Methodologies of Uncertainty: Philosophical Disagreement in the Economics of Climate Catastrophes

Corey Dethier

What probability distribution should we use when calculating the expected utility of different climate policies There's a substantial literature on this question in economics, where it is largely treated as empirical or technical—that is, the question is what distribution is justified by the empirical evidence and/or economic theory. The question has a largely overlooked methodological component, however—a component that concerns how (climate) economics should be carried out, rather than what the science tells us. Indeed, the major dispute in the literature is over precisely this aspect of the question: figures like Nordhaus and Weitzman disagree less about the evidence or the theory than they do about which possibilities we should consider when making political decisions—or offering economic advice—about climate change. There are two important implications. First, at least some of the economic literature misfires in attempting to treat the debate as open to empirical or technical resolution; a better path to progress on the question involves further investigating the policy recommendations that can be derived from the two positions. Second, the choice of discount rate is entangled with the choice of probability distribution: as both choices are responsive to the same normative reasons, we cannot evaluate the arguments in favour of a particular discount rate without considering the implications of those same arguments for the choice of distribution.

1. Introduction

What probability distribution should we use when calculating the expected utility of different climate policies? There's a substantial literature on this question in economics, where it is largely treated as empirical or technical—that is, the question is what distribution is justified by the empirical evidence and/or economic theory. The question has a largely overlooked methodological component, however—a component that concerns how (climate) economics should be carried out, rather than what the science tells us. Indeed, the major dispute in the literature is over precisely this aspect of the question: figures like Nordhaus and Weitzman disagree less about the evidence or the theory than they do about which possibilities we should consider when making political decisions—or offering economic advice—about climate change.

Here's the argument, roughly: In the economic literature, efforts to address climate change are considered investments; we're putting away money now to decrease the future harms. In making an investment, there are some events whose probabilities we can (at least roughly) estimate based on past experience and known features of the world. Other events that might conceivably occur—*Independence Day*-style alien invasions, for example—are so outside the bounds of any experience that we've had that we can't pin down their probabilities in the same way. Essentially, the debate concerns whether to include possible consequences of climate change that fall in this later category in our economic models. This question can't be answered—at least not in any direct way—by investigating the empirical world or examining more closely the economic models themselves. Doing so might shift particular events from the latter category to the former, but it doesn't address the general question. This more general question is a methodological one: which possibilities should we consider when building economic models? As a methodological question, it's open to influence from practical and ethical considerations—as we'll see, I think disagreements here are what primarily drives the debate.

The plan is as follows. I begin by outlining the basics of the relevant economic theory (sec. 2). I then contrast the positions of Weitzman and Nordhaus, the former of whom argues for the use of 'fat-tailed' probability distributions in modelling the economics of climate change and the latter of whom defends the 'thin-tailed' orthodoxy, arguing that this disagreement is neither empirical (sec. 3) nor a matter of economic theory (sec. 4). To offer a better diagnosis, I compare the debate found in the economic literature to a similar debate in epistemology over the appropriate response to 'higher-order uncertainty' (sec. 5). The analogy helps us see the methodological character of the economic debate: in the economic setting, as in the epistemic one, the question concerns which possibilities to 'count' when constructing a probability distribution; unlike what's true in epistemology, by contrast, in the economic setting we're interested finding the best probability distribution to use rather than in determining which ones are rational. Hence a methodological disagreement rather than a (purely) epistemic one.

I end by sketching lessons for both the economic debate itself and the related ethical debate surrounding the discount rate. Regarding the former, I argue that at least some of the economic literature in the debate misfires in attempting to treat the disagreement as open to empirical or technical resolution and suggest that the best path to progress involves investigating potential differences in the policy recommendations that can be derived from the two positions (sec. 6). With regard to the ethical debate about discounting, I show that the choice of discount rate is entangled with the choice of probability distributions: as both choices are responsive to the same normative reasons, we cannot evaluate the arguments in favour of a particular discount rate without considering the implications of those same arguments for the choice of distribution (sec. 7).

2. The Economic Background

This first section lays out the basic theory employed in economic modelling of climate change and motivates the central question: what probability distribution should we use when calculating the expected utility of different climate policies?

Consider a very simple economic problem: we're trying to determine whether we should invest some of our present wealth. The cost of investment is that we'll be able to 'consume' less at present, where 'consumption' is here understood broadly to represent all of the goods and services that we use. The benefit is that we'll be able to consume more in the future. The question then boils down to how we value the change in future consumption relative to the change in present consumption; the higher the ratio, the more we ought to invest.

To a first approximation, we can treat the purely economic aspects of acting to mitigate climate change as an investment problem of exactly this kind.[1] The idea is simple: climate change represents a future economic burden—a foreseeable cost that we will have to pay—and investing now can lower that cost. Economically speaking, then, we can view acting to address climate change as an investment: paying a cost now to receive a benefit in the future.

Understanding climate change as an investment problem allows us to represent it using a simple economic model that formalizes the intuitive sketch given above. In our model, an agent's total welfare W is given by the sum of their current utility $\big(\mu(C_t)\big)$ and their discounted expected future utility $\Big(\beta \cdot \mathbf{E}\big[\mu(C_f)\big]\Big)$, where in both cases utility is a function of their consumption (C_t or C_f) and where the discount factor β captures how much more the agent values present utility relative to future utility:

$$W = \mu(C_t) + \beta \cdot \mathbf{E}\big[\mu(C_f)\big].$$

The choice that the agent faces is to exchange some current consumption (and thus utility) for future consumption. Since a rational agent should be willing to make any exchange that results in equal welfare, their marginal willingness to make an exchange is given by:

$$M = \beta \frac{\mathbf{E}\big[\mu'(C_f)\big]}{\mu'(C_t)}.$$

Which is just to say that their willingness to exchange a small amount of consumption today for a small amount of consumption in the future is a function of three variables:

[1] Purely economic because there may be important ethical or political aspects to addressing climate change that cannot be captured by this kind of model; for arguments aling these lines, see, for example, (Hartzell-Nichols [2017]; Jamieson [2014]).

first, how much they value a small change in present consumption $\left(\mu'(C_t)\right)$; second, how much they discount future utility relative to present utility (β); and, third, how they distribute their confidence over potential future marginal utility rates $\left(\mathbf{E}\left[\mu'(C_f)\right]\right)$.

Our focus in this article is on the last of these variables, the expected marginal future utility. For present purposes, what makes this variable interesting is the following observation. It's standard to assume that there's a decreasing marginal return on consumption: a small increase in the amount of consumption matters much more when you have nothing than when you're already extremely wealthy. Formally, that means that $\mu'(C)^{-1}$ is logarithmic, and that as you approach zero consumption, the marginal utility of a small increase in consumption becomes infinite. The behaviour of the curve as consumption approaches zero is widely seen as unrealistic: the marginal utility of a small increase in consumption should go up as consumption goes down, but shouldn't go to infinity. Still, the assumption of logarithmic behaviour is considered a good approximation generally and the behaviour near zero consumption is unproblematic so long as the probability of being in a zero-consumption scenario is infinitesimal.[2]

There's a solid consensus within climate science that the truly catastrophic scenarios—the scenarios where civilization collapses and thus that might reasonably be represented as 'zero consumption' in the model—are highly unlikely. But highly unlikely is not the same as infinitesmally unlikely. Introducing a small but not infinitesmal probability of a zero-consumption scenario into the model has the effect of massively increasing the expected future value of a small change in consumption $\left(\mathbf{E}\left[\mu'(C_f)\right]\right)$ relative to the present value $\left(\mu'(C_t)\right)$—indeed, absent other assumptions it has the effect of driving the former to be infinitely larger than the latter. When we plug the resulting numbers back into the model, the outcome is that we should be willing to pay any price today for a marginal reduction in the expected consequences of climate change. So: employing a probability distribution that assigns non-infinitesmal probability to zero-consumption scenarios has a dramatic effect on the results given by the model and (thus) on the policy that it recommends (though see sections 4 and 6).

