

# Against “Possibilist” Interpretations of Climate Models

Corey Dethier (please contact before citing or quoting)

## Abstract

---

Climate scientists frequently employ (groups of) heavily idealized models. How should these models be interpreted? A number of philosophers have suggested that a *possibilist* interpretation might be preferable, where this entails interpreting climate models as standing in for possible scenarios that could occur, but not as providing any sort of information about how probable those scenarios are. The present paper argues that possibilism is (a) undermotivated by the philosophical and empirical arguments that have been advanced in the literature, (b) incompatible with successful practices in the science, and (c) liable to present a less accurate picture of the current state of research and/or uncertainty than probabilistic alternatives. There are good arguments to be had about how precisely to interpret climate models but our starting point should be that the models provide evidence relevant to the evaluation of hypotheses concerning the actual world in at least some cases.

---

## 0 Introduction

Climate scientists frequently employ groups or “ensembles” of heavily idealized models. How should these models be interpreted? To make this question more approachable, consider a standard case in which a simulation of some climate scenario is run on a model that represents the global climate. This simulation generates an output  $P$ . We can distinguish between two broad approaches to interpreting what this result indicates:

*Probabilistic*: the result tells us something about the probability that  $P$  holds in the actual world.

*Possibilistic*: the result tells us only that  $P$  is a “real possibility.”<sup>1</sup>

The dominant approach in climate science is the probabilist one: climate scientists regularly use groups of models to generate (precise) probability distributions in a manner that is incompatible with taking the results of model simulations to tell us nothing about probabilities. In the philosophical literature, by contrast, the probabilistic approach has few defenders and a large number of detractors.<sup>2</sup>

Some philosophers have suggested that the possibilist framework offers a superior alternative. Betz (2007), for instance, argues that *in the context of policy advice*, climate scientists should not present precise probabilistic claims about the future. More recently, Betz (2015) argues for possibilism in the context of model-generated predictions more generally and Katzav (2014) extends the view to the context of model evaluation, arguing that the primary focus of studies that compare the performance of models to empirical data should be on showing that models are “real possibilities,” not on determining how well they represent the actual world. Finally, Katzav et al. urge that probability density functions “should not be used in the climate context” (Katzav et al. 2021, 2) and tentatively advocate for a possibilistic approach instead.

Given these arguments, it makes sense to ask whether possibilism can replace probabilism; that is, whether it is a good *general* framework for the interpretation of climate models. The answer is “no”: while it may be that possibilistic approaches are preferable in some cases, the arguments that have been offered in the literature do not support the generalization to all cases. Furthermore, the possibilist approach is both incompatible with successful practices in the science and liable to present a less accurate picture of the current state of research than the probabilistic alternative in some realistic scenarios. As such, while there are good arguments to be had about how precisely to interpret climate models, our starting point should be that the models provide evidence relevant to the evaluation of hypotheses concerning the actual world in at least some cases.

---

<sup>1</sup>We should really attach an “assuming the internal validity of the simulation” qualifier here—we’re interested in the cases where the simulation has successfully been carried out according to the standards of the field—but since the rider applies in both cases, we can drop it for present purposes.

<sup>2</sup>Betz (2007, 2015), Katzav (2014), and Katzav et al. (2021) argue against probabilism and for possibilism. Parker (2010) and Winsberg (2018) raise problems for probabilism without rejecting it entirely.

# 1 Possibilism and individual models

Extant arguments for possibilism all operate by way of rejecting the (opposed) probabilist interpretation. What we'll see in the next two sections is that none of these arguments motivate rejecting probabilism in general. At worst, they motivate a case-by-case evaluation of the costs and benefits of the two approaches.

As is widely appreciated, climate models are imperfect representations of our actual climate. Furthermore, ensembles of climate models do not behave like random samples centered on the true climate. Both of these facts are established by empirical work on the subject (Knutti et al. 2010); both are also predictable based on our knowledge of how models are constructed. The upshot is that both individual models and groups of models must be understood as idealized representations in the sense that there are some aspects of their target that they distort. Proponents of possibilism have used the idealized character of climate models to argue that possibilism should be preferred to probabilism. Betz (2015) and Katzav (2014), for instance, suggest that we're not justified in treating a climate model as providing information about the actual world when we know that said model includes assumptions that are false.

