Epistemic and Objective Possibility in Science

Ylwa Sjölin Wirling

Till Grüne-Yanoff

Abstract

Scientists regularly make possibility claims. While philosophers of science are well aware of the distinction between epistemic and objective notions of possibility, we believe that they often fail to apply this distinction in their analyses of scientific practices that employ modal concepts. We argue that heeding this distinction will help further progress in current debates in the philosophy of science, as it shows that the debaters talk about different things, rather than disagree on the same issue. We first discuss how the two notions differ with respect to their epistemology and show that these differences are sometimes ignored in the philosophy of science. We then revisit four current philosophy of science debates about modelling, that are framed in modal terms, to showcase how the distinction significantly clarifies these debates and thereby and helps advance them.

1. Introduction

Scientists, across a wide range of disciplines, make and seek to justify possibility claims. Recently, philosophers of science focusing on scientific modelling have begun to pay more attention to such practices (e.g., Batterman & Rice [2014], Gelfert [2016], Grüne-Yanoff [2013], Massimi [2019]). However, a claim of the form 'It is possible that *p*' can be read in multiple ways. In particular, it can be read as marking either the *objective* or the *epistemic* possibility of the relevant proposition *p*. While most philosophers of science will be familiar with the conceptual difference between these two, we believe that they often fail to apply this distinction in their analyses of scientific practices that centrally involve modal concepts – as witnessed for example in (Massimi [2019]), where a criterion apparently derived from concepts of epistemic possibility is claimed to justify modal knowledge of 'objective' and 'causal' possibility.

Neglecting this distinction, we believe, has contributed to some longstanding disputes between philosophers of science. We revisit four such debates to showcase how better heeding the epistemic/objective distinction helps resolve them. These are: (i) whether *how-possibly explanations* of actual phenomena are just incomplete how-actually explanations, or have a *sui generis* epistemic function (Bokulich [2014], Forber [2010]); (ii) whether models used to refute necessity claims must represent actual targets or not (Grüne-Yanoff [2009], Fumagalli [2016]); (iii) whether constraint-based reasoning can provide non-causal explanations, or only assists heuristically in the search for an explanation (Craver [2007], Greene & Jones [2016]); and (iv) whether the many idealisations of the members of so-called climate model ensembles prevent an interpretation of these ensembles as representing the spread of possibility or not (Betz [2015], Katzav [2014]). We argue that differentiating different senses of 'possible' helps to clarify what is at issue in these disputes and improve our understanding of the positions on each side. As they stand, without making the distinction explicit, the different sides risks talking past each other. What will emerge is a picture where apparently conflicting positions on one and the same issue turn out to be focusing on different things. This should help advance these four longstanding arguments between philosophers of science. Furthermore, the distinction was not specifically developed for clarifying these four debates. Instead, that it matters for them makes us more confident that the epistemic/objective distinction might be relevant for a wide range of modal issues relating to scientific modelling.

The paper is structured as follows: In section 2, we rehash the distinction between objective and epistemic possibility and show how these notions differ with respect to both subject matter and epistemology. We also discuss a recent analysis of modal modelling that disregards this distinction. Section 3 discusses the four disputes and showcases how the objective/epistemic distinction helps clarify what is at issue in them. Section 4 concludes.

2. The Distinction

Claims about what is epistemically possible are claims about what is compatible with a given body of knowledge or justified beliefs. To say that some proposition *p* is epistemically possible, is roughly to say that we cannot, given what we know, rule out that *p* is true (Chalmers [2011], pp. 60-61, Edgington [2004], pp. 5-6, Hacking [1967], Huemer [2007], Vetter [2015], p. 216, Weatherson and Egan [2011], p.1).

The truth-value of a given claim of epistemic possibility can vary with a number of factors. First, it depends on *whose* corpus of knowledge, at what time, is the relevant one. Second, it depends on what the relation between *p* and the corpus should be, in order for *p* to not be 'ruled out', e.g., is logical compatibility enough, or is more required? Third, it depends on exactly what counts as being part of a given corpus. This turns on questions such as how much and how strong evidence one must have for *p* in order for *p* to be part of the corpus, and on whether only true beliefs can be part of the corpus. [1]

We assume that a notion of epistemic possibility of relevance to science will not depend on the beliefs of any individual scientist. Instead, the relevant corpus would be something like the shared body of acknowledged scientific evidence and scientifically established facts in a scientific community.[2] The second and third moving part of an epistemic possibility invoke more complex issues, and philosophers who work on epistemic modality disagree over how they are best fixed.[3] We will assume the standard, *negative* characterisation of the relevant relation between *p* and corpus, i.e. it is enough for epistemic possibility that *p* is *not* thus-and-so (e.g., 'known to be false' or 'ruled out by what is known').[4] Beyond that, we recognise that what standards gatekeep the corpus might differ between and even within disciplines, thus leading to differences in the truth-value of a given claim of epistemic possibility.

A natural way to think about objective (sometimes 'alethic') possibility is as expressing something about the world – in particular, about how the world could be. The notion of objective possibility makes best sense in light of the assumption that many things are only *contingently* the way they are: the world could have been different from how it is, and that there is more than one way the world can be in the future, even if there is just one way it *will* be. Objective possibility comes in different kinds, depending on the facts that restrict or determine it. Philosophers have paid most attention to *metaphysical* possibility, but outside of philosophy the relevant sense of possibility is

[1] We are not interested in notions of epistemic possibility strongly associated with what cannot be ruled out *a priori,* e.g., what Chalmers ([2011], p. 63) calls 'deep epistemic possibility'.

[2] This leaves open the possibility that in some scientific fields there are multiple, mutually exclusive corpora that generate distinct spaces of epistemic possibilities. This is desirable since we want to allow the relevant notion of evidence to go beyond the empirically observed, and thus include theories, some of which might not be universally accepted in a community.

[3] Moore ([1959]) defines epistemic possibility in relation to what is *known*, whereas Hacking ([1967]) and DeRose ([1991]) urge that the corpus should also include what we 'could easily come to know'. The issue of factivity is also contentious: is epistemic possibility about what is not excluded by *what we really know*, or by what we *take ourselves to know*?