Hence the question: what probability distribution should we use when calculating the expected utility of different climate policies? Or, slightly more precisely, should we use a probability distribution with 'thin tails'—in which case the zero-consumption scenarios have an infinitesmal probability—or ones with 'fat' (sometimes: 'thick') tails, in which case the probability of zero-consumption scenarios is still small but is large enough to throw off the expected utility calculations in the manner described in the last paragraph?[3]

[2] That's not quite right; what matters is actually the relationship between the slope of the probability distribution and the slope of the marginal return on consumption (see Pindyck [2011]). For present purposes, though, we can focus on the simpler condition.

[3] As one reviewer rightly stressed to me, there is another way that we could view this debate: not over what probability distribution to use, but over what events to distribute probability over. In some ways, at least on the reading that I'll be offering, this is actually a more telling way of understanding

3. The Debate over Fat Tails

The most famous defence of the use of fat tails is offered by Weitzman ([2007], [2009a],[2009b], [2011], [2012a], [2013]). The thin-tailed position is more orthodox, but Nordhaus ([2007a], [2011], [2012], [2013]) is one of its most prominent and most explicit proponents. In this section, we'll examine the arguments that each offer in favour of their preferred distribution; as we'll see, this disagreement is not—at least not primarily—over the empirical facts.[4]

It's helpful to begin with what looks like an empirical disagreement. Weitzman motivates preferring fat-tailed distributions on the grounds that there are scenarios, deemed possible but unlikely by climate scientists, in which certain feedback effects and/or discontinuous events (what have come to be called 'tipping points') would lead to dramatic amplification of global warming. Here's a representative passage:

> A more remote (but even more vivid) possibility, which in principle should also be included, is heat-induced releases of the even-vaster offshore deposits of CH_4 trapped in the form of hydrates (aka clathrates)—which has a decidedly non-zero probability over the long run of having destabilized methane seep into the atmosphere if water temperatures over the continental shelves warm just slightly. The amount of CH_4 involved is huge, although it is not precisely known. Most estimates place the carbon-equivalent content of methane hydrate deposits at about the same order of magnitude as all fossil fuels combined. Over the long run, a CH_4 outgassing-amplifier process could potentially precipitate a disastrous strong-positive-feedback warming. (Weitzman [2009a], p. 7)

On first pass, then, Weitzman's argument is that we should assign some non-infinitesmal probability to zero-consumption scenarios because there are potential mechanisms such as the outgassing of CH_4 that—if activated—would lead to massive amounts of warming. And massive amounts of warming would have catastrophic effects on human life.

Nordhaus appears to disagree. As he explicitly recognizes, his choice to use thin-tailed distributions amounts to the assumption 'that there are no genuinely catastrophic outcomes that would wipe out the human species or destroy the fabric of human civilization' (Nordhaus [2007a], p. 33). In a later critique of Weitzman, he echoes the same sentiment, arguing that Weitzman's characterization of various disastrous scenarios as true zero-consumption outcomes is unconvincing—true zero-consumption scenarios should not even be considered remote possibilities (Nordhaus [2011], pp. 252–53).[5] On first inspec-

the disagreement. It isn't how the extant debate has been framed, however, so I'll stick with the formulation in the main text.

[4] A useful overview of the economic literature as of 2013 is (Millner [2013]); for an introduction aimed at a more general audience, see (Wagner and Weitzman [2015]).

[5] What counts as a 'true' zero-consumption outcome is a semi-arbitrary modelling decision. My own view is that the Weitzman's argument doesn't really turn on exactly how we represent catastrophically bad scenarios in the model; see the next section for details.

tion, therefore, Weitzman and Nordhaus simply disagree about whether the empirical data support assigning non-infinitesmal probability to zero-consumption scenarios.

Close inspection, however, reveals that the apparent empirical disagreement is not nearly as dramatic as it initially appears. In the two paragraphs following the above-quoted sentence, Nordhaus goes on to stress how much uncertainty there is surrounding catastrophic outcomes, particularly in the long term. His position is not that such outcomes are truly infinitesimally unlikely so much as it is that we don't know precisely how unlikely they are, which inhibits making any sort of informed judgement about their probability (see also Nordhaus [2013], pp. 59, 66). Similarly, at no point does Weitzman insist that the zero-consumption outcomes are even moderately likely. He agrees with Nordhaus that the best characterization of the situation is that we don't know how unlikely such scenarios are (Weitzman [2009b], p. 5); his concern is really that they might be more than infinitesimally unlikely and his view is that our models should account for this possibility.

In other words, the disagreement between Weitzman and Nordhaus concerns how we should respond to uncertainty about the probability of zero-consumption outcomes, with the latter arguing in favour of ignoring them given that they are at the very least highly unlikely and the former arguing in favour of including them in the analysis. Looking more closely at the arguments that Weitzman gives for his position—particularly in (Weitzman [2009b])—reinforces this conclusion. Recall, the question here is what probability distribution we should use when calculating the expected marginal future utility. That is, the probabilities that we're concerned with are the probabilities of various future consumption scenarios. For present purposes, we can think of future consumption as a function of temperature and temperature as a function of CO_2 concentrations. We know (or can stipulate) the relevant CO_2 levels; what we want is to assign probabilities to different future consumption scenarios.

A simplified version of Weitzman's argument is as follows. Suppose that we know the relationship between temperature and consumption.[6] Then the only problem is to determine the relationship between temperature and specific CO_2 scenarios. There is, of course, uncertainty about the true relationship between these two variables. One way to proceed here is to use past temperature readings to generate a single probability distribution that captures this uncertainty. We could then treat this probability distribution as known when calculating the probability of different consumption scenarios—in effect, this means behaving as though there's no uncertainty about whether a particular event has (or should be assigned) a probability of x or y.

[6] In the present context, where we're concerned primarily with disastrously high temperature concentrations, this assumption isn't wholly unreasonable. In other contexts, however, it's extremely dangerous: as McLaughlin ([2024]) stresses, the idea that 'we'll' survive so long as we can avoid scenarios where temperatures alone kill us is entirely unjustified.

The problem with this approach is that past temperatures aren't sufficient to fix a single probability distribution for the relationship between CO_2 and temperature. Part of the reason for this is that there isn't a single simple mechanism relating the two variables but instead a host of complex and interrelated mechanisms. The Intergovernmental Panel on Climate Change (IPCC) uses a summary variable called equilibrium climate sensitivity (ECS) to capture the effect of all of these different processes. Since some of the mechanisms are well understood, there are aspects of ECS that are relatively well constrained. For instance, the lower bound on ECS is determined almost entirely by the greenhouse effect and both this lower bound and the 'best guess' range have remained largely steady since the 1970s. By contrast, some of the mechanisms are relatively poorly understood—such as the potential outgassing of stored CH_4 mentioned earlier—and these poorly understood mechanisms have the potential to have a substantial positive feedback effect on temperature. The upshot is that there's far more uncertainty about the upper end of possible values for ECS than on other aspects. As the IPCC put the point in the same year that Weitzman published his first article on the subject:

