

# Arguing for and interpreting epistemic possibilities in climate science

## Joel Katzav

Email: [j.katzav@uq.edu.au](mailto:j.katzav@uq.edu.au)

Forthcoming in *European Journal for Philosophy of Science*.

<https://doi.org/10.1007/s13194-025-00690-0>

### Abstract

Recent work on the epistemology of climate science includes arguments that are against probabilistic representations of uncertainty about climate and for possibilistic ones as well as some development and use of the latter. I reinstate these arguments, partly by rebutting Corey Dethier's recent challenge to them and partly by arguing that they remain effective against recent improvements to probabilistic representations. Recognising, however, that the case for possibilistic representations can be undermined by problematic interpretations of epistemic possibilities, I set out criteria of adequacy for such interpretations in the climate context while arguing for a preferred interpretation. I criticise the appropriateness of standard interpretations, according to which a proposition is epistemically possible if and only if it is not recognised to be excluded by what is known, as well as some other prominent non-probabilistic interpretations. So too, I criticise interpretations of epistemic possibilities in terms of upper probabilities. I conclude that an interpretation of epistemic possibilities as possibilities that are consistent with knowledge that approximates the basic way things are is preferable to the other available interpretations.

**Keywords:** Uncertainty in Science; Uncertainty in Climate Science; Possibility Theory

### 1. Introduction

Work on the possibilistic epistemology of climate science has focused primarily on how to evaluate uncertainty about the predictions of state-of-the-art climate models, that is, state-of-the-art global climate models and Earth system models. In turn, most of the work with this focus argues that uncertainty about state-of-the-art models' predictions should typically be expressed in terms of epistemic possibilities rather than probabilistically (Stainforth et al. 2007; Betz 2010 and 2015; Katzav et al. 2012; Katzav et al. 2021; Le Cozannet et al. 2017). I, more specifically, argue that uncertainty assessment of these predictions should typically be of the extent to which they are epistemically possible and that, in some cases where they are

epistemically possible, the possibilities should further be ranked as to how remote they are (Katzav 2014; Katzav et al. 2021).

Criticism of the case for preferring possibilistic over probabilistic interpretations of climate model predictions is provided by Dethier (Dethier 2023). In addition, the case requires reconsideration given recent, more sophisticated probabilistic approaches to generating probabilistic representations of uncertainty about future climate. Such approaches involve, for example, the relaxation by the Intergovernmental Panel on Climate Change (IPCC) sixth assessment report (AR6) of the assumption that all state-of-the-art climate models are equally good (Lee et al. 2021).

In the course of criticising the case for the possibilistic approach, Dethier also criticises the approach itself. Further criticism of the approach is offered by Frame et al. (Frame et al. 2007) and Frigg et al. (Frigg et al. 2015), each of whom note the need for a specification of which model predictions are to be identified with epistemic possibilities and thus as those that should be taken seriously in practice, that is, in contexts other than the context of discovery. Betz (2015) and I (2014; 2023), however, offer interpretations of epistemic possibilities that specify how a prediction must relate to existing knowledge in order to count as epistemically possible. Additional, relevant interpretations of epistemic possibilities can be found in the literature on the semantics of modals (Przyjemski 2017) and the literature in logic and computing (Dubois and Prade 2015). In presenting my interpretation, I have argued that it ought to be preferred over Betz's when it comes to representing uncertainty about climate (Katzav 2023). My argument and Betz's original argument (Betz 2015) for his own position comprise what limited argument there is by philosophers about how to interpret epistemic possibilities in climate science.

Beyond discussion of the case for and against the possibilistic approach to representing uncertainty about climate, there is some actual application of it. The most prominent interpretation in the logic and computing literature identifies epistemic possibilities with upper probabilities. This interpretation has been applied in assessing the output of state-of-the-art climate models (Le Cozannet et al. 2017; Tebaldi et al. 2021). So too, I argue that my notion of epistemic possibility matches the one actually used as part of informal evaluations of the predictions of climate models in the context of discovery and in contexts where discovery and application overlap (Katzav 2023).

My aim here is to reinstate the case against probabilistic and for possibilistic assessments of our uncertainty about the predictions of state-of-the-art climate models. I aim to do so partly through a rebuttal of the objections to this case. According to Dethier, the arguments against probabilistic assessments tell us that the predictions of climate models and climate model ensembles should not be interpreted probabilistically because the models involved are idealised. On his view, such arguments fail because idealisation is not itself a reason to suppose that models are unreliable. I will show, however, that this misrepresents the arguments Dethier discusses. The relevant parts of these arguments tell us that since we typically lack a reliable way of assessing what state-of-the-art climate model empirical successes and failures imply about the accuracy of uncertain model assumptions and predictions, probabilistic interpretations of this uncertainty typically threaten to be unreliable to an unquantifiable extent. When properly understood, these arguments are good as they stand. Dethier further claims that the arguments supporting possibilism do not recognise that probabilistic but not possibilistic assessment of our uncertainty is suited to assessing extremes of uncertainty distributions, where uncertainty is high. I reject this claim and, in doing so, point to and fill out a further argument that is found in the literature, one according to which probabilities are not suited to representing ignorance and partial ignorance.

My support for the case for possibilistic representations of climate uncertainty is completed in two further steps. First, I consider the extent to which the most sophisticated improvements to probabilistic representations of uncertainty in the climate context avoid the challenges to such representations. Second, I am more explicit than the existing literature about why the challenges to probabilistic representations are avoided by formal, possibilistic representations of uncertainty.

I recognise, however, that the case for possibilistic representations of climate uncertainty and the suggestion that epistemic possibilities are decision relevant can be undermined by problematic interpretations of epistemic possibilities. Thus, in addition to restating and defending the case for a possibilistic approach to representing our uncertainty about climate, I develop and apply criteria of adequacy for interpreting epistemic possibilities in the climate context. My existing case for preferring my notion of epistemic possibilities over Betz's rests on the requirement that such notions should not make assessing uncertainty too demanding to be practicable. I here examine existing interpretations from the philosophy of climate modelling and the semantics literature in light of this requirement but also in light of two more requirements. One of these is the requirement that an interpretation should be such that predictions are taken to be epistemically possible only if they overlap with the real potentialities of the climate system. The final requirement, which I apply to the view that epistemic possibilities are to be identified with upper probabilities, is that interpretations of epistemic possibilities should avoid reintroducing the problems faced by probabilistic interpretations of uncertainty. I conclude that my interpretation is preferable to available alternatives.<sup>1</sup>

---

<sup>1</sup> I will not, in what follows, consider issues relating to possibility-based decision making. For a discussion of these that is sympathetic to the possibilist approach see Roussos (2025).

One could try to extend the existing case for my interpretation of epistemic possibilities by appealing to the argument that it best fits the actual use of epistemic possibilities in climate science. But the present paper aims to answer the more basic question of which interpretation ought to be adopted in the first place. Considering how to interpret epistemic possibilities matters because this question is relevant to how possibilistic assessments are carried out, including to determining which possibilities ought to be taken seriously in practice. The possibilistic approach itself deserves further development not only because of the case for its preferability over probabilistic approaches but because of the use of possibilistic ones.

I begin, in section 2, by responding to Dethier's criticism of the arguments for probabilistic approaches to representing uncertainty about climate. I then, in section 3, consider whether recent developments in probabilistic approaches adequately respond to the arguments. In section 4, I introduce possibilistic representations of uncertainty and explain their advantages over probabilistic ones in the climate context. In sections 5 and 6, I discuss how to interpret epistemic possibilities. Section 5 focuses on interpretations from the philosophy of climate modelling and the semantics literature while section 6 focuses on interpretations that identify epistemic possibilities with upper probabilities. My conclusion is in section 7.

## **2. Problems for probabilism in climate science**

Dethier identifies two arguments for the application of possibilistic approaches to uncertainty in the climate context. He takes the first of these to be focused on the assessment of individual models. According to Dethier, Betz (2015) and I (2014) "suggest that we're not justified in treating a climate model as providing information about the actual world when we know that said model includes false assumptions" (Dethier 2023, p. 1418). To this, Dethier objects that the mere fact that certain models are idealised is not an obstacle to supposing that

they are reliable and, further, that there is nothing specific to climate models that leads to supposing that they are unreliable because of their idealisations. Thus, on his objection, if we are to argue that climate models or climate model ensembles provide us with unreliable probabilistic estimates of uncertainty, we will have to do so on a case-by-case basis, identifying in each case special reasons for distrusting the estimates (Dethier 2023, pp. 1418-1419).