If Betz and Katzav are right, however, that would make climate models the exception, not the rule. The received wisdom in philosophy of science is that all models are idealized in some respect. It's unquestionable that some models provide us with the information about the actual world, however. For example: point-mass models of the solar system can provide knowledge about the velocities of stellar bodies even though they massively misrepresent their densities. So if the argument offered by Betz and Katzav is to work, it must be because there's something about the idealizations involved in climate models in particular that makes it so that we cannot trust their outputs qua representations of the actual world.

There are a few candidates for distinguishing factors. First, it might be that there's empirical literature that indicates that climate models are generally untrustworthy. There is in fact substantial empirical literature on the accuracy of climate models, but it would be a serious misreading to interpret that literature as wholly negative. On contrary, the papers cited by proponents of possibilism show that the models are in fact quite trustworthy with respect to *some* aspects of the climate, and groups of climate models are in general even better (Knutti et al. 2010). Indeed, ensembles often outperform what we would expect given what we know about their idealized character, leading climate scientists to search for explanations for their "surprising" success (Annan and Hargreaves 2011). Nor is all of the evidence here fully backwards-

looking; recent work by Hausfather et al. (2020) indicates that climate models dating as far back as the 1970s have been quite good at predicting the relationship between increases in greenhouse gas concentration and temperature. In other words, what the empirical literature indicates is not that we cannot trust climate models to tell us about the actual world but instead that (a) we should trust climate models more in some cases and less in others and (b) we should generally trust ensembles more than individual models (though how much more varies with context).

Of course, another potentially differentiating factor is that there might be some reason why we cannot use the empirical evidence to evaluate model accuracy in this way. For example, Katzav (2014) and Lenhard and Winsberg (2010) raise the spectre of holism in arguing that models are essentially black boxes and so empirical success cannot be attributed to particular elements of the climate models. Katzav (2014) further argues that in many cases we lack independent access to the relevant empirical facts—we don't really have a good, model-independent means of estimating the internal variability of the climate, to pick just one example—and so we cannot test how accurate the models are with respect to those variables. And a number of authors have rightly pointed out that climate science is a domain in which the extrapolation from past performance is dubious. At worst, however, these points undermine the trust we should have in the models in specific cases; they don't indicate that the models never license conclusions about the actual world. To support the latter view, the empirical literature would need to show *both* that the predictions generated by models are generally untrustworthy *and* that we cannot typically distinguish between the good cases and the bad cases. But neither thesis is warranted by the evidence on hand.

The position motivated by the empirical literature therefore seems to be a moderate one: in at least some cases, we should take a given climate model to provide us with information about the actual world. It's conceivable that in other cases, the most that a climate model will warrant is a possibility claim—that a simulation run on the model delivers a result  $P$  indicates only that  $P$  is a “real possibility.” But the idealized character of climate models doesn't give us any reason to think that this is generally the case. As the data gathered by Hausfather et al. (2020) makes clear, if you were previously indifferent over a wide range of possible values for the relationship between greenhouse gases and temperature, learning the results of model-based simulations should lead you to adopt a narrower band of confidence around those results (compare Parker 2020). As such, the flaws in individual models can't be said to motivate adopting possibilism as a general framework.

## 2 Possibilism and multiple models

Until recently, proponents of possibilism have focused their arguments on individual models, arguing that these cannot be said to represent the actual world (at least not in the cases that they’re considering). In a recent paper, by contrast, Katzav et al. (2021) take a different approach, arguing that whatever the status of individual models, we’re not justified in trusting the probability distributions generated by groups, or “ensembles,” of models. After all, groups of models share idealizations and don’t behave like random samples. Even if an individual climate model has something to say about the actual world, probabilities generated by treating extant groups of them as some kind of sample are liable to distort the “evidential” or “objective” probabilities that we in some sense ought to have. As a consequence, employing the probabilities generated by an ensemble in the context of decision-making is likely to lead to bad decisions. A number of authors, including Parker and Risbey (2015) and Winsberg (2018), take this argument to cast doubt on the probabilistic interpretation of climate models; Katzav et al. (2021) go further and tentatively suggest that possibilism is one alternative that might be preferable.<sup>3</sup>