> Most of the results confirm that climate sensitivity is very unlikely below 1.5°C. The upper bound is more difficult to constrain because of a nonlinear relationship between climate sensitivity and the observed transient response, and is further hampered by the limited length of the observational record and uncertainties in the observations, which are particularly large for ocean heat uptake and for the magnitude of the aerosol radiative forcing. Studies that take all the important known uncertainties in observed historical trends into account cannot rule out the possibility that the climate sensitivity exceeds 4.5°C, although such high values are consistently found to be less likely than values of around 2.0°C to 3.5°C. (Solomon et al. [2007], p. 798)

The science has advanced in the intervening years, though the IPCC still assigns lower confidence to the upper 'very likely' bound of 5.0°C than to either the lower bound or the 'likely' envelope (Masson-Delmotte et al. [2021], p. 1006) .[7]

Given this uncertainty about ECS, Weitzman argues, we cannot treat the relationship between CO_2 and temperature as known. What does this mean? Weitzman asks us to imagine the best-case scenario: the relationship between the two variables can be characterized by a normal distribution. As noted above, treating the relationship as known involves estimating the 'true' normal distribution, that is, we assume that the true distribution has the same (or the same up to standard corrections) variance as the observed distribution. Subsequently using that distribution—plugging it in when calculating the probability distribution for expected consumption—amounts to ignoring

[7] For an outline of the IPCC's use of probabilistic terminology and the relationship between 'confidence' and 'likelihood' judgements, see (Mastrandrea et al. [2010]); for a philosophical discussion, see (Dethier [2023]). See also (Steel et al. [2023]), which argues that 'climate collapse' is a real risk at much lower temperatures than those considered by Weitzman.

the possibility that our estimate is wrong. To reiterate the main point: it amounts to treating our best guess at the distribution as something we know, rather than something that we're uncertain about.

Treating the distribution as unknown, by contrast, means treating the observed relationship between CO_2 and temperature as a sample from the true distribution. Things work slightly differently in classical and Bayesian approaches here, but the lesson is the same: if we allow there to be some uncertainty about whether features of the sample (particularly the sample variance) actually line up with the true population, the resulting estimate for the true distribution will be a t distribution rather than a normal one. That is, building the possibility that we're wrong about the true distribution into our estimate results in a distribution with a different shape than ignoring that possibility. And, crucially, a t distribution has much fatter tails than a normal one.

The point of this excursion into Weitzman's reasoning is to emphasize that while the argument does have empirical foundations—it's thoroughly based in the results from climate science—the key contention is not that the true probability of particular catastrophic scenarios is non-zero. On the contrary, the central contention of the argument is that we're uncertain about what probability to assign these scenarios and therefore that we should assign them non-zero probability when calculating the expected marginal future utility. And as we've already seen, this is precisely where Nordhaus disagrees: at least as of 2007, his view was that our inability to pin down a probability for these (highly unlikely) scenarios was a reason not to assign them a non-zero probability (Nordhaus [2007a], p. 33).

4. A Matter of Economic Theory?

Recall the question: what probability distribution should we use when calculating the expected utility of different climate policies? The takeaway from the last section is that the economic debate over this question is not primarily a matter of empirical disagreement. Nor, however, is it primarily a matter of a technical disagreement concerning economic theory—or so I'll argue in this section.

Just as there were reasons to think that the debate is primarily empirical, there are reasons to think that it is primarily technical. Weitzman's articles on the subject are themselves highly technical at various points, and much of the subsequent economic literature has focused on these technicalities. Anthoff and Tol ([2022], p. 885), for example, characterize Weitzman's 'dismal theorem' as the claim that 'the expected net present value of a stock problem with a stochastic growth rate with unknown variance is unbounded'. The expected net present value of a stock problem is not something we can (directly) observe. While the boundaries are blurry here, this claim is a matter of economic theory and is open to investigation by (for example) mathematical models in the same way that other elements of economic theory are. Indeed, this is how both Anthoff and Tol and quite a few other economists approach the problem.

Similarly, much of the subsequent debate in the literature concerns various assumptions of Weitzman's model and the extent to which his formal results do (or don't) hold when those assumptions are weakened. Nordhaus ([2011], p. 250) again offers a characteristic example. He characterizes Weitzman's argument as relying three assumptions: fat-tailed distributions, high risk aversion, and potentially unbounded costs. As in the last example, we may not be able to demonstrate that (say) costs are potentially unbounded empirically, but there are technical arguments to be given concerning the inclusion of such assumptions in the model. Other work, by Arrow ([2009]) and Cato ([2020]), likens Weitzman's results to the St. Petersburg paradox, which might suggest the disagreement here concerns technical questions about decision theory that are raised by the paradox.

In fact, I think that much of economic literature does take the debate to be a technical one, at least in principle answerable by the clever application of models or other mathematical tools (an important exception is Millner [2013]). This reading simply doesn't square with Weitzman's original argument, however. As we've seen, Weitzman argues in favour of employing a probability distribution that assigns non-infinitesmal probability to zero-consumption scenarios in which the marginal utility of consumption of utility is infinite. Within the context of the assumptions laid on in section 2, this choice causes the expected utility of a marginal change to future consumption to go to infinity as well, and thus implies—at least in the context of the model—that an agent should be willing to pay infinite amounts of money to receive a marginal increase in future consumption.

Taken literally, this result is ridiculous. And the model that generates the result is consciously and explicitly unrealistic. The obvious conclusion is that we should not take the result literally—a point that Weitzman reiterates repeatedly from his first article on the subject onwards. In particular, while zero-consumption scenarios and infinite utilities drive the shocking mathematical results, they aren't central to what (Weitzman [2009b], p. 11) sees as the main takeaway: 'One can easily remove the $+\infty$ in [the model], but one cannot so easily remove the underlying economic problem [...] The take-away message here is that reasonable attempts to constrict the length or the fatness of the 'bad' tail (or to modify the utility function) still can leave us with uncomfortably big numbers whose exact value depends nonrobustly upon artificial constraints or parameters that we really do not understand'. As we've seen, there are scenarios that are ignored by standard economic approaches because we have no grounds for assigning them a precise probability. Weitzman's point is that these scenarios are largely 'bad' and thus including them in our calculations has a dramatic effect on the results—as illustrated by the simple model in which everything goes to infinity. The model is unrealistic, and so introducing technical restrictions on how 'bad' a possibility can be (Weitzman's 'artificial constraints' and 'parameters') is liable to improve it in various ways. But we have no grounds for particular restrictions of this kind, and the exact character of the results will depend heavily on which restrictions we choose to employ.

As Weitzman sees it, in other words, the point of his critique is not the technical results themselves—what Millner ([2013], p. 318) characterizes as an 'abstract parable'—but rather that these results reveal important deficiencies in the tools and methods that economists choose to employ in representing climate change. Which means that treating Weitzman's critique and the questions raised by it as technical economic questions largely misses the mark. While there are many worthwhile reasons to examine whether current economic models exhibit fat-tailed behaviour as Anthoff and Tol ([2022]) do, such an investigation cannot help us evaluate Weitzman's criticisms given that said criticisms are aimed at the assumptions underwriting those economic models. Similarly, while it may well be relevant that Weitzman's formal results depend on particular assumptions as stressed by Nordhaus ([2011]) and many others, the formal results are only intended to be illustrative or motivating—Weitzman certainly doesn't want to endorse them as anything like a correct representation of how we ought to approach climate change. At minimum, therefore, criticizing the modelling assumptions undermines his arguments only if those assumptions are crucial not just for accurate representation but for the illustrating function as well.