[4] In contrast, Przyjemski ([2017]) argues that we need to distinguish strong and weak epistemic possibility, where *p* is epistemically possible in the weak sense just in case it is compatible with the evidence, and *p* is epistemically possible in the strong sense just in case we have some evidence *for p*. Strong epistemic possibility is an example of a *positive* characterisation where being epistemically possible requires something *more* from our body of knowledge than mere non-exclusion (see also Adler [2002], pp. 103-33, Malcolm [1963], p. 31). This is a step towards a more fine-grained picture of the conceptual space of epistemic modality. Although we are sympathetic to the development of a more nuanced account in this spirit, we set the issue aside here.

typically more restricted. Nomological (sometimes 'physical' or 'natural') possibility is one prominent kind as far as science is concerned: *p* is nomologically possible just in case *p* is objectively possible given the actual laws of physics. But in many (including scientific) contexts, the relevant sense of possibility is even more restricted, e.g., biological, practical, or 'easy' [5] possibility: What, more exactly, restricts at least the two latter notions clearly vary with context and human interest, but they are nonetheless objective: whether *p* is e.g., easily possible depends on whether *p* is a way the world could be, given the relevant facts. For some representative definitions of objective possibility in contrast to epistemic possibility, see e.g. (Deutsch ([1990], p. 752, Gendler and Hawthorne [2002], pp. 3-6, Lange [2009], ch. 2, Williamson [2016]). Exactly what makes an objective modal claim true is subject to extensive debate and we do not take a stance on this. The important thing, for current purposes, is that the truth of an objective possibility claim is independent of humans' epistemic situation.

Many possibility claims are true both when read as epistemic and as objective. But clearly, the two often come apart. They can come apart due to our knowledge of actuality, since many known actual truths are only contingently true. For instance, that Norway voted to join the EU in 1994 is not epistemically possible – we know that they did not. They can also come apart due to *ignorance:* it is epistemically possible that Goldbach's conjecture is true, and it is epistemically possible that it is false – at present we can't tell which. But if true (or false), it is necessarily true (or false), so whatever the right answer is, the other option is *not* objectively possible.

As should now be evident, objective and epistemic possibility claims *are claims about different things:* about the world, and about what we (don't) know, respectively. In light of that, it is unsurprising that they also have different epistemologies.

One is justified in accepting 'It is epistemically possible that *p*' just in case one is justified in believing that *p* is not ruled out by the corpus of evidence. While there is no need to suppose that this is a trivial matter, nor that we could not sometimes be mistaken, our epistemic access to these facts is relatively unmysterious. There is nothing strange about the idea that scientists have access to the corpus of scientific knowledge from which further research proceeds, and that they can figure out whether *p* is compatible with that corpus – even if doing so may be cumbersome. In contrast, the epistemology of objective possibility claims is a vexed issue. Ignoring the controversial idea that humans have something like direct access – by e.g., intuition (Bealer [2002]) – to how the world can possibly be, we have to infer the possibility of *p* from the evidence that *is* accessible to

[5] *p* is easily possible if *p* is in accordance with roughly how things (contingently) actually are (Strohminger and Yli-Vakkuri [2019], Kment [2014]). For example, at the time of writing, it is easily possible that there will be a hard Brexit, but not that Tony Blair replaces Angela Merkel as *Bundeskanzler* of Germany. Note that many easy possibilities can be counterfactual.

us (empirically and with the aid of ampliative inferences). Thus, in order to make a justified objective possibility claim, one needs to refer to evidence that plausibly indicates that so-and-so *is* a way the world could be. Exactly what sorts of facts one can draw on, and how the inferences need to go, is extensively discussed in *modal epistemology*. We stay neutral here[6], but note that what evidence is required partly depends on what sense of objective possibility one has in mind.

2.1. Three consequences

The picture just sketched should be familiar and relatively uncontroversial. In this subsection, we highlight three epistemologically relevant differences between epistemic and objective possibility implied by it.

First, in determining objective possibility, we normally disregard a lot of what we know to be the case – we engage only a *proper subset* of the knowledge corpus. For instance, if we are interested in whether *p* is nomologically possible, certain evidence strongly indicating that actually not-*p* will often be irrelevant, because many objective possibilities are counterfactual. In contrast, finding out whether *p* is epistemically possible involves checking how *p* relates to the *full* current corpus of evidence.

Second, ignorance on a given issue is no hinderance to justified claims of epistemic possibility, but obstructs justified claims of objective possibility. Just as one is not justified in believing that it will rain tomorrow, simply in virtue of lacking any evidence about tomorrow's weather, one is not justified in believing that life on Mars is objectively possible simply in virtue of having no evidence speaking against the possibility. Claims of epistemic possibility, in contrast, are claims about agents' current epistemic situation, so one *is* justified in believing that it is epistemically possible that it rains tomorrow or that there is life on Mars, with reference to one's ignorance.

Third and relatedly, we can make claims about the epistemic possibility (or not) of objective possibilities, e.g., 'it is epistemically possible that *p* is objectively possible'. It is important to note that justification for such a nested objective possibility claim does *not* automatically translate into justification for the 'first-order' objective possibility claim. Consider:

A. It is epistemically possible that *p,*
B. It is epistemically possible that *p* is objectively possible,
C. It is objectively possible that *p*,

[6] Notably we take 'evidence that plausibly indicates that so-and-so *is* a way the world could be' to be compatible both with *necessity-first* and *possibility-first* approaches to modal knowledge. See (Hale [2002], Fischer [2016], Roca-Royes [2017]) for discussion of this distinction.

Justification for A is also justification for B: if the objective possibility of *p* were ruled out by my evidence, its actuality would be ruled out too.[7] So if life on Mars is epistemically possible, it is also epistemically possible that it is objectively possible. However, justification for B does *not* automatically constitute justification for C: if I am sufficiently ignorant about relevant matters and recognise that my evidence does not speak to whether *p* is objectively possible, I am justified in claiming that it is epistemically possible that *p* is objectively possible. But given my ignorance, my evidence does not allow me to rule out that *p* is objectively *im*possible. So, both the objective possibility and the objective impossibility of *p* are epistemically possible. And it is precisely because of this that I am *not* justified in making any first-order claims about the objective possibility (or impossibility) of *p*: I do not have enough justified beliefs that bear on the issue.[8]

2.2. The distinction in current philosophy of science

Again, what has been said so far is neither new nor controversial. Most will agree that interpreting 'It is possible that *p*' as expressing either an epistemic or objective possibility matters for whether or not the claim is justified. Yet, philosophers of science analysing scientific modelling practices in terms of possibility have not been very clear whether they take the relevant possibilities to be epistemic or objective. Of course, even if the terms 'epistemic' and 'objective' (or 'alethic') are not used, the context will often let on what is at issue. But this is not always the case.[9] Moreover, even those who take care to be explicit about kinds of possibilities sometimes fail to see the relevance of the epistemic/objective distinction.