While the argument just given might show that the probabilities generated by extant ensembles misrepresent in various ways, it doesn’t show that we shouldn’t use them, let alone that we shouldn’t use them in general. To begin, it’s worth reiterating the main point from the last section: that a probability distribution involves some sort of misrepresentation is not sufficient to motivate abandoning it. We accept misrepresentations in the form of idealizations and abstractions throughout the sciences, and there’s no obvious reason why probability distributions should be exceptions to the general rule (compare Dethier 2022). Even if we accept that extant ensembles misrepresent, they may nevertheless be our best option for representing the “true” probabilities (compare Katzav and Parker 2015). So, for example, if we have to choose between making a decision based solely on the point-value estimate given by a single model or one based on the probability distribution generated by an ensemble, we should in general prefer the latter even if we expect it to misrepresent the “true” probabilities—after all, both representations can be thought of as probability distributions, with the difference being that the former assigns a confidence of 1 to the output of a single model while the latter distributes confidence more equitably over the available options. So even establishing that the probability distributions in question *grossly* misrepresent the true proba-

---

<sup>3</sup>The other alternative that they discuss involves employing imprecise probability distributions. At least at face value, this alternative is a version of the probabilist framework, and so not a target of this critiques found in this paper.

bilities only gives us a defeasible reason to reject them—we still have to ask about the relative costs and benefits of various alternatives.

Notably, the possibilist alternative is not free of misrepresentation. Once we accept that individual models provide information about the actual world, the possibilist option is just as much guaranteed to misrepresent the true state of our knowledge as any probabilistic approach. Where probabilism errs towards overstating our knowledge, the possibilism errs towards understanding it (compare Risbey 2007). After all, as its proponents explicitly acknowledge, the possibilist interpretation allows us to say nothing about the outer limits of the “real possibilities” on the basis of modeling results, meaning that we have to ignore whatever information climate models provide about the limits of what is really possible. It’s plausible that even if climate models don’t warrant precise probabilistic judgments concerning the actual world, they tell us something about which scenarios are more and less likely. Which approach should be preferred thus depends at least on whether errors due to over-precision or under-precision are more worrying in a given context.

Following the arguments offered by Betz (2007) and Parker and Risbey (2015), we might think that over-precision is generally more worrying in the setting of policy advice. Plausibly, scientists should only give advice within the scope of their expertise, and this principle tells against over-reaching in a way that it doesn’t tell against more conservative approaches. There are reasons to doubt this principle, but even if we accept it, there are many contexts other than the contexts of decision-making and policy advice where the calculus should be expected to be different. It may be the case, for example, that for the purposes of decision-making climate scientists should prefer to only make firm judgments about what is and isn’t possible, but that for all other purposes it’s worthwhile to adopt a more fine-grained probabilistic approach. This imagined position could be motivated by making an analogy with rounding to a particular significant digit. Scientists typically round their measurements to avoid adopting hypotheses that are more precise than is warranted—exactly the same motivation that Parker and Risbey (2015) have appealed to in criticizing the use of probabilities in climate science. But rounding is a step that occurs only at the end of the measurement process; it’s a serious methodological error to round results at every stage, because repeatedly rounding can lead to estimates that are significantly different from an entirely un-rounded estimate. Similarly, we might think that moving from the precision offered by precise probabilities to a more coarse-grained possibilistic representation is ultimately what’s demanded by our evidence while simultaneously thinking that it’s a mistake to coarse-grain in this manner at any point other than the end of the calculation. Notably, this is essentially the approach adopted by the

IPCC; see, e.g., IPCC ([in press](#)).

The upshot of the foregoing is that arguments found in the literature do not motivate adopting possibilism as a general framework. At most, these arguments establish that whether we should adopt a probabilistic interpretation of climate models depends on the costs and benefits of the approach relative to its competitors. In what follows, we'll see that possibilism is in fact generally worse than probabilism on at least two axes: it declares large chunks of successful science to be unmotivated and is liable to be less accurate in practice than the probabilistic alternative.

### 3 Possibilism and successful practice

Proponents of possibilism have largely focused on the future forecasting. But neither climate models nor the probabilities that they generate are used solely in that context. On the contrary, complex climate models are relied on throughout the science, and they're consistently understood as providing information about the way the world actually is. This section borrows an example from Dethier ([2022](#)) to illustrate the point and then argues that there's no principled distinction between the use of climate models in forecasting and other uses of climate models for the possibilist to fall back on.