I do not mean to imply that there are not technical disagreements to be had here or that there are no possible technically interesting criticisms of Weitzman's argument. For example, Weitzman ([2009b], pp. 16–17) suggests that one practical implication of his arguments is that the Monte Carlo-style simulations typically carried out in more complex economic models are liable to under-represent catastrophic scenarios (see also Tol [2003]). If it could be demonstrated that policy recommendations were robust when these scenarios were 'properly' represented in the Monte Carlo simulations, that result would indicate that there's something special about Weitzman's modelling assumptions that doesn't extend to more realistic applications. In so doing, it would undercut the ability of the simple model to perform the illustrative role that it plays in Weitzman's arguments. And, thus, it would seriously undermine the force of his criticism. To my knowledge, however, no such study has been carried out.

At present, in other words, there is not much interesting technical disagreement to be found in the economic debate—at least, there isn't enough technical disagreement to account for either the existence or prominence of the debate in the economic literature. Perhaps we should take this result to simply indicate that the debate has misfired. I think that would be a mistake, however, and in the next section I'll argue that there is an interesting disagreement to be found here—but it is a methodological disagreement rather than an empirical or technical one.

5. Diagnosing the Disagreement

In the economic literature, the debate surveyed above is often characterized as a disagreement about how to make decisions under conditions of 'deep uncertainty', or uncertainty about what probability to assign to various scenarios. (Deep uncertainty is sometimes called 'great', 'severe', 'Knightian', or 'Keynesian' uncertainty or simply 'uncertainty' and opposed to 'risk', where risk occurs when probabilities are known.) To better understand the present debate, therefore, we might benefit from examining the philosophical literature on the same subject.[8]

On the other hand, while it's certainly right to say that the debate is about decision making under uncertainty in some sense, this description fails to capture the precise nature of the disagreement. As we've seen, the primary question at issue is which probability distribution to assign under conditions of deep uncertainty. We wouldn't normally characterize two expected utility maximizers with different probability distributions as disagreeing about 'how to make decisions', and that's essentially the nature of the disagreement that we find here. As such, a more promising comparison is to be found in that area of philosophy concerned with the choice of probability distributions—namely, epistemology.

When we examine the epistemic literature, we find that a similar problem has been taken up under the heading of 'higher-order uncertainty' or uncertainty about what beliefs or credences to adopt.[9] The standard motivating example in this literature goes something like this. An agent has some evidence that supports a particular belief that q. They then receive new evidence, which does not support $\neg q$ but does indicate that the original evidence—and/or the agent's evaluation of it—is not trustworthy in some other way. (Standard examples include evidence that the agent has ingested a drug that impairs their ability to carry out simple addition.) This evidence is 'higher-order' because it is relevant to the agent's attitudes about q but not relevant to q itself. The question is then what attitudes the agent should adopt in face of the new evidence.

Some philosophers hold that the new evidence should make the agent uncertain that they should believe that q, but that this uncertainty should not—or, at least, should not always—affect their belief that q itself. This position requires endorsing the rationality of what's called 'epistemic akrasia'—believing q while also believing that one should not believe that q or that one's evidence does not support q—a condition that many other philosophers take to be inherently irrational. Dorst ([2020]) calls this first group 'splitters'; 'bridgers', by contrast, aim to find principles for incorporating the higher-

[8] There's been an increase in philosophical interest in this question in recent years due in part to climate change and in part to the development of imprecise probability measures; see, for example, (Rinard [2015]; Bradley [2017]; Helgeson [2022]).

[9] For a survey, see (Horowitz [2022]); my treatment draws heavily from (Dorst [2019], [2020]).

order evidence into the first-order evaluation. The problem for bridgers is that intuitive bridging principles are prone to counter-examples where applying the principle forces the agent into either scepticism or inconsistency.

There are important parallels between the questions addressed in this literature and those that motivate the economic one. The disagreement between bridgers and splitters is a disagreement about how to respond in the face of (rational) higher-order uncertainty: what should an agent believe about q when they have good reason to believe that their evidence concerning q is partial, untrustworthy, or deficient in some other way? Splitters, like Nordhaus, advocate adopting those beliefs best supported by the relevant evidence: if the relevant evidence does in fact support q, then q is what the agent should believe—even if they're rationally uncertain that they should believe q because they have good reasons to doubt (their evaluation of) the relevant evidence. By contrast, bridgers advocate for positions like Weitzman's: when faced with higher-order uncertainty, the beliefs that an agent adopts should reflect or incorporate that uncertainty.

Of course, while Weitzman and Nordhaus are engaged in an epistemic debate in some important sense, they're not concerned with what agents should believe, let alone what it is rational for agents to believe. Instead, they're concerned with the much more practical question of what probability distributions we should use when calculating expected utilities, and specifically which ones we should use when calculating expected utilities in the context of policy decisions concerning (for example) carbon taxes.

This difference in contexts drives the economic question apart from the questions asked in epistemology. For one thing, whereas many philosophers want to exclude cases where the relevant source of uncertainty is a mere lack of evidence from the discussion, these cases are paradigmatic for the debate in economics. The reason why is simple. Philosophers are concerned with what individual agents should do, and it is (thus) appropriate for them to appeal to agent's subjective priors to dissolve cases where the (sole) source of uncertainty is the lack of evidence. The same move is unjustified in the context of the debate between Weitzman and Nordhaus, who are concerned with providing expert advice regarding public policy. It would be inappropriate for an expert to present conclusions that depended overly on the choice of subjective prior as though they were the results of scientific investigation or determined by scientific evidence. As a consequence, while the epistemologist's problem arguably arises only when they encounter a very specific kind of evidence, the economist's can arise whenever the empirical evidence fails to pin down a single probability distribution.

For another, the debate between Weitzman and Nordhaus turns on practical considerations in a way that the debate in epistemology does not. For instance, Nordhaus argues that individual modellers are inevitably constrained by their own perspective and so cannot possibly hope to account for all of the different uncertainties and factors that should be considered in making climate policy (Nordhaus [2007a], p. 32). On his view, one reason that economic modellers should not attempt to build every uncertainty into their models is that 'individual scholars' are less well suited to accounting for the full range of uncertainties than 'committees and panels' of people with diverse areas of expertise (Nordhaus [2007a], p. 33). Similarly, though Nordhaus never to my knowledge puts the argument in these terms, one can easily extract from his work a criticism of Weitzman that the latter is allowing the perfect to be the enemy of the good: in deciding what policy to make today, we should work with our current 'best guesses' about the climate, on the understanding that as more information comes in, we'll revise and adjust our policies to account for the inevitable 'surprises' that it brings (Nordhaus [2013], p. 66).[10]

The debate between Nordhaus and Weitzman is thus akin to the debate in epistemology in some respects and distinct from it in others. On the one hand, like the debate in epistemology, the economic debate concerns the appropriate way to represent the world as it is—rather than how to represent the world as it should be or what rules agents should follow when making decisions. More precisely, it concerns the question of how to incorporate higher-order uncertainty into our representation of the world. Or, in slightly different terminology, both debates concern whether we should 'count' certain possibilities when we are evaluating which representation to use.

On the other hand, the economic disagreement is methodological rather than (purely) epistemic in that it concerns the appropriate practices to employ in pursuing the epistemic products of science—products like the probability distributions deployed in calculating the expected utility of different political decisions.[11] Whereas epistemology is (often) taken to exclude considerations of feasibility or practicality, practicality is a quintessential methodological concern: a method of testing that relies on infeasibly strong experimental controls isn't a good method, whatever its epistemic credentials.

