

Why simpler computer simulation models can be epistemically better for informing decisions

Casey Helgeson^a, Vivek Srikrishnan^a, Klaus Keller^{a,b}, Nancy Tuana^c

draft manuscript, July 3, 2019

ABSTRACT. For computer simulation models to usefully inform climate risk management decisions, uncertainties in model projections must be explored and characterized. Because doing so requires running the model many times over, and because computing resources are finite, uncertainty assessment is more feasible using models that need less computer processor time. Such models are generally simpler in the sense of being more idealized, or less realistic. So modelers face a trade-off between realism and extent of uncertainty quantification. Seeing this trade-off for the important epistemic issue that it is requires a shift in perspective from the established simplicity literature in philosophy of science.

1 Introduction

Computer simulation models are now essential tools in many scientific fields, and a rapidly-expanding philosophical literature examines a host of accompanying methodological and epistemological questions about their roles and uses (e.g., Beisbart and Saam, 2019; Frigg and Reiss, 2009; Grüne-Yanoff and Weirich, 2010; Jebeile, 2017; Weisberg, 2013; Winsberg, 2010, 2018). Climate science is one such field (Edwards, 2001), and questions about the interpretation and reliability of the simulation modes used to understand, attribute, and predict climate change have received considerable attention (e.g., Frigg et al., 2013, 2015; Lloyd, 2010, 2015; Lloyd and Winsberg, 2018; Oreskes et al., 2010; Parker, 2011, 2013; Petersen, 2012; Steele and Werndl, 2016; Thompson et al., 2016; Vezér, 2016).

One conspicuous feature of scientific discourse about the simulation models used in climate science, and in environmental modeling more broadly, is the attention given to where a model lies on a spectrum from simple to complex (e.g., Jakeman et al., 2006; McGuffie and Henderson-Sellers, 2001; Smith et al., 2014). While this attention to model complexity has

^aEarth and Environmental Systems Institute, Pennsylvania State University

^bDepartment of Geosciences, Pennsylvania State University

^cDepartment of Philosophy, Pennsylvania State University

informed some of the philosophical discourse on simulation modeling (e.g., Parker, 2010), its relevance for the literature on simplicity in science remains largely unexplored.

This literature on simplicity addresses whether and why *simpler* theories (or hypotheses, models, etc.) might be—other things being equal—*better* than complex ones. Different ways of unpacking “simpler” and “better” yield a diversity of specific theses, with correspondingly different justifications (see Baker, 2016; Sober, 2015, for details and history). A number of distinctly modern variants rest on mathematical theorems tying well-defined notions of simplicity to benefits such as *predictive accuracy* (Akaike, 1973; Forster and Sober, 1994), *reliability* (Vapnik, 1998; Harman and Kulkarni, 2007), and *efficient inquiry* (Kelly, 2004, 2007). Arguably more domain-specific appeals to parsimony include instances in phylogenetics (see Sober, 1988, 2015) and animal cognition (Sober, 2009, 2015; Clatterbuck, 2015).

Here we discuss a notion of simplicity drawn from scientific discourse on environmental simulation modeling and expound its importance in the context of climate risk management. The new idea that we bring to the simplicity literature is that simplicity benefits the assessment of uncertainty in the model’s predictions. The short explanation for this is that quantifying uncertainty in the predictions of computer simulation models requires running the model many times over using different inputs; simpler models allow this to be done more thoroughly and more rigorously because they use less computer processor time. (The quantification of uncertainty in light of present knowledge and available data should be clearly distinguished from the reduction of uncertainty that may occur as knowledge and data accumulate over time. We address the former, not the latter.)

While complexity obstructs uncertainty quantification, complex models may behave more like the real-world system, especially when pushed into Anthropocene conditions. So there is a trade-off between a model’s capacity to realistically represent the system and its capacity to tell us how confident it is in its predictions. Both are desirable from a purely scientific or epistemic perspective as well as for their contributions to the model’s utility in climate risk management. Whether simpler is better in any given case depends on details that go beyond the scope of this paper, but the critical importance of uncertainty assessment for addressing climate risks (e.g., Reilly et al., 2001; Smith and Stern, 2011) is why simpler models can be epistemically better for informing decisions.

In what follows, we introduce the relevant notion of simplicity and a way to measure it (§2). We then explain the link from simplicity to uncertainty quantification (§3), arguing that through this link, simplicity becomes *epistemically* relevant to model choice and model development (§4). We briefly discuss the resulting trade-off, highlighting the roles of non-epistemic values and high-impact, low-probability outcomes in mediating the importance of uncertainty assessment for climate risk management (§5).

2 Simplicity and run time

Environmental simulation models populate a spectrum from simple to complex, and attention to a model's position on this spectrum is a pervasive feature of both published research and everyday scientific discourse. All computational models idealize their target systems by neglecting less important processes and by discretizing a spatially and temporally continuous reality. What makes a model more *complex* is: explicit representation of more processes and feedbacks thought to operate in the real-world system (or more detailed representation of them), and/or greater resolution in the discretization (i.e., smaller grid size or shorter time step). Greater complexity can allow for a more realistic depiction of the target system, while simpler models must work with a more idealized picture.

Realistic depictions can provide benefits but come at a cost: complex models demand more computer processor time. Here we use the length of time needed to run the simulation model on a computer, or the model's *run time*, as a proxy for model complexity. Run time is, of course, a processor-dependent measure: a faster computer processor can run the same program in less time. But this will not hamper our discussion since we are ultimately concerned with between-model comparisons that can be relativized to fixed hardware without loss of import.^a Moreover, differences in processor speed are in practice relatively small (in the neighborhood of times two) when compared with differences in run time across models (factors of tens to trillions). Run time also depends on the timespan simulated. But this too is something we can hold constant across models in order to compare apples to apples.

Another feature of run time as a measure of simplicity is that it applies to models understood in the most concrete sense. Run time is not a feature of calculations understood abstractly but of a specific piece of computer code written in a specific programming language (and run on a specific machine). A consequence of this is that two pieces of code with meaningfully different run times can instantiate what is, in some sense, the same model. What that means for our discussion is that the trade-off we examine applies most unyieldingly to computationally efficient programming; inefficiently-coded models can, to a point, be sped up without sacrificing realism.

Run time can be contrasted with other concepts already implicated in discussions of simplicity's role in science, for example, the number of adjustable parameters in a hypothesis. Run time quantifies the amount of calculating needed to use the model, while number of parameters concerns the model's plasticity in the face of observations. Simulation models contain adjustable parameters, but their number is often poorly defined, since quantities appearing in the computer code might be fixed in advance for one application but allowed to

^aAn alternative, processor-independent way to quantify computational expense might be to count the number of floating-point operations (FLOPs) required by the model's compiled program. But this is not a common metric; modelers typically know their model's run time but not its FLOP count, and the processor-hour is the usual unit in which scientific computing resources are allocated and purchased. So we continue with run time.

vary in the next. Run time is typically insensitive to the choice of how many parameters to treat as adjustable in a given application.^b

To make the discussion more concrete (and point readers to further details) we introduce a small collection of models that together illustrate the simplicity–complexity spectrum in environmental simulation modeling. We choose from models that have been used to investigate the contribution to sea-level rise from the Antarctic ice sheet (AIS), a key source of uncertainty about future sea-level rise on time scales of decades to centuries (Bakker, Louchard, and Keller, 2017; Bakker, Wong, et al., 2017; DeConto and Pollard 2016).

The *Danish Antarctic Ice Sheet* model (DAIS) (Shaffer, 2014; Ruckert et al., 2017) is the simplest of four models we consider here. It represents the AIS as a perfect half-spheroid resting on a shallow cone of land—a highly *aggregate* representation in the sense that just a few numbers summarize a vast and varied landscape (the AIS is larger than the United States). DAIS represents several key processes governing ice mass balance, including snow accumulation and melting at contact surfaces with air, water, and land.

