



# Climate Uncertainty, Real Possibilities and the Precautionary Principle

Jeroen Hopster<sup>1,2</sup>

Received: 30 October 2020 / Accepted: 2 September 2021  
© The Author(s) 2021

## Abstract

A challenge faced by defenders of the precautionary principle is to clarify when the evidence that a harmful event might occur suffices to regard this prospect as a real possibility. Plausible versions of the principle must articulate some epistemic threshold, or *de minimis* requirement, which specifies when precautionary measures are justified. Critics have argued that formulating such a threshold is problematic in the context of the precautionary principle. First, this is because the precautionary principle appears to be ambiguous about the distinction between risk and uncertainty: should the principle merely be invoked when evidential probabilities are absent, or also when probabilities have low epistemic credentials? Secondly, defenders of the precautionary principle face an aggregation puzzle: in judging whether or not the *de minimis* requirement has been met, how should first-order evidential probabilities and their second-order epistemic standing be aggregated? This article argues that the ambiguity can be resolved, and the epistemological puzzle can be solved. Focusing on decisions in the context of climate uncertainty, I advance a version of the precautionary principle that serves as a plausible decision rule, to be adopted in situations where its main alternative—cost–benefit analysis—does not deliver.

## 1 Introduction

The precautionary principle (henceforth: PP) is subject to substantial disagreement. Critics question its unity, coherence, and non-triviality (Sunstein, 2005); supporters argue that each of these can be vindicated (Steel, 2013, 2014). But controversies aside, about one issue all scholars working on PP seem to agree: precaution is only warranted if a threat of harm constitutes a *realistic possibility*, rather than a far-fetched fantasy (e.g. Betz, 2010; Carter & Peterson, 2015; Gardiner, 2006;

---

✉ Jeroen Hopster  
j.k.g.hopster@utwente.nl

<sup>1</sup> University of Graz, Graz, Austria

<sup>2</sup> University of Twente, Enschede, The Netherlands

Hartzell-Nichols, 2017). We should not take precautionary measures in the face of any dreamed-up catastrophe, however fanciful. Instead, envisioned doom scenarios should surpass a minimal level of plausibility to warrant such measures.

This plausibility requirement goes by different names. Shue (2010, p. 549) designates it as PP's "anti-paranoia requirement". Gardiner (2006, pp. 52–53) calls for restricting PP to "realistic outcomes". Carter and Peterson (2015, p. 8) speak of a "*de minimis* requirement", from the legal principle *de minimis non curat lex*—the court should not concern itself with trifles. The underlying ideas are the same and well established in the literature on PP: any plausible version of the principle should ignore sufficiently improbable risks, to avoid precautionary paranoia.

A plausibility requirement is not unique to PP: *all* normative approaches to decision-making under risk or uncertainty face the question of what level of evidence suffices for a call to action. However, the question of what possibilities we should regard as satisfying the minimal epistemic threshold for being regarded as realistic—i.e. the question of how the *de minimis* requirement should be operationalized—seems particularly pressing in the context of PP, for two reasons. First, existing defenses of PP tend to leave the relation between real possibilities and evidential probabilities unclear. Some defenders of PP argue that the principle should only be invoked in contexts where evidential probabilities cannot be calculated, but where we nonetheless have reason to believe that a given outcome constitutes a real possibility (e.g. Gardiner, 2006; Shue, 2010; Steel, 2013). But this raises a problem regarding the *de minimis* requirement: if we cannot rely on evidential probabilities, then what are the grounds for identifying realistic possibilities? To operationalize PP as a decision rule, the muddled relation between real possibilities and evidential probabilities requires clarification.

Secondly, as Carter and Peterson (2015) argue, in the face of the *de minimis* requirement PP faces an aggregation puzzle. One way to formulate this puzzle is by distinguishing between the first-order probabilities associated with a given piece of evidence and the second-order epistemic credentials of this evidence. Consider an example from climate science, which will be the reference point for case-studies in this paper. Assuming a given level of anthropogenic forcing on the climate system, a climate projection might indicate that the probability that the West Antarctic ice sheet will start to collapse by 2050 is 18%. However, this first-order probability is itself subject to uncertainty: not only is the actual level of anthropogenic forcing up until 2050 uncertain, but the model on which the probability assessment is based might be inadequate. How should first-order probabilities and their second-order evaluations be combined, to judge whether the *de minimis* requirement has been satisfied? Carter and Peterson argue that it is far from obvious how this should be done in the context of PP, and that this saddles defenders of PP with a serious challenge—a challenge that has not yet been resolved.<sup>1</sup>

The aim of this article is to address and resolve both unclaritys. I will argue that epistemic challenges notwithstanding, PP constitutes a plausible decision-rule to

<sup>1</sup> Steglich-Petersen (2015) attempts to do so but fails to resolve the original puzzle, as Carter and Peterson (2016) point out in reply. The solution I offer in Sect. 5, I submit, *does* solve the puzzle.

adopt in the face of realistic possibilities of harm that are couched in substantial evidential uncertainty. To make this case I first outline in Sect. 2 how PP can be understood as a decision rule and clarify how the *de minimis* requirement relates to it. In Sect. 3 I contrast PP with cost–benefit analysis (henceforth: CBA) and argue that the former may be preferable to the latter, though only in circumscribed decision-contexts: CBA is preferable as a decision-rule in contexts that are epistemically transparent, whereas PP is superior in decision-contexts that are epistemically opaque but do involve real possibilities of substantial harm. In Sect. 4 I address the abovementioned unclarity about PP, by elucidating how real possibilities relate to evidential probabilities. This sets the stage for addressing Carter and Peterson’s challenge in Sect. 5: how can PP solve the aggregation puzzle, to arrive at an all-things-considered-judgment of whether the *de minimis* requirement is satisfied? Sect. 6 concludes by offering suggestions for future scholarship on real possibilities in relation to climate uncertainty.