The similarity between the economic and epistemic debates indicates that the economic question, like the epistemic, is normative—it's about how we should proceed—and thus not easily amenable to empirical investigation or standard modelling techniques. In epistemology, where what's at issue is precisely what scenarios 'count' for the evaluation of the appropriateness of the resulting beliefs, we cannot use mathematical techniques to simply settle the debate by (say) demonstrating that one option leads to more accurate beliefs that the other (Dethier [2022]). Similarly, we cannot settle the economic question by appealing to standard economic tools such as (for example) the expected utility of the different distributions, as what's at issue is how different possibilities figure into

[10] A third difference, less interesting for present purposes, is that both Nordhaus and Weitzman are more concerned with solving this particular problem than with identifying a rule to apply to all cases that involve uncertainty of this kind.

[11] The boundaries between methodology in this sense and epistemology are blurry—as made clear by recent debates over 'inquiry' (for example, Falbo [2023]; Friedman [2020])—but drawing the distinction in terms of concern with communal products like scientific evidence as opposed to mental attitudes like belief is a good first pass (see Staley and Cobb [2011]).

those calculations. What we need are normative grounds for excluding or including those possibilities.

The contrast between the economic and epistemic debates is also telling, however, in that it indicates that the normative considerations that are relevant to the economic debate are unlikely to be identical to those that relevant to the epistemic one. In the economic context, what's needed are arguments for why certain possibilities should or shouldn't be counted in political decision making. These arguments might be more purely epistemic—that is, we can't rule this or that possibility out—but they also might be practical, political, even ethical. We've already seen both Weitzman and Nordhaus making arguments of this sort, the former explicitly appealing to precautionary principles in motivating his choice of probability function (Weitzman [2007], p. 719, [2009b], p. 12) and the latter to feasibility arguments for preferring an achievable good option over an out-of-reach perfect one (Nordhaus [2007a], p. 33, [2013], p. 66).

To summarize what we've learned, then: the debate between Nordhaus and Weitzman is (primarily) methodological in that it concerns a normative question about how climate economics should be carried out. Specifically, it concerns what possibilities should get counted when evaluating climate policies. And, as such, the debate is unlikely to be resolved by either purely empirical or purely epistemic considerations—practical and political realities are at least as relevant as what the empirical evidence, economic theory, or epistemic first principles have to say.

6. Re-orienting the Economic Debate

What's gained by properly diagnosing this debate as methodological? We've already seen (in sec. 4) that some of the extant economic debate misfires by attempting to treat what is ultimately a philosophical disagreement as a matter of economic theory. This is not true of the entirety of the literature, however. In particular, a number of Weitzman's critics have pressed the point that there's a gap between Weitzman's results and actionable policy recommendations. In this section, I'll argue that this gap deserves more focus—both as a potential objection to Weitzman and insofar as closing it provides a better understanding of the implications of the debate.

Let's start with the objection. Notice that the model given in section 2 makes no mention of policies or of the differential effects thereof: from the perspective of the model, we could spend our infinite amounts of money on arresting climate change or on building a giant bouncy castle. So long as the effect is a marginal increase in future consumption, the model doesn't say anything in favour of either alternative. In fact, the problem is deeper: the model doesn't necessarily recommend action at all. As Nordhaus ([2011], p. 251) puts it: 'there is no particular relationship between the impact of fat tails on expected utility and the impact of fat tails on the optimal policy. For example, it might well occur that the outcome of an asteroid collision has very fat tails; but if we cannot prevent collisions, then policy is tail-irrelevant'. In other words: absent reason to think that some policies have differential effects on the size of the tail, the move from thin- to fat-tailed distributions vastly complicates the mathematics while making essentially no difference to policy.

Of course, we do have reasons to think that some policies have differential effects on the size of the tail. Policies that result in lower CO_2 emissions can be expected to thin the tail relative to other policies that result in higher ones. The reasoning here is simple. We're worried about mechanisms that lead to discontinuities in the relationship between CO_2 and temperature such as the possibility of large-scale CH_4 outgassing, which would serve as a major feedback effect. We don't know exactly how much CO_2 we would have to emit to trigger these mechanisms, but if they trigger, things will be much worse than if they don't. Our uncertainty about whether (and when and how) these mechanisms will fire is what generates the fat tail in the probability distribution over future consumption scenarios. Keeping CO_2 low reduces the probability that these mechanisms fire and thus serves to thin the tail. Essentially this reasoning drives researcehrs like Hwang et al. ([2016]) and Weitzman ([2013]) to conclude that employing a fat-tailed distribution motivates adopting a much more stringent carbon tax than is normally derived from thin-tailed distributions.

While fine as a first pass, the simple argument of the prior paragraph faces a substantial problem, namely, the existence of other mechanisms that could result in zero-consumption scenarios. This problem has typically been raised to stress the potential opportunity costs associated with devoting large sums of money to climate change; the more we devote to addressing climate change, the less we have to stave off nuclear war or the relevant technological apocalypse of the month (see, as usual, Nordhaus [2011], pp. 254–55). There is empirical disagreement to be had about which of these catastrophes is most likely, but Weitzman's fat-tailed approach suggests that we should treat all of them extremely seriously (Weitzman [2009b], p. 14). Why is climate change special?

The potential opportunity costs of spending money on climate change rather than some other catastrophic threat fail to capture the full scope of the problem. To my eyes, the more interesting concern is that it's plausible that the value of 'thinning the tail' by decreasing the probability of a climate catastrophe depends in part on how probable catastrophes are in general. For example, reducing the risk of climate catastrophe in 200 years from (say) 1% to 0% is extremely valuable if we know that there are no other sources of catastrophe that will threaten human civilization within the same time span, but would plausibly be much less valuable if we were very confident that civilization would be destroyed by some other mechanism.[12] Taking Weitzman's arguments seriously means

[12]Thorstad ([2023]) shows that the idea expressed here—namely, that there's an inverse correlation

employing a thicker tail not just for climate change catastrophes but for all catastrophes—indeed, may involve a much thicker tail in cases like (say) a Skynet-style AI that are far more speculative than the potential catastrophes associated with climate change.

To see why this matters, an example will be helpful. Millner ([2013], pp. 312–14) shows that we can extend the simple model given in section 2 to account for policies with differential tail effects. The rough idea is that a policy is optimal only if it thins the tail to the point where the expected marginal utility of an additional investment is non-infinite. Imagine that for large values of T, the temperature distribution behaves as $T^{-f(a)}$. Here, $f(a)$ is a non-decreasing function that measures the 'thinness' of the tail of the probability distribution for future temperatures given some policy intervention of 'size' a. So: the bigger $f(a)$, the more the intervention serves to thin the tail, the better off we are. Millner's result is that if we let γ measure the slope of the damage function for climate change and η measure the elasticity of marginal utility (essentially the degree of risk aversion), the tail will be sufficiently thin when:

$$f(a) > 1 + \gamma(\eta - 1).$$

Of course, none of $f(a)$, γ, or η is easy to fix. But we can estimate the character of the tail distribution for a policy that would serve to stabilize CO_2 concentrations at twice pre-industrial levels using ECS. Employing also Nordhaus's relatively conservative estimate for γ, Millner shows that such a policy can be optimal only if $\eta < 1.5$—which is an extremely significant constraint on policy.

