau, (ed.), *Oxford Studies in Metaethics*, Vol. 10, Oxford: Oxford University Press, pp. 53–79.

- Kahneman, D. (2011), *Thinking, Fast and Slow*, London: Penguin.
- Knobe, J. (2006), "The concept of intentional action: a case study in the uses of folk psychology," *Philosophical Studies*, 130(2): 203–31.
- Lenman, J. (2000), "Consequentialism and cluelessness," *Philosophy & Public Affairs*, 29(4): 342–70.
- Parfit, D. (1984), *Reasons and Persons*, Oxford: Oxford University Press.
- Parfit, D. (2011), *On What Matters*, Vol. 1, Oxford: Oxford University Press.
- Sepielli, A. (2009), "What to do when you don't know what to do," in R. Shafer-Landau, (ed.), *Oxford Studies in Metaethics*, Vol. 4, Oxford: Oxford University Press, pp. 5–28.
- Shinar, D. (2012), "Safety and mobility of vulnerable road users: pedestrians, bicyclists, and motorcyclists," *Accident: Analysis and Prevention*, 44(1): 1–2.
- Smart, J. J. C. (1973), "An outline of a system of utilitarian ethics," in J. J. C. Smart and B. Williams (eds.), *Utilitarianism: For and Against*, Cambridge: Cambridge University Press, pp. 1–74.
- Stern, N. (2007), *The Economics of Climate Change: The Stern Review*, Cambridge and New York: Cambridge University Press.
- Stern, N. (2014), "Ethics, equity and the economics of climate change, Paper 1: Science and Philosophy," *Economics and Philosophy*, 30(03): 397–444.
- Stern, N. (2015), *Why Are We Waiting?*, Cambridge, MA: MIT Press.
- Timpe, K. (2009), "Review: *Living with Uncertainty: The Moral Significance of Ignorance*," *Notre Dame Philosophical Reviews*.
- Tiwari, G. (2015), "The safety of vulnerable road users: the challenge for twenty-first century," *International Journal of Injury Control and Safety Promotion*, 22(2): 93–4.
- Zimmerman, M. (2008), *Living with Uncertainty: The Moral Significance of Ignorance*, Cambridge: Cambridge University Press.
- Zimmerman, M. (2014), *Ignorance and Moral Obligation*, Oxford: Oxford University Press.