Epistemic possibilities in climate science: lessons from some recent research in the context of discovery

By Joel Katzav (2023)

European Journal for Philosophy of Science 13(4)

https://doi.org/10.1007/s13194-023-00560-7

Abstract

A number of authors, including me, have argued that the output of our most complex climate models, that is, of global climate models and Earth system models, should be assessed possibilistically. Worries about the viability of doing so have also been expressed. I examine the assessment of the output of relatively simple climate models in the context of discovery and point out that this assessment is of epistemic possibilities. At the same time, I show that the concept of epistemic possibility used in the relevant studies does not fit available analyses of this concept. Moreover, I provide an alternative analysis that does fit the studies and broad climate modelling practices as well as meshes with my existing view that climate model assessment should typically be of real possibilities. On my analysis, to assert that a proposition is epistemically possible is to assert that it is not known to be false and is consistent with at least approximate knowledge of the basic way things are. I, finally, consider some of the implications of my discussion for available possibilistic views of climate model assessment and for worries about such views. I conclude that my view helps to address worries about such assessment and permits using the full range of climate models in it.

Keywords: uncertainty, epistemic possibility, scientific discovery, climate modelling, sea-level rise

1. Introduction

A number of climate scientists and philosophers of science have criticised probabilistic assessments of our uncertainty about the output of our most complex climate models, that is, of global climate models and Earth system models, and have suggested that it is preferable to present our uncertainty about such output in possibilistic terms (Stainforth 2008; Betz 2010; Katzav 2014; Le Cozannet et al. 2017; Katzav et al. 2021). In the philosophy of science, Gregor Betz has argued that assessing whether the outputs of climate models are

epistemic possibilities is preferable to assigning probabilities to them, though possibilistic assessments of the outputs of even the most sophisticated models are themselves, on his view, subject to a serious challenge (2010; 2015). I, also in the philosophy of science, agree that the outputs of our most complex climate models should typically be assessed as to whether they are epistemically possible, though I have not explicitly offered an analysis of epistemic possibility. I also propose that they be assessed as to whether they represent real possibilities and offer an analysis of real possibility. On my view, when we take outputs to be epistemically possible, we ought also to be able to suppose that they represent real possibilities (Katzav et al. 2021). Here, real possibilities are conceived of as a species of objective possibility, that is, as a species of modal property of the mind-independent world rather than, like epistemic possibility, of our knowledge.

There is, however, no substantive philosophical discussion of the possibilistic assessment of the output of relatively simple climate models. Nor has there been discussion of the applicability in this context of available analyses of epistemic possibility found in the semantics literature or in the literature on climate models. Such discussion would contribute to developing and evaluating analyses of epistemic possibility as well as possibilistic approaches to climate model assessment. Doing this is not merely of theoretical value. Climate research is often policy relevant and, as we will see, some such policy-relevant research is already assessed possibilistically.

I examine the possibilistic assessment of the output of relatively simple climate models in the context of discovery, that is, in a context in which hypotheses are being pursued but it is premature to decide on their acceptance. I argue that analyses of epistemic possibility which fit this context differ from those which have been developed in general

research into the semantics of natural languages as well as from Betz's analysis. I also offer an analysis which does fit this context and meshes well with my analysis of real possibility. Finally, I consider some of the implications of my discussion for existing views about possibilistic assessment of climate model output and for further consideration of the viability of such assessment.

I begin by presenting (section 2.1) standard, general analyses of assertions of epistemic possibility, that is, of assertions that typically have the form 'It is possible that *p*' (or of that of a cognate). Standard analyses share the view, call it 'the ignorance view', that such assertions are basically expressions of ignorance; each such assertion supposedly states that the truth of the proposition it is about is not (recognised to be) excluded by the propositions in some set of epistemically privileged propositions. On the most common variant of this view, *p* is epistemically possible if and only if *p*'s truth is not excluded by what is known. Nonstandard analyses of assertions of epistemic possibility share the view, call it 'the strong possibility view', that such assertions tell us that the total available evidence provides some/non-negligible support for the truth of the propositions they are about or, at least, that some evidence supports the propositions' truth, and their truth is not defeated by any available evidence. On one variant of the strong possibility view, for example, *p* is epistemically possible if and only if *p* has a non-negligible probability given the total evidence.

I go on (section 2.2) to present possibilistic views from the philosophy of climate science literature. I present Betz's view, according to which *p* is epistemically possible if and only if *p* is consistent with everything we know, as well as present his challenge to possibilistic assessment of climate models. I also present my view of real possibilities and

develop a corresponding notion of epistemic possibility. On my view, *p* is epistemically possible if and only if it is not known that *p* is false and *p* is consistent with at least approximate knowledge of the basic way things are in the domain *p* is about.

The various analyses of epistemic possibility are examined in light of two case studies of climate science research in the context of discovery. My first case study is of research into the possibility that marine ice-cliff instability in Antarctica will substantially impact sea-level rise over the coming centuries. We will see (section 3) that this research contains a pattern of argumentation that appears to be about whether certain future sea levels are epistemically possible but that, on the simple ignorance view, must be problematically interpreted as being about something else. As a result, the simple ignorance view fails to predict the pattern of argumentation. The same case study undermines the strong possibility view by suggesting (section 4) that epistemic possibility assertions are sometimes based solely on a portion of available, relevant evidence. Similarly, this case study undermines Betz's analysis by indicating that demonstrating consistency with less than all of our knowledge is often taken to suffice to establish epistemic possibilities. The study also suggests that non-negligible evidence for an hypothesis is not generally sufficient for asserting that it is epistemically possible, which we will see challenges attempts to save the strong possibility view and Betz's view by weakening them.

My second case study (section 4) focuses on projections of the end of the current interglacial. This study further supports the view that non-negligible evidence is not generally sufficient for asserting an epistemic possibility. We will see that finding non-negligible evidence for an hypothesis according to which the current interglacial will end within some timespan is not generally taken by climate scientists to be a sufficient basis for

asserting that the hypothesis is epistemically possible, even when the hypothesis is otherwise not implausible. My second case study also illustrates when evidence does warrant asserting epistemic possibilities, information that is used (section 5) to support my analysis of epistemic possibilities. After providing this support, I respond briefly to a number of potential worries about my analysis.

I conclude (section 6) with a discussion of the implications of my study for existing views about possibilistic assessment of climate model output. I argue that Betz's reservations about possibilistic assessment of climate model output are based on an unmotivated notion of epistemic possibility that does not fit established possibilistic practice. I also argue that my analysis of real possibility, and the complementary analysis of epistemic possibility, fit with the viability of possibilistic assessment of the output of climate models in general.