The *Building blocks for Relevant Ice and Climate Knowledge* model (BRICK) (Bakker et al., 2017; Wong et al., 2017), couples a slightly-expanded version of DAIS with similarly aggregate models of global atmosphere and ocean temperature, thermal expansion of ocean water, and contributions to sea-level from other land ice (glaciers and the Greenland ice sheet). Compared to DAIS, BRICK represents an additional AIS process (marine ice sheet instability), as well as a number of interactions with other elements of the global climate system, including feedback between sea level change and AIS behavior.

The Pennsylvania State University 3-D ice sheet-shelf model (PSU3D) (Pollard and DeConto, 2012; DeConto and Pollard, 2016) includes fewer global-scale interconnections than BRICK but a much richer representation of both AIS processes and local ocean and atmosphere interactions. PSU3D is a *spatially resolved* model of AIS in the sense that spatial variation in ice thickness and underlying topography are explicitly represented and incorporated into AIS dynamics. PSU3D represents many additional AIS processes (beyond DAIS), including ice flow through deformation and sliding, marine ice cliff instability, and ice shelf calving.

The last and most complex model we use to illustrate the simplicity spectrum is the *Community Earth System Model* (CESM) (Hurrell et al., 2013; Lipscomb and Sacks, 2013; Lenaerts et al., 2016), incorporating spatially resolved atmosphere, ocean, land surface, and sea ice components and allowing global ocean and atmosphere circulation to interact with the AIS. While dynamic (two-way) coupling of CESM with a full AIS model is still under

^bStrictly speaking, a model does not have a single run time but a range of run times, one for each parameter choice. Still, we can continue to speak sensibly of run time as a property of the model, since different parameter choices typically result in similar run times so long as time step and resolution are not treated as parameters. Because these two features contribute to our notion of simplicity, and because they are generally held fixed during calibration and projection, we see them as part of what defines the model. It's a question of how you individuate models; on our approach, same model entails same time step and resolution.

development (Lipscomb, 2017, 2018), recent work (Lenaerts et al., 2016) uses a static ice sheet surface topography to investigate *one aspect* of future AIS dynamics, the *surface mass balance* (net change in ice mass due to precipitation, sublimation, and surface melt).

The resolution and run time of the models described above are provided in Table 1. Figure 1 summarizes and illustrates the models with special attention to key differences in complexity that account for their different run times.

model	resolution, AIS (km)	resolution, atmosphere (km)	approximate run time (min)	reference
DAIS	n/a	n/a	.001	Ruckert et al. (2017)
BRICK	n/a	n/a	.1	Wong et al. (2017)
PSU3D	10	40 (regional)	25,000	DeConto and Pollard (2016)
CESM	110*	110 (global)	2,000,000,000	Lenaerts et al. (2016)

Table 1: Resolution and approximate run time of four simulation models. Run times are for 240,000-year hindcasts, which enable model calibration to incorporate key paleoclimate data. Model configurations are as per the reference. *Refers to resolution of CESM's land surface component used to calculate surface mass balance. Based on Ruckert et al. (2017); Wong et al. (2017); DeConto and Pollard (2016); Bakker et al. (2016); Lenaerts et al. (2016); UCAR (2016), and personal communication with Kelsey Ruckert, Tony Wong, and Rob Fuller.

3 Epistemic relevance for model choice

Having introduced a notion of simplicity and a way to measure it, we now turn to the benefits enabled by this kind of simplicity. The proximate consequence of using a simpler model is that shorter run time allows for more model runs. That, in turn, has consequences for what one can learn from the model. But we begin with the proximate step.

Given a computing budget of some number of processor-hours, a simple calculation of computing budget divided by run time yields a theoretical maximum for the number of times the model can be run. Figure 2a displays this reciprocal relationship for two example computing budgets. Each point along such a curve corresponds to a different model choice: moving from left to right, one trades away model complexity (run time) in exchange for more runs. (Because such plots become hard to read with larger numbers, it will be helpful to use a logarithmic scale on the axes; we introduce this visualization in Figure 2b.)

As per Figure 2, run time and a computing budget determine how many model runs can be carried out. This run limit in turn constrains what methods can be employed at key stages of a modeling study, including calibration and projection. Model *calibration* is a process of tuning,

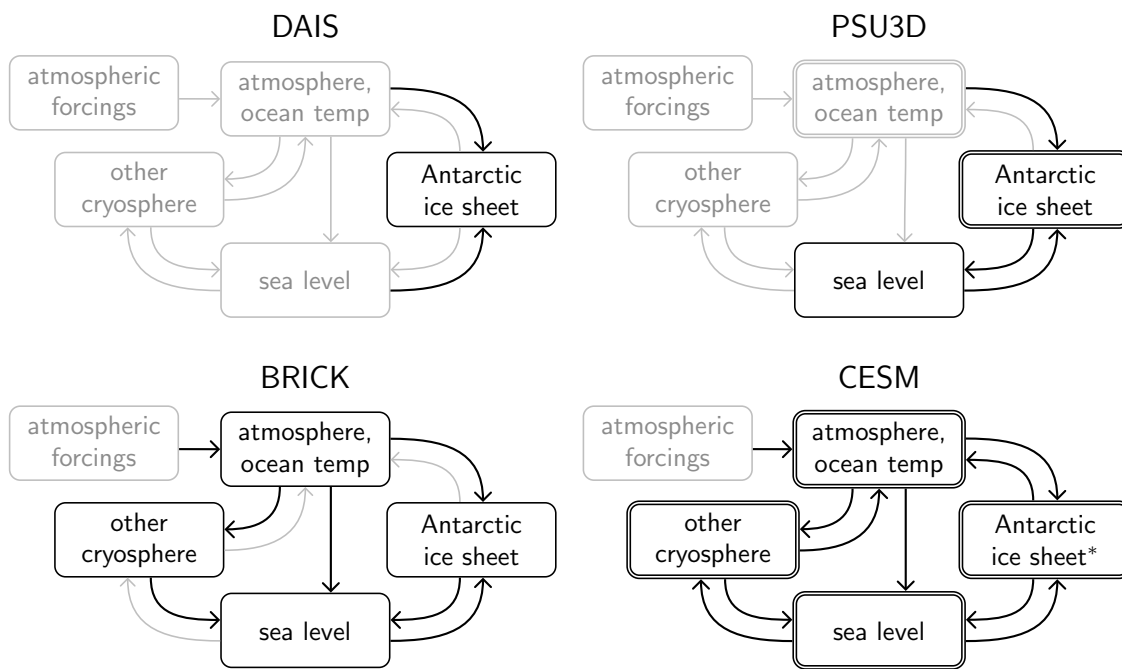
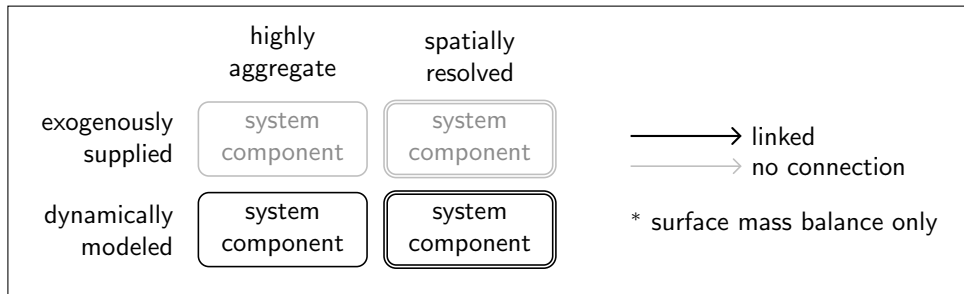


Figure 1: Schematic diagrams illustrating four environmental simulation models discussed in the main text. Based on configurations of DAIS, BRICK, PSU3D and CESM in Ruckert et al. (2017), Wong et al. (2017), DeConto and Pollard (2016) (Pliocene simulation), and Lenaerts et al. (2016), respectively. (Different configurations of the same model may correspond to somewhat different depictions within the visual schema used here.) CESM includes many additional system components not pictured.