## 2 The Precautionary Principle as a Decision Rule

Any treatment of PP faces the foundational question of clarifying how the principle should be understood. This is no trivial matter, as the principle has been framed in various ways. Important statements of PP have been given in legal and political documents, including the canonical statement in article 15 of the 1992 *Rio Declaration on Environment and Development*:

where there are threats of serious or irreversible damage, lack of full scientific certainty shall be not used as a reason for postponing cost-effective measures to prevent environmental degradation

as well as the 1998 *Wingspread Statement on the Precautionary Principle*:

when an activity raises threats of harm to human health or the environment, precautionary measures should be taken even if some cause and effect relationships are not fully established scientifically.

In philosophical discussions, on the other hand, PP is typically abstracted from applied contexts and treated as a general principle. Yet there is substantial variety in philosophical interpretations over what *kind* of principle PP might be taken to be. PP has been explicated as a midlevel moral principle (Sandin & Peterson, 2019) and as a procedural approach to risk-assessment (Goklany, 2001); it has been associated with the epistemic counsel of reversing the burden of proof (Birch, 2017) and with the normative decision rule of maximin (Gardiner, 2006). In the context of climate ethics PP is commonly associated with frameworks for decision-making (Hartzell-Nichols, 2017). Even understood as an approach to or rule for decision-making, however, several different statements of PP might be given (e.g. Koplin et al., 2020; Selgelid, 2016).

In this paper I will operationalize PP as a rule for decision-making, which is meant to avoid particularly bad outcomes in situations of uncertainty. I argue that

this principle can be vindicated, though only in a circumscribed set of decision contexts: PP is plausibly endorsed in decision-contexts that involve real possibilities of substantial harm, which are epistemically opaque. Before discussing the merits of the principle—and defending it from recent objections—in this section I specify how PP can be understood as a decision-rule. I do so by largely following the work of Steel (2013, 2014), who suggests that, notwithstanding appearances to the contrary, there is in fact substantial unity to existing formulations of PP. Inductively generalizing from existing statements, Steel presents PP as a decision-rule made up of three components. In the present exposition I add a fourth component to this generalization—the rule of ‘inverse linkage’—which is commonly associated with PP and helps to shed light on the role of the *de minimis* requirement.

The first and core component of PP is the so-called ‘decision tripod’ (Carter & Peterson, 2015; Hartzell-Nichols, 2017; Manson, 2002), which consists of the following three conditions:

- A *damage condition* (D), which specifies a threat of harm, or anticipated catastrophic outcome, which should be avoided.
- An *epistemic condition* (E), which serves to substantiate that the evidential probability that this outcome will occur is non-negligible, or that there are good epistemic grounds to take the threat seriously (i.e. that the *de minimis* requirement is satisfied).
- A *suggested remedy* (R), which prescribes or recommends the measures that should be taken in order to avoid the catastrophe, or to reduce its risk.<sup>2</sup>

These three conditions jointly constitute the decision rule of PP: if an envisioned outcome is regarded as damaging (D), and the prospect that damage will materialize lest precautionary measures be taken is sufficiently plausible (E), then precautionary measures—i.e. the suggested remedy—should be prescribed (R). The tripod can serve to generate more distinct versions of PP, tailored to specific contexts (Manson, 2002; Steel, 2013). Both D, E and R admit of degrees: the anticipated harm may be more or less severe, the evidence related to its occurrence may be more or less substantial, and the suggested remedy may be more or less effective. Precisely how these three conditions are specified is case-dependent. Indeed, this flexibility partly

<sup>2</sup> The remedy (R) can be formulated either to highlight the costs of acting or the costs of inaction (cf. Steglich-Petersen 2015). The former emphasis is appropriate where the introduction of novel risky technologies is concerned: if undertaking action, e.g. by conducting a risky technological experiment, seems harmful, then PP will counsel to prohibit action in the face of dangerous unknowns (cf. Taleb et al., 2014; Wingspread Statement). The latter emphasis is appropriate in the face of a threat towards which societies are presently heading, lest precautionary action be taken (cf. Rio Declaration). Anthropogenic climate change fits the latter description: action is needed to reduce global levels of CO<sub>2</sub> emissions, to reduce potentially catastrophic risks. Notice, however, that nothing of substance hinges on whether R is formulated in terms of action or inaction. Ultimately such formulations are interchangeable; they depend on how an issue is framed. For instance, the extent to which humanity is contributing to atmospheric concentrations of CO<sub>2</sub> might be framed as an unprecedented experiment, with enormous risks of harm. Adopting PP, one might argue that the proper remedy consists of *preventing* action in the face of these risks: i.e. the reckless experiment should be halted, by sharply reducing global CO<sub>2</sub> emissions.

explains the multiplicity of existing formulations of PP. Yet the basic skeleton of the tripod remains: at minimum, there should be *some* anticipated harm, *some* grounds for thinking that this harm will occur in the absence of the suggested remedy, and *some* grounds for thinking that the suggested remedy can mitigate the harm.

Steel (2013, 2014) extends PP with two further components, which set constraints on the decision-tripod and its application. A first extension is what he calls the rule of *proportionality*. The general idea behind this constraint is that precautionary measures should be calibrated to the degree of uncertainty and the seriousness of the consequences feared. Steel specifies proportionality in terms of what he calls ‘consistency’ and ‘efficiency’: the cure should not be worse than the disease and the negative side-effects of precautionary measures should be kept to the minimum. A second extension of the decision tripod is what Steel calls the *meta-precautionary principle*, which sets a general constraint on the kinds of decision rules that policy-makers should adopt. The meta-rule says that scientific uncertainty should not lead to paralysis in decision-making, in the face of a threat of serious harm (Steel, 2013). This is, for instance, how the emphasis of the Rio Declaration can be understood, which highlights that:

(...) lack of full scientific certainty shall be not used as a reason for postponing cost-effective measures to prevent environmental degradation.

On Steel’s account, the decision tripod, together with these two extensions, constitute the three components of PP.