Dethier takes proponents of possibilistic interpretations of uncertainty to have, more recently, offered a further argument for their position. He identifies this argument with a standard argument in the climate science literature that concerns what has been the preferred approach to producing probabilistic predictions of future climate. The approach produces such predictions using ensembles of state-of-the-art climate models, especially the ensembles of the Coupled Model Intercomparison Project (CMIP) (IPCC 2013; Lee et al. 2021). The range of the predictions of these ensembles is often taken to be an indication of the range of possibilities that we ought to take seriously. Agreement between ensemble models about a prediction is taken to indicate an increased probability for the prediction. Thus, for example, the IPCC fifth assessment report (AR5) chapter on short-term predictions of future climate tells us that, with one exception, “the spread among the CMIP5 models is used as a simple, but crude, measure of uncertainty. The extent of agreement between the CMIP5 projections provides rough guidance about the likelihood of a particular outcome” (Kirtman et al. 2013, p. 982). Similar statements are made in the AR5 chapter on long-term predictions and in AR6 (e.g., Collins et al. 2013, p. 1038; Lee et al. 2021, p. 568). On the basis of these considerations about ensemble spread and inter-model agreement, the runs of models in

model ensembles are collectively transformed into partial or full probability distributions.<sup>2,3</sup> However, the ensembles are ensembles of opportunity, that is, they comprise models that are not designed to be part of a systematic exploration of our uncertainty about the climate system but instead individually to provide as good a representation of climate as feasible. The worry is that model agreement in such ensembles reflects shared assumptions and aims but not increased evidence. So too, model disagreement will not span the full range of uncertainty. As a result, it is concluded, probabilistic approaches lead to overconfidence about where to place the bulk of the probability mass of predicted variables as well as related neglect of extremes (for better and for worse), because they are assigned low probability (Tebaldi and Knutti 2007; Baumberger et al. 2017; Katzav et al. 2021; Lee et al. 2021).

According to Dethier, the appeal to ensembles of opportunity tells us that because the probability distributions about future climate which are based on the output of such ensembles are idealised, these distributions are not the probabilities we ought to have. In response, he repeats the claim that the fact that a model is an idealisation does not suffice to make the model unreliable. A model of uncertainty might still be the best model we have and, as such, might still be the model on which we rely (2023, p. 1420).

Dethier, however, misrepresents my arguments against probabilistic interpretations of uncertainty about climate. I do not argue that since models are idealised, they are not reliable. In my most general argument (Katzav 2014), I am concerned with what the target of climate model assessment should be when we aim to learn something new about climate from a state-of-the-art model or an ensemble of such models. As part of this argument, I do draw a

---

<sup>2</sup> The transformation involves fitting continuous statistical distributions (e.g., Gaussian ones) to the runs of ensemble models and then ignoring the details of the shapes of these distributions and associating with their ranges (e.g., the 5-95% range) uncertainty ranges (e.g., a likely or very likely range) (Collins et al. 2013, p. 1055).

<sup>3</sup> Dethier focuses on a different approach to interpreting the output of models than the IPCC's one (Dethier 2022). I focus on the IPCC's approach because it is my focus in earlier work (2014), and my goal is to evaluate my arguments. I will, however, consider Dethier's preferred approach below.

conclusion from the fact that such models are known to be false. But the conclusion is just the trivial one that, in such cases, models' truth as such cannot be confirmed and thus that the target of confirmation assessment needs to be confirming something other than models in their entirety. Similarly, I point out that since the conjunctions of assumptions upon which the empirical successes of state-of-the-art climate models depend include assumptions that are known to be false, we cannot aim to confirm the truth of such conjunctions (2014, p. 230). Indeed, I go on to consider further ways of specifying confirmable assumptions or results of these models despite idealisation (2014, pp. 230-235).

The main challenge that I see for assessment of state-of-the-art climate models that is feasible and useful is that, because we typically have no reliable way of determining what the implications of these models' empirical successes and failures are for the accuracy of the models' uncertain assumptions and predictions, taking some set of these to be confirmed by successes typically threatens to be (and sometimes is known to be) an unreliable way of learning about the climate system. As a special case, taking model successes to confirm probabilistic interpretations of climate model predictions threatens to be unreliable. The reason we typically have no reliable way of determining the implications of the models' successes and failures is that the models themselves are the main tool for arriving at quantitative judgements about the climate system. As a result, if we want to know how some known model empirical success or limitation bears on a models' uncertain assumptions or predictions, it is to the models themselves that we need to turn. And yet, it is precisely the models that we are trying to assess. So, we are stuck.

Dethier neither outlines nor responds to this argument. The closest that he comes to doing so is in responding to the suggestion that there might be empirical evidence supporting the thesis that the models should not be trusted. His response is to cite literature on the empirical successes of the models and to assert that it shows that the models should be trusted

more for some purposes than others and that ensembles should be trusted more than individual models (2023, pp. 1418-1419). However, recognising the empirical successes of the models is the starting point of the possibilist argument. The challenge of the argument is to figure out whether the models are reliable given their empirical successes *and* failures. We can add (Katzav 2014, p. 235) that if we have a method that is, or threatens to be, unreliable and, like the ones being discussed here, is being used to inform important policy, we should replace or improve on the method. This applies even if the method is the best we have. It certainly applies if, as in the present case, there are alternative approaches to interpreting the predictions of climate models.

What Dethier takes to be my more recent appeal to standard worries about ensembles of opportunity is the just outlined argument applied to the special case of confirming probabilistic, including ensemble-based, predictions (Katzav et al. 2012; Katzav et al. 2021, pp. 7-11). In this formulation, the argument is that there is no reliable empirical, theoretical, or expert basis for estimating the consequences of known limitations involved in generating predictions about the climate by state-of-the-art models and that, therefore, probabilistic estimates of our uncertainty about these predictions are unreliable to an unquantifiable extent. In particular, we have no principled basis upon which to assess the reliability of ensembles of opportunity-based representations of our uncertainty, and there is reason to suspect that they substantially do misrepresent this uncertainty in policy-relevant ways and are thus unreliable. Relevant misrepresentations due to available CMIP ensembles include those identified in standard worries about them, namely those concerning the location of the bulk of the probability mass of predictions and the exploration of the extremes of our uncertainty.<sup>4</sup>

---

<sup>4</sup> My argument here should be distinguished from the standard worries. The standard worries are about the fact that the impacts of uncertainty about model parameters and structures on prediction ranges is not adequately explored by ensembles. My worry is about how ensembles are in principle

To illustrate the possibilist argument, consider a segment of it (Katzav et al. 2012). This segment observes that state-of-the-art climate models from about 2010 were able to simulate the collapse of the Atlantic meridional overturning circulation (AMOC) in the twenty-first century but only by assuming about ten times the influx of fresh water into the ocean than best estimates of what historically caused such collapse. The argument adds that, at the time, there was no well-supported theoretical or empirical estimate of how this limitation in models' representations of the climate system would affect their predictions for the twenty-first century though it clearly would be substantial. As a result, ensemble-based representations of such uncertainty were substantially incomplete in decision-relevant ways and were thus unreliable (Katzav et al. 2012, pp. 264-267). Not surprisingly, challenges to estimating uncertainty relating to AMOC collapse remain (van Westen et al. 2024). To respond to such an argument by saying that state-of-the-art models' representations of AMOC collapse are just idealisations is not to recognise the magnitude of the relevant errors, the absence of a principled understanding of their impacts on uncertainty beyond the knowledge that it is substantial, and the known significance of the impact for our future. It would be irresponsible to use an assessment that fails to take the possibility of AMOC weakening and collapse into account even if it were the only assessment we had. We would be obliged to develop a better one. To add that the models have substantial empirical successes in simulating various known climate phenomena, such as twentieth-century global mean surface temperature change, would be irrelevant given the specific worries about their ability to simulate changes in AMOC strength.

This response to Dethier might seem to make pressing his claim that arguments against probabilistic interpretations of climate model ensembles need to be developed on a

---

unable adequately to assist in estimating what their successes and limitations imply. That said, I think that Dethier also misrepresents the standard worry when he takes it merely to be about the use of idealisation. It too is about known, policy relevant limitations in representations of uncertainty.

case-by-case basis. However, the kinds of limitations appealed to in arguments for possibilistic interpretations of uncertainty suffice to sustain general worries about such interpretations. The worry about model ability to simulate AMOC collapse is intimately tied to a worry about limitations in simulating internal variability. There is no established quantitative theory of patterns of internal climate variability and their impacts that allows reliably estimating the impacts of substantial limitations in representation of such variability on climate predictions. Moreover, since internal variability is always only a component of observed climate, there is limited ability to estimate it in more empirical ways. Thus, we know that model predictions are typically unreliable to an unquantifiable extent when used as an indicator of where to locate the bulk of the probability mass regarding ensemble predictions as well as about where to locate the tails (Katzav 2014; Katzav et al. 2021, pp. 8-9). Similar considerations relate to the limited abilities of state-of-the-art models to represent sub-grid phenomena that affect cloud and aerosol processes (Katzav et al. 2021, pp. 8-9). For example, state-of-the-art climate models have been unable to resolve dynamical instabilities related to ocean currents, instabilities that affect atmospheric variables of interest—including, precipitation and wind—on the regional scale. As a result, one should be worried about whether model output can meaningfully capture information about the evolution of these variables on regional scales (Risbey and O’Kane 2011). Even the range of ensemble-based estimates of a key quantity such as (equilibrium) climate sensitivity may well substantially underestimate the actual possibilities for the future (Stainforth et al. 2005; Rowlands et al. 2012; Yamazaki et al. 2013).