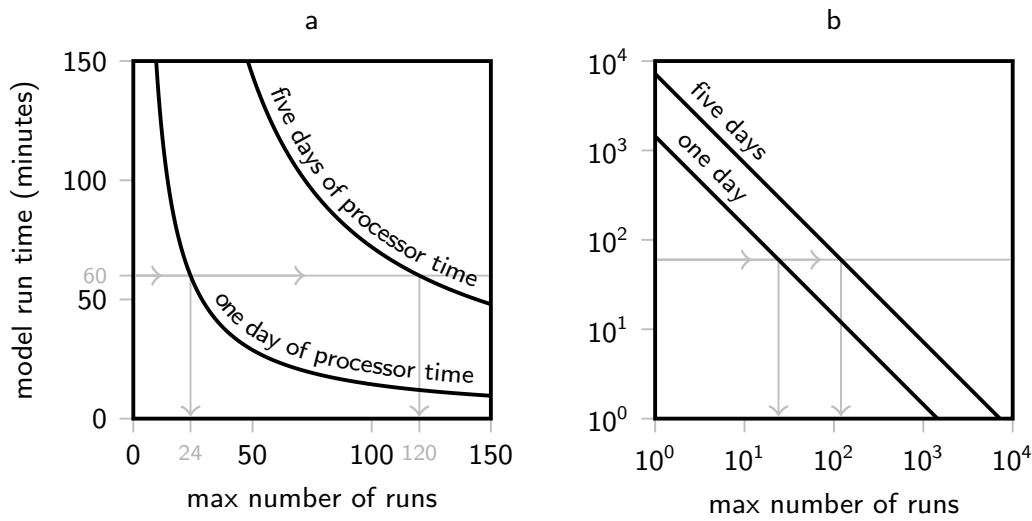


Figure 2: (a) Visual depiction of two computing budgets, illustrating the reciprocal relationship between run time and number of runs. Gray arrows illustrate how to read the figure: a model that runs in one hour (for example) can be run 24 times on a one-day budget and 120 times on a five-day budget. (b) The same figure, now plotted on logarithmic axes for a wider view (we use this format below to accommodate very large and very small numbers in one plot).

weighting, or otherwise constraining the values of adjustable parameters in order to make the model the best representation that it can be of the system under study. Subsequent *projection* involves running the calibrated model into the future to see what it foretells, conditional on assumptions about how required *exogenous* (supplied from outside the model) inputs will play out over the time frame in question.^c We discuss calibration and projection in turn, in each case detailing how the feasible number of model runs constrains the approach taken to these modeling tasks and what those constraints mean for the quantification of uncertainty.

3.1 Calibration

In the geosciences, model calibration typically aims to exploit both fit with data and prior knowledge about parameters (which often have a physical interpretation—see, e.g., Shaffer,

^cThough we used the word “prediction” in our broad-brush introduction, strictly speaking, we are concerned with projections. While projection is a subspecies of prediction construed broadly, the two are mutually exclusive on some finer-grained nomenclatures (Bray and von Storch, 2009; MacCracken, 2001). We use “projection” to emphasize the hypothetical nature of some assumptions—especially exogenous inputs such as future greenhouse gas emissions. Another way to put it is that by “projection” we mean a prediction conditional on certain inputs (like future emissions) about which the modeler explicitly makes no probability judgements.

2014; Pollard and DeConto, 2012). How this is done varies, and some methods require more model runs than others. To illustrate the dependence between model simplicity and calibration methods, we contrast three methods that together span the runs-required gamut. At one extreme lies *Markov Chain Monte Carlo* (MCMC), the gold standard in Bayesian model calibration (Bayesian inference is a natural fit for the task of integrating observations with prior knowledge). Near the middle is an approach called (somewhat confusingly) *precalibration*.^d At the other extreme lies *hand tuning*. We describe each below.

For the following discussion, assume a model that is deterministic and at each time step calculates the next system state as a function of the current state plus exogenous inputs (also called *forcings*) impinging on the system. Assume a set of historical observations, both of the forcings and of quantities corresponding to the model's state variables. Finally, assume a measure of fit between the time series of observations and the corresponding series of values produced by the model when driven by historical forcings (a *hindcast*).

Bayesian calibration begins with a prior probability distribution over all parameter combinations (the model's *parameter space*) and updates that prior in light of observations to arrive at a posterior distribution. In principle, the updating takes into account fit between those observations and every possible version of the model (every combination of parameter values). The posterior can therefore be calculated exactly (analytically) only where a suitably tractable mathematical formula maps parameter choices to hindcast–observation fit. This is not the case for computer simulation models, where fit can be assessed only by running the simulation and comparing the resulting hindcast with the observations. In this case, the posterior can still be approximated numerically, but this requires evaluating fit with observations (running the model each time) for a very large number of parameter choices. MCMC is a standard approach to numerically approximating Bayesian posterior distributions and requires tens of thousands to millions of model runs to implement (Kennedy and O'Hagan, 2001; Metropolis et al., 1953; van Ravenzwaaij et al., 2018; for application to ice sheet modeling, see Bakker et al., 2016; Ruckert et al., 2017).

In contrast to MCMC's thorough survey of parameter space, precalibration involves running the model at a smaller number of strategically sampled parameter-value combinations (e.g., 250, 1000, and 2000, in Srivier et al., 2012, Ruckert et al., 2017, and Edwards et al., 2011, respectively). Each of the resulting hindcasts is compared against observations using a streamlined, binary standard of fit—sorting the hindcasts into two classes: those *reasonably* similar to the observations and those not. The result is a dichotomous characterization of parameter choices as *plausible* or *implausible*, the latter unsuitable for use in projection.

The third method we highlight, hand tuning, encompasses a diverse set of practices that share some common features including varying parameters one at a time rather than jointly, using

^dThe approach can be applied as a preliminary step before further calibration (Edwards et al., 2011, §7), but here we discuss its use as a stand-alone method of calibration (as per Ruckert et al., 2017; Srivier et al., 2012).

different approaches to model–observation fit (or different observations altogether) for different parameters, calibrating sub-model components separately, a greater emphasis on expert assessment of parameter values, and a goal of identifying a single best parameter choice (for at least a majority of the parameters, sometimes for all). Hand tuning may not be clearly separated from model development, tends to be less transparent than the previous two approaches, and examines overall model behavior (and model–observation fit) for a relatively small number of parameter choices (very roughly, less than one hundred). Examples include Pollard and DeConto (2012) and Scheller et al. (2007); calibration of general circulation models (Meehl et al., 2007) and Earth system models such as CESM also generally fall into this category (Hourdin et al., 2017).

To summarize key points for present purposes, hand tuning samples the smallest number of possible parameter choices and aims to identify the best among them. Precalibration examines (on the order of) one hundred times more parameter choices and issues a dichotomous division of those into the plausible and implausible. MCMC examines (roughly) one thousand times more than *that* and exhaustively quantifies the relative plausibility of every possible parameter choice in the form of a probability distribution. Methods that demand more model runs do more to characterize uncertainty about parameter choices both by testing more possible values for the parameters and by furnishing a richer characterization of uncertainty about those values.

By limiting the feasible number of runs, model complexity undermines the characterization of parameter uncertainty. Figure 3 illustrates this point using the models from §2 and two example computing budgets. The lower diagonal line shows a ten-day budget (240 processor-hours). For researchers working on a single processor, this is a plausible limit on the computing time that can realistically be devoted to model calibration (in part since the calibration procedure might be repeated three to five times in the course of troubleshooting and replicating results). Points on or below this line are feasible on a 240 processor-hour budget. The figure shows that on this budget, DAIS can be calibrated by MCMC and BRICK by precalibration; PSU3D and CESM cannot be calibrated by any means.