Consider attribution studies, which are backwards-facing in that they aim to determine how responsible humans are for observed climate change. Many of the specific concerns raised by the proponents of possibilism don't apply in the attribution context—we're not extrapolating from present climate data to a future that we will know will be different, for instance—and though there remains substantial uncertainty about many specific details, the IPCC takes attribution studies to provide us “unequivacable” knowledge in at least some cases (IPCC [in press](#)).

Broadly speaking, attribution studies proceed as follows. Climate scientists collect substantial data on past changes to temperature and then run complex regressions to determine how much of the past temperature change can be attributed to CO<sub>2</sub> and how much to other factors such as the internal variability of the climate system. To run these regressions, they need a quantified understanding of how different factors affect the climate. So, for example, we consistently observe that while the planet as a whole is warming, the upper atmosphere is actually cooling. To determine how much of the observed warming is caused by CO<sub>2</sub> and how much by other factors, we need to know how these different factors affect the distribution of heat throughout the atmosphere. This information—what's sometimes called the “signature” or “fingerprint” of

a particular factor—is usually provided by climate models.

Simplifying and abstracting substantially, the resulting regression equation looks like this:

$$Y = \sum_i^n \beta_i X_i + v_Y$$

where  $Y$  is the observed data;  $\beta_i$  and  $X_i$  are the percentage of the increase due to the  $i^{\text{th}}$  factor and the signature of that factor, respectively; and  $v_Y$  is the internal variability of the climate. Standard least squares algorithms are then used to estimate the  $\beta$  terms. The results indicate how much of observed warming a particular factor is responsible for; if the least squares analysis yields a result that  $\beta_{GHG} = .95$ , for example, that would indicate that greenhouse gases are responsible for 95% of observed warming.

On a possibilist reading, however, this conclusion is not warranted—in fact, a general possibilist view is required to say that attribution studies cannot tell us anything about the actual world. The reason why is that the  $X$  terms are typically derived from climate models.<sup>4</sup> So if climate models do not provide information about the actual world, then the  $X$  terms only represent “real possibilities” and the results for the  $\beta$  terms must be thought of as providing information only about which contributions to climate change are possible. It’s not clear whether the consistent possibilist is even capable of making qualitative claims about humanity’s contribution to climate change in the actual world. Perhaps there are other lines of evidence that would allow them to say more than just that “it’s a real possibility” that humanity has contributed to climate change, but certainly they’re not warranted in adopting stronger conclusions on the basis of the standard form of attribution studies employed in climate science.

Of course, the possibilist has a couple of potential responses here. First, they could make the retreat already suggested in the last section and argue that while climate models can provide information about the actual world, the information generated by ensembles should not be given a probabilistic reading. Second, they could argue that there’s something particularly special about forecasting that distinguishes it from backwards-facing attribution science. Neither rejoinder is successful.

On the first count, we’ve already seen that the position in question is an unhappy halfway home: it’s hard to motivate the idea that individual models

---

<sup>4</sup>Of course, it’s not just the  $X$  terms. The description of attribution studies given here is highly simplified relative to actual practice; examining the complexities makes it clear that climate models are heavily involved at every step in the process.



tell us about the actual world but when we group them together we are no longer justified in drawing conclusions about anything more than real possibilities. It's especially unmotivated in this particular case, however. For roughly the last twenty years, climate scientists have employed the probability distributions generated by ensembles either in place of the  $X$  terms (Huntingford et al. 2006) or as part more complicated Bayesian updating procedures (Schurer et al. 2018). So even if the possibilist can recover the basic approach adopted in attribution studies by retreating in this way, they're committed to rejecting the conclusions generated by more complex versions of the approach—and, notably, these more complex approaches consistently yield results that are more accurate and reliable than the simple method (Hannart, Ribes, and Naveau 2014; Schurer et al. 2018). This means of recovering the successful practice of climate scientists thus looks like a non-starter.

The other rejoinder is no more successful. Contrary to what we might expect, there's not really a practical bright line to be drawn in climate science between how models are used in forecasting the future and how they're used in analyzing the past. Attribution studies illustrate the point nicely: since Stott et al. (2006), it's been standard practice to use the results of attribution studies to estimate how greenhouse gases will affect temperatures in the future. The possibilist thus faces a dilemma. On one horn is the position that the resulting forecasts should be interpreted in the same possibilist manner as those directly generated by climate models. But then it's not clear what their position has to do with climate *models*—it seems as though what they're really advocating is simply skepticism about climate scientists' ability to predict the future. On the other horn, they can assimilate the forecasts generated in this manner to the successful practices of attribution science and allow for a probabilistic interpretation here as well. It's not clear what the motivation for the resulting position could possibly be, however: why is that forecasts in which the models play one role can be interpreted probabilistically but those in which the models play a different role can't be?