2. Analyses of epistemic possibility

2.1 Epistemic possibilities in the literature on the semantics of modals

My discussion will distinguish between two variants of the ignorance view. According to the first of these, which I will call 'the easy ignorance view', to assert that some proposition is epistemically possible is to assert that the proposition's truth is not (recognised to be) excluded by some class of epistemically privileged propositions, where the class might be that of known propositions or, alternatively, that of sufficiently justified propositions. Here are some (rough)[1] examples of analyses that fit the easy ignorance variant:

[1] I have for the most part bracketed nuances of analyses, such as relativisation to speaker, that are irrelevant to my concerns in this paper.

(a) *p* is possible if and only if it is not known that *p* is false.[2]

(b) *p* is possible if and only if it is not certain that *p* is false.[3]

(c) *p* is possible if and only if nothing known obviously entails not-*p.*[4]

According to the second variant of the ignorance view, which I will call 'the sustained ignorance view', an assertion of an epistemic possibility tells us that the truth of some proposition not only is not excluded by epistemically privileged propositions but will continue not to be excluded despite some relevant, further investigation. (Rough) examples of analyses that fit the sustained ignorance view are:

(d) *p* is possible if and only if *p* is not known to be false and further practicable investigation will not establish it to be false.[5]

(e) *p* is possible if and only if *p* is not known to be false by members of the relevant community and there is no relevant way in which they can establish that *p* is false (where it is the context of the assertion that *p* is possible which determines which community and ways are relevant).[6]

Opposed to the ignorance view, there is the is the strong possibility view. On this view, to assert that some proposition is epistemically possible is to assert that the total available evidence provides some/non-negligible support for the truth of the proposition or, at least, that there is some/non-negligible support for this truth and no evidence undermines it. Here are two examples of analyses that fit the strong possibility view:

[2] See, e.g., J. Hawthorn (2012, p. 493).
[3] See, e.g., G. E. Moore (1962, p. 184).
[4] See, e.g., J. Stanley (2005, pp. 127-128).
[5] See, e.g., I. Hacking (1967, p. 149).
[6] See, e.g., K. DeRose (1991, pp. 593-594).

(f) *p* is possible if and only if *p* has a non-negligible/nonzero probability on the total available evidence available.[7]

(g) *p* is possible if and only if *p* is supported by evidence that is undefeated, that is, not opposed or overridden by other evidence.[8]

Importantly for what follows, proponents of the strong possibility view require very little of evidence for it to establish epistemic possibilities. Any argument or evidence that has not been defeated will do. Przyjemski is explicit that any evidence, however weak, will do (2017, p. 190). Dougherty and Rysiew require that the evidence be non-negligible but assume that evidence for a claim is non-negligible if its negation is not recognised to be entailed by the evidence (2009, p. 127).

Much of the discussion of analyses of assertions of epistemic possibilities concerns which bodies of knowledge need to be considered in evaluating the truth of such assertions. This issue arises because assertions about what is epistemically possible are not generally explicit about which bodies of knowledge they are about while what is epistemically possible does vary with what is known. Discussions of analyses of assertions of epistemic possibilities also often focus on which analysis best explains the awkwardness involved in concessive knowledge attribution claims, that is, claims of the form 'I know that *p*, but *p* might be false'. My discussion brackets these issues and focuses on general challenges to ignorance and strong possibility views.

2.2 Epistemic Possibilities in the Philosophy of Climate Science

[7] See, e.g., T. Dougherty and P. Rysiew (2009, p. 127; 2011, p. 396).
[8] See, e.g., K. Przyjemski (2017, p. 190).

Betz proposes using epistemic possibilities rather than probabilities to capture our uncertainty regarding climate model-based claims about future climate. Moreover, he offers the only explicit analysis of epistemic possibilities in the philosophy of climate science literature (2010 and 2015). On his view,

(h) *p* is possible if and only if *p* is logically consistent with what is known.

Betz understands this analysis to have epistemically demanding implications. One needs to show that *p* is consistent with what is known to establish that *p* is epistemically possible. We can do this, Betz notes, by arguing for *p* on the basis of what we know to be true. But since climate models are idealised, and thus known to be false, they do not provide such arguments. So, the fact that their simulations lead to some result is not itself sufficient grounds for establishing that the result is consistent with all background knowledge. Betz, further, is sceptical about our ability to develop arguments that would, by taking into account the errors in the models, establish such consistency (Betz 2010 and 2015).

Similarly to Betz, I propose that it is typically preferable to capture our uncertainty about the output of complex climate models using epistemic possibilities rather than probabilities (Katzav et al. 2021, pp. 2-3 and 7).[9] I add, however, that epistemic possibilities should ideally appropriately map onto real possibilities, conceived of as a species of objective possibility which has the potential to be realised in the actual world. More specifically, if we are justified in judging that it is epistemically possible for some concrete state of affairs to obtain, then we ought to have justification for supposing that the state of affairs is a potentiality of some concrete part of the actual world, that is, an objective way that part might, or might have, come to be (Katzav et al. 2021, p. 7). So, for example, we

[9] The sense of 'uncertainty' used here is epistemic.

should aim to judge that it is epistemically possible that sea levels will reach a certain level when sea levels can be judged to have the potential to do so in current circumstances.

Sjölin Wirling and Grüne-Yanoff (2021) have a different reading of my position. They suggest that my view is just that climate models should be used to establish objective possibilities. Supposedly, I am not concerned with whether these objective possibilities might (in the objective sense) turn out to be actual. Thus, I need not be concerned with whether the possibilities established by models are also epistemically possible. Sjölin Wirling and Grüne-Yanoff, however, do not take into account my recent work, which, as noted above, is explicit that assessment of climate model output should be of whether it is epistemically possible. Further, in my earlier work, I emphasised that the objective possibilities that assessment should aim to establish are real possibilities. I took such possibilities to be ones that might (in the objective sense) turn out to be actual and that need to be taken into account in practical decisions (Katzav et al. 2012, p. 270). I thus was committed to taking climate model output to be epistemically possible. This is as it should be. A decision relevant possibilistic assessment of climate model output needs to inform us about our uncertainty about the output and so needs to involve assessing whether it is epistemically possible.

It will help to clarify my position further if I consider why we might suppose that assessing whether a claim is epistemically possible ought also to involve assessing whether it represents a real possibility. An analogy between probabilistic and possibilistic reasoning is helpful here. Plausibly (see, e.g., Pettigrew 2021), if our reliance on subjective probabilities about future states of affairs to guide our actions is in principle to be justifiable, these probabilities should appropriately correlate with the way things actually are, e.g., with

propensities or long-run frequencies. For example, perhaps taking the future occurrence of a type of state of affairs in specified circumstances to have a 50% subjective probability should correlate with the type having an objective propensity to occur 50% of the time in similar circumstances.[10] Analogously, if our reliance on a judgement that a type of state of affairs is epistemically possible in certain circumstances is to be justifiable, it is plausible that the type should be a candidate real possibility in those and similar circumstances. For the less this is so, that is, the more we take possibilities that (objectively) cannot be actual into account in our reasoning, the less our reasoning is useful.