Consider for instance a recent paper by Michela Massimi ([2019]). Massimi is concerned with identifying two novel functions of exploratory modelling that are characterised by their epistemic function: they deliver knowledge of possibilities. Unlike many others, she takes care to specify, and distinguish between, two *kinds* of possibility that exploratory models may afford knowledge of, and she also endeavors to spell out in virtue of what models can afford such possibility knowledge. However, in doing so, she blurs the epistemic/objective distinction.

Massimi claims that hypothetical-exploratory models afford knowledge of 'objective possibilities' whereas fictional-exploratory models afford knowledge of 'causal possibilities'. Both of these possibility-notions appear to be objective, in the sense that we have been using it above. The main difference between Massimi's 'objective' and 'causal' possibility is that the latter are

---

[7] Conversely, justification for B does not automatically constitute justification for A: one's evidence may leave open whether *p* is objectively possibility while it conclusively indicates that actually, not-*p*.

[8] Of course, if I *have* evidence for C, I also have evidence for B, in the uninteresting sense that B is part of the corpus (hence epistemically necessary in the relevant sense, and hence possible – this is reminiscent of the uninteresting way in which evidence for '*p* is actual' is evidence for '*p* is possible').

[9] For a review of recent work on modal modelling, see (Gelfert [2019]). Most of the reviewed papers do not address the distinction between epistemic and objective possibilities at all.

possibilities pertaining to some actual phenomenon (e.g., what mechanism(s) could possibly cause electromagnetic induction) whereas the former concerns the possibility of the very existence of certain (hypothetical) things (e.g., could a particle with a set of properties P_i possibly exist).

Exploratory models can, according to Massimi, afford possibility knowledge in virtue of involving a particular mode of restricted imagination, which she calls 'physical conceivability'. She defines it as follows:

> *p* is physically conceivable for an epistemic subject S (or an epistemic community C) if S's (or C's) imagining that *p* not only complies with the state of knowledge and conceptual resources of S (or C) but it is also consistent with the laws of nature known by S (or C) ([2019], p. 872).

That is, physical conceivability requires (i) compliance with C's state of knowledge and conceptual resources, and (ii) consistency with the laws of nature known by C.

A number of things about this definition are puzzling. First, (i) ensures that if *p* is physically conceivable for C, *p* is epistemically possible for S. But physical conceivability is supposed to be a guide to objective possibility, so there appears to be a mismatch between the types of modality that characterise claim and justification respectively. It might be objected that Massimi's notion of physical conceivability as a guide to objective possibility might be tailor-made to the case studies she considers, where scientists are evidently not interested in any counterfactual objective possibilities, but only in those objective possibilities that are also epistemically possible (i.e., might be actual for all we know). In light of that, it makes sense to allow as conceivable only that which complies with the full corpus of knowledge, i.e., the first epistemological difference between epistemic and objective possibility highlighted in 2.1 is irrelevant here. Maybe so (although one may then certainly question the usefulness of the notion of physical conceivability in elucidating instances of modal modelling more generally, since as we shall see in the next section, scientists who use models to gauge possibilities are often also interested in counterfactual possibilities). However, it is not obvious that this helps. The second epistemological difference highlighted in 2.1 made clear that justified epistemic possibility claims are generally compatible with ignorance on relevant issues, whereas justified objective possibility claims are not. Presumably physical conceivability can only be evidence of objective possibility if our current state of knowledge contains evidence which bears on the prospective possibility of *p*. Physical conceivability, as defined by Massimi, does not ensure this. Of course, physical conceivability of *p* may justify the claim that it is epistemically possible that *p* is objectively possible, but as per the third epistemological difference highlighted in 2.1 this does not translate into guaranteed justification for the objective possibility claim itself.

Another puzzling thing about the definition is that (i) appears to imply (ii): C's knowledge of laws is presumably included in the 'state of knowledge', so why the double requirement? Our best guess is that (ii) is supposed to pull in the direction of objective possibility, i.e., that knowledge of the laws of nature, specifically, is evidence for claims of objective possibility. Of course, that seems quite right. But what will be 'consistent with the laws of nature known by C' depends on what C knows about the laws of nature. If C knows very little about natural laws, *p* may easily be physically conceivable despite being in conflict with the laws of nature. Of course, this doesn't detract from the epistemic possibility of the objective (i.e., nomological) possibility of *p*, or prevent agents in C from being justified in believing that it is epistemically possible that *p* is nomologically possible – but this does not automatically translate into justification for the objective, nomological possibility of *p*.

In sum, the emerging literature on modal aspects of scientific modelling rarely distinguishes different types of possibility claims inferred from models. Where it does – as in Massimi's case – the distinction between epistemic and objective possibility is disregarded. We believe that this theoretical neglect obstructs proper evaluation of claims about the epistemic functions of exploratory models. To support this claim, we now proceed to show that heeding the epistemic/objective distinction indeed is relevant for a number of current debates in the philosophy of science.

3. The Relevance of the Distinction for Science

As has by now been extensively documented (see Gelfert [2019] for an overview), scientists often consider possibility in their work, and arguably both epistemic and objective modalities are relevant to science. Of course, philosophers of science discuss these (overtly or implicitly) modal scientific practices, and often they disagree on how the practices are to be understood or reconstructed. In this section we present four debates from the philosophy of science that deal in possibility, but where there is a general lack of specification as regards the relevant sense of 'possible'. In particular, the epistemic/objective distinction is conspicuously absent from these debates, despite the fact that, as we will presently argue, what is at stake in these disagreements can be significantly clarified if the distinction is brought to bear on the scientific practices in question.

3.1. How-possibly explanation

One prominent way in which possibility claims figure in the philosophy of science literature is in so-called *how-possibly* explanations (HPEs). An HPE tells us how a phenomenon possibly could have come about. Philosophers of science disagree over whether HPEs form a *sui generis* kind of explanation (e.g., Forber [2010]), or whether they are merely incomplete stages towards the right how-actually explanation (e.g., Bokulich [2014], Brandon [1990]). From the latter perspective, HPE can be subsumed under ones' preferred account of explanation; the only conceptual addition needed is to determine in which way HPEs are incomplete. From the *sui generis* perspective, in contrast, HPE contribute epistemically in a substantially different way from how-actually explanations. Therefore, additional conceptual resources that account for how HPE generates understanding are required.[10] Avoiding that would be desirable for unificatory and simplicity reasons, so the disagreement centers on whether all HPE practices can be subsumed under standard explanatory accounts – and if not, what distinguishes the non-subsumable from the subsumable ones (see Verreault-Julien [2019] for a useful overview of this debate).