Note that the exact contents of the tripod are open to substantive debate (cf. Bogner, 2011; Gardiner, 2006). Some PP adherents frame the damage condition (D) in terms of catastrophic outcomes. But what makes an outcome catastrophic? Hartzell-Nichols (2017, p. 46) designates catastrophes as “outcomes in which millions of people could suffer severely”. But surely, we might consider alternative specifications. Does a loss of species-diversity count as a catastrophe? Does a sea-level rise of five meters by the end of this millennium? What outcomes we regard as harmful, and how grave we take this harm to be, is a matter of ethical debate. PP does not preclude such debate, nor does it settle it. Instead, PP is better thought of outlining a decision rule that is conditional on the outcome of normative debate. The contribution of ethicists and epistemologists will be specifically important in specifying which standards D and E should satisfy: what makes an outcome damaging, and which amount of evidence should be in place to take a threat seriously?

With this exposition in place, let me propose one further extension of the tripod that merits explicit discussion in the present context, because it is regularly appealed to by defenders of PP and because it helps to illuminate the role of the *de minimis* requirement in the principle. This is the rule of *inverse linkage*: the greater the envisioned catastrophe, the less evidence is needed to warrant precautionary action. Conversely, the lesser the catastrophe, the more evidence is needed to warrant precaution. Hence, if we hold the envisioned remedy (R) fixed, then D and E are inversely linked: if the magnitude of the impact is enormous, then only little evidence is needed to justify precautionary action (Carter & Peterson, 2015).

Prima facie, this rule might seem to conflict with the decision tripod as outlined above, which specifies the damage condition (D) and the epistemic condition (E)

independently, thereby indicating that the evidential probability that an outcome will occur should *not* hinge on the anticipated damage. Indeed, linkage of the epistemic condition and the damage condition has often been raised as an objection to PP (e.g. Sunstein, 2005), and for good reasons, or so it seems. After all, such linkage seems to entail that only very little evidence is needed to warrant precaution with respect to choice X when the envisioned catastrophe is immense—say the gradual extinction of all species on our planet after centuries of suffering—even if the evidence that choice X will have this detrimental impact is incredibly weak. But a decision rule along these lines seems misguided: if the relevant evidence is indeed incredibly weak, then a precautionary stance is uncalled for, irrespective of the envisioned impact. Indeed, examples along these lines invoke the very worry of precautionary paranoia that an epistemically credible version of PP is meant to rule out.

However, the rule of *inverse linkage* is *not* meant to trump the anti-paranoia requirement—or what we have called the *de minimis* requirement. Instead, the *de minimis* requirement has lexical priority over the *inverse linkage* of D and E: such linkage comes into play only *after* this requirement has been fulfilled. That is, evidential standards may be lowered in the face of impending harm, *provided* that the minimal evidential threshold has been satisfied in the first place.<sup>3</sup> By way of example, consider proposals to develop technologies for solar radiation management, which have raised the worry among opponents that catastrophic harm to human health or the environment might ensue. Invoking PP, opponents should first inquire whether the *de minimis* requirement is satisfied: are there indeed good epistemic grounds for taking this threat seriously? If this is the case, then inverse linkage comes into play: even limited evidence that the worry of catastrophic harm is justified should suffice to forestall the technology's implementation. That is, if the anticipated harm (D) is indeed catastrophic, then the standards to satisfy E should be low, such that R can easily be triggered. This is, for instance, how the emphasis of the Wingspread Statement could be interpreted:

when an activity raises threats of harm to human health or the environment [i.e.: D is substantial], precautionary measures should be taken [i.e.: R should be triggered] even if some cause and effect relationships are not fully established scientifically [i.e.: it might be justified to lower the standard for E, compared to the default scientific standard].

Similar reasoning might also be employed to shift the burden of proof (Powell, 2010; Steglich-Petersen, 2015). For instance, given a suspected risk of substantial harm, it might be argued that proponents of a novel technology must provide positive evidence of its safety, over and above standard procedures. Hence, a heavier

<sup>3</sup> It has been argued against PP that it makes decisions overly sensitive to the bounds of what probabilities count as realistic, and insensitive to any further probabilities (Peterson 2006, p. 597; Roser 2017, p. 1402). This criticism, however, is misguided. PP is sensitive to *all* evidential probabilities. The realism criterion does not demarcate a point beyond which further evidence is arbitrarily ignored, but a reasonable distinction between what evidence *does*, and what evidence *does not* suffice to take precautionary measures. The boundary itself may be vague and contested, but there are paradigm examples on either side.

burden of proof might be placed on proponents of solar radiation management technology to show that the technology is *safe*, rather than on regulators to show that it is *unsafe*.

Does this not suggest that PP is overly restrictive? Consider the fear, common among contemporary conspiracy theorists, that deploying the 5G telecommunications network will have detrimental health effects. With the enormous health risks that are at stake, does the rule of inverse linkage not imply that evidential standards should be lowered, and precautionary measures are warranted? That implication would appear to make PP quite implausible—after all, there does not seem to be a credible scientific basis underlying these health concerns. But again, the rule of inverse linkage only comes into play after the *de minimis* requirement has been satisfied, and regarding the envisioned health risks of the 5G network this does not seem to be the case. The *de minimis* requirement serves to prevent precautionary paranoia.

To sum up, we have identified four components of PP:

- *Decision tripod*: if an envisioned outcome is regarded as damaging (D), and the prospect that damage will materialize lest precautionary measures be taken is sufficiently plausible (E), then precautionary measures—i.e. the suggested remedy—should be prescribed (R).
- *Rule of proportionality*: the aggressiveness of precautionary measures should correspond to the plausibility and severity of the threat.
- *Meta-precautionary principle*: uncertainty should not lead to inaction.
- *Inverse linkage*: holding R fixed, D and E are inversely linked: the magnitude of one depends on the magnitude of the other. As a result, depending on how catastrophic an envisioned outcome is, evidential standards may be lowered and the burden of proof may be shifted, provided that the *de minimis* requirement is still satisfied.

Each of these components appears to be quite reasonable. At first face, then, this fourfold PP constitutes a plausible decision rule.