But while I am committed to supposing that assessments of epistemic and real possibilities go hand in hand, I have only offered an analysis of real possibility. I, accordingly, will propose what is plausibly an appropriate, corresponding analysis of epistemic possibility. Let me explain my analysis of real possibility before introducing the corresponding notion of epistemic possibility. On my view (2014, p. 235), a state of affairs in a target domain is a real possibility if and only if it is not known that it does not obtain and its obtaining is compatible with the basic way things are in the domain (2014, p. 236). The basic way things are in a domain is taken to comprise the circumstances in the domain and its laws and/or mechanisms. A state of affairs is compatible with the basic way things are in a domain if something *like* the circumstances obtaining when the state of affairs is supposed

[10] It is contentious whether probabilities should be interpreted in objective terms in the context of climate predictions, but the concern here is how to justify the reliance on probabilities rather than how to interpret them. Justification seems to be required if we are to answer the question why we ought to rely on probabilistic climate predictions.

to come to obtain would, with the help of something *like* the domain's laws and/or mechanisms, bring the state of affairs about (2014, p. 236).[11]

Importantly, my inclusion of an epistemic component in my analysis of real possibility–the requirement that real possibilities not be known to be non-actual–does not mean that real possibilities are not objective. The requirement that possibilities be compatible with the basic way things are specifies a class of objective possibilities. Selecting those that are not known to be non-actual from the class merely restricts our attention to objective possibilities that are still decision relevant.

My analysis of epistemic possibilities results from substituting, in my analysis of real possibility, a consistency condition for the ontological condition of compatibility with the basic way things are. I propose that

> (NEWE) *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.[12]

It is plausible that, as per my earlier recommendation, if we justify the claim that some state of affairs is epistemically possible in the sense of NEWE, we are also justifying the claim that the state of affairs is a real possibility in my specified sense. Justifying the claim that *p* is epistemically possible in the sense of NEWE involves showing not only that *p* is not known to be false but that *p* is consistent with approximate knowledge of the basic way things are in the domain *p* is about, including not just of circumstances but also of laws and/or

[11] The characterisation of real possibility I use here permits real possibilities to be in the past, present or future. I thus recognise that we are interested in past, present and future real possibilities. In an earlier discussion (Katzav et al. 2012, p. 270) my focus was solely on future real possibilities. Further, note that I am not here offering a necessary condition for being a real possibility but only a sufficient one.

[12] Buckwalter et al. (2020) review and support the case for the view that knowledge can be approximate.

mechanisms. But, in the present science-focused context, we know that *p* is about a state of affairs which, if it is actual, is brought about by the laws and/or mechanisms in the relevant domain. So, establishing the consistency of *p* with approximate knowledge of these laws and/or mechanisms must involve showing that the state of affairs *p* is about can somehow be brought about by what the approximate knowledge describes, that is, by something like the actual laws and/or mechanisms in the relevant domain. The idea here is just that, since we assume that states of affairs are somehow brought about, demonstrating consistency of our claims about states of affairs in a domain with our assumptions about how the domain develops requires showing that this development can lead to the states of affairs.

3. The context of discovery and ignorance views

Let me turn to evaluating views of epistemic possibilities from the semantics literature, starting with ignorance views. Floating ice shelves sometimes have a role in holding back the flow of ice from grounded ice sheets in the Antarctic. As the climate warms, however, the floating ice shelves might be lost, thus exposing tall, unsupported grounded ice cliffs. The tall ice cliffs, in turn, might collapse, exposing further cliffs and, in doing so, lead to runaway ice-sheet retreat (Clerc et al. 2019). DeConto and Pollard (2016) use a three-dimensional ice sheet–ice shelf model in order to support the possible truth of RETREAT, the hypothesis that, partly through the marine ice-cliff instability (MICI) mechanism just described, the Antarctic will contribute more than one meter of sea-level rise by the year 2100. The selling point of their study was that their simulations were more realistic along several dimensions. The simulations took into account how atmospheric warming–and not just the ocean

warming taken into account in previous studies–might affect buttressing ice cliffs and better represented Pliocene sea-levels. The authors note, in drawing conclusions, that

> given uncertainties in model initial conditions, simplified hybrid ice dynamics, parameterized sub-ice melt, calving, structural ice-margin failure, and the ancient sea-level estimates used in our Large Ensemble analysis, the rates of ice loss simulated here should not be viewed as actual predictions, but rather as possible envelopes of behaviour…that include processes not previously considered at the continental scale (2016, p. 596).

The upper bound of the envelope of possibilities referred to in this quote is neatly stated in the paper's abstract: "Antarctica has the potential to contribute more than a metre of sea-level rise by 2100 and more than 15 metres by 2500, if emissions continue unabated" (2016 p. 591).

The notion of possibility being used by DeConto and Pollard in their talk about an envelope of possibilities is epistemic. As is standard in current research into the future of climate, they are explicitly concerned with the uncertainty associated with their claims about future climate and explore this uncertainty in a variety of ways, including using an ensemble of varying simulations to provide a range of results (2016, p. 594). The envelope of possibilities is the way in which the study encapsulates this exploration and can only represent uncertainty if it is understood epistemically. The concern with epistemic possibility, further, can be seen in the contrast DeConto and Pollard make between predicting climate and specifying an envelope of possibilities. Uncertainty arising from model limitations, they note, is too great to allow prediction. Instead, the less ambitious aim of specifying an envelope of possibilities is adopted. If specifying an envelope of possibilities

is to address the difficulties associated with representing uncertainty about predictions, the envelope must capture uncertainty and thus be epistemic. Indeed, the epistemic use of 'envelope of possibilities' is standard and explicitly acknowledged in climate science as a way of representing uncertainty (see Stainforth et al. (2007) for an influential discussion of this notion within the context of climate science).[13]

However, according to the easy ignorance view, it seems that DeConto and Pollard were entitled to assert that RETREAT is epistemically possible without providing arguments for their assertion. For the upper bound of estimates of end of the century sea-level rise that did not take MICI into account was, at the time of the publication of their paper, 0.45m (Pattyn and Morlighem 2020, p. 1334). Moreover, MICI was a relatively neglected hypothetical source of sea level rise. Nobody had ruled out the possibility that it would contribute more than 0.5m of sea-level rise by the end of the century (DeConto and Pollard 2016; Edwards et al. 2019). It was thus accepted that RETREAT was not excluded by what is known (or believed) by the scientific community. This, of course, means that, on the easy ignorance view, it was already generally agreed that RETREAT might be true.

The literature on the easy ignorance view does recognise that assertions of epistemic possibilities are often accompanied by arguments. Moreover, this literature includes two strategies for explaining the presence of these arguments. According to one strategy, such arguments can be supposed merely to reflect the pragmatics of assertions of epistemic possibilities. Supposedly, assertions of epistemic possibilities are often used, with the help of the arguments associated with them, to highlight the possibilities they are about as ones

[13] Perhaps DeConto and Pollard are, like me, concerned with epistemic possibilities that appropriately map onto real possibilities. I think this is correct and is supported by some of the language in their paper, such as the use of 'potentiality' in the already provided quote about the upper bound of Antarctica's contribution to sea-levels. But this is tangential to my discussion here. For present purposes, what matters is that DeConto and Pollard are concerned with epistemic possibilities.

that should not be ignored (see, e.g., von Fintel and Gillies 2008, p. 83). This supposition is implausible in the context of the research into RETREAT.

