Evidence and Knowledge from Computer Simulation

Wendy S. Parker¹

Received: 28 April 2018 / Accepted: 3 April 2020 / Published online: 3 May 2020
© The Author(s) 2020

Abstract

Can computer simulation results be evidence for hypotheses about real-world systems and phenomena? If so, what sort of evidence? Can we gain genuinely new knowledge of the world via simulation? I argue that evidence from computer simulation is aptly characterized as *higher-order evidence*: it is evidence that other evidence regarding a hypothesis about the world has been collected. Insofar as particular epistemic agents do not have this other evidence, it is possible that they will gain genuinely new knowledge of the world via simulation. I illustrate with examples inspired by uses of simulation in meteorology and astrophysics.

1 Introduction

Computer simulations are ubiquitous in science and engineering today. They frequently serve as replacements for traditional experimentation and observation. Philosophers of science, however, disagree about the extent to which computer simulation, as a methodological practice, really is like experimentation and observation, with similar epistemic powers (see e.g. Barberousse et al. 2009; Beisbart 2012, 2018; Giere 2009; Humphreys 2013; Lusk 2016; Morrison 2009; Parker 2009, 2017; Winsberg 2010). Those who emphasize the differences have argued that, while computer simulation can provide us with knowledge of the world in the way that argumentation does—by helping us to recognize the implications of our existing knowledge—it cannot provide us with knowledge that goes beyond this; observation and traditional experimentation, by contrast, are capable of doing so, because in these practices we “gather experience” (Beisbart 2012, p. 425). On this view, observation and experimentation have the potential to provide us with what might be called *genuinely new knowledge* of the world, while computer simulation does not.

A closely-related but little-discussed question is whether computer simulation, like experimentation and observation, is capable of providing *evidence* for hypotheses about real-world systems and phenomena (hereafter, ‘hypotheses about the

✉ Wendy S. Parker
wendy.parker@durham.ac.uk

¹ Department of Philosophy, Durham University, 50 Old Elvet, Durham DH1 3HN, UK

world’). Intuitions pull in both directions. On one hand, weather forecasts produced by today’s computer models seem to provide good grounds for revising our beliefs about tomorrow’s weather. And scientists in a wide range of fields—chemistry, physics, meteorology, hydrology, sociology, medicine and more—at least sometimes seem to treat simulation results as evidence for claims about the world. On the other hand, if computer simulation is at bottom an attempt to calculate the implications of a set of modelling assumptions, then simulation results, like results derived from a scientific theory by hand, seem to be predictions rather than evidence; they are the kind of thing we might seek evidence *for*, by making observations of the real-world systems and phenomena referenced in those predictions.

Here I attempt to reconcile these competing intuitions, while also challenging the view that simulation is a methodology incapable of providing genuinely new knowledge of the world. I do this, first, by arguing that evidence from computer simulation is aptly characterized as *higher-order evidence*: it is evidence that there exists other evidence for a hypothesis h about the world. It is then a short step to the conclusion that it is possible for epistemic agents to gain genuinely new knowledge of the world via simulation: insofar as particular agents either do not *have* the other evidence e whose existence is indicated by simulation results, or else lack background knowledge in virtue of which e can be interpreted as bearing on h , they can obtain new information regarding h from those simulation results which, in some cases, might be sufficient for them to come to *know* that h . Thus, while it is true that simulation models do not provide information about the world that goes beyond that which is already implicit in their assumptions, particular epistemic agents—including even scientists and engineers using simulation models—might still gain genuinely new knowledge of the world via simulation.

The discussion proceeds as follows. In Sect. 2, I introduce some basic terminology related to computer simulation. In Sect. 3, I argue for the main theses of the paper: that computer simulation results can be evidence for hypotheses about the world, that such evidence is aptly characterized as higher-order evidence, and that it is possible to gain genuinely new knowledge of the world via simulation. In Sect. 4, I show that it is plausible that scientists today do sometimes obtain evidence and knowledge from computer simulation, with the help of two extended examples from meteorology and astrophysics. In Sect. 5, I consider what is special about computer simulation as a source of evidence. Finally, Sect. 6 offers some concluding remarks.