We believe that this unresolved debate can be significantly clarified by introducing the epistemic/objective distinction into the HPE debate. Consider Bokulich's ([2014]) argument that HPEs are an incomplete form of how-actually explanations, using the example of the tiger bush. Tiger bush is the phenomenon where vegetation in semi-arid areas grow in stripes, separated by barren areas, forming a pattern reminiscent of that in the tiger's fur. Scientists do not know exactly what causes this self-organising pattern formation. But they construct models that are supposed to provide *possible explanations* – e.g., Turing models, kernel models, differential flow models – of the tiger bush phenomenon. From the study of this modelling practice, Bokulich draws the conclusion that HPEs and how-actually explanations are 'not in fact two distinct and exhaustive categories, but rather two poles of a spectrum' (p. 334). In particular, HPEs are incomplete in the sense that they are not yet empirically confirmed. As Bokulich puts it, 'how-possibly explanations are explanations that, though not known to be the case, do not conflict with known facts' (ibid.). But with the accumulation of more empirical evidence, some of the HPEs will be culled, i.e. scientists will rule out this or that mechanism as *not* in fact responsible for producing the phenomenon, while other HPEs are further supported or not yet ruled out. That is, they move along the spectrum towards the ultimate goal of a how-actually explanation.[11]

Not all cases of HPE can be understood in this way, and this motivates the *sui generis* claim. In biomimetic chemistry, for example, researchers ask whether natural DNA could be replaced with a new, size-expanded geometry, while retaining all the functions that natural DNA

---

[10] There seems to be a consensus today that earlier proposals along the lines of (Dray [1957]) do not work. More recent discussions (Forber [2010], Rohwer and Rice [2013], Cuffaro [2015]) do not offer full accounts of how *sui generis* HPE generate understanding. Instead, they take their starting point from the observation of distinct scientific practices, and argue that their explanation requires the assumption of sui generis HPE.

[11] Bokulich ([2014], pp. 334-35) agrees with Brandon on this, but also stresses that the interplay between HPE and HAEs is more complicated, because explanations are subject to different levels of abstraction, so that scientists may have settled on a comparatively abstract HAE of *A,* and at the same time entertain different, more detailed HPEs of *A*.

has in nature's genetic system (Krueger et al. [2007], Lynch et al. [2006]).[12] Researchers synthesise such XNA and explore its functional properties. Some of them conclude (e.g., Eschenmoser [1999]) on the basis of this research that such alternative systems *could have* existed – the evolution of life could have been based on them, either instead of or in addition to DNA. But these scientists know full well that *in fact* evolution of life was based on RNA/DNA. That is, the HPE of life, supported by this research, *does* 'conflict with known facts'. The epistemic contribution of such an HPE cannot be explained as being steps towards the how-actually explanation, but requires separate methodological evaluation.

The epistemic/objective distinction illuminates this debate. If an HPE of a phenomenon *A* is supposed to provide an *epistemically* possible cause of *A*, it *is* reasonable to think that HPEs are just stages towards a how-actually explanation, and that its epistemic contribution can be subsumed under a standard account. In contrast, if an HPE of *A* is supposed to provide an *objectively* possible but non-actual cause of *A*, the practice of providing HPEs is indeed *sui generis* and not reducible to an instrument in the quest for the ultimate how-actually explanation. In the XNA case, scientists' claims about the functional requirements of genetic systems – and about them being satisfied by certain properties of these expanded DNAs – describe states that are, if their argument holds, *objectively* possible. Such an HPE provides us with a piece of modal knowledge that is interesting in itself, and it may support certain counterfactuals important to our scientific understanding of the relevant phenomena. However, it does not directly contribute to an explanation of any actual phenomena. Thus, if such an HPE makes any epistemic contribution at all, it must be *sui generis*.

The distinction also shows that the two camps are not necessarily in conflict. They both describe scientific practices that have legitimate claim to being called explaining how-possibly: they are concerned with different, but equally potentially relevant, senses of 'possibly'. The story told by Bokulich makes sense in contexts when scientists want to know what the *actual* mechanism behind some phenomenon is, so by way of modelling they try to map the causes of e.g., tiger bush that are possible *relative to the evidence available to them* – i.e., epistemically possible. In other contexts, such as with the case of XNA, the aim is rather to support objective (sometimes counterfactual) possibility claims.

Consequently, distinguishing between epistemic and objective HPEs firstly provides a way of making sense of the otherwise somewhat opaque claim that HPEs are *sui generis*, by cashing it out in terms of aiming to establish objective possibility claims. Secondly, it dissolves some of the

[12] For philosophical analyses of synthetic biology, see (Knuuttila & Loettgers [2013] and Koskinen [2017]).

disagreement, since it clarifies that the object contention divides into two different practices that each can legitimately be called explaining how-possibly.

3.2. Refuting necessity claims

Questions of whether certain states of affairs are necessary or contingent are among science's most prominent issues. Does all matter necessarily contain quarks? Is DNA necessary for the evolution of cells? Is geographic clustering a necessary condition for social cooperation? In order to find out, scientists investigate possibilities, and in doing so they often employ models. But it remains controversial whether models must be (perhaps partial or approximate) representations of actual targets or not, in order to be relevant to refuting necessity claims.

Those who argue that they must, follow models-as-representations accounts (e.g., Weisberg 2013) in arguing that models must be adequate representations of actual targets in order to teach us about the world: '*some* world-linking relation must hold between a model and real-world target if consideration of this model is to prompt learning about such targets' (Fumagalli [2016], p. 450; cf. Fumagalli [2015]), where 'world-linking relations' are to be understood as common adequacy conditions like 'similarity, isomorphism, resemblance … between such models and targets' ([2016], p. 434), and 'learning' is indicated as a justified change in one's belief – in particular one's confidence in necessity or impossibility hypotheses (Grüne-Yanoff [2009], p. 81).