The takeaway from this section is that the probabilistic interpretation of climate models is much more deeply intertwined with successful climate science than the literature typically acknowledges. Applied consistently, possibilism would require us to completely rewrite not just future-facing climate science, but all of climate science, including those areas deemed highly successful by the scientists themselves. As we've just seen, while the possibilist has potential avenues for resisting these applications, the resulting positions are hard to motivate.

## 4 Possibilism and accuracy

As stressed above, there’s no reason to think that the possibilist interpretation of climate models is *guaranteed* to be a more accurate representation of the current state of the evidence than a probabilistic one. On the contrary, just as the probabilistic representation risks overstating the evidence, the possibilist one risks understating it. In this sense, the two positions are analogous—as is often true, the question is simply which kind of error is more worrying.

There is an important sense, however, in which the possibilist interpretation is liable to be less *informative* than the probabilist one. Recall that there are empirical studies that examine how accurately groups of models represent present-day climate targets (see, e.g., Knutti et al. 2010). The principal finding of these studies is that extant ensembles under-sample from the extremes relative to your standard normal distribution centered on the truth—there’s less variance in the sample than we would expect there to be, indicating that a probability distribution generated on the assumption that the models are normally distributed is liable to underestimate the probability of extreme outcomes.<sup>5</sup>

Unsurprisingly, these findings are widely cited by critics of the probabilist interpretation. But the same results undermine possibilist approaches for essentially the same reason: if the set of models available under-samples from the extremes, then the set of “real possibilities” that are represented is unlikely to include extreme scenarios. To illustrate, suppose that the sample were normally distributed. Then you would need 14 models to have even a 50% chance of getting a result with a “true” probability of 5% or less. Given that we know that extant ensembles are *not* normally distributed, we should expect the numbers to be even higher, meaning that—on the possibilist interpretation—we should expect extant ensembles not to tell us anything about extreme cases. In aiming to avoid misrepresenting the probability of extreme scenarios, possibilism takes the quietist route and refuses to say anything about them at all.

Quietism isn’t the only option here. One of the great advantages of the probabilist approach is that it is flexible: we don’t have to assume that extant ensembles are normally distributed around the truth. Indeed, in the last decade, it’s become increasingly common for climate scientists to adopt a different interpretation of the models motivated by the same empirical literature

---

<sup>5</sup>This situation seems to have reversed in the last few years (Tokarska et al. 2020). If an ensemble over-samples from the extremes rather than under-samples, however, that doesn’t change the main point of this section, which is that probabilism allows scientists to respond to known biases.

cited above—in particular, they tend to assume that the “truth” behaves like a single sample from the same population as the models, an assumption that fits better with the empirical data and that (if anything) leads to an over-sampling from extremes (Annan and Hargreaves 2011; Sedláček and Knutti 2013).

Of course, as stressed by Katzav et al. (2021), correctly estimating the biases in extant ensembles is a tricky empirical task. The point here is not that these corrections are easy or even successful, but that probabilism opens the door for this kind of response to known biases. Furthermore, while the information that the probabilist approach provides with respect to the extreme scenarios may be flawed in some ways, the tails of the probability distribution do provide some information about what we can expect in these cases. That’s no guarantee that the resulting probabilities will be perfectly accurate, but it’s something. The possibilist interpretation, by contrast, doesn’t provide any information about these scenarios at all: it requires climate scientists to draw *no* conclusions about the extremes on the basis of extant models, and it has no way of taking into account any of the empirical information derived from studies like Knutti et al. (2010). Insofar as we’re worried about scenarios that we think are unlikely but not impossible, therefore, we should prefer the probabilist approach unless we think that the resulting information is more likely than not to be misleading in the relevant context. This extreme view is entirely unmotivated by the arguments surveyed above, however, at least in its general form: that extant probability distributions are more precise than is warranted does not mean that they are likely to be so inaccurate that we cannot draw any conclusions from them.