2 Computer Simulation

It will be helpful to first clarify some terminology related to computer simulation. In what follows, a *computer simulation model* is a computer program that is designed to iteratively solve a set of dynamical modelling equations, either exactly or approximately, following a particular algorithm.¹ The models that will be the focus of this

¹ One might object, as an anonymous referee did, that it is wrong to say that the computer program is a “model”, since a program is a list of instructions while a model is not. We of course can choose to reserve the label “computer simulation model” for the set of equations that the computer solves during

paper are ones that have particular real-world targets, i.e. particular real-world systems (such as earth's climate system) or particular instances of a real-world phenomenon (such as a set of financial crises that occurred in the twentieth century). A *computer simulation* occurs when the program that constitutes the computer simulation model is successfully executed by a computer; starting from an initial assignment of values to the variables of the modelling equations (i.e. initial conditions), the computer calculates values for a sequence of later times by repeatedly solving, or estimating solutions to, the dynamical equations. The computer's calculated values for those times are the *solutions* of the simulation.

In the simplest cases, the *results* of a simulation are the records of some or all of those solutions, stored in the computer's physical memory or on some external storage device. In other cases, solutions undergo additional processing after the simulation is completed, e.g. to correct for known model biases; the results associated with the simulation will then be records of some or all of the outputs of that additional processing. In either case, results are interpreted as claims about the modelled system, e.g. that it will or would have particular properties at a particular time. Not infrequently, simulation studies involve multiple simulations whose results are interpreted or analyzed together, as we will see in an example in Sect. 4.

A *computer simulation procedure* is a procedure for generating results using a computer simulation model. It includes (at least) steps for obtaining one or more sets of initial conditions for the model, for launching one or more simulations, and for mapping some or all of the solutions of the simulation(s) to claims about the world. For example, to generate weather forecasts for particular cities, a national meteorological office might follow a procedure that involves estimating initial conditions with the help of recent weather observations, launching a simulation of the atmosphere with those initial conditions, and mapping (e.g. with the help of bias corrections and other steps) some of the solutions generated during the simulation to estimates of temperature at particular geographical locations. A physical implementation of the various components of the procedure—the initial condition estimation, simulation and mapping algorithms implemented on particular digital computers—constitutes a *computer simulation system*.

A computer simulation system is *reliable* in a domain of application if, and to the extent that, the majority of results that it would produce in that domain are true (or are accurate enough, given an agent's tolerance for error), when interpreted as claims about the world. What is usually desired in practice is that a simulation system is *highly reliable*, i.e., that it would almost always produce accurate-enough results in the domain of application of interest. Note that a domain of application can be quite narrow. For instance, one might be interested in the reliability of a weather forecasting system for predicting next-day high temperatures for London Heathrow airport only, and for just the next calendar year. The system might be quite reliable when it comes to predicting these temperatures to within some acceptable margin of error,

Footnote 1 (continued)

a simulation. It is not uncommon in scientific practice, however, to use the term “computer simulation model” to refer to the computer program that incorporates such equations, as I do here.

but not nearly so reliable when the domain of application encompasses next-day high temperatures for e.g. all major European airports over the coming five years.² Note also that, whether a simulation system is reliable in a domain of application does not depend on how many times it is *actually* used. What matters is that, if it *were* used sufficiently many times to answer questions within that domain, then the majority of its results *would* meet some specified standard of accuracy.

3 Evidence and Knowledge from Computer Simulation

With these preliminaries, we now turn to three key questions: Can results from computer simulations constitute *evidence* for hypotheses about the world? If so, what *kind* of evidence? Can agents gain *genuinely new knowledge* of the world via simulation? In this section, I answer these questions in general terms, with very brief examples. In Sect. 4, I provide extended illustrations inspired by scientific practice.