Now, consider the example of Axelrod et al. ([2002]), who investigated whether social networks with distant links can perform as well as networks based on local geography. The authors motivate their investigation by noting the widespread concern that 'broadening patterns of inter-action among people … would destroy the basis of community' (ibid., p. 341). They interpret this concern as the belief that geographic clustering is a *necessary* condition for social cooperation (and that the decrease of local interaction thus would imply an unravelling of social cooperation). To reject this necessity claim, the authors employ a highly stylised Prisoner's Dilemma game model. In the model, stable cooperation with high average population payoffs can be achieved with a variety of network structures that determine how agents select neighbours to interact with. From this, the authors conclude that geographic clustering is not necessary for social cooperation *in the real world*: 'a persistent random network can support cooperation as well as a geographic network' (ibid., p. 345). The authors thus claim that considering this model should reduce one's confidence in the claim that geographic clustering is a necessary condition for social cooperation. But the model is *not* a representation of any particular real-world target; Axelrod et al. do not make any effort to justify the many idealisations of their model as compatible with an adequate representation, in fact, they don't bother identifying any actual target – probably because it didn't exist at the time: mass social networking arose only after this article was written.[13] Axelrod et al. purportedly learn about necessity claims concerning the actual world from models that did not adequately represent actual targets.

Critics of course might maintain that this is too bad for Axelrod and colleagues, and that the technical feasibility provided by the machinery of agent-based simulations does not amount to a possibility relevant to refuting necessity claims as the one discussed here. We believe that they are mistaken, and that this disagreement rests on a misconception of the relevant possibilities. Axelrod et al. were not aiming to identify an epistemic possibility – how actual cooperation was possibly generated, given current knowledge. They knew that geographic clustering was still the dominant cause of cooperation. They saw the increased 'ability of electronic communication' (ibid., p. 341) as a possible future development that might undermine this dominance. Thus, they were interested in the effect of currently non-actual possibilities on cooperation – that is, objective possibilities. This relates to their main motive, the worry that modernisation dynamics would destroy community, based on the belief that geographic clustering was necessary for cooperation. This necessity claim went beyond actuality, including non-actual future developments. Rejecting it therefore also required considering such non-actual possibilities, which is what they claim their model provided. Satisfying this goal not only allows for models not adequately representing actual targets; *it requires them*. Unless the above critics deny that the inferences exemplified by Axelrod et al. constitute learning about the actual world, their claims are mistaken.

Consequently, distinguishing between epistemic and objective possibility claims clarifies what models are required to investigate necessity claims. In particular, if the focus is on an objective necessity claim, the models for investigating it must support objective possibility claims. These may of course be actual, but very many and in this context equally relevant objective possibilities will be counterfactual. Models representing them thus cannot and should not be required to adequately represent actual targets. If, however, the focus is on epistemic possibility, then requiring adequate representation of actual targets seems justified. This shows once again that introducing the distinctions between epistemic and objective possibility can dissolve apparent conflict on one issue into positions on substantially different practices.

[13] The first mass social networking site, the South Korean *Cyworld*, was launched in 2001 Cf. (Boyd & Ellison [2007], p. 214): 'While people were already flocking to the Internet [in 2000], most did not have extended networks of friends who were online.'

3.3. Constraint-based reasoning

Practices of *constraint-based reasoning* are widespread in biology and the engineering sciences. Its core idea is to describe a system by a set of constraints that characterise its possible behaviours, but typically are not detailed enough to allow making precise predictions (Orth et al. [2010]). While some constraints are physically imposed upon the structure and function of natural systems, constraint-based reasoning employs formal constraints - general principles, often conceptualised in mathematical terms, often aiming to approximate the natural constraints present for those systems. Such formal constraints help characterise a system's dependency in a way that limits and affords a certain set of possibilities for the models of that system. Some philosophers have argued that certain practices of constraint-based reasoning are explanatory, despite or even because these practices often neglect relevant mechanistic and causal details (Green and Jones [2016]). But whether or not non-causal or non-mechanistic accounts identify genuine explanations remains controversial.

Those who do consider such practices explanatory argue that they facilitate reflection about whether certain system architectures are necessary or sufficient for a certain function to be performed, why certain dynamic or structural properties are present, and through that facilitate higher-level reflections on types of system organisation (Green and Jones [2016], p. 355). In their view, the reflections that constraint-based reasoning facilitates constitutes a separate mode of explanation. In contrast, authors who deny the existence of such non-mechanistic or non-causal explanations have typically accepted that scientists engage in constraint-based reasoning, but have maintained that it serves only a heuristic, not an explanatory function (e.g., Craver [2007], Matthiessen [2017]). They instead explain the usefulness of constraint-based reasoning through its provision of templates or mechanism schemes, which help both with ordering the search space for some actual causal process in a system, and narrowing it down with incoming new evidence. Such *constraint-based search* can be explanatory only to the extent that it helps identify the relevant mechanistic detail of the actual system.

We believe that this apparent conflict can be better understood if the different notions of possibility referred to in the discussion are made more explicit. Scientists employ constraint-based reasoning in order to identify possible behaviours of a system; that much both defenders and deniers of non-causal explanations agree on. They don't agree on what functions the possibilities serve. Deniers of non-causal explanations like Craver ([2007]) argue that any attempt to explain by pointing to possibilities identified through constraint-based reasoning would fail, because such explanations do not satisfy the fundamental actuality requirement of explanation: a set of possible mechanisms might or might not include the actual mechanism, but explanation requires identifying the actual one. However, 'mechanism schemes' and the like are useful tools in the quest for this actual mechanism, because they describe *how it could look, given what we know*. This usefulness tracks an epistemic notion of possibility: only mechanisms consistent with current knowledge can assist in identifying the actual mechanism that will eventually do the explanation. Non-actual objective possibilities, in contrast, could not assist in this quest.

Defenders of non-causal explanations, on the other hand, while they do not deny the function just described, are not primarily concerned with epistemic possibility. The additional function they claim for constrained-based reasoning turns on something quite different. Green and Jones ([2016]), for example, argue that these models help identify sufficiency or necessity relations between dynamic, structural or organisational properties on the one hand and functional properties on the other.[14] This arguably requires identifying the possibility of such relations *in the world*, not in terms of compatibility with our knowledge. Consequently, if their argument is to succeed (which we remain uncommitted on), then it requires objective possibility.

In light of this, we can see that defenders and deniers of non-causal explanations agree that constraint-based reasoning can provide information about possibilities, but they interpret these possibilities differently and *therefore* they disagree about what functions these possibilities can serve. As it turns out, both might well be correct, because the notion of constraint-based reasoning itself remains vague as to what possibilities it produces, much like the notion of how-possibly explanation, as discussed in 3.1 above. If the 'set of constraints that characterize its possible behaviors' is interpreted as *consistency with* current knowledge, then constraint-based reasoning yields epistemic possibilities, and critics are right in observing that mere compatibility with current knowledge may come too cheap to serve any justificatory or explanatory role. However, if it is instead interpreted as *evidential support* from suitably delineated background knowledge or laws, then it might well yield knowledge of objective possibilities, which as the defenders rightly point out may well be of explanatory relevance. Thus, an analysis of the possibility concept helps dissolve the apparent conflict over *the* function of constraint-based reasoning into two separate views on it different and potentially equally legitimate functions.