It was clear at the time of DeConto and Pollard's study not just that RETREAT was not known to be false but that knowing whether it and some of the hypotheses used in arguing for it were false would not be feasible in the study or in the years immediately following its publication in 2016. So, if assertions of epistemic possibility are mere expressions of ignorance, the semantics of such assertions will play no role in explaining assertions about the possibility or otherwise of RETREAT in the context of recent investigations about it. Similarly, the semantics of assertions of epistemic possibilities will play no role in explaining the assertions of epistemic possibilities used in arguing for or against the possible truth of RETREAT. Even questions about whether RETREAT and related, similarly uncertain hypotheses are possible would have to be taken to be about whether they should be taken seriously rather than, as their semantics suggests, about what is epistemically possible.

RETREAT's truth, and the truth of some of the hypotheses used in its evaluation, were not to be expected to be decided in the years immediately following the study because, as noted above, prior research had largely focused on causes of ice-sheet retreat other than MICI. Moreover, DeConto and Pollard's study was just a first step into quantitively examining the potential effects of MICI (DeConto and Pollard 2016). As Edwards et al. point out in a 2019 follow-up study,

> very little is known about whether initial cliff collapse would lead to a positive feedback effect (that is, MICI), how such a feedback effect would vary in different locations, the consequent rate of ice wastage and how long it would last (Edwards et al. 2019, p. 61).

Not surprisingly, even though subsequent studies *do* undermine confidence in RETREAT's possibility, a 2020 review of the relevant research tells us no more than that predictions based on MICI are "very hard to verify" (Colledge 2020 p. 6). Another review from the same year, is similarly sceptical about MICI and its effects, but also does not rule these out, stating that

> the hypothesis of marine ice cliff instability….has emerged….postulating that ice cliffs become unstable and collapse if higher than ~90 m above sea level, facilitating the rapid retreat of ice sheets. This process may have been important in Antarctica during past warms periods (Pattyn and Morlighem 2020, p. 1332).

After noting the results of Clerc et al., the review states that "it remains unclear how rapidly an ice cliff would retreat as a function of its height" (Pattyn and Morlighem 2020, p. 1333). Similarly, Kopp et al. note that the "potential for MICI" is shrouded under deeper uncertainty than the potential for the instability of the Antarctic sea-ice sheet from marine ice-sheet instability (2019, pp. 1240-1241).

Clerc et al.'s response (2019) to DeConto and Pollard illustrates how, if we accept the easy ignorance view, the semantics of assertions of epistemic possibilities will be supposed to be irrelevant to understanding the chains of reasoning in research into RETREAT. Clerc et al. note that studies earlier than theirs were based on the assumption, derived from studies of exposed, static ice cliffs, that a cliff taller than ~90 meters in height would collapse. They add, however, that it is not clear whether this assumption applies to ice cliffs that are being formed as a result of the loss of buttressing ice and accordingly model the response of cliffs to loss of buttressing ice as a function of the speed of the loss. Further, observations of the collapse of the Larsen B ice shelf allow them to estimate that "removal of a buttressing ice

shelf prior to the potential initiation of runaway cliff failure may take days to weeks" (2019, p. 12115). At these timescales, Clerc et al.'s model indicates that collapse requires cliffs taller than ~540 meters. Thus, we should expect that, if Clerc et al.'s results were taken into account in modelling the impacts of ice-cliff collapse, we would get a prediction of less loss of ice-sheet mass than predicted by DeConto and Pollard. Now, Clerc et al.'s hypothesis about the possible timescales for collapse problematizes DeConto and Pollard's assumption about collapse heights and thus undermines the support offered for the possibility of RETREAT. But, on the easy ignorance view, Clerc at al.'s hypothesis is irrelevant to DeConto and Pollard's argument. DeConto and Pollard would not deny that the relevant timescales are an open issue. So, on the easy ignorance view, the semantics of assertions of epistemic possibilities fail to explain why Clerc et al.'s assertion about collapse timescales is part of a critique of the possible truth of RETREAT.

Consider the second way in which proponents of the easy ignorance view might try to account for the arguments about the possibility of MICI-related effects. Discourse about epistemic possibilities includes the grading of such possibilities as to how remote they are. One can understand this grading in terms of degrees of plausibility, that is, degrees to which possibilities are supported by evidence or argument. The more plausible a possibility is, the less remote it supposedly is.[14] Now, versions of the ignorance view can be extended so as to include a semantics for assertions of graded possibilities (Kratzer 1991; Rudin 2016). So, if the arguments about MICI and its effects are not about whether they might occur but about

[14] Proponents of the ignorance view can distinguish between 'plausible' and 'epistemically possible' because they think epistemic possibilities require no supporting arguments. In section 4, I will be pointing out that 'plausible' in the sense just specified is in practice distinguished from 'epistemically possible' because establishing that something is epistemically possible is supposed to require more than just having some supporting argument for it.

how plausible they are, the arguments may yet be explained by an extended version of the ignorance view.

However, the debate about MICI is explicitly about its bare possibility (at least in the modern era) as well as the possibility that it might, in the not too distant future, occur with worrying results. Thus, for example, Gilford et al., which includes DeConto and Pollard as authors and continues to explore MICI, concludes that "this study can neither confirm nor exclude MICI as a primary driver of AIS [Antarctic ice-sheet] mass loss" (2020, p. 16). Similarly, Pattyn and Morlighem write, about whether MICI occurs at all, "the existence of MICI remains controversial today" (2020, p. 1333). Nicholls et al. (2021, p. 9) write that "there is still considerable debate as to whether the MICI mechanism could occur on the timescale and over the spatial extent required" to yield sea-level changes such as those predicted by DeConto and Pollard. Given that it really matters how certain we are about the occurrence of MICI in the not too distant future, including whether we are confident enough to say that its occurrence is plausible rather than possible but not yet established to be plausible, it would be unreasonable to insist that the authors under consideration are concerned with plausibility but are presenting their concern as being about what is possible.

The sustained ignorance view does suggest that assertions of epistemic possibilities will be associated with arguments that aim to establish that the ignorance about the possibilities will withstand some degree of investigation. However, such arguments are in fact not the kinds of arguments which are offered in favour of the possibility that RETREAT is true. As we have seen, at the time of DeConto and Pollard's study, there was no dispute about whether RETREAT would feasibly come to be known to be false in the immediately following years. Thus, DeConto and Pollard should not be thought of as arguing for the thesis that there was at the time no practical way of knowing or establishing that RETREAT is

false. This leaves open the possibility that DeConto and Pollard were arguing that future, then impracticable, research would not reveal that RETREAT is false. But DeConto and Pollard offer no support for such a strong claim. Their study includes no arguments for thinking that their conclusions will be robust to further developments in the modelling of MICI or of other sources of Antarctic ice-sheet instability. On the contrary, they recognise that future modelling should account for a variety of further processes relevant to the occurrence and impacts of MICI (DeConto and Pollard 2016, p. 596).