3.1 Evidence from Simulation

Many accounts of evidence have been developed in philosophy of science and in epistemology (see e.g. Achinstein 2014; Kelly 2016). Some accounts make evidence relative to what an agent believes or knows, while others do not. Some allow that fact *e* is evidence for hypothesis *h* if it provides *some* support for *h*, while others require that *e* provides *strong* support for *h* if it is to count as evidence. And so on. In order to keep the analysis here manageable, I consider the evidential status of simulation results from the perspective of just two accounts of evidence. I do not mean to endorse either of these accounts as the best available, though each has something to recommend it; my aim is to show that simulation results can be evidence for hypotheses about the world under quite different accounts of evidence.³

The first account is non-agent-relative and requires strong support: some fact (i.e. true proposition) *e* is *evidence_i* for hypothesis *h* just to the extent that *e* is a good indication of *h*. Some *e* is a good indication of *h* just to the extent that it is likely that *e* would be the case if *h* were true and not likely that *e* would be the case if *h* were false. Thus, the fact that (*e*) [a person has a particular constellation of unusual medical symptoms] might be evidence_i that (*h*) [the person has a particular disease], because those symptoms are highly likely to occur among people who have the disease and not likely to occur otherwise. On this view, even before such a disease is known to exist—indeed even if nobody ever discovers it—the symptoms are

² The reliability attributed to a process generating a particular result—say, a weather forecast for London—can vary, depending on how the process is conceptualized (e.g. as one generating forecasts for London, or for London in the winter, or for London when the recent weather has been such-and-such, etc.). This is a well-known problem faced by process reliabilism in epistemology (see e.g. Conee and Feldman 1998; see Hájek 2007 for a general discussion of the reference class problem in philosophy of science). In what follows, I assume that a reasonable domain of application is specified, picking out a description of the simulation procedure as one generating results for that domain.

³ These two accounts also correspond to two different types of evidence discussed by Kelly (2016), which he calls *indicator* and *normative* evidence, respectively.

evidence of its presence. In the extreme, e is a perfect indicator of h . In a genuinely deterministic system, for instance, the fact that (e) [the system is in state s at time t_1] guarantees (and thus counts as particularly strong evidence, or perhaps even proof) that (h) [it will be in state s' at a later time t_2].

By contrast, a standard Bayesian account makes evidence relative to an epistemic situation and requires only that e provides some support for hypothesis h : fact e is *evidence_B* for h for agent A just in case learning e warrants an increase in the probability that A assigns to h , given A 's background knowledge (see e.g. Howson and Urbach 2007). Whether learning e warrants such an increase in probability is determined by Bayes' Theorem: $p(h|e) = p(e|h) \times p(h)/p(e)$, where these probabilities reflect A 's degrees of belief, given her background knowledge. Thus, if agent A 's background knowledge includes the fact that a given medical symptom is more likely to occur if one has a particular disease than if one does not, then the fact that (e) [she has that symptom] might be *evidence_B* for her that (h) [she has the disease]. By contrast, for a different agent A' , who does not have this background knowledge about the disease, the fact that (e) [he has that symptom] might not be *evidence_B* that (h) [he has the disease].

On either of these accounts, it is possible for computer simulation results to be evidence for hypotheses about the world. If a simulation system is used in a domain of application in which it will very reliably give results that are accurate to within a specified margin of error ϵ , then its result R , attributing to the target the feature r , can be *evidence_i* for (h) [the target has feature $r(\pm \epsilon)$].⁴ Similarly, R can be *evidence_B* for that same h for an epistemic agent who is confident that the simulation system has such reliability. For example, if layperson Jane learns from a forecast comparison website that a particular weather forecasting system very reliably predicts the next-day's temperature in her city to within a few degrees Celsius, then her learning that (e) [today the forecast system predicted 22 °C for tomorrow] could warrant an increase in the probability that she assigns to (and thus be *evidence_B* for) the hypothesis that (h) [it will be within a few degrees of 22 °C tomorrow].

3.2 What Sort of Evidence?

When a computer simulation result is evidence for a hypothesis h about the world, what sort of evidence is it? Is it different in kind from evidence obtained by observing and experimenting on real-world systems and phenomena? I suggest that it is generally different in kind, and that the distinction between first-order and higher-order evidence can help to articulate the difference.