3.4. Possibilistic climate modelling

The Intergovernmental Panel on Climate Change (IPCC) bases its climate projection on *ensembles* of climate models. The ensembles, called the *Climate Model Intercomparison Project* (CMIP), are

[14] '[W]hile the possibility space for constraints used as heuristics ranges over potential mechanisms, the possibility space for constraint-based explanations ranges over potential variations among structural aspects of biological systems, given specified functional requirements' (Green & Jones [2016], p. 359).

collections of models developed by a number of certified modelling groups across the world.[15] CMIP5 is the model ensemble for the IPCC's Fifth Assessment Report (AR5), released in 2013, consisting of more than 50 models developed by more than 20 modelling groups. CMIP6 began in 2013. By 2018 it had endorsed 23 Model Intercomparison Projects, involving 33 modelling groups. It will be published with the IPCC 6th Assessment Report Working Group I in 2022.

The models in each CMIP are incompatible in the sense that they represent the physical processes of the climate system in mutually incompatible ways. Yet they are not regarded as mutually exclusive, but as complementing each other (cf. Parker [2006], [2013]); the ensemble is not merely seen as a reservoir of models that might be individually useful, rather the ensemble *as a whole* is taken to serve some epistemic function. Yet philosophers, and climate scientists, disagree over how to interpret this function and its constraints.

One prominent interpretation of model ensembles (not only in climate science) is the *multiple model idealisation* account (Levins [1966], Weisberg [2013], pp. 103-5). The basic idea is that modellers are always forced to make simplifying and idealising assumptions about their target. But because there are different ways how one can idealise, because these idealisations support different epistemic model desiderata (e.g., precision, accuracy, simplicity, compatibility with certain theories, etc.) and because the influence of these idealisations on model outcomes is often uncertain, scientists often collect those differently idealising models of the same target in an ensemble.

An alternative account of climate model is the *possibilistic* interpretation. On this view, the model results can be taken as representing future possibilities for the world: 'The model simulations are therefore taken as *possibilities* for future real-world climate and as such of potential value to society' (Stainforth et al. [2007], p. 2155, our emphasis). Furthermore, the models are not just providing information about individual possibilities, but the ensemble helps mapping the *spread* or *range* of possible future climate change (Knutti et al. [2010], p. 5; Stainforth et al. [2007], p. 2159). The possibilistic interpretation is an alternative to the multiple model idealisation account, because it accounts for the variation in the model ensemble in a different way. It is not the plurality of model desiderata that drive different idealisations, but rather the extend of the possibility space that determines the constitution of the model ensemble according to the possibilistic interpretation.

Of course, this re-interpretation does not do away with the fact that models in the ensemble are typically idealised. But unlike the multiple model idealisation account, which

[15] More specifically, each model in the CMIP is a global coupled ocean-atmosphere general circulation model (GCM) that takes as input different CO2 emission scenarios, and gives as output the projected change in annual mean surface air temperature, typically from the late 20th century to the middle 21st century.

considers each ensemble member an (idealised) representation of the actual target, the possibilistic account considers each member a (potentially idealised) representation of a possibility. Philosophers of science have disagreed what implications idealisations have for the possibilistic interpretation. For instance, Betz ([2015]) worries that because all ensemble members are idealised, climate ensembles do not support conclusions about possibilities and their ranges at all. To remedy this, he develops a framework that is supposed to isolate the model idealisation from the possibilities they represent. In contrast, Katzav ([2014]) is much more optimistic about, and sees less complication besetting, the idea that climate projections indicate possibilities, irrespective of their idealisations. We believe that their apparent disagreement here rests on a difference in how the possibilities supported by climate models are understood, and that the distinction between epistemic and objective possibilities brings this out.[16]

In particular, we believe that Betz assumes the relevant possibility concept to be epistemic, and for this reason is worried about model idealisations. Taking the relevant sense of possibility to be 'serious possibility', he writes: 'P is seriously possible if and only if P is consistent with the entire body of background knowledge K' (Betz [2015], p. 195). Both the consistency requirement, as well as the inclusion of the entire body of K, indicate that Betz follows an epistemic notion of possibility here. Under such an epistemic interpretation, model ensembles' function consists in mapping scientists' uncertainty at a given point in time. The spread of the CMIP model results shows what they know (expressed by the margins of this spread and its trend in time) as well as what they don't know (the space within this spread). As long as this possibility space is constrained by knowledge, the true model must lie within it – even though scientists lack evidence to determine which one it is. Idealisation threatens such a function. Idealised models attribute properties to targets that modellers know to be inaccurate. They are thus inconsistent with current knowledge, and do not represent epistemic possibilities. If models are supposed to represent epistemic possibilities, it is thus of utmost importance to deal with this problem of model idealisation.

Katzav, in contrast, seems to endorse an objective notion of possibility. He takes the relevant sense to be 'real possibility' (Katzav [2014], p. 236; cf. also Katzav et al. [2012]). Katzav refers to Deutsch ([1999]) here, who offers a restricted notion of objective possibility that holds fixed the actual history of the world up to a certain point. According to Katzav, to be a real possibility involves being 'compatible with the basic way things are in the target domain over the [relevant] period of time' ([2014], p. 236). Importantly, the 'mere absence of knowledge that something is not the case does not make something a real possibility' (ibid.). Both the rejection of

[16] Parker ([2018]) observes in passing that their disagreement is at least partly due to disagreement on the relevant notion of 'possibility', but offers no analysis what these notions could be.

the mere consistency requirement, as well as the insistence on support from a constrained set of facts ('the basic way things are') indicate that Katzav follows an objective notion of possibility. Because of this, he is not as worried about idealisation as Betz is.

Objective possibilities are often counterfactual, and many idealizations can be interpreted as counterfactual states of affairs. That modellers have chosen them for reasons of tractability or simplicity does not speak against this: they might be the simplest conceivable possibility that yield a certain property, as with e.g., Axelrod and colleagues' 2*2 PD game with a random network. The only exception are *impossible* idealisations, e.g., zero friction, infinite population, or no transaction costs. Some idealisations in scientific models are impossible in this way, but not all. Therefore, the function of providing objective possibilities is less threatened by model idealisation than the function of providing epistemic possibilities.