Millner, like Weitzman, is working with a toy model; his aim is to show that introducing fat tails can lead to distinctive policy recommendations. Insofar as our goal is to actually determine the policy ramifications of Weitzman's arguments, however, simply introducing fat tails into $f(a)$ is insufficient. The same reasoning offered by Weitzman also motivates complicating the model by incorporating uncertainty with respect to γ and allowing η to vary with the degree of background risk. The problem here is not just a lack of the philosophical virtue of consistency on the part economic modellers. Instead, it's that attending to possibilities that cannot be assigned a precise probability in some contexts but not in others is liable to skew the analysis by making some policy proposals appear better than they would in either a consistently thin-tailed or a consistently fat-tailed setting. That is: regardless of which side of the debate we find ourselves on, we should be suspicious of economic models that build in climate catastrophes that cannot be assigned a precise probability but that ignore other scenarios whose probabilities are

between the probability of 'existential risk' and the value of working to address said risk—holds for a variety of models that measure value by total utility. On the other hand, a venerable tradition in economics holds that risk aversion should be proper: taking on one undesirable risk shouldn't make another more desirable (see, for example, Pratt [1987]). And empirical studies (for example, Guiso and Paiella [2008]) indicate that agents tend to be more risk averse in the presence of other risks. I take it that squaring these various attitudes towards risk is an open problem.

similarly unknown. The implications of Weitzman's critique simply cannot be evaluated absent models that approach uncertainties about the proper probability assignments more systematically.

If I'm right, in other words, one upshot of my diagnosis is that the economic debate should be re-oriented away from attempts to 'test' Weitzman's view by way of the same economic models that he critiques and towards the construction of models that allow us to systematically evaluate the policy implications of the two different views. The re-orientation would serve both to explore the most significant criticism of Weitzman's position and to better indicate what consequences (if any) follow from 'counting' scenarios with uncertain probabilities when making policy decisions.

Notice the importance of the diagnosis for this recommendation. Were the debate primarily empirical or technical, the policy implications of the two approaches would be at best beside the point; the question would simply then be which probability distribution is best supported by the evidence. Whether the two distributions lead to different policy recommendations is irrelevant to this question, and choosing one over the other on the grounds that we prefer its implications looks dangerously close to motivated reasoning (cf. Frank [2019]). Policy implications are similarly irrelevant in the context of purely epistemic debates about rationality, and for the same reason. In either case, it would be bizarre to recommend resolving the debate by examining the practical payoffs.

In the context of a methodological disagreement, by contrast, the recommendation makes much more sense. As we noted above, practical and ethical considerations such as the feasibility of different methods are quintessential grounds for methodological decisions. If we were able to show that adopting Weitzman's preferred approach rather than Nordhaus's led to no important difference in policy recommendations, that would make no difference to the epistemic standing of either approach, but it would give us very good methodological reason to prefer Nordhaus's given that it is much simpler, and thus less costly in both time and money. If the debate is methodological, therefore, determining the policy implications of Weitzman's proposal is imperative in a way that it otherwise would not be.

7. Implications for the Discounting Debate

As we saw in section 2, the marginal willingness to make an investment depends on the discount factor, β. In the famous report that bears his name, Stern ([2007]) cites normative considerations to argue in favour of almost no discounting. Various critics—including both Nordhaus ([2007b]) and Weitzman ([2007])—argued that Stern's normative arguments are inappropriate, as discount factors should be decided empirically (or democratically), rather than by the normative views of scientists. This economic debate engendered a substantial philosophical literature—or, really, re-ignited an already existing discussion about how we should value future utility relevant to present utility.[13] To date, however, the philosophical discussion has almost entirely been conducted independently of the economic debate about probability distributions.[14] In this final section, I argue that it shouldn't be: we can't separate the problem of discounting from the methodological questions discussed in this article.

In order to understand the connections between the choice of probability distribution and the choice of discount rate, we need to dive a little bit further into the technical details. As discussed above, discount factor β captures how we value consumption at some future time, t_i, relative to consumption today. This factor is a function of the discount rate, r:

$$\beta = \exp\Big(- \int_0^{t_i} r(t)dt \Big).$$

For our purposes, the crucial feature of this relationship is that the smaller the discount rate, the closer the discount factor is to one, and thus the less we ultimately discount future consumption. The discount rate, in turn, is typically taken to be a function of three quantities: the 'pure time preference', δ, the elasticity of marginal utility, η, and the growth rate, g:

$$r = \delta + \eta g. \tag{1}$$

Of these three quantities, g and to a slightly lesser extent η are typically taken to be open to empirical estimation. The heart of the (philosophical) debate is over δ, the rate of pure time preference. Philosophers (for example, Parfit [1984]) and economists (Stern [2007]) who favour low values argue that any significant pure time preference is impermissible on the grounds that it's wrong to weight some people's lives over others based solely on when they were born. (According to this line of thinking, a small pure time preference may be permissible to account for the possibility of human extinction.) Philosophers (Heath [2021]) and economists (Arrow [1999]) who favour higher values argue that truly valuing the future at the same level as the present is overly demanding on the present given the high probability that the number of people currently existing is much smaller than the number of people who will exist.

The methodological questions discussed in this article have a variety of potential consequences for the debate over pure time preference. The alleged consequence that has

[13] Mintz-Woo ([2021]) offers an introduction to the philosophical literature; for a discussion of the technicalities involved, see (Greaves [2017]). Additionally, note that empirical surveys by Drupp et al. ([2018]) and Nesje et al. ([2023]) indicate that there is relatively little difference between the average opinions of philosophers and economists on the subject.

[14] So far as I'm aware, (Hartzell-Nichols [2017]) offers the only extended discussion of the economic debate over probability distributions in this context, and does so in service not of arguing for a particular discounting rate but instead of arguing for rejecting cost–benefit analysis altogether. Though see also (Broome [2010]).

received the most attention is Weitzman's claim that the debate about pure time preference is much less important than the choice of probability distribution. Here's a sketch: As Weitzman stresses, equation 1 is actually a special case—the special case where the growth rate is a deterministic constant. If the growth rate is instead a normally distributed random variable, $g \sim \mathcal{N}(\mu, s^2)$, equation 1 becomes:

$$r = \delta + \eta\mu - \frac{1}{2}\eta^2 s^2.$$

This is just to say that the discount rate depends on the variability of growth.[15] Weitzman then argues that the final term of equation 7 has a much larger effect on r than δ does when we account the possibility of climate catastrophe (cf. Broome [2010]; Dietz [2011]).

Suppose that Weitzman is right. Then disagreements about the discount rate (largely) wash out. That's important for policy-making purposes, but does not—so far as I can see—have any particularly interesting implications for the philosophical debate over discounting. That is, I take it that the ethical questions that arise in the debate over δ are interesting on their own merits, regardless of whether the different possible values of δ result in different climate policies.

The more interesting connections between the methodological questions discussed in this article and the discounting debate are more subtle. Consider, for example, an argument by Mintz-Woo ([2019]) for a positive δ. As Mintz-Woo notes, real-world decision models universally exclude a wide variety of presumed-to-be-exogenous factors and outcomes from consideration: climate models, for example, don't calculate probabilities and utilities of an asteroid hitting the earth or Facebook permanently going offline because these scenarios are presumed to be unaffected by the decisions in the model. Mintz-Woo's argument is essentially that just as the possibility of human extinction justifies some small pure time preference, so too does the possibility of certain model-independent outcomes: an asteroid hitting the earth between now and the realization of this or that consequence of climate change alters how we should value that consequence. Since the probability of some interfering event of this kind grows as time increases, we are justified in discounting outcomes further in the future more than events that are more immediate. Hence a positive δ.

The methodological debate rehearsed above has (at least) two important implications for Mintz-Woo's suggestion. First, taking Weitzman's arguments regarding uncertainty seriously would suggest even further discounting than Mintz-Woo's initial argument supports: the kind of events that would seriously interfere with climate abatement tend to involve deep uncertainty, and thus by Weitzman's reasoning the confidence distributions should plausibly have fatter tails than our 'best guesses' would suggest.

[15]The basic proof of this result can be found in (Weitzman [2009b], pp. 6–7); for a more extensive and general discussion, see (Weitzman [2012b]).