With access to a high-performance computing cluster, much larger computing budgets can be contemplated. 400,000 processor-hours is a routine high-performance computing allocation in 2019 for research supported by awards from the United States National Science Foundation (UCAR, 2019). Since 400,000 hours would occupy a single processor for forty-six years, such a budget can be properly exploited only where the computing workload can be *parallelized* (split between multiple processors that run in parallel). Spread over 1,000 processors, 400,000 hours lasts a little over two weeks. While precalibration is easily parallelized, the most widely-used algorithm for implementing MCMC (Metropolis et al., 1953; Robert and Casella, 1999) requires that model runs be executed serially. There are, however, a number of kindred approaches to numerically approximating a Bayesian posterior, some of which can be substantially parallelized (Lee et al., 2019, and references therein).

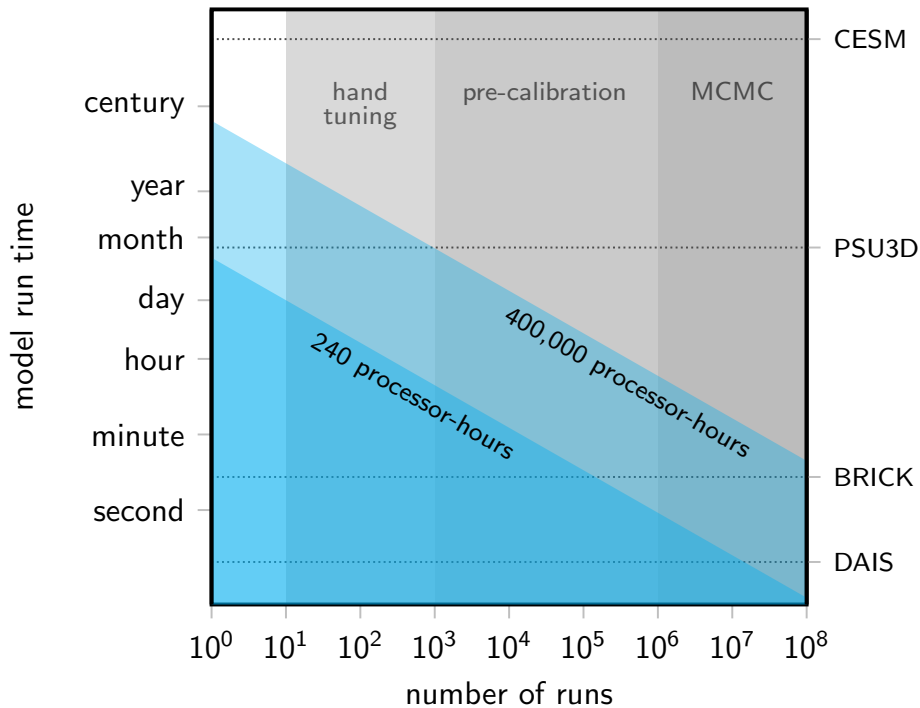


Figure 3: Example computing budgets compared with approximate run times for four simulation models (Table 1). Left boundaries of the shaded columns show approximate minimum run requirements for calibration methods discussed in the main text (supposing ~ 10 parameters are calibrated). The region below a computing-budget diagonal shows which calibration methods are feasible for each model on that budget. Figure expands on Bakker et al. (2016).

The upper diagonal in Figure 3 shows a 400,000-hour budget. With those resources, BRICK can easily be calibrated by numerical Bayesian methods and PSU3D moves to the edge of precalibration territory. A single hindcast using CESM is still well out of reach. Parallelization and large computing clusters substantially shift the goalposts, but they cannot dissolve the fundamental trade-off between model complexity and uncertainty quantification.

3.2 Projection

To the degree that parameter uncertainty has been characterized during calibration, it can then be propagated into projections. The most frugal approach to projection would be a single model run looking forward into the future. This can provide a best guess about future system behavior but does not offer any characterization of uncertainty around that guess. To do that requires additional runs using alternate, also-plausible parameter choices to generate

correspondingly also-plausible projections. A collection of different parameter choices leads to an *ensemble* of projections that can collectively characterize how uncertainty in parameter values translates to uncertainty in future system behavior.

The characterization of parameter uncertainty provided by MCMC (or other Bayesian numerical methods) allows for projection ensembles that are interpretable probabilistically (Lee et al., 2019; Ruckert et al., 2017; Wong et al., 2017). Precalibration allows for a dichotomous grading of plausibility in projected futures. Hand tuning offers little information about parameter uncertainty that could be propagated into a projection ensemble.

The size of the ensemble (i.e., number of projection runs) needed for high-fidelity propagation of characterized parameter uncertainty varies depending on many particulars, including the number of parameters calibrated. (Numbers in the thousands are typical, e.g., Ruckert et al., 2017, Wong et al., 2017.) By limiting the feasible number of runs, complexity can preclude projection ensembles of sufficient size. In this way, model complexity constrains not only the characterization of parameter uncertainty, but also its propagation into projections.

Moreover, parameter values are not the only uncertain inputs into model projections. Incorporating additional sources of uncertainty requires expanding the projection ensemble. Where uncertainty about *initial conditions* makes a meaningful difference to model projections, these conditions can be varied across ensemble members (e.g. Daron and Stainforth, 2013; Deser et al., 2014; Sriver et al., 2015). *Exogenous forcings* are often deeply uncertain and treated using a *scenarios* approach (Carlsen et al., 2016; Schwartz, 1996), multiplying the projection ensemble by the number of scenarios used (three to five is typical). The *model structure* (assumptions built into the model regardless of parameter choice) can also be questioned. Expanding the model structure (Draper, 1995) or explicitly characterizing model discrepancy (Brynjarsdóttir and O'Hagan, 2014) adds new parameters that, in turn, raise the run demands of both calibration and projection. Alternatively, repeating the entire work flow (calibration and projection) with several different models multiplies the total runs by the number of models used (supposing comparable run times).

The overall message on model runs and projection is that the more thoroughly one wishes to characterize uncertainty in projections, the larger the required ensemble. Specifically, more thorough characterization of uncertainty means that more of the assumptions built into the modeling have been questioned, with the consequences of questioning (varying) those assumptions having been propagated into projected system behavior. Jointly addressing multiple sources of uncertainty can lead to very large ensembles (e.g., ten million model runs, Wong and Keller, 2017).^e

^eThe computing demands of projection relative to calibration vary from one modeling exercise to the next, and depend not only on the number of runs used at each stage, but also the length of those runs. Regarding the AIS, sparse data and processes operating on geological time scales motivate calibration hindcasts of at least hundreds of millennia. The length of projections, on the other hand, often reflects policy planning horizons, and may extend only a century or two into the future. Since a model's run time is generally proportional to the

We have discussed two key phases in simulation modeling studies: calibration and projection. In each phase, the approach taken and the results obtainable are strongly constrained by a model's run time.^f The absolute numbers of runs needed for thorough characterization of known uncertainties can be very large, easily outstripping realistic computing budgets for more complex models. A fixed budget can therefore enforce a harsh trade-off between model complexity (and the realism it enables), and characterization of uncertainty in model projections. Another way to put it is that complexity limits what can be learned from the model. For this reason, simplicity is an important scientific, or epistemic consideration in model choice and model development.

4 Epistemic and non-epistemic benefits

We have argued that simplicity, measured via run time, is *epistemically* relevant to model choice. But some readers may be drawn to another framing of the issue, on which the value of this kind of simplicity is in fact not epistemic but merely *practical* (and therefore outside the traditional focus of the simplicity literature). When comparing the consequences of different model choices, one must hold something else fixed in order to structure the comparison. We hold the computing budget fixed, in which case simpler models enable better uncertainty quantification—a recognizably epistemic upshot. But you can make a different sort of comparison by holding something else fixed. For example, set in stone the desired approach to calibration and projection (including the number and length of model runs needed). Now the benefit of simplicity shows up in the processor time required to complete the envisioned research—which may sound like a practical matter rather than an epistemic one.