To conclude, idealisations often contradict current knowledge and thus pose a substantial problem for epistemic possibility. Objective possibility, in contrast, because it concerns often counterfactual situations, is not similarly affected by all idealisation. The distinction thus helps to show that those who claim that model idealisation undermines the possibilistic account of model ensembles and those who argue that it doesn't can justify their respective positions with different notions of possibility.

3.5 Summary

We reviewed four topics in which possibilities play a prominent role (table 1, first column). For each topic, we identified an unresolved philosophical debate (table 1, second column) and showed that the disagreement at least partially rested on different interpretations of the possibility notions involved. By distinguishing between epistemic and objective possibility for each case, we showed that the conflict could be significantly clarified, and the risk of participants talking past each other reduced. This improved understanding of the debates should, in turn, provide good grounds for fruitfully advancing these debates beyond their current state. To be clear, we do not claim to have once and for all resolved these debates. Rather, we wanted to showcase how more care in distinguishing between these different notions of possibility can be useful for a range of issues in philosophy of science.

Table 1

Topic	Debate	Analysis result
How-possibly explanation (HPE)	Are HPEs more than defective how-actual explanations?	Epistemic HPEs are merely defective how-actual explanations, but objective HPEs are not.
Necessity claims	Can models that do not represent actual targets refute necessity claims?	Rejecting epistemic necessity claims might require adequate representation of actual targets. Rejecting objective necessity claims require objective possibility, objective possibility requires models of non-actual targets.
Non-causal explanations	Does constraint-based reasoning provide non-causal explanation?	Epistemic possibilities might assist in constructing but do not provide causal explanation; objective possibility might offer non-causal explanation; constraint-based reasoning provides both.
Model ensembles	Does idealisation in model ensembles prevent the identification of possibilities?	Idealisations often contradict current knowledge and thus pose a substantial problem for epistemic possibility. Objective possibility, because it often concerns counterfactual situations, is not similarly affected by all idealisation.

4. Conclusions

Epistemic possibilities and objective possibilities concern different things, and as a result differ importantly with respect to their epistemology. Philosophers of science often do not pay sufficient attention to the work that this distinction can do in analyses of modelling practices that deal in possibility claims. We have argued that explicitly introducing the distinction into particular philosophy of science debates resolves confusion and thereby fosters progress on the relevant issues, using four examples: how-possibly explanation, refutation of necessity claims on the basis of models, the possibility of non-causal constraint-based explanation, and interpretation of climate model ensembles. The distinction was obviously not constructed for the purpose of analysing any of these cases. Instead, we consider it a fortuitous result that each case illustrates the importance of the epistemic/objective distinction for the evaluation of scientific inference practices. Furthermore, precisely because of this independence, we believe that the distinction captures something fundamental about modal modelling practices, and are confident that other modal issues relating to scientific modelling could benefit in similar ways from heeding it.

Acknowledgements

We would like to thank Philippe Verreault-Julien, the audience at the Joint Stockholm/Uppsala Seminar in Philosophy of Science, and several anonymous referees for their helpful comments. Both authors acknowledge funding from the Swedish Research Council, grants no. 2018-01353 and no. 2019-00635.

Ylwa Sjölin Wirling
Department of Philosophy, Linguistics, and Theory of Science
University of Gothenburg
Gothenburg, Sweden

Department of Philosophy
University of Manchester
Manchester, United Kingdom

ylwa.wirling@gmail.com

Till Grüne-Yanoff
Divison of Philosophy
KTH Royal Institute of Technology
Stockholm, Sweden

gryne@kth.se

References

Adler, J. [2002]: *Belief's Own Ethics*. Cambridge, MA: Bradford Books/MIT Press.

Axelrod, R., Ford, G. R., Riolo, R. L., and Cohen, M.D. [2002]: 'Beyond Geography: Cooperation with Persistent Links in the Absence of Clustered Neighborhoods', *Personality and Social Psychology Review* **6 (4)**, pp. 341-46.

Batterman, R. and Rice, C. [2014]: 'Minimal Model Explanations', *Philosophy of Science*, **81(3)**, pp. 349-76.

Bealer, G. [2002]: 'Modal Epistemology and the Rationalist Renaissance'. In T. Gendler and J. Hawthorne *(Eds.), Conceivability and Possibility*, pp. 71-125, Oxford: Oxford University Press.

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 for Philosophy of Science,* **5(2),** pp. 191-215.

Bokulich, A. [2014]: 'How the tiger bush got its stripes: "how possibly" vs. "how actually" model explanations', *The Monist*, **97(3)**, pp. 321-38.

Boyd, D. M. and N. B. Ellison. [2007]: 'Social network sites: Definition, history, and scholarship', *Journal of computer-mediated Communication*, **13(1)**, pp. 210-30.

Brandon, R. N. [1990]: *Adapatation and Environment.* Princeton: Princeton University Press

Chalmers, D. [2011]: 'The nature of epistemic space', In A. Egan and B. Weatherson *(eds.), Epistemic Modality*, pp. 60-107, New York: Oxford University Press

Craver, C. [2007]: *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*, Oxford: Clarendon Press.

Cuffaro, M. E. [2015]: 'How-possibly explanations in (quantum) computer science' *Philosophy of Science*, **82(5)**, pp. 737-48.

DeRose, K. [1991]: 'Epistemic possibilities', *Philosophical Review,* **100(4)**, pp. 581-605.

Deutsch, H. [1990]: 'Real possibility', *Noûs,* **24(5)**, 751-55.

Dray, W. H. [1957]: *Laws and explanation in history*. Oxford: Clarendon Press.

Edgington, D. [2004]: 'Two kinds of possibility', *Supplement to the Proceedings of the Aristotelian Society,* **78(1)**, pp. 1-22.

Eschenmoser, A. [1999]: 'Chemical etiology of nucleic acid structure', *Science*, **284(5423)**, pp. 2118-24.

Fischer, B. [2016]: 'Hale on the Architecture or Modal Knowledge', *Analytic Philosophy,* **57(1)**, pp. 76-89.

Forber, P. [2010]: 'Confirmation and explaining how possible', *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences,* **41(1)**, pp. 32-40.

Fumagalli, R. [2015]: 'No learning from minimal models', *Philosophy of Science,* **46(2)**, pp. 798-809.