Those criticising DeConto and Pollard are not in a fundamentally different position when it comes to developing arguments about what future research into RETREAT might find. Critics accept that RETREAT cannot be refuted given current research options. Moreover, their lines of criticism rest on assertions of epistemic possibilities which are precarious in much the same way as was DeConto and Pollard's assertion in 2016 that RETREAT might be true. When, for example, Clerc et al. assert that initiation of runaway cliff failure may take days to weeks, their evidence comprises, as we have seen, only the collapse of the Larsen B ice shelf, and they do not attempt to argue that their claim will be sustained in future research. Thus, the sustained ignorance view does not leave us with any more of a substantial account of discourse about epistemic possibilities in the context of research into RETREAT and of related claims than does the easy ignorance view.

In summary, if we accept the ignorance or sustained ignorance views, we should conclude that the semantics of assertions of epistemic possibilities is irrelevant to research into the possible truth of RETREAT and into that of related, uncertain hypotheses. These views fail to predict the arguments that aim to support RETREAT as well as the attempts to undermine these arguments' premises. The extent to which this failure is a worry depends on whether the semantics of assertions of epistemic possibilities must be supposed to be

similarly irrelevant to much other research in the context of discovery. Pervasive irrelevance might lead us to suspect that, at least in the context of discovery, the easy ignorance and sustained ignorance views fail to capture the semantics of assertions of epistemic possibilities. After all, it seems plausible that major chunks of research that are straightforwardly presented as being about such possibilities play some role in fixing the semantics in question.

4. The context of discovery, the strong possibility view and Betz's view

Model-based research such as the research into whether RETREAT might be true is also a problem for the strong possibility view. According to the strong possibility view, recall, if a claim is epistemically possible, the claim must be supported given all available evidence, or, at least, there must be no available evidence that undermines it. But while modelling studies are used to draw conclusions about long-term sea level rise, such studies are only able to consider some of the available evidence that is relevant to their conclusions and thus are not in a position to make claims about what, all things considered, available evidence implies about their conclusions. Thus, for example, the limited availability of computational power means that climate scientists are not able fully to estimate the implications of what they know about the coupling of ice, ocean and atmosphere for estimates of the contribution of ice loss in Antarctica to long-term sea level rise. Yet there is reason to think that these implications are substantial (Pattyn and Morlighem 2020, p. 1335). Those

investigating Antarctica's contribution to long-term sea level rise are thus not yet able to make claims about what all the evidence relevant to their investigations implies.[15]

Betz's analysis is similarly problematic. Betz, recall, requires that we show the consistency of a possibility with all available knowledge for the possibility to be taken to be an epistemic one. Betz thus implies that establishing an epistemic possibility involves considering its consistency with all of the implications of available knowledge. As we have seen, this is not feasible in practice.

One might suggest, in response, that the views just considered can easily be modified to accommodate the observation that model-based studies need not take all relevant evidence into account before drawing conclusions about what might be the case. We can propose a modified strong possibility view according to which

(MSP*) p* is possible if and only if there is some/non-negligible support for *p* and this support is not known to be defeated by available evidence.

Similarly, one could modify Betz's proposal so that it only ties epistemic possibilities to establishing consistency with some available knowledge. One can propose that

(MBETZ) *p* is possible if and only if *p* is consistent with some knowledge.

However, in modelling studies, the existence of non-negligible evidence for an hypothesis is not generally a sufficient condition for taking the hypothesis to be epistemically possible even when the hypothesis is otherwise plausible. In other terms, sometimes a level of substantial relevant evidence that is above the level of non-negligible

[15] Here, and in what follows, I follow climate science research in treating knowledge as evidence. More in particular, knowledge of ice, ocean and atmosphere coupling that climate models cannot currently take into account is recognised to be evidence here. Although we cannot calculate what this impact might be and thus do not have quantitative knowledge of it, we do have qualitative knowledge of it. In other terms, we do recognise in an imprecise way how it can enhance/mitigate effects.

evidence for an hypothesis is required before taking the hypothesis to be epistemically possible. This is a problem for MSP, since it tells us that any argument or evidence that is undefeated will do. Similarly, MBETZ is challenged, since it tells us that showing consistency with some knowledge suffices to establish epistemic possibilities. Any argument for an hypothesis, based on however limited knowledge, can be used to show consistency of the knowledge with the hypothesis.

The reception of DeConto and Pollard's model-based evidence for RETREAT already indicates that non-negligible evidence for an hypothesis is not sufficient for taking it to be epistemically possible. While DeConto and Pollard aimed to justify the claim that REATREAT is possible, others responded by denying that this had been achieved and, as we have seen, the weight of opinion seems to favour the more sceptical position. As a result, whether RETREAT is possible remained an open issue in the scientific community. Yet the reception of DeConto and Pollard's study, specifically the fact that respondents did not dismiss its evidential weight but rather viewed it as worthy of further study, makes it clear that the evidence for RETREAT was not taken to be negligible by the parties to the debate about it.

DeConto and Pollard do appear to be less demanding than their critics about what is needed in order to establish that RETREAT is epistemically possible. Nevertheless, the disagreement with these critics concerns whether a suitable amount of evidence above non-negligible evidence has been provided not about whether more than non-negligible evidence is needed. DeConto and Pollard emphasise that they are taking claims such as RETREAT to be epistemically possible because their study represents a substantial advance in the provision of evidence for MICI's possible effects. They already had relevant evidence from an earlier study of theirs with the same model but merely took that evidence to be *suggestive* of epistemic possibilities in the vicinity of RETREAT. In the earlier study, the

model was not used to simulate future climate but only paleoclimate. Nevertheless, its simulations were, because they had to appeal to MICI to successfully simulate 20 or more meter past sea-level rise, taken to "suggest that East Antarctic subglacial basins may be more vulnerable than in most previous models" (Pollard et al. 2015, p. 117). The extra vulnerability of the East Antarctic, in turn, enhanced worries about future sea levels. Additional suggestive evidence was provided before the 2016 study by observed increased seaward flow of ice due to loss of buttressing in a number of glaciers (2016, pp. 591-592).

Looking at attempts to project the end of the current interglacial also supports the view that non-negligible evidence does not generally suffice to establish epistemic possibilities and, as we will see in the next section, helps to indicate when evidence sometimes is good enough to justify epistemic possibilities. Early attempts to project the end of the current interglacial did not explicitly take the impact of CO_2 forcing into account and thus implicitly assumed that this impact would continue as it was in the past.[16] Some of these attempts were based on statistical analyses of palaeontological evidence and some on simple dynamical models (Berger and Loutre 2002). Simple dynamical models are physical models that comprise a small number of differential equations for interdependent variables of a target system. The variables represent key quantities of the system and the equations represent physically plausible interdependencies between these quantities. Simple dynamical models, further, do not represent the overall structure of their target systems but only aspects of it.