Analogous contrasting perspectives can be applied to the issue of *cognitive benefits* (such as ease of use), which are routinely dismissed as *non-epistemic* advantages of simplicity (e.g., Baker, 2016; Douglas, 2017; Kelly and Mayo-Wilson, 2010; Sober, 2015). One way to reach this dismissive conclusion is to assume a fixed research plan detailing the concrete steps to be taken within a research project (analogous to fixing the desired approach to calibration and projection). On this sort of comparison, the benefit of employing a simple, easy-to-use theory rather than a complex and burdensome one appears to be getting the proposed work done faster or with less effort (a seemingly non-epistemic benefit). On the other hand, a different sort of comparison can be made by supposing a fixed cognitive-effort “budget,” in which case easier use translates to more research completed, and as a result, more knowledge (or greater fulfillment of some epistemic value or other). Both perspectives are valid, each isolating and

number of simulated years, in this case one calibration run (hindcast) needs a thousand times more processor time than one projection run.

^fAnother important modeling activity (that we lack space to discuss, but where a similar lesson applies) is *sensitivity testing* (Banks, 1993; Sobol, 2001; Wong and Keller, 2017); also see the notion of *relevant dominant uncertainty* (Smith and Petersen, 2014).

illuminating one aspect of a bigger-picture bundle of trade-offs.⁸

For comparison, it is worth noting that the benefits of other, well-discussed notions of simplicity also admit of multiple framings, where one perspective highlights an epistemic upshot and another highlights a practical one. Notions of simplicity that concern the *flexibility* of a model or hypothesis get their epistemic relevance as a result of viewing the choice between simple and complex models while holding fixed the *quantity of data available*. AIC scores (see Forster and Sober, 1994), for example, give advice about which statistical model will yield more accurate out-of-sample predictions after fitting to data. AIC does this by rewarding fit with data while penalizing flexibility (number of parameters). But the influence of the parameter penalty decreases as the number of data increase, so the more data one has, the less simplicity matters. This means that if we instead hold fixed the goal of some desired degree of predictive accuracy, the benefit of simplicity will now show up in the quantity of data needed to achieve that goal—or more to the point, the time and expense of obtaining those data. As before, making a different sort of comparison pivots an epistemic consideration into one more naturally viewed as non-epistemic.

The fixed-data perspective is often *salient* because obtaining more data can be costly, slow, or otherwise impractical (and because the division of scientific labor often divorces statistical analyses from data gathering). But developments in the nature of scientific research have made our fixed-budget perspective equally salient. The growth of scientific computing has shifted work from brains to computer processors where it is more easily quantified and tracked. The complexity of computational simulation models has shadowed the exponential growth of computing power, massively increasing the calculating required to answer even simple questions using a model. At the same time, run-hungry statistical computing methods for calibration and projection of these simulations multiply the “cognitive effort” advantage of simpler models thousands to millions of times over, all within the scope of a single study or publishable unit of research. As a result, the trade-offs illuminated by contrasting modeling options on a fixed computing budget are now critical to a full understanding of the epistemology of simulation modeling.

5 Purpose and values

Simplicity facilitates uncertainty quantification, but complex models can be more realistic and may behave more like the real-world system. How much complexity is the right amount? The

⁸Acknowledging different ways of making such comparisons (by holding different things fixed) may help clarify a disagreement on whether using streamlined, less-reliable scientific methods in resource-constrained regulatory contexts illustrates non-epistemic factors intruding on method choice (Elliott and McKaughan, 2014; Steel, 2016a). If the resource budget is framed as a part of the give and take, then yes, epistemic considerations have been traded against non-epistemic. But if that budget is viewed as a fixed, external constraint, then the methods trade-off pits quantity against quality of scientific results, both of which are recognizably epistemic.

question raises challenging scientific and technical issues requiring deep, case-by-case integration of geoscience, statistics, computing, and numerical approximation (issues that go far beyond the scope of this paper). But equally important is the general qualitative point that, like other aspects of model evaluation (Addor and Melsen, 2019; Haasnoot et al., 2014; Parker, 2009), much depends on the *purpose* of the modeling exercise. Simulation modeling to improve scientific understanding, for example, may demand realism and benefit little from uncertainty quantification. Informing decisions, on the other hand, often demands attention to uncertainty (Keller and Nicholas, 2015; Rougier and Crucifix, 2018; Smith and Stern, 2011).

Broadly speaking, risk assessment involves contemplating what outcomes might occur, how likely each is, and how bad each would be (Kaplan and Garrick, 1981). These components jointly characterize the risk associated with a given course of action. Thinking in terms of *probability* and *cost*, for example, risk might be expressed as *expected cost*. This is not to say that risk management requires probabilities (Dessai and Hulme, 2004; Lempert et al., 2013; Weaver et al., 2013), only that some sense of the plausibility of different outcomes is needed to assess and manage risk (and that probability estimates are a common medium for this). Because simplicity enables the required uncertainty quantification while complexity impedes it, the simplicity–complexity dimension of model choice strongly influences a model's adequacy for the purpose of supporting decisions.

In climate risk management specifically, the importance of simplicity is magnified by the role of high-impact, low-probability outcomes. The limits that complexity places on uncertainty quantification are particularly unfavorable to estimating the chances of *extreme* possibilities, or what are referred to (in probability terms) as the *tails* of a distribution (e.g., Sriver et al., 2012; Wong and Keller, 2017; Lee et al., 2019). But since these extreme outcomes (e.g., large and/or rapid sea-level rise) are also the most dangerous and costly, estimating their probability can be central to managing risks, and relatively small changes to their estimated probability can have an outsize impact on risk calculations and management strategies.

A study by Wong et al. (2017) serves to illustrate these points. The authors use a relatively simple model of the AIS and other contributors to sea-level rise (BRICK, §2), allowing for rigorous quantification of parameter uncertainty via MCMC, followed by multiple large ensembles to propagate that uncertainty into local sea-level rise projections for the city of New Orleans under each of several forcing (greenhouse gas concentration) scenarios. The resulting projections (plus other inputs and assumptions) allow for estimation of a site-specific, economically optimal levee height (such that building any higher costs more than the flood damage it would be expected to prevent) for each concentration scenario.

The simplicity of BRICK also enabled Wong et al. to characterize some model uncertainty by repeating the entire workflow for two different model configurations: one with and one without an additional (poorly-understood but potentially important) mechanism of ice sheet behavior labelled *fast dynamics* (DeConto and Pollard, 2016; Pollard et al., 2015). Focussing

on just one of the city's five levee rings and assuming a business-as-usual greenhouse gas scenario (RCP8.5, van Vuuren et al., 2011), the authors quantify the impact of this model uncertainty by confronting the base-case model's economically optimal levee with projections from the fast-dynamics model configuration. The result is an increase in estimated annual chance of flooding (seawater overtopping the levee) of one-half of one-tenth of a percent (from eight in ten thousand to thirteen in ten thousand). This seemingly small change adds \$175 million in expected flood damage between now and 2100. A levee 25 centimeters higher could prevent much of that damage, with estimated net savings of \$53 million.

To underline the key points of the illustration: simplicity can contribute to a model's adequacy-for-purpose by enabling quantification and propagation of parameter uncertainty into projections, estimation of probabilities for high-impact, low-probability outcomes, and characterization of deeper uncertainties (e.g., model structure, forcing scenario) by spelling out how alternative assumptions impact management strategies. Where complexity undermines such modeling activities, the model's adequacy-for-purpose suffers.