### 3 Adopting PP or CBA? A Criterial Approach

Even if PP is a plausible decision rule, it may be of little use if it is subordinate to another plausible decision rules in its vicinity. More specifically, it has been objected that PP is merely a muddled version of CBA (Goklany, 2001; Sunstein, 2001). According to CBA, it is rational to take whichever course of action has the highest sum of benefits minus costs, adjusted for the probability that these will be realized.<sup>4</sup> PP, too, involves a probability adjusted weighting of costs and benefits. However, in CBA these costs and benefits are quantified precisely, whereas PP

<sup>4</sup> CBA admits of different versions (Boardman et al., 2018). When applied under conditions of risk, CBA is sometimes framed in terms of expected value theory (e.g. Broome 2012), or in terms of expected utility theory (e.g. Koplín et al., 2020).



proponents rely on what seem to be rather vague and interpretable criteria, such as the *de minimis* requirement. If it turns out that decisions generated by CBA do not clearly differ from decisions generated by PP, but that CBA provides a more straightforward and precise method to generate them, then we would be better off to rely on CBA, or so it seems.

Now, I submit that in a circumscribed set of contexts CBA is indeed more straightforwardly applicable than PP. Consider situations involving a small but determinate probability of a specific catastrophic event. If both the probability and the outcome are determinate—we are quite certain about what to expect, and how to value it—then so is the expected value of the action we decide to pursue. Colloquially, we might still say that in seeking to avoid catastrophic outcomes with a low probability, we are relying on precautionary measures, but doing so is not unique to PP. If the harm of a low probability risk is sufficiently grave, then CBA will counsel to avoid it just as well. If, on the other hand, the risk of harm is insufficiently grave and substantial benefits are at stake, then CBA will not counsel to avoid it, but *neither does PP*. After all, PP is only triggered if the *de minimis* condition (E) as well as the damage condition (D) are met.

Following Knight's (1921) distinction, decision contexts which involve determinate probabilities are typically labeled contexts of *risk*, which should be distinguished from contexts of *uncertainty*. The defining aspect of contexts of risk is that decision-makers have access to well-defined probabilities: the risk of a given outcome is understood as the probability of its occurrence times its negative consequences. In contexts of uncertainty, however, decision-makers do not have access to well-defined probabilities, although they do have some positive knowledge about the prospect that a given outcome will occur. Probabilities may be vague, or perhaps it is unclear how the knowledge that decision-makers possess can be framed in probabilistic terms. Even so, in conditions of uncertainty decision-makers do possess at least *some* knowledge that is relevant to decision-making. Hence, uncertainty should not only be distinguished from risk, but also from the epistemic state of pure *ignorance*.

CBA seems to be more practicable than PP in contexts of risk: it allows for more clarity and expression of quantitative detail than PP. Proponents of PP, however, typically argue that the same does not hold for contexts of uncertainty, which lack determinate probabilities, but do involve a threat of substantial harm. Such characteristics frequently apply to policy-decisions concerning climate change, especially where local variations and long-term impacts are concerned (e.g. Sutton, 2019). Returning to our previous example, while it has been well-established that melting land-ice is likely to provide a major contribution to future sea-level rise, there is significant uncertainty about the rate and timing with which ice sheets melt. It is particularly difficult to anticipate the timing of potential threshold effects, such as the collapse of the West Antarctic marine ice sheet, which could raise global sea-levels by approximately 3.3 m in the long run (Bamber et al., 2009). Importantly, the process resulting in collapse is likely to be irreversible once a tipping point has been crossed. Hence, there appears to be a real possibility of substantial harm, but there is substantial uncertainty as well, as the rate of melting and the exact tipping point are difficult to pin down.



Such conditions, which involve potentially catastrophic events that should be taken seriously on scientific grounds (they constitute real possibilities), but the probabilities of which are themselves subject to substantial uncertainty, confront CBA with a dilemma. Let's call real possibilities of this kind *U-events*. The dilemma, for CBA-proponents, is whether to incorporate U-events in their calculations. On the one hand, if they *do* take U-events into account, then CBA cannot deliver precise recommendations. After all, U-events do not come along with precise probabilities; therefore, the balance of expected benefits and costs remains unclear.<sup>5</sup> As a result, relying on CBA can lead to paralysis in decision-making (Steel, 2013). Since this renders CBA unhelpful as a decision rule, the only viable alternative for CBA-proponents is to take on the second horn of the dilemma and ignore U-events. In doing so, however, CBA ignores decision-relevant information. Even worse, it does so in a way that seems decidedly immoral—namely, by ignoring outcomes that involve existential risks (Bostrom, 2013). In practice, some economists applying a version of CBA, such as Nobel laureate William Nordhaus, try to escape the dilemma by relying on best guesses (Hartzell-Nichols, 2017, ch. 4). But this simply brings us back to the dilemma: to be incorporated into CBA, these best guesses need to be quantified. This is problematic given the uncertainty at issue, whereby quantification can easily amount to false precision. Moreover, it is likely that false precision still comes at the cost of neglecting the most catastrophic outcomes, whose systemic and all-encompassing effects involve too many unknown variables. Hence, CBA seems to have an inherent tendency to shy away from 'fat tailed' risks of catastrophic harm (Weitzman, 2011).

Proponents of PP, by contrast, resolve the dilemma in a different manner. U-events should not be ignored, nor should uncertainty lead to inaction. This is exactly why the decision-making tripod is extended with the aforementioned meta-clause that uncertainty should not lead to paralysis. What is more, rather than risking false precision, it would be much more appropriate to make the genuine scientific uncertainty that comes along with U-events explicit and the focal point of decision-making. PP does just this, by focusing on the question of how E and D should be specified, both of which are normative questions. Subsequently, moral, political, and scientific debate over whether these conditions are satisfied should serve to break the deadlock and to avoid paralysis. Such debate is more likely to lead to clear policy advice. Of course, as with any normative issue, it

<sup>5</sup> Some proponents of CBA might want to resist this characterization. Employing expected value analysis, they might argue that it is possible to generate a weighted account of the different probabilities that are at play, adjusted for their respective epistemic credentials. For instance, John Broome (2012), who defends expected value theory, grants that to work out expected values in the context of climate change, we need much more detailed information about probabilities than climate scientists can typically provide. But he also maintains that "[w]hile we are waiting for it to be supplied by scientists, we have to do the best we can. On the basis of what evidence we can muster, we have to assign each event a probability as well as we can. In any case, the lack of firm probabilities is not a reason to give up expected value theory." (Broome 2012, p. 129) Defenders of PP should agree with Broome that "we have to do the best we can" in terms of assigning probabilities to the best of our knowledge. But contrary to Broome, they maintain that doing our best is not good enough to stick with CBA as a practicable decision rule: the unclarity of the probabilities involved will either paralyze CBA or lead to false precision.

is possible that no clear consensus will be reached, such that PP, too, ultimately leads to paralysis. But if so, then paralysis should not be diagnosed as a failure of economic quantification, but rather as a failure of reaching moral or political agreement—which seems entirely appropriate, given the nature of the issues that give rise to U-events, such as the issue of climate change (Gardiner, 2011).