If we are worried about climate model-based probabilistic predictions which rest on problematic assumptions about internal variability, sub-grid processes that are relevant to regional predictions in general, and climate sensitivity, we have a general worry about model-based predictions. If there are predictions of interest to which these worries do not apply, this

will require argument (Risbey and O’Kane 2011). Internal variability, for example, is generally a component of developments in the climate system. Such variability matters to global and regional climate predictions of key climatic quantities, including mean temperature trends, temperature extremes, sea-level rise, precipitation levels, and other quantities, on decadal, multi-decadal, and centennial scales (Deser et al. 2020; Lee et al. 2021; Meccia et al. 2023). It is only in special cases, such as estimates of climate sensitivity, that uncertainty about internal variability will arguably not matter. So, worries about underestimating uncertainty due to internal variability generally translate into worries about our predictions.

In addition to arguing that the possibilist is wrong to object to probabilistic approaches on the ground that models and model-based uncertainty distributions are idealised, Dethier objects that the possibilist’s reluctance to idealise leads to an unrecognised advantage for probabilistic approaches to uncertainty. On his view, the possibilist must suppose that ensembles of opportunity cannot assist us in representing the extremes of our uncertainty, that is, representing partial or complete ignorance about modelling results. This is because the possibilist aims to avoid misrepresenting extremes and, as a result, “refuses to say anything about them at all” (2023, p. 1424). At the same time, Dethier observes, one can at least fit a full probability distribution to the output of ensembles of sophisticated models, and the tails of the distribution will provide us with information about extremes. Here, Dethier follows some of the climate science literature and recommends that in constructing the distributions for the outputs of CMIP ensembles, we assume that each output is drawn from the same sample as the truth. As he puts it, we should assume “that the “truth” behaves like a single sample from the same population as the models” (2023, p. 1424).

Yet the possibilist is not objecting to idealisation as such and so can fit a continuous distribution with long tails to the discrete predictions of an ensemble of models. More

importantly, the shared, policy-relevant limitations of models within ensembles mean that it is a mistake to assume that their predictions are drawn from the same family as relevant climatic events. For example, fitting a probability distribution to CMIP models' predictions about future climate on the assumption that they are drawn from the same family as the truth ignores shared, policy-relevant biases relating to simulating AMOC weakening and collapse. In this case, model output is not from the same sample as the truth (Katzav et al. 2012, pp. 266-267).<sup>5</sup> In order to assess the impacts of these limitations on the predictions of CMIP ensembles and thus to be better placed to interpret what they mean for our uncertainty about future climate, we have to draw on knowledge of processes not adequately represented in the models. Doing so includes relying on information concerning processes about which there is near ignorance and choosing representations of uncertainty that are suitable for representing such ignorance.

This brings us to a second challenge to probabilistic representations of uncertainty about future climate. Representations of near ignorance are needed when the concern is with possibilities that we do not know are real potentialities of the climate system while there is some preliminary evidence that suggests that they are (Katzav et al. 2021, p. 14). The straightforward probabilistic approach to representing near ignorance about potentialities specifies a low probability that is a threshold for being taken seriously. Alternatively, one might specify a range of probabilities, e.g., from some specified low value to some case-by-case specified higher value. But neither of these options adequately represents our uncertainty

---

<sup>5</sup> Related points are made in the climate modelling literature. For example, it is argued that even if we assume that model output is from the same distribution as the truth when it comes to current observations, it is not clear how this assumption can be made about predictions of future climate (Sanderson and Knutti 2012). More forcefully, it is argued that this assumption fails to recognise how shared biases among ensemble members impact predictions (Sansom et al. 2021).

about predictions that should not be taken seriously or that stand at the threshold for being taken seriously.

Assigning a low probability to predictions of phenomena where evidence is too preliminary to establish that they are of real potentialities neglects recognising uncertainty about them, specifically uncertainty about what the evidence shows and how it will evolve in the future. This can be seen in that the uncertainty we would assign to a kind of event that is not clearly a real potentiality would have to be identified with the uncertainty assigned to some events with well known, low frequencies despite the extra uncertainty associated with the former. In the former case, it is unclear what probability to assign because available evidence is provisional and perhaps contradictory in what it suggests. In the latter case, assigning a probability is comparatively straightforward. Thus, for example, we can reasonably assign a low probability to a prediction of a supervolcanic eruption by 2100 because we know such eruptions have a low frequency. On the other hand, that marine ice-cliff instability in Antarctica will contribute substantially to sea-level rise by 2100 is not yet known to represent a real potentiality of the climate system; relevant evidence is too limited, conflicting, and expected to evolve substantially (Pattyn and Morlighem 2020; Morlighem et al. 2024). So, to take specific estimates of marine ice-cliff instability contributions to sea levels to have a low probability would be to ignore some of our uncertainty about these contributions. Even assigning a range of low probabilities would be problematic here, since we are not in a position sufficiently to constrain our estimates of how relevant knowledge will evolve and thus to provide an informative upper bound for the range.

Climate scientists aim to represent ignorance about the climate system when producing prior probability distributions for quantities, such as climate sensitivity, that are to be updated in light of data. In such cases, there is a need to represent the uncertainty we would have had about values of quantities of interest if we did not have empirical evidence of

their relative probabilities (Katzav et al. 2012, pp. 268-269). Standard practice is to use the principle of indifference and to assign equal probabilities to values. But we do not know whether equiprobable assignments of priors will yield reliable predictions upon updating. Moreover, we have no principled way of selecting outcome spaces under ignorance. We get different posterior probability distributions for climate sensitivity, for example, if the outcome space is climate sensitivity and if it is climate sensitivity and feedbacks (Katzav et al. 2012, p. 261).

### **3. Probabilistic responses in climate science to problems with probabilism**

Unlike Dethier, climate scientists have not dismissed the challenges to ensemble-based probabilistic predictions and to probabilistic prediction more broadly. AR5 aimed partially to address the worry that the CMIP5 ensemble does not capture the full range of our uncertainty by reducing confidence in the ensemble's ranges of predictions. 90% of each range, specifically the 5-95% interval, was only taken to be the likely range ( $\leq 66\%$ ). AR6 continued to do this regarding some key quantities, such as those relating to precipitation (Lee et al. 2021).

In some cases, AR5 used an ensemble of models to generate a likely range of predictions and then used information about processes not represented in the ensemble to explore phenomena outside of that range. Thus, for example, the report used an ensemble comprising the CMIP5 models and process models of ice sheets and glaciers to predict that, under representative concentration pathway (RCP) 8.5, it is at least likely that global mean sea-level rise will be between 0.52 and 0.98m by 2100. They used information about less understood processes that were not captured in this ensemble, especially about marine ice-sheet instability, to explore further possibilities. They judged that a rise of a few further 10ths of a meter would not be exceeded by 2100 but added that the probability for exceeding the likely range could not be precisely assessed (Church et al. 2013, pp. 1178-1180). Similarly,

AR6 relied on a variety of ensembles of opportunity, including the CMIP6 ensemble, to estimate contributions of different sources of sea-level rise to total sea-level rise by 2100. The contributions estimated by ensembles were summed up to provide a likely range for sea-level rise by this year. Scenarios that draw on information about processes not captured by the ensembles were then used to provide information about unlikely ranges of such rise, though no further information about the uncertainty of the scenarios was offered. The processes contributing to extremes in this case included marine ice-sheet instability and marine ice-cliff instability (Fox-Kemper et al. 2021, pp. 1295-1307).

The AR5 approach of taking the 5-95% interval to be what is, at least, likely is an improvement over taking CMIP ensemble predictions at face value but by itself leaves us with probability distributions that provide no information outside the likely range. This limits decision making. Further, the distributions still rest on the problematic assumption that model-agreement implies increased probability. As for AR5 and AR6's use of evidence to draw conclusions about the extremes of our uncertainty about sea-level rise, this evidence is not used to estimate the uncertainty about the extremes. More explicitly, no further specification of uncertainty is associated with the extra 10ths of a meter of rise that AR5 entertains, so that we are left with an uncorrected, and thus problematic, likely range with the addition of some information about what the unlikely range might involve. Similarly, no further specification of uncertainty is associated with the scenarios of sea-level rise considered in AR6, so that once again we are left with an unaltered probability distribution.

Another improvement to ensemble-based generation of probability distributions is the recent relaxation of the longtime assumption that each model in an ensemble has equal weight in producing probability distributions. Instead of assuming that confidence in a prediction increases solely as a function of the number of models that produce the prediction, it is possible to reduce the weight of the contributions of models to confidence in a prediction

when the models have a similar history of development. As a result, the double counting of evidence for predictions that results when similar models in an ensemble provide similar predictions will be reduced (Kuma et al. 2023). Weighting models, however, can only reduce some of the overconfidence in areas where clustering occurs. For example, it does not assist in addressing worries about shared limitations across available state-of-the-art models.