The broad-brush purpose of supporting climate risk management can be analyzed further in any particular instance to reveal specific non-epistemic concerns such as protecting livelihoods, preserving culture, and saving money and lives (Besette et al., 2017; CPRAL, 2017). By judging models in light of purpose while also viewing these motivating values as a part of that purpose, the simplicity–complexity dimension of model choice can be seen as a *coupled ethical–epistemic problem* (Tuana, 2013, 2017; Vezér et al., 2018) in which motivations and trade-offs encompass both epistemic and ethical values.

The prospect of ethical values motivating model choice may raise concerns about such values overstepping their proper role in science, and at this point our discussion links up with a large literature on ethical (or more broadly, non-epistemic) values in science (Douglas, 2009; Elliott, 2017), a portion of which addresses climate science specifically (e.g., Betz, 2013; Intemann, 2015; Parker, 2014; Steel, 2016b; Steele, 2012; Winsberg, 2012). Here we can only note this connection, leaving further exploration of the topic for future work.

6 Conclusion

Discussions of simplicity's role in scientific method and reasoning have often recognized a loose notion of cognitive benefit—or benefit in terms of cognitive effort—associated with simple theories or models. Yet this aspect of simplicity has largely escaped attention, at least in philosophical literature, either because the advantage is seen as self-evident and trivial, or because the upshot is judged a matter of convenience, not epistemology.

This convenience-not-epistemology verdict is a natural consequence of the practice (common in much traditional philosophy of science) of attending to formal relationships between theory

and data while idealizing away the messy human element in science. But for today's computer simulation models, the "effort" required to operate the model—now understood in terms of computing resources, not cognitive burden—is too consequential to neglect. Computing demands sharply constrain how a model can be used and what can be learned from it.

We have used the run time of a simulation model as a measure of the model's complexity: simple models run faster and complex models run more slowly. The importance of run time to the epistemology of computer simulation can be seen clearly by adopting what we have called a fixed-budget perspective: compare what can be achieved with a simpler model to what can be achieved with a more complex one on the same computing budget. On such a comparison, simplicity facilitates quantification of parameter uncertainty and propagation of this and other sources of uncertainty into model projections, including estimates of chances for low-probability, high-impact outcomes.

How much one values these benefits is a further question that is tied up with the purpose of a modeling activity. One purpose for which uncertainty assessment can be critical is informing climate risk management. One specific example is managing flood risk in coastal communities facing sea-level rise, but there are, of course, many others (e.g., Butler et al., 2014; Keller and Nicholas, 2015; Hoegh-Guldberg et al., 2018).

None of this takes away from the important purposes served by very complex—and maximally realistic—environmental simulation models, including advancing understanding of processes and their interactions across multiple scales, and expanding the range of model structures that can be explored by the research community as a whole. Our discussion highlights the high stakes and harsh trade-offs inherent in model choice and model development—and the central role of simplicity in prioritizing the various scientific and social benefits gleaned from environmental simulation modeling.

ACKNOWLEDGEMENTS. The authors thank Rob Fuller, Ben Lee, Kelsey Ruckert, and Tony Wong for discussions that informed the paper and Hayley Clatterbuck, Kevin Elliott, Ben Lee, and Elliott Sober for comments on a draft. This work was supported by the National Science Foundation through the Network for Sustainable Climate Risk Management (SCRiM) under NSF cooperative agreement GEO-1240507.

References

- Addor, N. and L. Melsen (2019). Legacy, rather than adequacy, drives the selection of hydrological models. *Water Resources Research* 55(1), 378–390.
- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In B. N. Petrov and F. Csaki (Eds.), *Proceedings of the 2nd International Symposium on Information Theory*, pp. 267–281. Budapest: Akademiai Kiado.
- Baker, A. (2016). Simplicity. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2016 ed.). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2016/entries/simplicity/> (accessed July 1, 2019).

- Bakker, A. M., P. J. Applegate, and K. Keller (2016). A simple, physically motivated model of sea-level contributions from the Greenland ice sheet in response to temperature changes. *Environmental Modelling & Software* 83, 27–35.
- Bakker, A. M., D. Louchard, and K. Keller (2017). Sources and implications of deep uncertainties surrounding sea-level projections. *Climatic change* 140(3-4), 339–347.
- Bakker, A. M., T. E. Wong, K. L. Ruckert, and K. Keller (2017). Sea-level projections representing the deeply uncertain contribution of the West Antarctic ice sheet. *Scientific Reports* 7(1), 3880.
- Bankes, S. (1993). Exploratory modeling for policy analysis. *Operations research* 41(3), 435–449.
- Beisbart, C. and J. J. Saam (Eds.) (2019). *Computer Simulation Validation: Fundamental Concepts, Methodological Frameworks, and Philosophical Perspectives*. Springer.
- Bessette, D. L., L. A. Mayer, B. Cwik, M. Vezér, K. Keller, R. J. Lempert, and N. Tuana (2017). Building a values-informed mental model for New Orleans climate risk management. *Risk Analysis* 37(10), 1993–2004.
- Betz, G. (2013). In defence of the value free ideal. *European Journal for Philosophy of Science* 3(2), 207–220.
- Bray, D. and H. von Storch (2009). “Prediction” or “projection”? The nomenclature of climate science. *Science Communication* 30(4), 534–543.
- Brynjarsdóttir, J. and A. O’Hagan (2014). Learning about physical parameters: The importance of model discrepancy. *Inverse Problems* 30(11), 114007.
- Butler, M. P., P. M. Reed, K. Fisher-Vanden, K. Keller, and T. Wagener (2014). Inaction and climate stabilization uncertainties lead to severe economic risks. *Climatic change* 127(3-4), 463–474.
- Carlsen, H., R. Lempert, P. Wikman-Svahn, and V. Schweizer (2016). Choosing small sets of policy-relevant scenarios by combining vulnerability and diversity approaches. *Environmental Modelling & Software* 84, 155–164.
- Clatterbuck, H. (2015). Chimpanzee mindreading and the value of parsimonious mental models. *Mind & Language* 30(4), 414–436.
- CPRAL (2017). Louisiana’s comprehensive master plan for a sustainable coast. Technical report, Coastal Protection and Restoration Authority of Louisiana (CPRAL). Baton Rouge, LA.
- Daron, J. D. and D. A. Stainforth (2013). On predicting climate under climate change. *Environmental Research Letters* 8(3), 034021.
- DeConto, R. M. and D. Pollard (2016). Contribution of Antarctica to past and future sea-level rise. *Nature* 531(7596), 591–597.
- Deser, C., A. S. Phillips, M. A. Alexander, and B. V. Smoliak (2014). Projecting North American climate over the next 50 years: Uncertainty due to internal variability. *Journal of Climate* 27(6), 2271–2296.
- Dessai, S. and M. Hulme (2004). Does climate adaptation policy need probabilities? *Climate policy* 4(2), 107–128.
- Douglas, H. (2009). *Science, policy, and the value-free ideal*. University of Pittsburgh Press.
- Douglas, H. (2017). Why inductive risk requires values in science. In K. Elliott and D. Steel (Eds.), *Current Controversies in Values and Science*. Routledge.
- Draper, D. (1995). Assessment and propagation of model uncertainty. *Journal of the Royal Statistical Society: Series B (Methodological)* 57(1), 45–70.
- Edwards, N. R., D. Cameron, and J. Rougier (2011). Precalibrating an intermediate complexity climate model. *Climate Dynamics* 37(7), 1469–1482.
- Edwards, P. N. (2001). Representing the global atmosphere: Computer models, data, and knowledge about climate change. In C. A. Miller and P. N. Edwards (Eds.), *Changing the atmosphere: Expert knowledge and environmental governance*. Mit Press.
- Elliott, K. C. (2017). *A tapestry of values: An introduction to values in science*. Oxford University Press.
- Elliott, K. C. and D. J. McKaughan (2014). Nonepistemic values and the multiple goals of science. *Philosophy of Science* 81(1), 1–21.
- Forster, M. and E. Sober (1994). How to tell when simpler, more unified, or less ad hoc theories will provide more accurate predictions. *The British Journal for the Philosophy of Science* 45(1), 1–35.
- Frigg, R. and J. Reiss (2009). The philosophy of simulation: Hot new issues or same old stew?