## 5 Conclusion

In the recent literature, it has become increasingly common for philosophers of climate science to advocate for a “possibilist” interpretation of climate models. What we’ve seen in the foregoing is that even if a possibilist interpretation has some worthwhile domains of application, it can’t replace the probabilistic one as a general framework. The explicit philosophical arguments advanced by the defenders of possibilism fail to show that the position is preferable in general. Moreover, consistent application of the possibilist interpretation requires us to reject successful applications of climate science and we should furthermore expect that the possibilist approach is less informative than a probabilistic one in at least some cases.

Ultimately, the possibilist interpretation involves a kind of quietism—a familiar “we don’t know everything about the probability of various scenarios,

so we can't say anything on the subject." There are cases where the evidence provided by climate models is so flimsy that this kind of quietism is well-motivated. Taken as a general framework, however, possibilism goes much further insofar as it demands that climate scientists always remain silent about probabilities. The general form of the position is a radical one and, given the poverty of the arguments in its favor, we should prefer instead the assumption that in at least some cases, groups of climate models provide us with probabilistic information about the actual world.

## References

- Annan, James D. and Julia C. Hargreaves (2011). Understanding the CMIP3 Model Ensemble. *Journal of Climate* 24: 4529–38.
- Betz, Gregor (2007). Probabilities in Climate Policy Advice: A Critical Comment. *Climatic Change* 85.1-2: 1–9.
- (2015). Are Climate Models Credible Worlds? Prospects and Limitations of Possibilistic Climate Prediction. *European Journal for Philosophy of Science* 5.2: 191–215.
- Dethier, Corey (2022). When is an Ensemble Like a Sample? 'Model-Based' Inferences in Climate Modeling. *Synthese* 200.52: 1–20.
- Hannart, Alexis, Aurélien Ribes, and Phillippe Naveau (2014). Optimal Fingerprinting under Multiple Sources of Uncertainty. *Geophysical Research Letters* 41: 1261–68.
- Hausfather, Zeke et al. (2020). Evaluating the Performance of Past Climate Model Projections. *Geophysical Research Letters* 1 (47).
- Huntingford, Chris et al. (2006). Incorporating Model Uncertainty Into Attribution of Observed Temperature Change. *Geophysical Research Letters* 33.L05710: 1–4.
- IPCC (in press). *Climate Change 2021: The Physical Science Basis*. Ed. by Valérie Masson-Delmotte et al. Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge: Cambridge University Press.
- Katzav, Joel (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* 46: 228–38.
- Katzav, Joel and Wendy S. Parker (2015). The Future of Climate Modeling. *Climatic Change* 132.4: 475–87.

- Katzav, Joel et al. (2021). On the Appropriate and Inappropriate Uses of Probability Distributions in Climate Projections, and Some Alternatives. *Climatic Change* 169.15: 1–20.
- Knutti, Reto et al. (2010). Challenges in Combining Projections from Multiple Climate Models. *Journal of Climate* 25.10: 2739–58.
- Lenhard, Johannes and Eric Winsberg (2010). Holism, Entrenchment, and the Future of Climate Model Pluralism. *Studies in History and Philosophy of Science Part B* 41.3: 253–62.
- Parker, Wendy S. (2010). Whose Probabilities? Predicting Climate Change with Ensembles of Models. *Philosophy of Science* 77.5: 985–97.
- (2020). Evidence and Knowledge from Computer Simulation. *Erkenntnis* (online first).
- Parker, Wendy S. and James S. Risbey (2015). False Precision, Surprise and Improved Uncertainty Assessment. *Philosophical Transactions of the Royal Society Part A* 373.3055: 20140453.
- Risbey, James S. (2007). Subjective Elements in Climate Policy Advice. *Climatic Change* 85.1: 11–17.
- Schurer, Andrew P. et al. (2018). Estimating the Transient Climate Response from Observed Warming. *Journal of Climate* 31.20: 8645–63.
- Sedláček, Jan and Reto Knutti (2013). Evidence for External Forcing on 20th-century Climate from Combined Ocean-atmosphere Warming Patterns. *Geophysical Research Letters* 29.20: 1–5.
- Stott, Peter A. et al. (2006). Observational Constraints on Past Attributable Warming and Predictions of Future Global Warming. *Journal of Climate* 19.13: 3055–69.
- Tokarska, Katarzyna B. et al. (2020). Observational Constraints on the Effective Climate Sensitivity from the Historical Period. *Environmental Research Letters* 15.3: 1–13.
- Winsberg, Eric (2018). *Philosophy and Climate Science*. Cambridge: Cambridge University Press.