It is also possible to use model ranking to account for differential model successes, that is, cases when some state-of-the-art models in an ensemble are in some ways less empirically successful than others in the ensemble. In such cases, the claim sometimes is that we should reduce the weight we give to the predictions of the less successful models in generating probability distributions in comparison with the weight given to predictions of other models. Most notably, the CMIP6 ensemble includes models that predict relatively high climate sensitivities but that do relatively poorly at simulating recent observed global mean temperature changes. Partly as a result and partly because the high sensitivities are supposed to be very unlikely given more empirical methods of estimating climate sensitivity, AR6 gave the predictions of these models less weight in generating probability distributions of future global mean surface temperature change (Eyring et al. 2021; Hausfather et al 2022). It turns out, however, that improved simulation of recent observed global mean temperature changes need not correlate with improved predictions of other climatic quantities, such as regional precipitation change. As a result, when to reduce the weight of the predictions of ‘hot’ models is not straightforward (McDonnell et al. 2024). More importantly for present purposes, it is unclear why down weighting predicted extremes is wise when it comes to ensembles about which we are worried that they do not adequately explore extremes. Further, such weighting underestimates the challenges to probabilistic approaches to assessing our uncertainty about ensemble outputs. If, as I have argued, our worries about the unreliability of this approach are due to an inability to assess the implications of model successes and

failures for predictions, we ought to resist the temptation to downgrade our confidence in predictions of hot CMIP6 models merely on the grounds that they are less successful in some ways. This leaves us with the suggestion that the weight of some models in generating probability distributions should be reduced because there is empirical evidence that the model predictions are very unlikely. Whether such a suggestion should be accepted depends on whether more data-driven methods do justify probabilistic predictions, a question I now consider.

There have been increasingly sophisticated attempts to rely on more data-driven methods in assessing our uncertainty about future climate. Most notably perhaps, AR6 provided such an assessment of climate sensitivity and used it to inform projections of other quantities, including globally averaged air surface temperature increases and global mean sea level rise (Forster et al. 2021; Lee et al. 2021). Following Sherwood et al. (2020), climate sensitivity was estimated using variety of evidence considerations, including appeals to evidence from feedback analysis/process studies, the instrumental record, paleoclimate data, and emergent constraints. Prior probability distribution functions for climate sensitivity were updated in light of each source of evidence and resulting posterior distributions were used to create a single probability distribution. The resulting probability distribution was then used to constrain ensemble-based estimates of global mean surface temperature trends and of some other key climatic quantities. This approach substantially reduces reliance on state-of-the-art climate models, though, as Sherwood et al. recognise, it brings to the fore ongoing issues relating to prior selection. Indeed, on my view (Katzav et al. 2021, p. 11), the underlying worry that we lack the theoretical and empirical basis for principled probabilistic assessments of our predictions applies to this more data based, probabilistic assessment of uncertainty. Since we lack the background theory needed to delimit the space of hypotheses that might explain the various bodies of data used to constrain equilibrium climate sensitivity, we are

not in a position to determine whether these bodies of data might be equally, or better, explained by other hypotheses than those which are preferred and thus in a position to use existing ones to delimit our uncertainty probabilistically.

In summary, while there has been progress in developing probabilistic representations of our uncertainty of future climate, the basic challenge of deciding where to place the bulk of the probability mass of such distributions remains as does the challenge of representing the extremes of our uncertainty. In line with this, progress with probabilistic approaches does not purport fully to tackle the challenges to them. Thus, while we should acknowledge the improvements to probabilistic approaches to assessing our uncertainty, there is a need to further develop non-probabilistic approaches to doing so. The need also exists because possibilistic approaches are starting formally to be used in climate assessment that relies on state-of-the-art models and because it is being used informally in assessing uncertainty about the predictions of climate models more broadly.

#### **4. Possibilistic representations of uncertainty and some of their advantages**

Let me turn to providing some clarification about the possibilistic approach to representing uncertainty and how it avoids the just posed challenges for probabilistic ones. We have an informal understanding of reasoning about epistemic possibilities. When representing our uncertainty about claims concerning complex future events, we often hesitate to assign probabilities to the claims and, instead, represent our uncertainty with epistemic possibilities. Such representations are then used as a basis for decision making. In thus using them, we conceive of epistemic possibilities as possibilities that are to be taken seriously in practical contexts. Thus, after Brexit, some of us might hesitate to assign a probability to a judgement according to which other large European countries will leave the European Union (EU) but accept that further withdrawals are epistemically possible. If we do, we take this possibility seriously in practice. At the individual level, for example, an EU citizen living in an EU

country of which they are not citizens might acquire citizenship of that country. At the collective level, we might act in ways that shore up support for the EU, or we might do the opposite, depending on our politics and situation.

Informally, we also distinguish between epistemic possibilities and possibilities that are not quite epistemically possible. While the former are possibilities upon which we ought to act in practice, it is at least unclear whether the latter are ones on which we should act. For example, we might be unclear whether it is possible for us to get a coveted job for which there are hundreds of applicants and, as a result, be unsure whether there is a point in applying for the job.

In formal discussions, epistemic possibilities are termed ‘full possibilities’ while possibilities that are not quite epistemically possible are termed ‘partial possibilities.’ A (quantitative) possibility distribution function,  $\pi(p)$ , over a set of propositions,  $S$ , can then be defined as follows (Dubois and Prade 2015, pp. 34-35):

- (1) For any  $p$  in  $S$ ,  $\pi(p) \in [0,1]$ , where  $\pi(p)=1$  means  $p$  is fully possible,  $\pi(p)=0$  means  $p$  is impossible and  $0<\pi(p)<1$  means  $p$  is partially possible.<sup>6</sup>
- (2) For at least one  $p$  in  $S$ ,  $\pi(p)=1$ .

Possibility and necessity measures are, in turn, defined on  $\pi(p)$ :

- (3) The degree to which a subset  $A$  of  $S$  is possible is given by the possibility measure

$$\Pi(A) = \sup_{(p \in A)} \pi(p)$$

- (4) The degree to which  $A$  is necessary is given by the necessity measure  $N(A) =$

$$\inf_{(p \notin A)} (1 - \pi(p))$$

---

<sup>6</sup> Qualitative possibility theory relaxes the assumption that possibilities are specified on the  $[0,1]$  interval. Qualitative possibility theory will not be further discussed here.

The possibility measure satisfies the following maxitivity property:

$$(5) \Pi(A \cup B) = \max(\Pi(A), \Pi(B))$$

(2) is a consistency condition. Roughly, if no proposition in  $S$  is taken to be fully possible, we recognise that further investigation might indicate that none of them is even partially possible and thus that there are possibilities outside of  $S$ . (3) and (5) involve the rejection of additivity and thus differentiate possibility distributions from probability distributions.

Possibility theory comes with its own conditionalization principle, which is the analogue of Bayes' principle in probability theory (Dubois 2006). For evidence,  $e$ , and an hypothesis,  $h$ , the possibilistic equation for updating on new evidence is:

$$(6) \Pi(h|e) = \Pi(e|h)\Pi(h)/\Pi(e)$$

However, in contrast with probability theory, priors in possibility are uninformative, so that (6) can be reduced to the following equation, where  $h^c$  is the complement of  $h$ :

$$(7) \Pi(h|e) = \min(1, \frac{\Pi(e|h)}{\Pi(e|h^c)})$$

One respect in which the formal side of possibility theory is not as developed as the informal side concerns the ranking of epistemic possibilities. Formal representations treat epistemic possibilities as equal while we sometimes informally distinguish between more and less remote ones. For example, now that Trump is in office for a second term, we might judge that the possibility that there will be a three-term presidency in the near future is less remote than it was.

That epistemic possibilities are not additive means that they are better suited to interpreting the typical predictions of climate models than are probabilities. If, for example, all the models in an ensemble of state-of-the-art models agree on a prediction and each model

is taken to establish that the prediction is fully possible, the non-additivity of possibilistic measures of uncertainty means that the prediction remains no more than fully possible. Similarly, if we test the robustness of a model result to changes in parameters, the non-additivity of possibilistic measures means that the result will be no more possible than it was assumed to be prior to the test. In this way, we avoid increasing/decreasing confidence in model predictions on the basis of model successes and failures the implications of which for predictions are unclear. In particular, we do not assume that inter-model agreement justifies increased confidence in shared predictions and thus avoid problematic specifications of the bulk of predictions' probability mass.

So too, partial possibilities allow us to compensate for the problematic clustering of the output of ensembles by accounting for knowledge relevant to our uncertainty about the extremes of distributions that is not captured by the ensembles. We can address the problem of distinguishing between uncertainty about possibilities that are known to be realised but occur rarely and those that, like marine ice-cliff instability induced sea-level rise in the twentieth century, are not known to be real. This, building on everyday practice, would involve reserving representations of full possibilities for possibilities that are to be taken seriously in practice while interpreting partial possibilities along a range from those that should not be taken seriously in any context, to those that should be taken seriously in research but not in practice and, finally, to those about which it is unclear whether they ought to be taken seriously in practice. In addition, an adequate notion of full possibility would at least assist in working out our full range of uncertainty by providing guidance about the extent to which knowledge, including model-based knowledge, used to generate climate predictions suffices to justify taking the predictions to be fully possible.

Finally, the problem of selecting prior uncertainty distributions for conditionalization can be bypassed by appealing to possibility theory. This is so because, recall, possibilistic

conditionalization can be prior independent. When, say, updating our possibilistic knowledge about climate sensitivity in light of empirical evidence, what we would be doing is akin to carrying out a likelihood test in probability theory.

### **5. Non-probabilistic interpretations of epistemic possibilities**

The above concludes my discussion of the case for the possibilistic approach and how the approach aims to remove specific threats to the reliability of probabilistic approaches. I now turn to ensuring that we have a notion of epistemic possibilities that does not undermine this case.