- Synthese* 169(3), 593–613.
- Frigg, R., L. A. Smith, and D. A. Stainforth (2013). The myopia of imperfect climate models: The case of UKCP09. *Philosophy of Science* 80(5), 886–897.
- Frigg, R., L. A. Smith, and D. A. Stainforth (2015). An assessment of the foundational assumptions in high-resolution climate projections: The case of UKCP09. *Synthese* 192(12), 3979–4008.
- Grüne-Yanoff, T. and P. Weirich (2010). The philosophy and epistemology of simulation: A review. *Simulation & Gaming* 41(1), 20–50.
- Haasnoot, M., W. Van Deursen, J. H. Guillaume, J. H. Kwakkel, E. van Beek, and H. Middelkoop (2014). Fit for purpose? Building and evaluating a fast, integrated model for exploring water policy pathways. *Environmental modelling & software* 60, 99–120.
- Harman, G. and S. Kulkarni (2007). *Reliable reasoning: Induction and statistical learning theory*. MIT Press.
- Hoegh-Guldberg, O., D. Jacob, M. Taylor, M. Bindi, S. Brown, I. Camilloni, A. Diedhiou, R. Djalante, K. Ebi, F. Engelbrecht, J. Guiot, Y. Hijikata, S. Mehrotra, A. Payne, S. Seneviratne, A. Thomas, R. Warren, , and G. Zhou (2018). Impacts of 1.5°C global warming on natural and human systems. In V. Masson-Delmotte, P. Zhai, H.-O. Pörtner, D. Roberts, J. Skea, P. Shukla, A. Pirani, W. Moufouma-Okia, C. Péan, R. Pidcock, S. Connors, J. Matthews, Y. Chen, X. Zhou, M. Gomis, E. Lonnoy, T. Maycock, M. Tignor, and T. Waterfield (Eds.), *Global Warming of 1.5°C. An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty*.
- Hourdin, F., T. Mauritsen, A. Gettelman, J.-C. Golaz, V. Balaji, Q. Duan, D. Folini, D. Ji, D. Klocke, Y. Qian, et al. (2017). The art and science of climate model tuning. *Bulletin of the American Meteorological Society* 98(3), 589–602.
- Hurrell, J. W., M. M. Holland, P. R. Gent, S. Ghan, J. E. Kay, P. J. Kushner, J.-F. Lamarque, W. G. Large, D. Lawrence, K. Lindsay, et al. (2013). The Community Earth System Model: A framework for collaborative research. *Bulletin of the American Meteorological Society* 94(9), 1339–1360.
- Intemann, K. (2015). Distinguishing between legitimate and illegitimate values in climate modeling. *European Journal for Philosophy of Science* 5(2), 217–232.
- Jakeman, A. J., R. A. Letcher, and J. P. Norton (2006). Ten iterative steps in development and evaluation of environmental models. *Environmental Modelling & Software* 21(5), 602–614.
- Jebeile, J. (2017). Computer simulation, experiment, and novelty. *International Studies in the Philosophy of Science* 31(4), 379–395.
- Kaplan, S. and B. J. Garrick (1981). On the quantitative definition of risk. *Risk analysis* 1(1), 11–27.
- Keller, K. and R. Nicholas (2015). Improving climate projections to better inform climate risk management. In *The Oxford Handbook of the Macroeconomics of Global Warming*, pp. 9–18. Oxford University Press, New York.
- Kelly, K. T. (2004). Justification as truth-finding efficiency: How Ockham’s Razor works. *Minds and Machines* 14(4), 485–505.
- Kelly, K. T. (2007). A new solution to the puzzle of simplicity. *Philosophy of Science* 74(5), 561–573.
- Kelly, K. T. and C. Mayo-Wilson (2010). Ockham efficiency theorem for stochastic empirical methods. *Journal of philosophical logic* 39(6), 679–712.
- Kennedy, M. C. and A. O’Hagan (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63(3), 425–464.
- Lee, B. S., M. Huran, R. Fuller, D. Pollard, and K. Keller (2019). A fast particle-based approach for calibrating a 3-D model of the Antarctic ice sheet. arXiv preprint arXiv:1903.10032.
- Lempert, R. J., S. W. Popper, D. G. Groves, N. Kalra, J. R. Fischbach, S. C. Bankes, B. P. Bryant, M. T. Collins, K. Keller, A. Hackbarth, et al. (2013). Making good decisions without predictions. Technical report, RAND corporation.
- Lenaerts, J. T., M. Vizcaino, J. Fyke, L. Van Kampenhout, and M. R. van den Broeke (2016). Present-day and future Antarctic ice sheet climate and surface mass balance in the Community Earth System Model. *Climate*

- Dynamics* 47(5-6), 1367–1381.
- Lipscomb, W. (2017). Steps toward modeling marine ice sheets in the Community Earth System Model. Technical Report LA-UR-17-21665, Los Alamos National Laboratory.
- Lipscomb, W. (2018). Ice sheet modeling and sea level rise. (lecture, NCAR CESM Sea Level Session, January 10, 2018).
- Lipscomb, W. and W. Sacks (2013). The CESM land ice model documentation and user's guide. Technical report, National Center for Atmospheric Research.
- Lloyd, E. A. (2010). Confirmation and robustness of climate models. *Philosophy of Science* 77(5), 971–984.
- Lloyd, E. A. (2015). Model robustness as a confirmatory virtue: The case of climate science. *Studies in History and Philosophy of Science Part A* 49, 58–68.
- Lloyd, E. A. and E. Winsberg (2018). *Climate Modelling: Philosophical and Conceptual Issues*. Springer.
- MacCracken, M. (2001, February). Prediction versus projection—forecast versus possibility. *WeatherZine* (Edition Number 26), <https://sciencepolicy.colorado.edu/zine/archives/1-29/26/guest.html> (accessed July 1, 2019).
- McGuffie, K. and A. Henderson-Sellers (2001). Forty years of numerical climate modelling. *International Journal of Climatology* 21(9), 1067–1109.
- Meehl, G. A., C. Covey, T. Delworth, M. Latif, B. McAvaney, J. F. Mitchell, R. J. Stouffer, and K. E. Taylor (2007). The WCRP CMIP3 multimodel dataset: A new era in climate change research. *Bulletin of the American Meteorological Society* 88(9), 1383–1394.
- Metropolis, N., A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller (1953). Equation of state calculations by fast computing machines. *The journal of chemical physics* 21(6), 1087–1092.
- Oreskes, N., D. A. Stainforth, and L. A. Smith (2010). Adaptation to global warming: Do climate models tell us what we need to know? *Philosophy of Science* 77(5), 1012–1028.
- Parker, W. (2014). Values and uncertainties in climate prediction, revisited. *Studies in History and Philosophy of Science Part A* 46, 24–30.
- Parker, W. S. (2009). Confirmation and adequacy-for-purpose in climate modelling. *Proceedings of the Aristotelian Society Supplementary Volume* 83(1), 233–249.
- Parker, W. S. (2010). Predicting weather and climate: Uncertainty, ensembles and probability. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 41(3), 263–272.
- Parker, W. S. (2011). When climate models agree: The significance of robust model predictions. *Philosophy of Science* 78(4), 579–600.
- Parker, W. S. (2013). Ensemble modeling, uncertainty and robust predictions. *Wiley Interdisciplinary Reviews: Climate Change* 4(3), 213–223.
- Petersen, A. C. (2012). *Simulating nature: A philosophical study of computer-simulation uncertainties and their role in climate science and policy advice*. CRC Press.
- Pollard, D. and R. DeConto (2012). Description of a hybrid ice sheet-shelf model, and application to Antarctica. *Geoscientific Model Development* 5(5), 1273–1295.
- Pollard, D., R. M. DeConto, and R. B. Alley (2015). Potential Antarctic ice sheet retreat driven by hydrofracturing and ice cliff failure. *Earth and Planetary Science Letters* 412, 112–121.
- Reilly, J., P. H. Stone, C. E. Forest, M. D. Webster, H. D. Jacoby, and R. G. Prinn (2001). Uncertainty and climate change assessments. *Science* 293(5529), 430–433.
- Robert, C. P. and G. Casella (1999). *Monte Carlo Statistical Methods*, Chapter 7, pp. 231–283. Springer.
- Rougier, J. and M. Crucifix (2018). Uncertainty in climate science and climate policy. In E. A. Lloyd and E. Winsberg (Eds.), *Climate Modelling: Philosophical and Conceptual Issues*, pp. 361–380. Springer.
- Ruckert, K. L., G. Shaffer, D. Pollard, Y. Guan, T. E. Wong, C. E. Forest, and K. Keller (2017). Assessing the impact of retreat mechanisms in a simple Antarctic ice sheet model using Bayesian calibration. *PLOS ONE* 12(1), e0170052.
- Scheller, R. M., J. B. Domingo, B. R. Sturtevant, J. S. Williams, A. Rudy, E. J. Gustafson, and D. J. Mladenoff (2007). Design, development, and application of LANDIS-II, a spatial landscape simulation model