Imbrie and Imbrie (1980) develop a simple (two-equation) dynamical model of glaciation cycles. Their model models variation in the volume of land ice as a function of

[16] A forcing is a net change in the radiative flux into the atmosphere, one that is due to some factor which, like anthropogenic CO_2, is external to the climate system.

orbital forcing. The model does not explicitly consider sources of variation in the climate system, including CO_2, with frequencies higher than every 19,000 years. Moreover, the model's four free parameters are calibrated to the historical record. A key conclusion drawn from the model's simulations is a prediction. As their abstract puts it, "this model predicts that the long-term cooling trend which began some 6000 years ago will continue for the next 23,000 years" (1980, p. 943). The conclusion in the body of the text is similar. They take "model output to indicate that orbital forcing will act over the next 23,000 years to continue the general cooling trend that began 6000 years ago" (1980, p. 951). Thus, Imbrie and Imbrie draw conclusions only about what their model implies and resist drawing any conclusion from their model about the actual climate system, including about its possible behaviour in the absence of sources of variation with frequencies higher than every 19,000 years. Nevertheless, the model does, according to the authors, provide useful information regarding its thesis about long-term sources of variation. Imbrie and Imbrie write that

> the output of this model compares favorably with the geological record of the past 250,000 years. Thus, the model's prediction for the next 100,000 years is useful as a basis for forecasting how climate would evolve at orbital frequencies in the absence of anthropogenic disturbance (1980, p. 943).

The idea here is that model output is a good starting point for drawing conclusions about future variation in ice mass in the absence of anthropogenic influences but one that requires supplementation by further considerations before such conclusions are drawn. So, although the model by itself does not justify accepting claims about the future variation in question, it does provide non-negligible evidence in considerations about such variation.

Since the early 1990s, attempts have been made to project the end of the current interglacial while explicitly taking CO_2 forcing into account (Berger and Loutre 2002). One early attempt is described by Saltzman et al. (1993). They, like Imbrie and Imbrie, make use of a simple dynamical model that is calibrated to the historical record. The four variables of Saltzman et al.'s model represent the mass of the northern ice sheets (excluding that of Greenland), bedrock depression under the ice sheets, concentrations of carbon dioxide in the atmosphere and global mean ocean temperature (1993, p. 1051). The model's predictions of ice-sheet mass describe trajectories of the climate state under the assumption that recent CO_2 levels persist into the future, trajectories in which the major ice age oscillations of the past few millions of years cease.

Saltzman et al. do not, however, use their model to draw something as strong as a claim about how the climate system might be in the future. They take the value of their work to be that it quantitatively illustrates a long-standing hypothesis according to which a change in atmospheric CO_2 concentrations can result in a major change to the behaviour of the climate system. As they put it,

> within the context of this example of a self-consistent (albeit non-unique) dynamical model of paleoclimatic change involving CO_2 as both a free and forced variable, we provide a quantitative demonstration of a physically plausible scenario illustrating potential effects of CO_2 on the possibility for ice age oscillations. Such a possibility has been conjectured at least since the time of Tyndall (1861) (1993, p. 1054).

Thus, it seems that, according to Saltzman et al., it is the quantitative nature of their illustration that makes it a non-negligible contribution to considering whether CO_2 might bring about a major change in the behaviour of the climate system. Offering a quantitative

illustration of a scenario brings us a step closer to seeing that, as has been conjectured, the scenario might occur.

Saying more about the attitudes of climate scientists to simple dynamical models in general and to those used in the above studies will help to support my suggestion that the models in the studies were, and were appropriately, being taken to provide non-negligible evidence. Later studies recognise that Imbrie and Imbrie's and Saltzman et al.'s models did, despite their limitations, make real contributions to evaluating claims about future glacial cycles (e.g., Berger and Loutre 2002). These later studies also note that the models' results led to using more realistic models to establish (at least) that the coming ice age might have been delayed by CO_2 emissions. Moreover, there are other contexts in which the simple dynamical models under consideration are taken to have non-negligible evidential value, though not individually to establish epistemic possibilities. Thus, a variety of simple dynamical models, including Imbrie and Imbrie's model, were collectively key in arguing that ice sheets have a role in explaining glacial cycles (Imbrie et al. 1993). Similarly, simple dynamical models, including Saltzman et al.'s model, were key in the revival of the hypothesis that atmospheric CO_2 is a driver of long-term climate change (Maasch et al. 2005, pp. 2146-2147).

In general, since simple dynamical models developed in paleoclimate studies represent merely plausible quantitative relations and largely accommodate, rather than provide out-of-sample predictions of, data, the models' predictive value is taken to be "rather limited" (Claussen et al. 2002, p. 580). Nevertheless, their theoretical and empirical basis means that they are taken to confer plausibility on their projections, without justifying acceptance of these projections. The models are also taken to make plausible hypotheses about possible causal processes (Claussen et al. 2002, p. 580; Maasch et al. 2005, p. 2147).

In summary, the strong possibility view is too demanding in requiring that all relevant evidence for an hypothesis be considered before we can take the hypothesis to be possibly true. Weakening the strong possibility view so that, as per MSP, it merely takes the existence of (not known to be defeated) support/non-negligible support for an hypothesis to be sufficient for taking that hypothesis to be possibly true results in a view of epistemic possibilities that is too weak. More than non-negligible support is at least sometimes needed here. Betz's view is also too demanding in requiring consistency of an hypothesis with all available knowledge before we can take the hypothesis to be epistemically possible. Weakening his view so that, as per MBETZ, it merely takes consistency of the hypothesis with some knowledge to be needed results in a position that is too weak. More than consistency with some knowledge is at least sometimes needed.

5. A novel analysis of epistemic possibility in the context of discovery

My discussion of research into RETREAT and associated hypotheses suggests that the easy ignorance and sustained ignorance views at least sometimes make the semantics of assertions of epistemic possibilities irrelevant to research programs which present themselves as being about what is epistemically possible. The greater the extent of such irrelevance in the context of discovery, the greater the challenge to the applicability of the views. Strong possibility views and Betz's view appear to be too demanding, given what is required in modelling studies. Modified versions of these views, that is, MSP and MBetz, appear to be insufficiently demanding in light of research into REREAT and research into the onset of the next glaciation.

This leaves the question of what view of epistemic possibility is suggested by my case studies. They suggest, first, that research into whether some hypothesis is possible occurs under the assumption that it is not known whether the hypothesis is false. Second, the studies suggest that establishing that an hypothesis might be true at least sometimes involves giving an argument for the hypothesis. Third, it seems that sufficient support for an epistemic possibility need not involve support on the basis of all the available evidence. Finally, it seems that sometimes more than merely non-negligible support is needed to establish an epistemic possibility.