Second, and by contrast, taking Weitzman's methodological arguments seriously motivates making other changes to the discounting framework that are likely to either obviate the changes suggested by Mintz-Woo or to render them unnecessary by incorporating the same uncertaintities in a different way. Indeed, prior to his work on fat tails, Weitzman ([1998]) himself gave an argument for modifying discount policies to account for the kinds of presumed-to-be-exogenous factors that motivate Mintz-Woo's position. Where Mintz-Woo suggests modifying δ, however, Weitzman argued that these factors motivate a discount rate that decreases with time.[16] Consider again equation 7. The final term in equation 7, $-(1/2)\eta^2 s^2$, can be thought of as capturing what Arrow et al. ([2014]) term a 'precautionary effect', essentially a damper on the degree of discounting that a rational agent will accept that arises due to risk aversion. Over time, the probability of exogenous shock of the kind discussed by Mintz-Woo grows, and so this third term becomes increasingly important. As it grows, the discount rate decreases, with the results that while we discount consumption at time t_{i+1} more than we discount consumption at time t_i, the difference in discount factors between t_{i+1} and t_i is smaller than the difference between t_i and t_{i-1}.

In other words, there's no way to evaluate the implications of Mintz-Woo's argument without engaging with the methodological questions raised in this article: not only the implications of the argument but the very question of whether we should build the relevant factors into δ (as opposed to some other element of the equation) depends on our methodological choices with respect to probability distributions.[17] More broadly, understanding the economic debate over probability distributions as methodological rather than empirical opens it to normative arguments: elements of the discounting framework other than δ are sensitive to normative considerations, if for no other reason than that our methodological choices influence (for example) whether to use equation 1 or 7. As a consequence, the debate over discounting becomes more complicated: if both the value of δ and the choice of probability assignment are sensitive to normative considerations, then the arguments that we offer for the former are liable to have implications for the latter as well. Mintz-Woo's argument is simply an example of this kind of interaction, where the same normative motivation—namely, the thought that our theory should account for the possibility of exogenous shocks to the economic system—can be incorporated in different ways that have (apparently) contradictory implications.

[16]The result is one half of what's called the 'Weitzman–Gollier puzzle' due to an (ultimately illusory) conflict with results by Gollier ([2004]). For an accessible introduction, see (Greaves [2017], pp. 416–17); for a generalized formal treatment that illustrates the compatibility of the two results, see (Gollier and Weitzman [2010]). Note that—unlike Weitzman's arguments regarding fat tails—this result is widely accepted within the economic literature; see, for example, (Arrow et al. [2014]).

[17]Gollier ([2013]) explores a number of other connections between declining discount rates and the discounting debate.

8. Conclusion

It is very unlikely that the world will see—oh—10°C of warming. How unlikely? We can't say. Global warming will be almost certainly be hugely 'costly'—costly in the sense of causing extraordinary amounts of death, misery, and suffering—regardless. But the 'costs' of temperature changes on the scale of 10°C are virtually inconceivable. Limited humans that we are, we simply don't have the resources to imagine what it would mean for some of the most populous regions of the world to become uninhabitable. Mass migration, certainly, on a previously unknown scale; the collapse of the world economy; famines, droughts, wars, and plagues that would otherwise be avoidable. All happening in a world where some states have nuclear capabilities and/or biological weapons.

The question at the heart of this article—what probability distribution should we use when calculating the expected utility of different different climate policies?—might seem empirical, to be determined entirely by the evidence available to us. What we've seen is that there are reasonable methodological disagreements to be had here, however, and that these disagreements turn on how to count scenarios like the ones where we encounter 10°C of warming. And it's important that this dispute is methodological. It means that to resolve it, we need philosophical tools, not (just) economic ones. It also means that we cannot continue to carry out ethical debates about how to discount the future without considering which possibilities we're discounting and how those possibilities are counted.

Acknowledgements

My thanks to Mathias Frisch, Paul Kelleher, Nathan Lackey, Alan Love, Dana Mathiessen, Nadia Ruiz, three referees, and audiences at the University of Minnesota and 2024 Central American Philosophy Association for comments on an earlier version of this article. Funding for this article was provided in part by the Deutsche Forschungsgemeinschaft (project no. 254954344/GRK2073) and in part by the National Science Foundation (grant no. 2042366).

Philosophy Department
University of Minnesota
Minneapolis, USA
corey.dethier@gmail.com

References

Anthoff, D. and Tol, R. S. J. [2022]: 'Testing the Dismal Theorem', *Journal of the Association of Environmental and Resource Economists*, **9**, pp. 885–920.

Arrow, K. J. [1999]: 'Discounting, Morality, and Gaming', in P. R. Portney and J. P. Weyant(*eds*), *Discounting and Intergenerational Equity*, Washington, DC: RFF Press, pp. 13–22.

Arrow, K. J. [2009]: 'A Note on Uncertainty and Discounting in Models of Economic Growth', *Journal of Risk and Uncertainty*, **38**, pp. 87–94.

Arrow, K. J., Cropper, M. L., Gollier, C., Groom, B., Heal, G. M., Newell, R. G., Nordhaus, W. D. et al. [2014]: 'Should Governments Use a Declining Discount Rate in Project Analysis?', *Review of Environmental Economics and Policy*, **8**, pp. 145–63.

Bradley, R. [2017]: *Decision Theory with a Human Face*, Cambridge: Cambridge University Press.

Broome, J. [2010]: 'The Most Important Thing about Climate Change', in J. Boston, A. Bradstock and D. Eng (*eds*), *Public Policy: Why Ethics Matters*, Canberra: ANU Press, pp. 101–16.

Cato, S. [2020]: 'From the St. Petersburg Paradox to the Dismal Theorem', *Environment and Development Economics*, **25**, pp. 423–32.

Dethier, C. [2022]: 'Accuracy, Probabilism, and the Insufficiency of the Alethic', *Philosophical Studies*, **179**, pp. 2285–301.

Dethier, C. [2023]: 'Interpreting the Probabilistic Language in IPCC Reports', *Ergo*, **10**, pp. 203–25.

Dietz, S. [2011]: 'High Impact, Low Probability? An Empirical Analysis of Risk in the Economics of Climate Change', *Climatic Change*, **108**, pp. 519–41.

Dorst, K. [2019]: 'Higher-Order Uncertainty', in M. Skipper and A. Steglich-Petersen(*eds*), *Higher-Order Evidence: New Essays*, Oxford: Oxford University Press, pp. 35–61.

Dorst, K. [2020]: 'Evidence: A Guide for the Uncertain', *Philosophy and Phenomenological Research*, **100**, pp. 586–632.

Drupp, M. A., Freeman. M. C., Groom, B. and Nesje, F. [2018]: 'Discounting Disentangled', *American Economic Journal: Economic Policy*, **10**, pp. 109–34.

Falbo, A. [2023]: 'Inquiring Minds Want to Improve', *Australasian Journal of Philosophy*, **101**, pp. 298–312.

Frank, D. M. [2019]: 'Ethics of the Scientist qua Policy Advisor: Inductive Risk, Uncertainty, and Catastrophe in Climate Economics', *Synthese*, **196**, pp. 3123–38.

Friedman, J. [2020]: ‘The Epistemic and the Zetetic’, *The Philosophical Review*, **129**, pp. 501–36.

Gollier, C. [2004]: ‘Maximizing the Expected Net Future Value as an Alternative Strategy to Gamma Discounting’, *Finance Research Letters*, **1**, pp. 85–89.