Fumagalli, R. [2016]: 'Why we cannot learn from minimal models', *Erkenntnis,* **80(3)**, pp. 433-55.

Gelfert, A. [2016]: *How to do science with models: A philosophical primer.* Cham, Switzerland: Springer International.

Gelfert, A. [2019]: 'Probing Possibilities: Toy Models, Minimal Models, and Exploratory Models', In M. Fontaine, C. Barés-Gómez, F. Salguero-Lamillar, L. Magnani and Á. Nepomuceno-Fernández *(eds), Model-Based Reasoning in Science and Technology*, pp. 3-19, Cham, Switzerland: Springer International.

Gendler, T. and Hawthorne, J. [2002]: 'Introduction: Conceivability and Possibility', In T. Gendler & J. Hawthorne *(eds.), Conceivability and Possibility,* pp. 1-70, Oxford: Clarendon Press

Green, S., & Jones, N. [2016]: 'Constraint-based reasoning for search and explanation: Strategies for understanding variation and patterns in biology', *Dialectica*, **70(3)**, pp. 343-74.

Grüne-Yanoff, T. [2009]: 'Learning from Minimal Economic Models', *Erkenntnis,* **70(1)**, pp. 81-99.

Grüne-Yanoff, T. [2013]: 'Appraising models non-respresentationally', *Philosophy of Science,* **80(5)**, pp. 850-61.

Hacking, I. [1967]: 'Possibility', *Philosophical Review,* **76(2)**, pp. 143-68.

Hale, B. [2002]: 'Knowledge of possibility and impossibility', *Proceedings of the Aristotelian Society*, **103,** pp. 1-20.

Huemer, M. [2007]: 'Epistemic Possibility', *Synthese,* **156(1)**, pp. 119-42.

Katzav, J. [2014]: 'The epistemology of climate models and some of its implications for climate science and the philosophy of science', *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics,* **46(2)**, pp. 228-38.

Katzav, J., Dijkstra, H. A., and 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**, pp. 258-76.

Kment, B. [2014]: *Modality and Explanatory Reasoning.* Oxford: Oxford University Press.

Knutti R., Furrer, R., Tebaldi, C., Cermak, J., and Meehl, G. A. [2010]: 'Challenges in Combining Projections from Multiple Climate Models', *Journal of Climate*, **23(10)**, pp. 2739-58.

Knuuttila, T. and Loettgers, A. [2013]: 'Synthetic modeling and mechanistic account: Material recombination and beyond', *Philosophy of Science*, **80**, pp. 874–85.

Koskinen, R. [2017]: 'Synthetic biology and the search for alternative genetic systems: Taking how-possibly models seriously', *European Journal for Philosophy of Science*, **7(3)**, pp. 493-506.

Krueger, A. T., Liu, H., Lee, A. H. and Kool, E. T. [2007]: 'Synthesis and properties of size-expanded DNAs: toward designed, functional genetic systemsä *Accounts of Chemical Research*, **40(2)**, pp. 141-50.

Lange, M. [2009]: *Laws and Lawmakers: Science, Metaphysics, and the Laws of Nature*. New York: Oxford University Press.

Levins, R. [1966]: 'The strategy of model building in population biology', *American Scientist,* **54(4)**, pp. 421-31.

Lynch S. R., Liu, H., Gao, J. and Kool, E. T. [2006]: 'Toward a designed, functioning genetic system with expanded-size base pairs: solution structure of the eight-base xDNA double helix', *Journal of the American Chemical Society*, **128(45)**, pp. 14704-11.

Malcolm, N. [1963]: 'The Verification Argument', In N. Malcolm, *Knowledge and Certainty,* pp. 1-57, Englewood Cliffs, NJ: Prentice-Hall, Inc.

Massimi, M. [2019]: 'Two kinds of exploratory models', *Philosophy of Science,* **86(5)**, pp. 869-81.

Matthiessen, D. [2017]: 'Mechanistic Explanation in Systems Biology: Cellular Networks', *The British Journal for the Philosophy of Science*, **68(1)**, pp. 1-25.

Moore, G. E. [1959]: 'Four Forms of Skepticism', In G. E. Moore, *Philosophical Papers*, pp. 196-225, London: Routledge.

Orth, J. D., Thiele, I. and Palsson, B. Ø. [2010]: 'What is flux balance analysis?', *Nature biotechnology*, **28(3)**, pp. 245-48.

Parker, W. [2006]: 'Understanding pluralism in climate modeling', *Foundations of Science,* **11(4)**, pp. 349-68.

Parker, W. [2013]: 'Ensemble Modeling, uncertainty and robust predictions', *Wiley Interdisciplinary Reviews: Climate Change,* **4(3)**, pp. 213-23.

Parker, W. [2018]: 'Climate Science', In E. N. Zalta *(Ed.), The Stanford Encyclopedia of Philosophy*, available at <https://plato.stanford.edu/entries/climate-science/>.

Przyjemski, K. [2017]: 'Strong epistemic possibility and evidentiality', *Topoi*, **36(1**), pp. 183-95.

Roca-Royes, S. [2017]: 'Similarity and Possibility: An epistemology of de re possibility for concrete entities', In B. Fischer and F. Leon *(eds.), Modal Epistemology After Rationalism*, pp. 221-45, Cham: Springer International.

Rohwer, Y. and Rice, C. [2013]: 'Hypothetical pattern idealization and explanatory models', *Philosophy of Science*, **80(3)**, pp. 334-55.

Strohminger, M. and Yli-Vakkuri, J. [2019]: 'Knowledge of objective modality', *Philosophical Studies,* **176(5)**, pp. 1155-75.

Stainforth, D. A., Allen, M. R., Tredger, E. R. and Smith, L. A. [2007]. 'Confidence, Uncertainty and Decision-support Relevance in Climate Predictions', *Philosophical Transactions of the Royal Society,* **365**, pp. 2145-61.

Verreault-Julien, P. [2019]: 'How could models possibly provide how-possibly explanations?' *Studies in History and Philosophy of Science Part A*, **73**, pp. 22-33.

Weatherson. B and Egan, A. [2011]: 'Epistemic Modals and Epistemic Modality', in A. Egan and B. Weatherson *(Eds.)*, *Epistemic Modality*, pp. 1-18, New York: Oxford University Press.

Weisberg, M. [2013]: *Simulation and Similarity*, Oxford: Oxford University Press.

Williamson, T. [2016]: 'Modal science', *Canadian Journal of Philosophy*, **46**, pp. 453-92.