To sum up, when deciding about how to act in the face of risks of harm under conditions of uncertainty, PP seems preferable over CBA, both for moral reasons and for reasons of clarity. PP seems preferable on moral grounds, because it does not shy away from considering realistic risks of catastrophe, even if these are themselves uncertain. PP seems preferable for reasons of clarity because it adequately identifies the key normative issues around which decisions regarding U-events revolve and raises the prospect of delivering a straightforward prescription for action, rather than resulting in paralysis. But, importantly, these practical advantages of PP over CBA only emerge when scientific probabilities are subject to substantial uncertainty. Hence, we arrive at what Gardiner (2006) calls a criterial approach to PP: PP might be preferred over CBA as a decision-rule, if and only if specific criteria are met.

#### 4 Real Possibilities and Evidential Probabilities

A challenge that has been raised for criterial accounts of PP is that they leave the relation between real possibilities and probabilities unclear. Consider Gardiner's (2006) account. Gardiner outlines four conditions that should be jointly met to legitimately invoke PP, or what he calls RCPP—the 'Rawlsian Core Precautionary Principle' (which is spelled out in terms of Rawls' maximin principle). Let's look at the two epistemic conditions of this version of PP. First, echoing Rawls (1999), Gardiner argues that deciding on the basis of RCPP requires that:

decision-makers either lack, or have reason to sharply discount, information about the probabilities of the possible outcomes of their actions. (Gardiner, 2006, p. 47)

I will call this the 'no probabilities criterion'. Second, Gardiner submits that:

the range of outcomes considered are in some appropriate sense "realistic," so that, for example, only credible threats are considered. (idem 2006, p. 51)

I will call this the 'realism criterion'.

As it stands, the 'no probabilities condition' is ambivalent between two options: probabilities are either *absent* or *unclear*. If we assume the former, then a tension arises: identifying real possibilities seems impossible in situations where probabilities are entirely absent. After all, if we cannot attach any meaningful probabilities to a given outcome, then it seems that we are also in no position to identify that outcome as a real possibility. Differently put, if we *can* identify an outcome as a real possibility, then there must be at least *some* evidence on the

basis of which we can justifiably do so. This evidence, in turn, could be described in terms of evidential probability.<sup>6</sup> Hence, in situations where we can identify real possibilities, evidential probabilities cannot be entirely absent (Roser, 2017).

A better interpretation of the ‘no probabilities condition’, then, is that probabilities are unclear, rather than absent. This conforms to the idea that for almost any outcome that is of scientific interest we can attach *some* probability to its occurrence (Powell, 2010). These probabilities might be too coarse-grained to provide much insight, or to be practicable in decision-making. Still, ascribing them is at least possible, albeit with low epistemic credentials. That is, if we have *at least some* evidence that X will occur, then we are also in a position to ascribe an evidential probability to X, although the epistemic credentials of this ascription may be rather low, depending on the quality of the evidence on which the statement rests. Moreover, as Roser (2017) argues, in principle there are myriad evidential grounds on which probabilities can be ascribed. For instance, in the context of climate change, the evidence of decision-makers ranges from things such as

background knowledge about the stability of natural systems under human influence (...), the tone of voice with which scientists speak about the dangers of climate change, the track record of science in forecasting long-term trends, the probabilities that scientists give in IPCC reports which in turn are based on empirical data and the general body of natural science, etc. (Roser, 2017, p. 1405)

Of course, some of these evidential sources are very weak, and not all of them may satisfy the *de minimis* requirement of realistic possibilities. But they are evidential sources nonetheless, and that holds for any context in which we are dealing with realistic possibilities: these are not contexts of *pure ignorance*, but contexts in which we possess at least some relevant evidence, which could in principle be framed in terms of (imprecise) evidential probabilities.

In sum, while probabilities are neither always necessary, nor always the preferred mode of expressing information that pertains to a body of evidence, such expression is nonetheless possible. The uncertainty that comes along with evidential probabilities can be expressed, for instance, in terms of a probability interval (Hansson, 2018). Scientists might claim, after integrating different models, that the probability that the West Antarctic marine ice sheet has already become unstable is, say, between 1 and 18%. This practice of assigning probability ranges is widespread in climate science. It is adopted, for instance, by the IPCC, which in its recent assessment reports employs a quantitative ‘likelihood scale’ that explicitly links specific probability ranges with specific statements of likelihood (a 99–100% probability

<sup>6</sup> Evidential probabilities express the degree of support for a proposition given by the evidence. For instance, if we consider the proposition that the collapse of the West Antarctic marine ice sheet ultimately becomes inevitable at a temperature level of 1.5–2 °C above pre-industrial levels (cf. Pattyn et al., 2018), then the evidential probability expresses to what degree the evidence supports this proposition.

range has the qualifier ‘virtually certain’; a 90–100% probability range is ‘very likely’, and so on; see Mastrandrea et al., 2010).<sup>7</sup> Hence, indeterminate probabilities often serve to designate likely ranges of outcomes.