NEWE fulfills all these requirements. It tells us that epistemic possibilities must not be known to be false. So too, it tells us that epistemic possibilities must be consistent with approximate knowledge of the basic way things are in the domains the possibilities are about, which means that establishing an epistemic possibility will require supporting arguments for its consistency with relevant, approximate knowledge. Since we must establish such consistency by showing how the possibility is supported by the body of knowledge, NEWE requires that we establish epistemic possibilities by arguing for them. Nor need the supporting case be based on all available evidence. NEWE permits making a case using knowledge that approximately captures the basic way things are. There is no requirement that all relevant knowledge, approximate or otherwise, be considered.

What remains to be seen is whether NEWE predicts that more than non-negligible evidence for a possibility is required before it is taken to be an epistemic one. I will argue for a stronger claim than this. NEWE provides the correct predictions about when climate model-based evidence supporting an hypothesis is strong enough to establish epistemic possibilities and when it is not strong enough to do so.

Global climate models, with their detailed representations of the structure of Earth's oceans, land surfaces, atmosphere and sea ice, as well as of how these components interact, are examples of models the simulations of which can approximately represent the basic way the climate system is (Katzav 2014, p. 236). These models' simulations can be taken to represent the basic way the climate system is because, while the models generating them are substantially idealised and incomplete, the models are nevertheless sufficiently complete to count as representing the laws and mechanisms of the climate system rather than just of an aspect of it. The justification for supposing that the models' simulations are approximately correct is that the models are built around principles from classical mechanics, including, e.g., continuity equations, the Navier-Stokes equations of motion and the first law of thermodynamics, in conjunction with a substantial amount of empirical knowledge specific to the climate system. The models are also subject to extensive empirical testing once constructed.

On the present construal of 'basic way things are', the justification needed to establish epistemic possibilities, as NEWE construes them, about some domain of investigation can be provided by theory and/or models that are known at least to approximately represent the causal laws and/or mechanisms of the domain. Global climate models, since they represent the basic way the climate system is, can be supposed to provide such justification for some projections relating to global climate, for example, by showing that something like twentieth century climate brings about the state of affairs the projection represents (Katzav 2014, pp. 236-237). This fits modelling practice. Global climate models are used to provide probabilistic predictions about future global climate, e.g., to tell us what the likely range of future global mean temperature is (Katzav et al. 2021). Whether we think such predictions are justified or not, that they are being provided means that the

models upon which they are based are being used to establish epistemic possibilities. If some future range of possibilities is being taken to be likely, it is being taken to be epistemically possible.

At the other end of the climate model spectrum, there are simple dynamical models. The relations posited by the simple dynamical models used in studying future glaciation, recall, are no more than plausible. Further, these models at best approximate an aspect of the structure of the climate system (Saltzman 1985 and 1988). The models are, accordingly, not the kinds of models that can be used to justify epistemic possibilities, according to NEWE. The studies which are taken to justify epistemic possibilities relating to glaciation rely on simulations with more complex models, specifically with Earth system models of intermediate complexity (EMICs). EMICs represent the structure of the climate system in a substantially cruder way than do global climate models but, like global climate models, are built around a substantial body of knowledge of the climate system and do approximately represent the structures of key components, including the atmosphere, the ocean and the ice sheets, along with how these interact (Claussen et al. 2002). These models' simulations can accordingly approximately represent the basic way the climate system is.[17] Moreover, these models are used to justify accepting conclusions about how the climate system might be and even about how it is likely to be. In some of the cases where EMICs are thus used (e.g., Ganopolski 2016), they are used to represent how components of the climate system might interact to bring about glaciation cycles, including a delay in the next cycle. In other cases (e.g., Archer and Ganopolski 2005), the parameters of simple models are calibrated so

[17] I here limit my discussion to full, as opposed to hemispheric, EMICS.

that they mimic the behaviour of EMICs. The simple models are then used to project future patterns of glaciation.

There are, of course, potential worries about my view. To begin with, it requires that epistemic possibilities be supported by argumentation and thus might be thought implausibly to suggest that it is not coherent to assert that a possibility is extremely unlikely. In response, cases such as those of research into RETREAT and into the end of the current interglacial illustrate that one can have an argument for an hypothesis and yet be unable to assess how likely or unlikely the hypothesis is. The arguments in these cases are offered against a background of such limited knowledge that it is arguably premature to judge hypotheses using qualifiers such as 'extremely likely' or 'extremely unlikely' (Katzav et al 2021). Another concern is whether the views described here fit with a uniform treatment of all modals, including physical and deontic ones. In response, note that our concern here is with a technical field in which vocabulary for handling uncertainty is explicitly developed and discussed with an eye towards informing policy. There is no a priori reason to expect that the semantics of epistemic modals in this context will more than resemble their semantics in the everyday contexts in which modals tend to be investigated and thus will fit in with an overall, fully uniform treatment of modals. Moreover, NEWE resembles available views. It can be seen as combining an ignorance requirement with a requirement, in the spirit of strong possibility views, that epistemic possibilities be supported by evidence. So even if we assume, for the sake of the argument, that ignorance and strong possibility views are acceptable in other contexts, we would not thereby have a reason to reject the proposals being put forward here about climate modelling research.

A final concern is whether my case studies uniquely select my thesis that NEWE is the notion of epistemic possibility that is involved in possibilistic claims about future

climate. In particular, even if it is plausible to assume that NEWE is the one used in the studies, other notions of epistemic possibility, e.g., the one specified by the ignorance view, might sometimes be used in the broader context of research about future climate. I suspect that the standard of justification for epistemic possibilities in climate prediction developed in response to its relatively specific characteristics, including the complexity of the system, the limitations and strengths of available data and theory, and policy relevance. Because of this, I suspect that one notion of epistemic possibility will tend to be used in justifying climate predictions. That said, I cannot develop my case further at this point and thus claim merely that my thesis about NEWE is preferable to the rivals that I have examined here. I do not claim to have offered a positive, general argument for its truth and acknowledge that, for example, hybrid views that allow that multiple notions of epistemic possibility are used in climate prediction are worthy of pursuit.

6. Concluding discussion

What lessons does my discussion of epistemic possibility in the context of discovery have for possibilistic epistemology? We have seen that, at least in climate science, possibilistic assessment is part of the assessment of climate models in the context of discovery. An important part of this assessment is determining whether certain possibilities are indeed epistemically possible. Relatively simple climate models are used in such investigations.

Betz's pessimism about whether even state-of-the art complex climate models can successfully establish epistemic possibilities is based on his view that epistemic possibilities are possibilities that are logically consistent with what is known. But this view is more demanding than available analysis of epistemic possibilities from the semantics literature.