Standard interpretations of epistemic possibilities are shared by the semantics literature and the logic and computing literature. These interpretations take epistemic possibilities to be those that are not recognised to be excluded by what is known (Dubois and Prade 2015; Przyjowski 2017). Such interpretations are not sufficiently nuanced. That a proposition is not recognised to be excluded by what is known is compatible with (complete) ignorance about whether the proposition is true. So, standard interpretations classify propositions about which we are ignorant as epistemic possibilities. This leads, assuming that epistemic possibilities are to be taken seriously, to taking propositions about which we are ignorant seriously. Yet doing so can leave us with too many possibilities to reason about in circumstances where decisions are needed. An interpretation of epistemic possibilities ought, accordingly, to distinguish between what is epistemically possible and what we are ignorant about.

Not distinguishing between what is epistemically possible and what we are ignorant about also leads to unreliable decision making. It allows focusing decision making on possibilities that are not real potentialities and thus that are irrelevant to attaining our goals. Just as we can argue that the appropriateness of our subjective probabilities can be assessed by how they map on to objective features of the world, e.g., frequencies of events, my

argument here is that the appropriateness of assignments of epistemic possibilities can be assessed by how they map on to objective features of the world, specifically events that have the potential to be actual. At a minimum, we should want the epistemic possibilities we are considering at least partially to overlap with relevant real potentialities.

Consider the possibility that a six-year-old neighbour has hacked the offline Minecraft world that your similarly aged child created and plays in the world while your family sleeps and the relevant computer is off. Not only that, but the neighbour does some amazing things with the world, including upgrading its block graphics (though, unfortunately, the world reverts to normal whenever family members enter it!). Such a possibility is not, as far as a six-year-old can tell, excluded by what they know about hacking or the abilities of other six-year-olds. Yet the child's relevant knowledge is sufficiently meagre that, in this case, the possibilities represent no actual, or even potential, states of affairs. So, the possibilities ought to be represented as ones about which there is ignorance and distinguished from something more substantial. After all, reasoning based on such possibilities would be problematic in practice. It would be intractable (since too many possibilities would need to be considered) and unreliable (since it would focus on possibilities none of which represent even potentially actual states of affairs).

Here is an example involving the thesis that anthropogenic CO<sub>2</sub> emissions are delaying the onset of the next glaciation. Since the early 1990s, good evidence has accumulated that this is a possibility that should be taken seriously (Berger and Loutre 2002). At least prior to the accumulation of this evidence, however, it would have been unwise to consider possibilities involving specific quantitative hypotheses about the delay, e.g., that on such and such assumptions about anthropogenic emissions, there will be a delay of roughly  $n \cdot 10000$  years (with  $n < 10$ ), seriously. This is so because quantitative estimates of effects on future glaciation then relied on extremely, and overly, simple statistical models and similarly

overly simple dynamical models. Here too, it seems plausible that knowledge was too meagre to speculate, never mind usefully speculate, about possibilities in practical matters and our epistemic situation would have been best characterised as close to ignorance. Moreover, here too, there are too many options about which to reason.

Still in the spirit of standard interpretations of epistemic possibilities are interpretations according to which  $p$  is epistemically possible if and only if  $p$  is not recognised to be excluded by what is known and some relevant, further practicable investigation will not show that  $p$  is false (e.g., DeRose 1991). Such interpretations are problematic in similar ways to standard interpretations. We need to distinguish between ignorance about a possibility and our having some relevant substantive knowledge about it, even when the possibility will withstand some further investigation. Moreover, we would be at risk of not being able to reason, never mind reason reliably, if we considered possibilities that would withstand practicable investigation even when we were ignorant about them. Consider once again predictions that anthropogenic emissions would delay future glaciation by various lengths of time. Such predictions were, in the 1970s and 1980s, not going to be shown to be false by then feasible research, yet, as we have seen, ought plausibly to have been classified as predictions about which we were close to ignorant rather than as ones that ought to have been taken seriously.

Nonstandard interpretations of epistemic possibilities from the semantics literature tell us that  $p$  is epistemically possible if and only if our total evidence provides some support for  $p$  (e.g., Dougherty and Rysiew 2009; Przyjemski 2017). This family of interpretations faces a challenge that is also faced by Betz's (2015) interpretation, from the philosophy of climate science literature. Betz tells us that  $p$  is epistemically possible if and only if  $p$  is consistent with everything we know. These interpretations would (Katzav 2023) stall possibilistic advice on future climate and thus are too demanding. A climate prediction is, even in the

most extensive prediction-relevant studies, only based on part of the relevant evidence and can only be tested for consistency with part of what is known. For there typically are (Betz 2015; Katzav et al. 2021) a variety of factors which are recognised to be, or recognised possibly to be, substantially relevant to the quantities we want to predict but the relevance of which cannot be quantified. This is due to limited data, theory, and computing power.

Still, the interpretation that epistemic possibilities require some supporting evidence if they are to be established might allow distinguishing between possibilities about which we are ignorant and possibilities which should be taken seriously. Unfortunately, the literature on such evidential interpretations seems to allow any evidence, no matter how limited, to establish epistemic possibilities. Dougherty and Rysiew, for example, suggest that not being known to contradict available evidence suffices to do so (2009, 127). As a result, these interpretations are not sufficiently distinguished from interpretations according to which epistemic possibilities are just possibilities that are not recognised to be incompatible with our knowledge.

An appropriate interpretation of epistemic possibilities for the climate context will, accordingly, need to constrain which possibilities are epistemic in a more substantive way than by requiring that they be supported by some (unspecified amount of) evidence. My interpretation of epistemic possibilities is the only available interpretation that fulfills this requirement while not being impracticable. On my interpretation,

*p* is possible if and only if it is not known that *p* is false and *p* is consistent with knowledge that at least approximately represents the basic way things are in the domain *p* is about (2023, p. 7).

Here, the basic way things are in a domain is constituted by “the circumstances in the domain and its laws and/or mechanisms” (2023, p. 7). Moreover, to show that a possibility is

consistent with knowledge that approximately represents the basic way things are in the domain the possibility is about, we need to show that the state of affairs described by the possibility can be brought about by something like the circumstances in the domain in conjunction with something like its laws and/or mechanisms (2023, p. 7). Thus, on my view, we can only establish epistemic possibilities by appealing to substantial bodies of knowledge.

In making my view more concrete, I have focused on Earth system models of intermediate complexity, models that do represent the four-dimensional structure of the climate system but in a cruder way than global climate models. I have argued that it was only when Earth system models of intermediate complexity started to capture delays in the onset of glaciation due to anthropogenic greenhouse gas emissions that climate scientists took model output to establish that such a delay is epistemically possible. Predictions of such delays by simpler models did not suffice to establish the possibility. I, accordingly, proposed that the Earth system models of intermediate complexity that simulated the delay can represent the basic way the climate system is at an appropriate level of approximation for establishing some epistemic possibilities (2023, p. 18).

There is a further component to my possibilistic approach, namely that if a possibility is held to be epistemically possible, we ought also to be in a position to suppose that it is a real possibility, where such possibilities are taken to correspond to real potentialities of the actual climate system (Katzav et al. 2012, p. 270; Katzav 2023, pp. 5-6). I thus accept the requirement that epistemic possibilities not be too permissive and propose that my interpretation meets this requirement.<sup>7</sup>

---

<sup>7</sup> Since it is the notion of epistemic possibility that is used to represent our uncertainty, how to understand my talk of real potentialities is not directly relevant to the present discussion. That said, the literature includes positions that allow fleshing this component of my position out. One option, for example, is to take potentialities to be dispositions, construed as irreducible universals, that cease to exist when actualised (McKittrick 2018). Another option is to identify real potentialities with all the

## 6. Possibilities and upper probabilities

The literature in logic and computing provides probabilistic interpretations of epistemic possibilities. Most commonly, in such interpretations, possibilities are associated with upper probabilities. A set of probability measures  $\mathbf{P}(\pi)$  can be specified as a function of a possibility measure  $\Pi(A)$  as follows:  $\forall A \subseteq S, P(A) \leq \Pi(A)$ . The possibility measure  $\Pi(A)$  then coincides with upper probability  $P^*(A)$ , where  $P^*(A) = \sup\{P(A), P \in \mathbf{P}(\pi)\}$ . The necessity measure  $N(A)$ , in turn, coincides with the lower probability  $P_*(A)$  where  $P_*(A) = \inf\{P(A), P \in \mathbf{P}(\pi)\}$  (Dubois and Prade 2015, p. 46).

Thus, Le Cozannet et al. (2017) aim to transform probabilistic estimates of twenty-first century sea levels into possibilistic ones and use the identification of possibilities with upper probabilities to achieve this. Recall that AR5 assesses that, under RCP 8.5, it is at least likely that global mean sea-level rise will be between 0.52 and 0.98m by 2100. Le Cozannet et al. (2017) assign possibility 1 to this range and, based on the range's supposed lower likelihood bound of 66%, a degree of possibility of 0.34 and lower to the remaining values of sea-level rise that they did not deem impossible. A minimal sea-level rise value of 0.38m with a degree of possibility 0.34 is assigned, on the assumption that anything lower than a linear continuation from present levels of sea-level rise is inconceivable under RCP 8.5 (2017, p. 3). They draw on available literature to provide three estimates of the "maximal value of sea-level rise" by 2100: 1.5m, 2m and 5m (2017, p. 4). They note that disagreement about what is epistemically possible increases with estimated levels and correspondingly reduce the degree of possibility assigned to each estimate. The three estimates, along with the estimated lower bound and full possibility range, thus yield three possibility distribution functions which disagree on the upper bound of sea-level rise. A weighted mean possibility

---

possibilities at a time that are not excluded by the concrete facts that exist at or before that time (Teller 2025).