One way of probabilistically expressing uncertainty about a given body of evidence is by distinguishing between first-order and second-order probabilities (Ord et al., 2010). First-order probabilities refer to the evidence itself (they might express, for instance, the probability of an outcome as given by a specific climate model), whereas second-order probabilities refer to the epistemic standing of the evidence (how reliable should we take the specific climate model to be?). Second-order probabilities might also be attached to probability ranges. Sticking with our example, scientists might claim that the first-order probability range that the West Antarctic ice sheet has already become unstable is 1% to 18%, alongside a second-order order probability which expresses the reliability of this claim—i.e. the probability that the assessment of the first-order probability is correct. Once again, such expressions are familiar from IPCC assessment reports, which do not only involve a quantitative likelihood scale, but also a qualitative scale which expresses the level of confidence in the validity of a finding, as a function of the nature of the evidence and the degree of expert agreement on the evidence (Mastrandrea et al., 2010). These qualitative expressions (‘low confidence’, ‘medium confidence’, etc.) could, in principle, be expressed in terms of second-order probabilities. With regard to criterial accounts of PP, whether or not the relevant evidence is presented in probabilistic terms is ultimately beside the point. What matters is whether this evidence is couched in substantial epistemic uncertainty.

## 5 Solving the Aggregation Puzzle

In the last section I have argued that, framed in terms of the distinction between first and second-order probabilities, PP should be adopted in situations where our second-order epistemic probabilities regarding the soundness of our first-order evidential probabilities is low.<sup>8</sup> But this leads to a further question pertinent to the epistemology of PP: how should these two probabilities be integrated, in making an all-things-considered judgement concerning the risk that a given threat will actualize? Carter and Peterson have recently argued that “it is far from obvious how the defender of the precautionary principle should combine the two types of

<sup>7</sup> Going by the IPCCs likelihood scale, this (fictitious) range of 1–18% would be qualified as ‘unlikely’. Arguably, however, this label does not communicate quite the right message to the public and to policy-makers. A more appropriate message, it seems, if this range were accurate, is that instability of the West Antarctic ice sheet constitutes a possibility that is very real. In its most recent Assessment Report (AR6), the IPCC has begun to place more emphasis on “low-likelihood, high-impact outcomes” of this sort.

<sup>8</sup> It might well be the case that one’s uncertainty about the correctness of a given first-order evidential probability is greater than the uncertainty the first-order probability itself expresses. That is, the probability estimate given by any scientific analysis might be dwarfed by the probability that the analysis itself is flawed (Ord et al., 2010). Therefore, to establish which outcomes constitute real possibilities, it is crucial to scrutinize second-order epistemic considerations.

probability.” (Carter & Peterson, 2015, p. 10) This seems problematic for PP, especially since one of the principal reasons for preferring PP over CBA in contexts of uncertainty is that the former can offer more clarity than the latter. If it turns out that PP cannot present a clear formula for handling uncertain probabilities, then this purported advantage disappears. Hence, we are faced with the puzzle of how first- and second-order probabilities should be aggregated, to arrive at an overall judgement of whether the *de minimis* requirement is satisfied. Let’s call this the aggregation puzzle.

Carter and Peterson (2015, 2016) frame the puzzle in terms of a dilemma for proponents of PP: either they aggregate the first- and second-order probability functions into a single measure of uncertainty, or they refrain from aggregating the two probabilities. Suppose that PP proponents take on the dilemma’s first horn. The obvious candidate aggregation rule would be to multiply first- and second-order probabilities. This rule entails that if a second-order probability is very small—i.e. if we should have very little confidence in the validity of a first-order probability—then the all-things-considered probability will be even smaller. Hence, the greater the uncertainty about our first-order probabilities, the less likely it is that the *de minimis* requirement will be satisfied, and that PP can be legitimately invoked. But according to Carter and Peterson (2015, p. 10) “[t]his is clearly the wrong conclusion. Intuitively, it would make more sense to apply the precautionary principle if the first-order probability is highly uncertain, which is the opposite of what the multiplicative rule suggests.” The rationale underlying their intuition is that if a first-order probability is uncertain, then there seems to be a real possibility that things might turn out *worse*. Therefore, plausibly, we should be *more*, rather than less, cautious in our decision-making.

If multiplying first- and second-order probabilities delivers the wrong result with respect to the *de minimis* requirement, then the first horn of the dilemma is unattractive. But the second horn—refrain from aggregating the two measures—is unattractive too. Multiplication constitutes the intuitive guideline for combining probabilities in PP and it is not clear whether another guideline is tenable. But if there are no clear guidelines for how second-order considerations should be incorporated into PP, then the principle itself is rendered unclear.

This apparent puzzle, I contend, can be solved. Pace Carter and Peterson, the correct way to aggregate first- and second-order probability functions is indeed to multiply them, in order to generate an overarching judgment of whether or not a given outcome constitutes a real possibility (i.e., of whether *de minimis* is satisfied). But crucially, this rule only applies if the outcome is held fixed and targets what might be realistically considered as the worst case. Under these conditions, the intuition articulated by Carter and Peterson that greater uncertainty warrants greater precaution misfires: if our first-order evidence suggests that a given outcome constitutes a realistic worst-case possibility, then the shakier this first-order evidence turns out to be, the less likely that *de minimis* is satisfied.

What explains the puzzlement of Carter and Peterson is that they treat the outcome to which the multiplication rule applies as variable. Their intuition that the multiplicative rule fails, after all, relies on there being a real possibility that things might turn out *worse*. Now, it is certainly true that if the outcome is variable, then

greater uncertainty about the standing of the first-order evidence pertaining to this outcome might justify *greater* precaution. Consider the proposition that the sea-level of the North Sea will rise with maximally 1 m during the twenty-first century. Judged by scientists' current first-order evidence, this constitutes approximately the worst-case possibility (van den Hurk et al., 2014). However, taking into account higher-order evidence about this first-order prediction, the all-things-considered realistic possibility envisioning the worst case should arguably employ a wider uncertainty margin. In this scenario, Carter and Peterson's intuition is appropriate: since the second-order probability implies that matters might turn out *worse*, we should be *more* cautious in our decision-making.

But if we focus on the worst-case realistic outcome taking into account both first- and second-order probabilities, and if we treat this this outcome as fixed, then the puzzle disappears. Say, for instance, that our all-things-considered estimate (i.e. the estimate based on our first-order evidence, calibrated by our second-order epistemic assessment) is that the North Sea will rise with maximally 1–2 m during the twenty-first century. Any envisioned outcome above this range (>2 m.) should be considered as too grandiose to take seriously. Now, if this is the case, then the multiplication rule *does* apply: the more solid the second-order epistemic standing of this worst-case estimate, the more likely it will be that the *de minimis* requirement is satisfied.