Gollier, C. [2013]: *Pricing the Planet’s Future: The Economics of Discounting in an Uncertain World*, Princeton, NJ: Princeton University Press.

Gollier, C. and Weitzman, M. L. [2010]: ‘How Should the Distant Future Be Discounted When Discount Rates Are Uncertain?’, *Economics Letters*, **107**, pp. 350–53.

Greaves, H. [2017]: ‘Discounting for Public Policy: A Survey’, *Economics and Philosophy*, **33**, pp. 391–439.

Guiso, L. and Paiella, M. [2008]: ‘Risk Aversion, Wealth, and Background Risk’, *Journal of the European Economic Association*, **6**, pp. 1109–50.

Hartzell-Nichols, L. [2017]: *A Climate of Risk: Precautionary Principles, Catastrophes, and Climate Change*, New York: Routledge.

Heath, J. [2021]: *Philosophical Foundations of Climate Change Policy*, Oxford: Oxford University Press.

Helgeson, C. [2022]: ‘Structuring Decisions under Deep Uncertainty’, *Topoi*, **39**, pp. 257–69.

Horowitz, S. [2022]: ‘Higher-Order Evidence’, in E. N. Zalta and U. Nodelman (*eds*), *Stanford Encyclopedia of Philosophy*, available at <plato.stanford.edu/archives/fall2022/entries/higher-order-evidence/>.

Hwang, I., Tol, R. S. J. and Hofkes, M. W. [2016]: ‘Fat-Tailed Risk about Climate Change and Climate Policy’, *Energy Policy*, **89**, pp. 25–35.

Jamieson, D. [2014]: *Reason in a Dark Time: Why the Struggle against Climate Change Failed—and What It Means for Our Future*, Oxford: Oxford University Press.

Masson-Delmotte, V., Zhai, P., Pirani, A., Connors, S. L., Péan, C., Berger, S., Caud, N. et al. [2021]: *Climate Change 2021—The Physical Science Basis: Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge: Cambridge University Press.

Mastrandrea, M. D., Field, C. B., Stocker, T. F., Edenhofer, O., Ebi, K. L., Frame, D. J., Held, H. et al. [2010]: ‘Guidance Note for Lead Authors of the IPCC Fifth Assessment Report on Consistent Treatment of Uncertainties’, Technical Report, Intergovernmental Panel on Climate Change, available at <www.ipcc.ch/site/assets/uploads/2017/08/AR5_Uncertainty_Guidance_Note.pdf>.

McLaughlin, A. [2024]: ‘Existential Risk, Climate Change, and Nonideal Justice’, *The Monist*, **107**, pp. 190–206.

Millner, A. [2013]: ‘On Welfare Frameworks and Catastrophic Climate Risks’, *Journal of Environmental Economics and Management*, **65**, pp. 310–25.

Mintz-Woo, K. [2019]: ‘Principled Utility Discounting under Risk’, *Moral Philosophy and Politics*, **6**, pp. 89–112.

Mintz-Woo, K. [2021]: ‘A Philosopher’s Guide to Discounting’, in M. B. Budolfson, T. McPherson and D. Plunkett(*eds*), *Philosophy and Climate Change*, Oxford: Oxford University Press, pp. 90–110.

Nesje, F., Drupp, M. A., Freeman, M. C. and Groom, B. [2023]: ‘Philosophers and Economists Agree on Climate Policy Paths but for Different Reasons’, *Nature Climate Change*, **13**, pp. 515–22.

Nordhaus, W. D. [2007a]: *A Question of Balance: Weighing the Options on Global Warming Policies*, New Haven, CT: Yale University Press.

Nordhaus, W. D. [2007b]: ‘Critical Assumptions in the Stern Review on Climate Change’, *Science*, **317**, pp. 201–2.

Nordhaus, W. D. [2011]: ‘The Economics of Tail Events with an Application to Climate Change’, *Review of Environmental Economics and Policy*, **5**, pp. 240–57.

Nordhaus, W. D. [2012]: ‘Economic Policy in the Face of Severe Tail Events’, *Journal of Public Economic Theory*, **14**, pp. 197–219.

Nordhaus, W. D. [2013]: *The Climate Casino: Risk, Uncertainty, and Economics for a Warming World*, New Haven, CT: Yale University Press.

Parfit, D. [1984]: *Reasons and Persons*, Oxford: Oxford University Press.

Pindyck, R. S. [2011]: ‘Fat Tails, Thin Tails, and Climate Change Policy’, *Review of Environmental Economics and Policy*, **5**, pp. 258–74.

Pratt, J. W. and Zeckhauser, R. J. [1987]: ‘Proper Risk Aversion’, *Econometrica*, **55**, pp. 143–54.

Rinard, S. [2015]: ‘A Decision Theory for Imprecise Probabilities’, *Philosopher’s Imprint*, **15**, available at <hdl.handle.net/2027/spo.3521354.0015.007>.

Solomon, S., Qin, D., Manning, M., Marquis, M., Averyt, K., Tignor, M. M. B., Miller, H. L. and Chen, Z. [2007]: *Climate Change 2007—The Physical Science Basis: Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge: Cambridge University Press.

Staley, K. W. and Cobb, A. [2011]: 'Internalist and Externalist Aspects of Justification in Scientific Inquiry', *Synthese*, **182**, pp. 475–92.

Steel, D., Mintz-Woo, K. and DesRoches, C. T. [2023]: 'Collapse, Social Tipping Dynamics, and Framing Climate Change', *Politics, Philosophy, and Economics*, **23**, pp. 230–51.

Stern, N. [2007]: *The Economics of Climate Change: The Stern Review*, Cambridge: Cambridge University Press.

Thorstad, D. [2023]: 'High Risk, Low Reward: A Challenge to the Astronomical Value of Existential Risk Mitigation', *Philosophy and Public Affairs*, **51**, pp. 373–412.

Tol, R. S. J. [2003]: 'Is the Uncertainty about Climate Change Too Large for Expected Cost–Benefit Analysis?', *Climatic Change*, **56**, pp. 265–89.

Wagner, G. and Weitzman, M. L. [2015]: *Climate Shock: The Economic Consequences of a Hotter Planet*, Princeton, NJ: Princeton University Press.

Weitzman, M. L. [1998]: 'Why the Far-Distant Future Should Be Discounted at Its Lowest Possible Rate', *Journal of Environmental Economics and Management*, **36**, pp. 201–8.

Weitzman, M. L. [2007]: 'A Review of the Stern Review on the Economics of Climate Change', *Review of Economics and Statistics*, **45**, pp. 703–24.

Weitzman, M. L. [2009a]: 'Additive Damages, Fat-Tailed Climate Dynamics, and Uncertain Discounting', *Economics*, **3**, pp. 1–29.

Weitzman, M. L. [2009b]: 'On Modelling and Interpreting the Economics of Catastrophic Climate Change', *Review of Economics and Statistics*, **91**, pp. 1–19.

Weitzman, M. L. [2011]: 'Fat-Tailed Uncertainty in the Economics of Catastrophic Climate Change', *Review of Environmental Economics and Policy*, **5**, pp. 275–92.

Weitzman, M. L. [2012a]: 'GHG Targets as Insurance against Catastrophic Climate Damages', *Journal of Public Economic Theory*, **14**, pp. 221–44.

Weitzman, M. L. [2012b]: 'The Ramsey Discounting Formula for a Hidden-State Stochastic Growth Process', *Environmental and Resource Economics*, **53**, pp. 309–21.

Weitzman, M. L. [2013]: 'Tail-Hedge Discounting and the Social Cost of Carbon', *Journal of Economic Literature*, **51**, pp. 873–82.