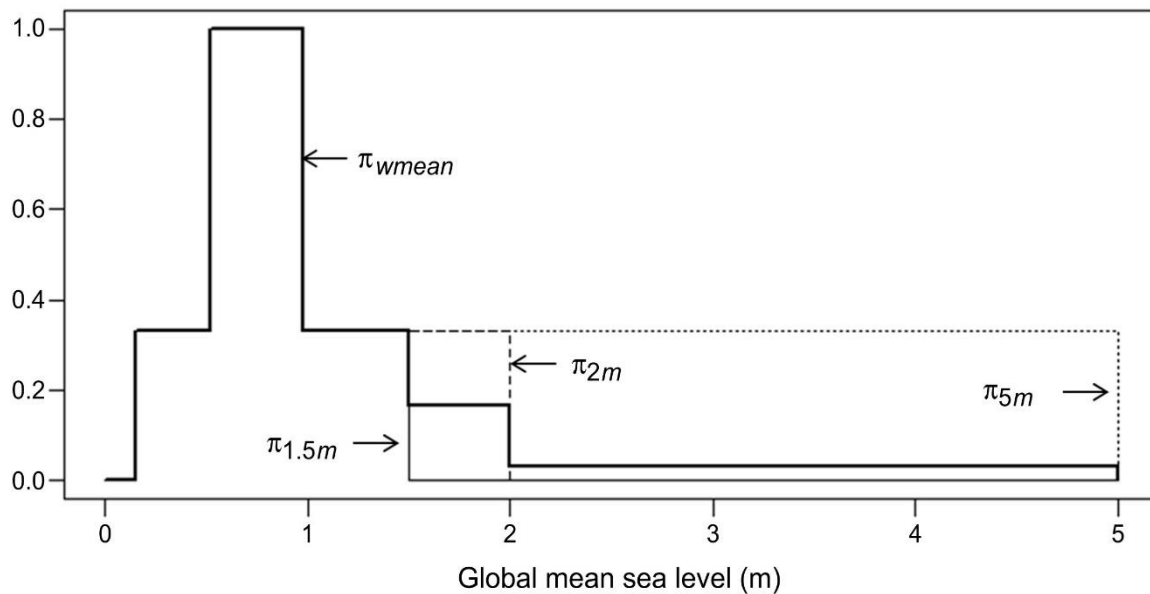
distribution function is generated from these three distributions (Fig. 1). This is done

according to the rule  $\pi_{wmean} = \frac{\sum_{i=1...N} w_i \pi^i(x)}{\sum_{i=1...n} w_i}$ , where the weights  $w_{1.5m}=0.5$ ,  $w_{2m}=0.4$  and

$w_{5m}=0.1$  are assigned to the distributions with upper estimates of 1.5m, 2m and 5m

respectively (2017, p. 5). The weighting used in this synthesis reflects the extent of

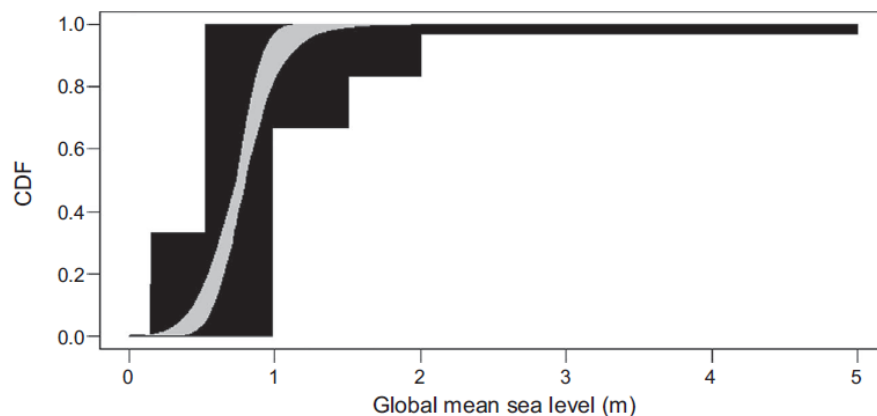
disagreement about physical upper limits.



**Fig. 1** Possibilistic interpretation of sea-level rise projections by 2100 for the climate change scenario RCP 8.5. (a) Possibility distributions corresponding to the three assumptions regarding the maximum sea-level rise by 2100 (1.5 m, 2m and 5m; fine black lines respectively solid, dashed and dotted) and aggregated possibility distribution (dark black line). Adapted from Le Cozannet et al., Figure 3 (2017, 6).

Following Dubois and Prade (2015), Le Cozannet et al. take  $\prod(-\infty, x)$  to be the upper bound for permitted cumulative probability distributions until any specified value of  $x$  and  $N(-\infty, x)$  to be the lower bound for permitted cumulative probability distributions (2017, p. 4). This means that, since  $\prod(-\infty, x)$  is in their study simply the measure of  $\pi_{wmean}(x)$  over  $[-\infty, x]$ ,  $\pi_{wmean}(x)$  is taken to coincide with the upper probability distribution for sea-level rise to 2100.

Le Cozannet et al. claim that their possibilistic representation of uncertainty about future sea-level rise “makes no assumption regarding the actual shape of probability density functions representing future sea-level rise” (2017, p. 5). But  $\pi_{wmean}(x)$  has been constructed so that, for  $x \leq 0.98\text{m}$ ,  $N(-\infty, x)$  is at least  $1 - 0.34$  (See the lower bound in Figure 2). So, permitted probability distributions will have a cumulative probability distribution of at least 66% for the interval  $[-\infty, 0.98]$ . Moreover, this probabilistic information is the only source the interpretation provides for understanding the interval while, as we have seen, part of what drives possibilistic approaches are worries about how to locate the bulk of the probability mass of assessments of uncertainty. Le Cozannet et al.’s study thus inherits the worries of probabilistic approaches by linking upper probabilities and possibilities.



**Fig. 2** Area of possible values taken by cumulative probability density functions bounded by (1; grey area) three probabilistic projections compliant with the AR5; (2; dark and grey areas) the p-box obtained in this study. Adapted from Le Cozannet et al. (2017).

This issue with Le Cozannet et al.’s reconstruction is a general one for probabilistic interpretations of epistemic possibilities. If we use even a rough probabilistic assumption about how probabilities are distributed to interpret our possibility distributions, we will be incorporating problematic aspects of probabilistic representations of our uncertainty into the

possibilistic representations. Accordingly, if we are to properly exploit possibilistic representations of our uncertainty, we ought to avoid interpreting them probabilistically. Of course, a possibilistic distribution will always bound a range of distributions that can be interpreted probabilistically. But there is no obligation on us to interpret bounded distributions in this way.<sup>8</sup>

I want to bring out three more problems with Le Cozannet et al.'s possibilistic study. The first of these arises because of the appeal to upper probabilities. More specifically, the problem arises because, when we construct and interpret a possibility distribution by bounding a probability distribution from above, we are committed to a probabilistic interpretation of partial possibilities, and such an interpretation is inconsistent with not taking partial possibilities seriously in practice. I will start by considering this issue in the case of Le Cozannet et al.'s specific choices about partial possibilities and then explain why the issue persists when other choices are made.

Le Cozannet et al. assign up to 34% of the probability mass of predictions of future sea-levels to predictions of more than 0.98m sea-level rise by 2100, while assigning no full possibilities to such levels. But to believe that some set of predictions has a probability of up to 34% is to have a belief that is consistent with believing that they have a probability of 34%; if we accept that predictions have a probability of up to 34%, we accept that we may well come to give the predictions a probability of 34%. So, however we interpret the

---

<sup>8</sup> It might be thought that the problem with Le Cozannet et al.'s study is its use of a probability distribution to construct a possibility distribution and not the interpretation of the possibility distribution. After all, the suggestion might be, reinterpreting their possibility distribution using my notion of epistemic possibility would still leave us with a distribution that is correlated with the problematic IPCC probability distribution. However, how we interpret epistemic possibilities is not independent of our use of probability distributions. If, say, we identify full possibilities with upper probabilities, we are obliged to constrain our possibility distribution in the way Le Cozannet et al. do and thus cannot avoid the challenges to probabilistic representations of uncertainty. If I am correct about epistemic possibilities, then we ought to reject IPCC distributions as misguided and construct possibilistic ones directly, with the help of my notion of epistemic possibilities.

assignment of up to 34% of a probability mass to some range of predictions, we must interpret it in a way that is consistent with believing the predictions have a 34% probability. The problem is that believing that a set of predictions has a 34% probability is inconsistent with supposing that the set is not fully possible. If there is a 34% chance that one or more events will occur, it is a full possibility that they will. Recall that the border between full and partial possibilities is the border between possibilities that are to be taken seriously in practice and those about which it is at least unclear whether they should be taken seriously in practice. Yet events with a 34% probability of occurring cannot be ignored in practical decision making.