Hence, it turns out that aggregation is no real puzzle for PP. Carter and Peterson's intuition that multiplication delivers the wrong result with respect to *de minimis* is misleading: it is an artifact of how Carter and Peterson frame the puzzle, rather than a deficit of PP. Reframed in terms of the decision-tripod, they treat the damage condition (D) as variable when considering the epistemic threshold (E) that should be satisfied to undertake action (R). But in the scenario under discussion, we are already considering the worst-case realistic outcome. In other words, it is not realistic that things will get any worse. Therefore, greater second-order uncertainty regarding the first-order probability that the damage condition will be satisfied should lead to a *less*, rather than a *more* precautionary stance.

The solution to the aggregation puzzle connects with a further lesson about how to handle imprecise probabilities, in the context of PP. Suppose that the probability that some catastrophic event will occur can only be specified in terms of a probability range of 0.001–10%. Given the wide extent of this range, what outcome should we anticipate? Some proponents of CBA might proceed by making an educated guess of which value in this range is most plausible and adopt this value in their calculations. Proponents of PP, by contrast, will argue that a strategy along the lines of 'maxiproability' (maximize the most probable outcome) is reckless, if we are dealing with potentially catastrophic outcomes (cf. Rendall, 2019; Weitzman, 2011). It is an example of wishful thinking, akin to the decision of not concluding a fire insurance, even after evidence has been presented that one lives in an area where wildfires are becoming increasingly common. For proponents of PP, by contrast, the relevant outcome to focus on is the upper bound of the realistic range. If, on scientific grounds, the whole range of 0.001–10% should be regarded as realistic, then it is only natural to treat a 10% probability of catastrophe as a worst-case realistic outcome. The epistemic credentials of this assessment should thereby of course be

considered: it makes a difference whether the evidence is solid or shaky. Pace Carter and Peterson, however, shaky evidence should not be taken to imply that things might still turn out worse, since we are already considering the worst-case realistic possibility.

## 6 Conclusion

I have argued that the PP can be understood as a plausible guide to decision-making under uncertainty, which is specifically useful when considering worst case outcomes that seem to constitute real possibilities—i.e. possibilities that are couched in substantial evidential uncertainty, but that satisfy the *de minimis* requirement nonetheless. It would be a mistake to think that PP should only be appealed to when evidential probabilities are entirely absent: *in principle* the evidence on the basis of which one assesses whether the *de minimis* requirement has been met can be framed in probabilistic terms, although the resultant probabilities may be rather imprecise and not very helpful for purposes of decision-making. Furthermore, I have argued that the task of aggregating first- and second-order probabilities constitutes no distinct challenge for PP. In principle we can come up with an all-things-considered probability estimate to assess whether the *de minimis* criterion has been met—though I hasten to add that in practice an alternative approach, not couched in probabilistic terms, might well be equally, or more, opportune.

Several second-order epistemic considerations are relevant to the task of identifying realistic possibilities (cf. Hansson, 2005). In the case of climate uncertainty these include, among other things, the question of whether there is a solid mechanistic understanding of the system under consideration (e.g. Shue, 2018), whether there are independent lines of evidence supporting an envisioned outcome (e.g. Winsberg, 2018), the level and quality of expert consensus (e.g. Oreskes, 2018), the question of whether surprises are to be expected in the system under consideration (e.g. Parker & Risbey, 2015), as well as the question of whether there are historical precedents, which may help us to calibrate our assessments of what the climate system is capable of (Woo, 2019). These and other guidelines help to ascertain whether the *de minimis* criterion has been met and to make an all-things-considered judgement of whether a given outcome constitutes a real possibility.

Important philosophical work remains to be done in outlining these guidelines in a more rigorous fashion and assessing their epistemic merits. Additionally, important work remains to be done on the question of how to *articulate* and *communicate* real possibilities (e.g. van der Bles et al., 2019), by coming up with the fine-grained distinctions necessary to navigate the knowledge space just above the level of ignorance, but involving substantial uncertainty. The sources of uncertainty in making predictions about future states of the climate are manifold (Hopster, 2021), but so are the sources of evidence that lift us from ignorance. For instance, we might take an outcome to constitute a realistic possibility because our best scientific models suggest its occurrence lies within a plausible range, or because analogous events have occurred in the past. But we might also be uncertain about evidential probabilities because we are still at the initial stage of our research, and only have preliminary



data at our disposal, with no independent lines of supporting evidence. These different epistemic conditions, in turn, might call for different kinds of precautionary remedies. For instance, if we only have preliminary findings at our disposal, then the remedy (R) of PP might gravitate towards conducting further research. But if outcomes are uncertain because of properties inherent to the system being studied, then further research is unlikely to diminish uncertainty in the short run, and precautionary resources should be allocated differently. Hence, the differences between the kinds of uncertainty at issue in PP can be decision-relevant: they might call for different precautionary responses. To develop a more fine-grained vocabulary that distinguishes between these different types of uncertainty, then, in a language suitable for decision-makers, is an important task for future scholarship.

**Acknowledgements** I thank Lukas Meyer, Thomas Pözlner, Norbert Paulo and Elias Moser for their helpful comments on an earlier version of this article.