On standard analyses, we can know that *p* is epistemically possible without extensive examination of *p*'s consistency with background knowledge. *p* need only not be known to be false, or something along these lines. NEWE merely adds to standard analyses that *p* needs to be consistent with not necessarily fully accurate knowledge comprising knowledge of the basic way things are in some domain; NEWE does not require consistency with everything we know about the basic way things are in a domain never mind with knowledge in general. Moreover, Betz provides no argument for such a requirement, and we have seen that it is not made in practice. It is thus unclear why climate scientists, or philosophers of science, ought to be worried that climate models cannot be used to establish epistemic possibilities when such possibilities are understood in the way Betz understands them.

My relative optimism about the use of climate models to establish epistemic possibilities, however, fits well with the investigation of relatively simple climate models carried out here. NEWE not only meshes well with my view that sophisticated climate models can be used to support epistemic possibilities but, as we have seen, fits with the practice of assessing simple and intermediate complexity climate models. Thus, my position that the output of the most complex climate models ought, in typical cases, to be assessed possibilistically can be supplemented with the view that such assessment is sometimes appropriate for the output of simpler climate models. At the moment, at least, there are no objections to such assessment of simpler models and its use in practice means that its feasibility has been demonstrated.

References

Archer, D., and A. Ganopolski (2005) “A movable trigger: Fossil fuel CO2 and the onset of the next glaciation”, Geochemistry, Geophysics, Geosystems, 6, Q05003, doi:10.1029/2004GC000891.

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*, 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.

Buckwalter, W. and Turri, J. (2020) “Knowledge, adequacy, and approximate truth,” *Consciousness and Cognition*, 83: 102950.

Claussen, M., Mysak, L., Weaver, A. et al. (2002) “Earth system models of intermediate complexity: closing the gap in the spectrum of climate system models,” *Climate Dynamics* 18: 579–586.

Colledge, N. R. (2020) “Long-term projections of sea-level rise from ice sheets,” *WIREs Climate Change*, 11: 1-21.

Clerc, F., Minchew, B. M., and Behn, M. D. (2019) “Marine Ice Cliff Instability Mitigated by Slow Removal of Ice Shelves”, *Geophysical Research Letters*, doi: 10.1029/2019GL084183.

DeConto, R. M. and Pollard, D. (2016) “Contribution of Antarctica to past and future sea-level rise”, *Nature*, 531: 591-597.

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

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

Dougherty, T. and Rysiew, P. (2011) “Clarity about concessive knowledge attributions: Reply to Dodd,” *Synthese*, 181(3): 395–403.

Edwards, T. L., Brandon, M. A.. Durand, G.. Edwards, N. R., Golledge, N. R.. Holden, P., Payne, A. J., Ritz, C. and Wernecke, A. (2019) “Revisiting Antarctic ice loss due to marine ice cliff instability. Nature”, *Nature*, 566: 58–64.

Ganopolski, A., Winkelmann, R. and Schellnhuber, H. (2016) “Critical insolation–CO2 relation for diagnosing past and future glacial inception,” *Nature*, 529: 200–203.

Gilford, D. M., Ashe, E. L., DeConto, R. M., Kopp, R. E., Pollard, D., & Rovere, A. (2020) “Could the Last Interglacial constrain projections of future Antarctic ice mass loss and sea-level rise?” *Journal of Geophysical Research: Earth Surface*, 125, e2019JF005418. https://doi.org/10.1029/2019JF005418

Hacking, I. (1967) “Possibility”, The Philosophical Review, 76(2): 143–168.

Hawthorn, J. (2012) “Knowledge and epistemic necessity,” Philosophical Studies, 158: 493-501.

Imbrie, J. and Imbrie, J. Z. (1980) “Modeling the climatic response to orbital variations”, Science, 207(4434): 943-953.

Imbrie, J., et al. (1993) “On the structure and origin of major glaciation cycles 2. The 100,000-year cycle,” Paleoceanography, 8(6): 699-735.

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

Kopp, R. E., Gilmore, E. A., Little, C. M., Lorenzo Trueba, J., Ramenzoni, V. C., Sweet, W. V. (2019). “Usable Science for Managing the Risks of Sea-Level Rise.” *Earth’s Future*, 7: 1235-1269.

Kratzer, A. (1991) “Modality.” In A. von Stechow and D. Wunderlich (eds), *Semantics: An International Handbook of Contemporary Research*. Berlin: de Gruyter.

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.

Maasch, K. A., Oglesby, R. J., and Fournier, A. (2005) “Barry Saltzman and the Theory of Climate”, *Journal of Climate*, 18(13): 2141-2150.

Moore, G. E. (1962) *Commonplace book: 1913-1953*, London and New York: Routledge.

Nicholls, R. J., Hanson, S. E., Lowe, J. A., Slangen, A. B. A., Wahl, T., Hinkel, J., Long, A. J. (2021) “Integrating new sea-level scenarios into coastal risk and adaptation assessments: An ongoing process,” *WIREs Climate Change*, 12:e706.

Pattyn, F. and Morlighem, M. (2020) “The uncertain future of the Antarctic ice sheet”, *Science*, 367: 1331-1335.

Pettigrew, R. (2021) “What is justified credence?” *Episteme*, 18(1): 16-30.

Pollard, D., DeConto, R. M. and Alley, R. B. (2015) “Potential Antarctic Ice Sheet retreat driven by hydrofracturing and ice cliff failure,” *Earth and Planetary Science Letters*, 412:112–121,

Przyjemski, K. (2017) “Strong epistemic possibility and evidentiality”, *Topoi*, 36: 183-195.

Rudin, D. (2016) “Deriving a Variable-Strength Might,” In *Proceedings of Sinn und Bedeutung*, 20: 587-603.

Saltzman, B. (1985) “Paleoclimatic modelling.” In A. D. Hecht (ed.) *Paleoclimate Analysis and Modelling*, pp. 341-396. Wiley: New York.

Saltzman, B. (1988) “Modelling the Slow Climate Attractor.” In Schlesinger, M. E. (ed.) *Physically-Based Modelling and Simulation of Climate and Climatic Change*. NATO ASI Series, vol. 243, 737-754. Springer: Dordrecht.

Saltzman, B., Maasch, K. A. and Verbitsky, M. Ya (1993) “Possible effects of anthropologically-increased CO2 on dynamics of climate: implications for ice age cycles”, *Geophysics Research Letters*, 20 (11): 1051-1054.

Stainforth, D. A., Downing, T. E., Washington, R., Lopez, A. & New, M. (2007) “Issues in the interpretation of climate model ensembles to inform decisions,” *Philosophical Transactions of the Royal Society A: Mathemtatics and Physics*, 365: 2163-2177.

Sjölin Wirling, Y. and Grüne-Yanoff, T. (2021) “Epistemic Possibility and Objective Possibility in Science,” *British Journal for the Philosophy of Science*, https://doi.org/10.1086/716925.

Stanley, J. (2005) “Fallibilism and concessive knowledge attributions”, *Analysis*, 65(2): 126-131.

von Fintel, K. and Gillies, S. (2008) “CIA leaks”, *The Philosophical Review*, 117(1): 77-98.