To try to remove this tension, it seems Le Cozannet et al. should have taken a lower range than up to 34% of the probability mass to be the full possibilities range. This would allow distinguishing between full and partial possibilities by assigning a probability range with a relatively low upper bound to the latter. There would, however, still be the threat of inconsistency. Consider the proposal that up to 3.4% of the probability mass, and thus all levels of sea-level rise above 2m, be classified as partially possible. Since the climate context is a context in which events with a lower than 3.4/100 frequency are hardly unusual, the belief that a kind of event occurs with this frequency must be consistent with it being partially possible. At the same time, it is plausible that taking a kind of event to be partially possible is inconsistent with supposing it is actual. Events that are actual should sometimes be taken seriously in practice, if only to decide that they can be lived with. So, from a possibilist perspective, such events should be taken to be full possibilities.

Even if we leave aside the requirement that partial possibilities not be relevant to practical decision making, the recognition that Le Cozannet et al. should have associated such possibilities with ranges of low probabilities emphasizes the problem of not being able to distinguish between low probability events which do occur and events the possibility of

which is unclear. It matters whether there is, say, a 1% chance of 2m sea-level rise or whether this is yet to be established to be epistemically possible. Finally, the appeal to upper probabilities provides no principled basis for deciding where the range of partial possibilities is to start; in deciding, we can appeal only to the vague intuition that partial possibilities must have relatively low probabilities if they are to be represented probabilistically.

Thus, the appeal to upper probabilities to interpret epistemic possibilities in uncertainty assessments means that our possibilistic assessments will inherit the challenges associated with probabilistic assessments concerning locating the bulk of the probability mass and representing partial ignorance. Worse, the appeal will result in an unprincipled delimitation of the range of partial possibilities. The first of these three worries is, recall, avoided by non-probabilistic interpretations of epistemic possibilities because possibilities are not additive. The second and third worries indicate that we need an interpretation of epistemic possibilities which, like non-probabilistic interpretations of epistemic possibilities in general, allow reserving partial possibilities for representations of partial ignorance and, like my interpretation, are also explicit about the standard of evidence needed to establish full possibilities while allowing that some possibilities do not meet this standard.

## **7. Conclusion**

The case for extending the use of possibilistic assessments about the climate system to assessing the predictions of state-of-the-art climate models stands as it is. The need to interpret these predictions possibilistically, as well as the existing practice of interpreting the predictions of models in the context of discovery possibilistically, makes the further development of possibilistic approaches a worthwhile goal. The present paper works towards this goal by considering how to interpret epistemic possibilities in the climate modelling context. We have seen that interpreting epistemic possibilities as possibilities that are not recognised to be, or that are known not to be, excluded by what is known, as well as equating

epistemic possibilities with those for which we have some evidential support, will not do. We need an interpretation which requires that substantial evidence support epistemic possibilities before they are taken seriously. My interpretation fulfills this requirement by telling us that establishing full possibilities involves showing that a possibility is consistent with knowledge approximating the basic way things are.

Interpretations of epistemic possibilities that correlate them with upper probabilities, further, inherit the worries about probabilistic approaches to climate prediction, including the difficulties concerning locating the bulks of probability masses and representing partial ignorance. The appeal to upper probabilities also faces the worry that it does not provide us with a principled distinction between partial and full possibilities. My interpretation of epistemic possibilities resolves these worries.

My discussion does leave open further questions for possibilistic representations of uncertainty, of which I mention two. To begin with, I have defended the thesis that interpreting epistemic possibilities as possibilities that are consistent with knowledge approximating the basic way things are allows representations of uncertainty that avoid the challenges to probabilistic ones. So too, I have argued that this interpretation is not so permissive as to allow just any climate model prediction to be epistemically possible nor so demanding as to make establishing epistemic possibilities unfeasible. This is not the same as arguing that possibilistic interpretations are reliable. I do not, to be more explicit, argue that the range of full possibilities of an appropriately developed possibility distribution of future climate will appropriately overlap with a range of real potentialities of the climate system and thus be a reliable representation of our uncertainty. While I cannot offer such an argument here, my view is that it is feasible to do so because this would only involve showing that the unreliability of a possibilistic approach is not a real possibility.

Finally, there is the question of how formally to represent uncertainty in cases where evidence really does support some full possibilities more than others. From an informal, possibilistic perspective, recall, we sometimes differentiate between full possibilities in terms of degrees of remoteness. Moreover, I propose that such differentiation is sometimes appropriate in the climate modelling context. Following this proposal, we can use informal assessments of remoteness to supplement formal possibilistic assessment or as part of informal ones. It would, however, be useful to have a formal treatment of degrees of remoteness of possibilities to use with available formal possibility theory.

**Conflicts of interest:** I have no conflicts of interest to declare.

## References

- Baumberg, C., Knutti, R, and Hirsch Hadorn, G. (2017) “Building confidence in climate model projections: an analysis of inferences from fit,” *WIREs Climate Change*, 8: e454. <https://doi.org/10.1002/wcc.454>
- Berger, A., and Loutre, M.-F. (2002) “An exceptionally long interglacial ahead?” *Science*, Vol. 297(5585): 1287-1288.
- Betz, G. (2010) “What's the worst case? The methodology of possibilistic prediction.” *Analyse & Kritik*, 32(1): 87–106.
- Betz, G. (2015) “Are climate models credible worlds? Prospects and limitations of possibilistic climate prediction,” *European Journal of Philosophy of Science*, 5, 191–215.
- Church, J. A., P. U. Clark, A. Cazenave, J. M. Gregory, S. Jevrejeva, A. Levermann, M. A. Merrifield, G. A. Milne, R. S. Nerem, P. D. Nunn, A. J. Payne, W. T. Pfeffer, D. Stammer and A. S. Unnikrishnan (2013) “Sea Level Change.” In *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* [Stocker, T. F., D. Qin, G. -K. Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P. M. Midgley (eds.)], pp. 1137-1216. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.
- Collins, M., R. Knutti, J. Arblaster, J. -L. Dufresne, T. Fichet, P. Friedlingstein, X. Gao, W. J. Gutowski, T. Johns, G. Krinner, M. Shongwe, C. Tebaldi, A. J. Weaver and M. Wehner, (2013) “Long-term Climate Change: Projections, Commitments and Irreversibility.” In *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the*

*Fifth Assessment Report of the Intergovernmental Panel on Climate Change* [Stocker, T.F., D. Qin, G. -K. Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P. M. Midgley (eds.)], pp. 1029-1136. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Deser, C., Lehner, F., Rodgers, K. B. *et al.* “Insights from Earth system model initial-condition large ensembles and future prospects,” *Nature Climate Change*, 10, 277–286 (2020). <https://doi.org/10.1038/s41558-020-0731-2>

DeRose, K. (1991) “Epistemic possibilities”, *The Philosophical Review*, 100(4): 581-605.

Dethier, C. (2022) “When is an ensemble like a sample? “Model-based” inferences in climate modeling,” *Synthese*, 200, 52. <https://doi.org/10.1007/s11229-022-03477-5>

Dethier C. (2023) “Against “Possibilist” Interpretations of Climate Models,” *Philosophy of Science*, 90(5):1417-1426. <https://doi:10.1017/psa.2023.6>

Dougherty, T., and Rysiew, P. (2009) “Fallibilism, Epistemic Possibility, and Concessive Knowledge Attributions,” *Philosophy and Phenomenological Research*, 78(1): 123-132.

Dubois, D. (2006) “Possibility theory and statistical reasoning,” *Computational Statistics & Data Analysis*, 51(1): 47-69. <https://doi.org/10.1016/j.csda.2006.04.015>

Dubois, D., and Prade, H. (2015) “Possibility Theory and Its Applications: Where Do We Stand?” In Kacprzyk, J., Pedrycz, W. (eds) *Springer Handbook of Computational Intelligence*. Springer Handbooks. Springer: Berlin, Heidelberg. [https://doi.org/10.1007/978-3-662-43505-2\\_3](https://doi.org/10.1007/978-3-662-43505-2_3)

Eyring, V., N. P. Gillett, K. M. Achuta Rao, R. Barimalala, M. Barreiro Parrillo, N. Bellouin, C. Cassou, P. J. Durack, Y. Kosaka, S. McGregor, S. Min, O. Morgenstern, and Y. Sun, (2021) “Human Influence on the Climate System.” In *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* [Masson-Delmotte, V., P. Zhai, A. Pirani, S. L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M. I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J. B. R. Matthews, T. K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, pp. 423–552. <https://doi:10.1017/9781009157896.005>

Forster, P., T. Storelvmo, K. Armour, W. Collins, J. -L. Dufresne, D. Frame, D. J. Lunt, T. Mauritsen, M. D. Palmer, M. Watanabe, M. Wild, and H. Zhang (2021) “The Earth’s Energy Budget, Climate Feedbacks, and Climate Sensitivity.” In *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* [Masson-Delmotte, V., P. Zhai, A. Pirani, S. L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M. I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J. B. R. Matthews, T. K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, pp. 923–1054. <https://doi:10.1017/9781009157896.009>.