**Funding** This work has been supported by an NWO Rubicon grant for the project *Climate Futures: Real Possibilities and the Ethics of Uncertainty*.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Bamber, J. L., Riva, R. E. M., Vermeersen, B. L. A., & LeBroq, A. M. (2009). Reassessment of the potential sea-level rise from a collapse of the West Antarctic ice sheet. *Science*, 324(5929), 901–903.
- Betz, G. (2010). What's the worst case? The methodology of possibilistic prediction. *Analyse Und Kritik*, 32(1), 87–106.
- Birch, J. (2017). Animal sentience and the precautionary principle. *Animal Sentience*, 2(16), 1–16.
- Boardman, A. E., Greenberg, D. H., Vining, A. R., & Weimer, D. L. (2018). *Cost-benefit analysis. Concepts and practice* (4th ed.). Cambridge University Press.
- Bognar, G. (2011). Can the maximin principle serve as a basis for climate change policy? *The Monist*, 94(3), 329–348.
- Bostrom, N. (2013). Existential risk prevention as global priority. *Global Policy*, 4, 15–31.
- Broome, J. (2012). *Climate matters: Ethics in a warming world*. W.W. Norton & Co.
- Carter, J. A., & Peterson, M. B. (2015). On the epistemology of the precautionary principle. *Erkenntnis*, 80(1), 1–13.
- Carter, J. A., & Peterson, M. B. (2016). On the epistemology of the precautionary principle: Reply to Steglich-Petersen. *Erkenntnis*, 81(2), 297–304.
- Gardiner, S. M. (2006). A core precautionary principle. *The Journal of Political Philosophy*, 14(1), 33–60.
- Gardiner, S. M. (2011). *A perfect moral storm*. Oxford University Press.
- Goklany, I. M. (2001). *The precautionary principle: A critical appraisal of environmental risk assessment*. Cato Institute.
- Hansson, S. O. (2005). The epistemology of technological risk. *Techné*, 9(2), 68–80.
- Hansson, S. O. (2018). Representing uncertainty. In S. O. Hansson & V. F. Hendricks (Eds.), *Introduction to formal philosophy* (pp. 387–400). Springer.
- Hartzell-Nichols, L. (2017). *A climate of risk. Precautionary principles, catastrophes and climate change*. Routledge.

- Hopster, J. (2021). Climate change, uncertainty and policy. In G. Pellegrino & M. Di Paolo (Eds.), *Handbook of the philosophy of climate change*. Springer.
- Knight, F. (1921). *Risk, uncertainty and profit*. Houghton Mifflin.
- Koplin, J. J., Gyngell, C., & Savulescu, J. (2020). Germline gene editing and the precautionary principle. *Bioethics*, 34(1), 49–59.
- Manson, N. A. (2002). Formulating the precautionary principle. *Environmental Ethics*, 24, 263–274.
- Mastrandrea, M. D., Field, C. B., Stocker, T. F., et al. (2010). Guidance note for lead authors of the IPCC fifth assessment report on consistent treatment of uncertainties. *Intergovernmental Panel on Climate Change (IPCC)*.
- Ord, T., Hillerbrand, R., & Sandberg, A. (2010). Probing the improbable: Methodological challenges for risks with low probabilities and high stakes. *Journal of Risk Research*, 13, 191–205.
- Oreskes, N. (2018). *Why trust science?* Princeton University Press.
- Parker, W. S., & Risbey, J. S. (2015). False precision, surprise and improved uncertainty assessment. *Philosophical Transactions Royal Society A*, 373, 20140453.
- Pattyn, F., Ritz, C., Hanna, E., et al. (2018). The Greenland and Antarctic ice sheets under 1.5 °C global warming. *Nature Climate Change*, 8, 1053–1061.
- Peterson, M. (2006). The precautionary principle is incoherent. *Risk Analysis*, 26(3), 595–601.
- Powell, R. (2010). What's the harm? An evolutionary theoretical critique of the precautionary principle. *Kennedy Institute of Ethics Journal*, 20(2), 181–206.
- Rawls, J. (1999). *A theory of justice*. Harvard University Press.
- Rendall, M. (2019). Discounting, climate change, and the ecological fallacy. *Ethics*, 129, 441–463.
- Roser, D. (2017). The irrelevance of the risk-uncertainty distinction. *Journal of Science and Engineering Ethics*, 23, 1387–1407.
- Sandin, P., & Peterson, M. (2019). Is the precautionary principle a midlevel principle? *Ethics, Policy & Environment*, 22(1), 34–48.
- Selgelid, M. (2016). Gain-of-function research: Ethical analysis. *Science and Engineering Ethics*, 22, 923–964.
- Shue, H. (2010). Deadly delays, saving opportunities: creating a more dangerous world? In S. M. Gardiner, S. Caney, D. Jamieson, & H. Shue (Eds.), *Climate ethics: Essential readings* (pp. 146–162). Oxford University Press.
- Shue, H. (2018). Climate Surprises: Risk transfer, negative transmissions, and the pivotal generation. Online at SSRN: <https://ssrn.com/abstract=3165064> or <https://doi.org/10.2139/ssrn.3165064>
- Steel, D. (2013). The precautionary principle and the dilemma objection. *Ethics, Policy & Environment*, 16(3), 321–340.
- Steel, D. (2014). *Philosophy and the precautionary principle*. Cambridge University Press.
- Steglich-Petersen, A. (2015). The epistemology of the precautionary principle: Two puzzles resolved. *Erkenntnis*, 80, 1013–1021.
- Sunstein, C. (2001). *Risk and reason*. Cambridge University Press.
- Sunstein, C. (2005). *Laws of fear*. Cambridge University Press.
- Sutton, R. T. (2019). Climate science needs to take risk assessment much more seriously. *Bulletin of the American Meteorological Society*, 100(9), 1637–1642.
- Taleb, N.N., Read, R., Douady, R., et al. (2014). The precautionary principle (with application to the genetic modification of organisms). Working paper; online via: <https://arxiv.org/pdf/1410.5787.pdf>
- van den Hurk, B., Siegmund S., Klein Tank, A., eds. (2014). *Climate change scenarios for the 21st Century: A Netherlands perspective*. Scientific Report WR2014-01, KNMI, De Bilt.
- van der Bles, A., van der Linden, S., Freeman, A., Mitchell, J., Galvao, A., Zaval, L., & Spiegelhalter, D. (2019). Communicating uncertainty about facts, numbers and science. *Royal Society Open Science*, 6, 181870.
- Weitzman, M. L. (2011). Fat-tailed uncertainty in the economics of catastrophic climate change. *Review of Environmental Economics and Policy*, 5, 275–292.
- Winsberg, E. B. (2018). *Philosophy and climate science*. Cambridge University Press.
- Woo, G. (2019). Downward counterfactual search for extreme events. *Frontiers in Earth Science*, 7, 340.