- with flexible temporal and spatial resolution. *Ecological Modelling* 201(3-4), 409–419.
- Schwartz, P. (1996). *The Art of the Long View: Planning in an uncertain world*. Currency-Doubleday, New York.
- Shaffer, G. (2014). Formulation, calibration and validation of the DAIS model (version 1), a simple Antarctic ice sheet model sensitive to variations of sea level and ocean subsurface temperature. *Geoscientific Model Development* 7(4), 1803–1818.
- Smith, L. A. and A. C. Petersen (2014). Variations on reliability: Connecting climate predictions to climate policy. In M. Boumans, G. Hon, and A. C. Petersen (Eds.), *Error and Uncertainty in Scientific Practice*, pp. 137–156. Pickering & Chatto Publishers.
- Smith, L. A. and N. Stern (2011). Uncertainty in science and its role in climate policy. *Philosophical transactions of the Royal Society A: Mathematical, physical and engineering sciences* 369(1956), 4818–4841.
- Smith, M. J., P. I. Palmer, D. W. Purves, M. C. Vanderwel, V. Lyutsarev, B. Calderhead, L. N. Joppa, C. M. Bishop, and S. Emmott (2014). Changing how Earth system modeling is done to provide more useful information for decision making, science, and society. *Bulletin of the American Meteorological Society* 95(9), 1453–1464.
- Sober, E. (1988). *Reconstructing the past: Parsimony, Evolution, and Inference*. MIT press.
- Sober, E. (2009). Parsimony and models of animal minds. In R. W. Lurz (Ed.), *The Philosophy of Animal Minds*, pp. 237–257. Cambridge University Press.
- Sober, E. (2015). *Ockham's Razors*. Cambridge University Press.
- Sobol, I. M. (2001). Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Mathematics and Computers in Simulation* 55(1-3), 271–280.
- Striver, R. L., C. E. Forest, and K. Keller (2015). Effects of initial conditions uncertainty on regional climate variability: An analysis using a low-resolution CESM ensemble. *Geophysical Research Letters* 42(13), 5468–5476.
- Striver, R. L., N. M. Urban, R. Olson, and K. Keller (2012). Toward a physically plausible upper bound of sea-level rise projections. *Climatic Change* 115(3-4), 893–902.
- Steel, D. (2016a). Accepting an epistemically inferior alternative? A comment on Elliott and McKaughan. *Philosophy of Science* 83(4), 606–612.
- Steel, D. (2016b). Climate change and second-order uncertainty: Defending a generalized, normative, and structural argument from inductive risk. *Perspectives on Science* 24(6), 696–721.
- Steele, K. (2012). The scientist qua policy advisor makes value judgments. *Philosophy of Science* 79(5), 893–904.
- Steele, K. and C. Werndl (2016). The diversity of model tuning practices in climate science. *Philosophy of Science* 83(5), 1133–1144.
- Thompson, E., R. Frigg, and C. Helgeson (2016). Expert judgement for climate change adaptation. *Philosophy of Science* 83(5), 1110–1121.
- Tuana, N. (2013). Embedding philosophers in the practices of science: Bringing humanities to the sciences. *Synthese* 190(11), 1955–1973.
- Tuana, N. (2017). Understanding coupled ethical-epistemic issues relevant to climate modeling and decision support science. In L. C. Gundersen (Ed.), *Scientific Integrity and Ethics in the Geosciences*, pp. 157–173. the American Geophysical Union and John Wiley and Sons, Inc.
- UCAR (2016). CESM 1.2 timing table. Technical report, University Corporation for Atmospheric Research. <http://www.cesm.ucar.edu/models/cesm1.2/timing/> (accessed July 1, 2019).
- UCAR (2019). University allocations. Technical report, University Corporation for Atmospheric Research. <https://www2.cisl.ucar.edu/user-support/allocations/university-allocations> (accessed July 1, 2019).
- van Ravenzwaaij, D., P. Cassey, and S. D. Brown (2018, Feb). A simple introduction to Markov chain Monte-Carlo sampling. *Psychonomic Bulletin & Review* 25(1), 143–154.
- van Vuuren, D. P., J. Edmonds, M. Kainuma, K. Riahi, A. Thomson, K. Hibbard, G. C. Hurtt, T. Kram, V. Krey, J.-F. Lamarque, et al. (2011). The representative concentration pathways: An overview. *Climatic change* 109(1-2), 5.

- Vapnik, V. (1998). *Statistical Learning Theory*. Wiley, New York.
- Vezér, M., A. Bakker, K. Keller, and N. Tuana (2018). Epistemic and ethical trade-offs in decision analytical modelling. *Climatic Change* 147(1), 1–10.
- Vezér, M. A. (2016). Computer models and the evidence of anthropogenic climate change: An epistemology of variety-of-evidence inferences and robustness analysis. *Studies in History and Philosophy of Science Part A* 56, 95–102.
- Weaver, C. P., R. J. Lempert, C. Brown, J. A. Hall, D. Revell, and D. Sarewitz (2013). Improving the contribution of climate model information to decision making: The value and demands of robust decision frameworks. *Wiley Interdisciplinary Reviews: Climate Change* 4(1), 39–60.
- Weisberg, M. (2013). *Simulation and Similarity: Using Models to Understand the World*. Oxford University Press.
- Winsberg, E. (2010). *Science in the age of computer simulation*. University of Chicago Press.
- Winsberg, E. (2012). Values and uncertainties in the predictions of global climate models. *Kennedy Institute of Ethics Journal* 22(2), 111–137.
- Winsberg, E. (2018). Computer simulations in science. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2018 ed.), <https://plato.stanford.edu/archives/sum2018/entries/simulations-science/>.
- Wong, T. E., A. M. Bakker, and K. Keller (2017). Impacts of Antarctic fast dynamics on sea-level projections and coastal flood defense. *Climatic Change* 144(2), 347–364.
- Wong, T. E., A. M. Bakker, K. Ruckert, P. Applegate, A. Slangen, and K. Keller (2017). BRICK v0.2, a simple, accessible, and transparent model framework for climate and regional sea-level projections. *Geoscientific Model Development* 10(7), 2741–2760.
- Wong, T. E. and K. Keller (2017). Deep uncertainty surrounding coastal flood risk projections: A case study for New Orleans. *Earth's Future* 5(10), 1015–1026.