Fox-Kemper, B., H. T. Hewitt, C. Xiao, G. Aðalgeirsdóttir, S. S. Drijfhout, T. L. Edwards, N. R. Golledge, M. Hemer, R. E. Kopp, G. Krinner, A. Mix, D. Notz, S. Nowicki, I. S. Nurhati, L. Ruiz, J. -B. Sallée, A. B. A. Slangen, and Y. Yu (2021) “Ocean, Cryosphere and Sea Level

Change.” In *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* [Masson-Delmotte, V., P. Zhai, A. Pirani, S. L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M. I. Gomis, M. Huang, K. Leitzell, E., Lonnoy, J. B. R. Matthews, T. K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, pp. 1211–1362.  
<https://doi.org/10.1017/9781009157896.011>.

Frame, D. J., Faull, N. E., Joshi, M. M., Allen, M. R. (2007) “Probabilistic climate forecasts and inductive problems,” *Philosophical Transactions of the Royal Society A*, 365: 1971-1992.  
<https://doi.org/10.1098/rsta.2007.2069>

Frigg, R., Thompson, E., Werndl, C. (2015) “Philosophy of Climate Science Part II: Modelling Climate Change,” *Philosophy Compass*, 10(12): 965-977.  
<https://doi.org/10.1111/phc3.12297>

Hausfather, Z., K. Marvel, G. A. Schmidt, J. W. Nielsen-Gammon, and M. Zelinka (2022) “Climate simulations: Recognize the 'hot model' problem,” *Nature*, 605: 26-29.  
<https://doi.org/10.1038/d41586-022-01192-2>.

IPCC (2013) “Summary for Policymakers.” In *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* [Stocker, T. F., D. Qin, G. -K. Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P. M. Midgley (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Katzav, J., H. A. Dijkstra, and A. T. J. (Jos) de Laat (2012) “Assessing climate model projections: State of the art and philosophical reflections,” *Studies in History and Philosophy of Modern Physics*, 43(4): pp. 258-276.

Katzav, J. (2014) “The epistemology of climate models and some of its implications for climate science and the philosophy of science.” *Studies in the History and Philosophy of Modern Physics*, 46(B): 228–238.

Katzav, J., Thompson, E. L., Risbey, J., Stainforth, D. A., Bradley, S., and Frisch, M. (2021) “On the appropriate and inappropriate uses of probability distributions in climate projections and some alternatives,” *Climatic Change*, 169: 15. <https://doi.org/10.1007/s10584-021-03267-x>.

Katzav, J. (2023) “Epistemic possibilities in climate science: lessons from some recent research in the context of discovery,” *European Journal for Philosophy of Science*, 13: 57.  
<https://doi.org/10.1007/s13194-023-00560-7>.

Kirtman, B., S. B. Power, J. A. Adedoyin, G. J. Boer, R. Bojariu, I. Camilloni, F. J. Doblas-Reyes, A. M. Fiore, M. Kimoto, G. A. Meehl, M. Prather, A. Sarr, C. Schär, R. Sutton, G. J. van Oldenborgh, G. Vecchi and H. J. Wang (2013) “Near-term Climate Change: Projections and Predictability.” In *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* [Stocker, T. F., D. Qin, G. -K. Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P. M. Midgley (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, pp. 953-1028.

- Kuma, P., Bender, F. A.-M., and Jönsson, A. R. (2023) “Climate model code genealogy and its relation to climate feedbacks and sensitivity,” *Journal of Advances in Modeling Earth Systems*, 15, e2022MS003588. <https://doi.org/10.1029/2022MS003588>
- Le Cozannet, G., Manceau, J.-C., and Rohmer, J. (2017) “Bounding probabilistic sea-level projections within the framework of the possibility theory,” *Environmental Research Letters* 12:014012.
- Lee, J.-Y., J. Marotzke, G. Bala, L. Cao, S. Corti, J. P. Dunne, F. Engelbrecht, E. Fischer, J. C. Fyfe, C. Jones, A. Maycock, J. Mutemi, O. Ndiaye, S. Panickal, and T. Zhou (2021) *Future Global Climate: Scenario-Based Projections and Near Term Information*. In *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* [Masson-Delmotte, V., P. Zhai, A. Pirani, S. L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M. I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J. B. R. Matthews, T. K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou (eds.)]. Cambridge University Press, Cambridge and New York, pp. 553–672. <https://doi.org/10.1017/9781009157896.006>.
- McDonnell, A., Bauer, A. M., and Proistosescu, C. (2024) “To what extent does discounting ‘hot’ climate models improve the predictive skill of climate model ensembles?” *Earth's Future*, 12, e2024EF004844. <https://doi.org/10.1029/2024EF004844>
- McKittrick, J. (2018) "Real Potential." Faculty Publications - Department of Philosophy. 52. <https://digitalcommons.unl.edu/philosfacpub/52>
- Meccia, V. L., Fuentes-Franco, R., Davini, P. *et al.* (2023) “Internal multi-centennial variability of the Atlantic Meridional Overturning Circulation simulated by EC-Earth3,” *Climate Dynamics*, 60: 3695–3712. <https://doi.org/10.1007/s00382-022-06534-4>
- Morlighem, M. *et al.* (2024) “The West Antarctic Ice Sheet may not be vulnerable to marine ice cliff instability during the 21st century,” *Science Advances* 10, ead07794. <https://www.science.org/doi/10.1126/sciadv.ado7794>
- Pattyn, F., and Morlighem, M. (2020) “The uncertain future of the Antarctic ice sheet”, *Science*, 367: 1331-1335.
- Przyjemski, K. (2017) “Strong epistemic possibility and evidentiality”, *Topoi*, 36: 183-195.
- Sherwood, S. C, Webb, M. J., Annan, J. D., Armour, K. C., Forster, P. M., Hargreaves, J. C. *et al.* (2020) “An assessment of Earth’s climate sensitivity using multiple lines of evidence,” *Reviews of Geophysics*, 58:e2019RG000678. <https://doi.org/10.1029/2019RG000678>
- Risbey, J. S., and O’Kane, T.J. (2011) “Sources of knowledge and ignorance in climate research,” *Climatic Change*, 108, 755. <https://doi.org/10.1007/s10584-011-0186-6>
- Roussos, J. (2025) “Modeling Climate Possibilities.” In Knuuttila, T., Grüne-Yanoff, T., Koskinen, R., & Sjölin Wirling, Y., eds., *Modeling the Possible: Perspectives from Philosophy of Science* (1st ed.), pp. 196-220. Routledge. <https://doi.org/10.4324/9781003342816-14>

- Rowlands, D., Frame, D., Ackerley, D. et al. (2012) “Broad range of 2050 warming from an observationally constrained large climate model ensemble,” *Nature Geosciences*, 5: 256–260. <https://doi.org/10.1038/ngeo1430>
- Sanderson, B. M., and R. Knutti (2012) “On the interpretation of constrained climate model ensembles,” *Geophysical Research Letters*, 39, L16708. <https://doi.org/10.1029/2012GL052665>.
- Sansom, P. G., Stephenson, D. B., & Bracegirdle, T. J. (2021) “On Constraining Projections of Future Climate Using Observations and Simulations From Multiple Climate Models,” *Journal of the American Statistical Association*, 116(534): 546–557. <https://doi.org/10.1080/01621459.2020.1851696>
- Stainforth, D. A., Aina, T., Christensen, C. et al. (2005) “Uncertainty in predictions of the climate response to rising levels of greenhouse gases,” *Nature*, 433: 403–406. <https://doi.org/10.1038/nature03301>
- Stainforth, D. A., Downing, T. E., Washington, R., Lopez, A. and New, M. (2007) “Issues in the interpretation of climate model ensembles to inform decisions,” *Philosophical Transactions of the Royal Society A: Mathematics and Physics*, 365: 2163–2177.
- Tebaldi C, and Knutti R. (2007) “The use of the multi-model ensemble in probabilistic climate projections,” *Philosophical Transactions of the Royal Society A*, 365:2053–2075.
- Tebaldi, C., Ranasinghe, R, Vousdoukas, M. et al. (2021) “Extreme sea levels at different global warming levels,” *Nature Climate Change*, 11: 746–751. <https://doi.org/10.1038/s41558-021-01127-1>
- Teller, P. (2025) “Actual Possibility.” In Knuuttila, T., Grüne-Yanoff, T., Koskinen, R., & Sjölin Wirling, Y., eds., *Modeling the Possible: Perspectives from Philosophy of Science* (1st ed.), pp. 66–81. Routledge. <https://doi.org/10.4324/9781003342816-5>
- Van Westen, R, M., Kliphuis, M., Dijkstra, H. A. (2024) “Physics-based early warning signal shows that AMOC is on tipping course,” *Science Advances*, 10, eadk1189. <https://www.science.org/doi/full/10.1126/sciadv.adk1189>
- Yamazaki, K., et al. (2013) “Obtaining diverse behaviors in a climate model without the use of flux adjustments,” *Journal of Geophysical Research Atmospheres*, 118: 2781–2793. <https://doi.org/10.1002/jgrd.50304>