First-order evidence regarding a proposition or hypothesis h is often characterized, roughly, as evidence that bears directly on h (Kelly 2016). The fact that (e) [there are drops of water hitting the window of my office] is first-order *evidence_i* that (h) [it is raining outside]. *Higher-order evidence* is "evidence about the existence,

⁴ This is analogous to the way in which a result R from a traditional instrument (e.g. a thermometer) is evidence for the hypothesis that the value of the measured parameter T (e.g. temperature) is within a given range, e.g. $T = R (\pm \epsilon)$, where ϵ reflects the limited accuracy and/or precision of the measuring instrument.

merits, or significance of a body of evidence” (Feldman 2009, p. 304; see also Christensen 2010). The fact that (e') [the expert pathologist tells a patient that the tumour in his liver is benign] is evidence that (h') [there exists first-order medical evidence that (h) [the patient’s tumour is benign]], though the patient might not know what that first-order evidence is. Higher-order evidence of this sort—evidence that there exists some first-order evidence for a hypothesis h —is often itself evidence for h : if the pathologist is highly reliable in determining whether tumours of the liver are benign in light of the medical evidence he collects, then his testimony in the case of the present patient can be evidence_i that (h) [the patient’s tumour is benign].⁵ In effect, the pathologist’s testimony serves as a proxy for the relevant medical evidence, which the patient either is unaware of or is not in a position to interpret.⁶

Evidence for a hypothesis about some real-world system that comes from observing or experimenting on that system is often aptly characterized as first-order evidence. For example, facts about water collected in rain gauges can be first-order evidence for hypotheses about rainfall; the results of randomized controlled trials can be first-order evidence for hypotheses about drug efficacy in a population; reports of bird sightings in an area can be first-order evidence for hypotheses about changing migration patterns; and so on.⁷ By contrast, when evidence for a hypothesis h about the world comes from computer simulation, it is often aptly characterized as higher-order evidence: it is *evidence that there is other evidence for h* , which has already been collected. This already-collected evidence e underwrites the choice of dynamical equations and/or initial conditions for the simulation that produced the result. In the simplest case, the salient, already-collected evidence e consists in observed facts about the state of the modelled system at a given time, which informed the choice of initial conditions for the simulation: these facts e are evidence_i for a hypothesis h about the state of that system at some other time, given that various regularities (i.e. law-like relationships) obtain; they are evidence_B for h for any agent A who has knowledge of those regularities, or of a close enough approximation to them. In the case of weather forecasting, for instance, the already-collected evidence e might

⁵ As Tal and Comesaña (2017) and others have discussed, not all evidence of evidence for a proposition p is itself evidence for p . However, the evidence of evidence for p that, according to the discussion below, is sometimes provided by computer simulation seems to be an unproblematic case.

⁶ Upon reflection, the distinction between first-order and higher-order evidence as presented above is not clear-cut. While the fact that the dial on my bathroom scale points to 60 kg seems to be first-order evidence regarding my current weight, one might argue that it is also higher-order evidence regarding my weight: it is evidence that there exists some other evidence regarding my weight, causally upstream from the pointer (e.g. the fact that a plate within the scale is displaced by a certain amount or that some other mechanical aspect of the scale was in a particular state). On this way of thinking, higher-order evidence of the existence variety—i.e. evidence that there exists other evidence for h —is ubiquitous. There are various possible responses. One could add conditions to further limit what counts as higher-order evidence. Or one could argue that some facts seem *only* aptly characterized as higher-order evidence for a hypothesis h ; facts about the results of inference or derivation from other facts—as in the case of the pathologists’ testimony—would be prime candidates.

⁷ There are exceptions. Increasingly, computational tools are embedded in experiments and observations; in some of these cases, it might be more apt to characterize the results of the experiments and observations as higher-order evidence, with the inputs to the computations understood as first-order evidence.

consist of observed facts about recent weather conditions, which inform the initial conditions of the forecast simulation. This case is discussed further in Sect. 4.

3.3 Genuinely New Knowledge from Simulation

Analyses of the concept of knowledge are at least as numerous as accounts of evidence (see e.g. Ichikawa and Steup 2018). I do not adopt any particular account of knowledge in this discussion, with the aim of accommodating various views. For purposes of illustration, I will sometimes employ simple accounts, e.g. knowledge as justified true belief, or as warranted credence above some threshold, or as true belief produced via a sufficiently reliable process, etc. It should become clear in what follows that, under a number of different accounts of knowledge, including more sophisticated accounts, it will be possible in principle for an epistemic agent to come to know that h via computer simulation.

More important for the present analysis is what it means for knowledge to be genuinely new. I assume that, if agent A obtains evidence e for hypothesis h , and as a consequence comes to *know* that h , then his knowledge of h is *genuinely new knowledge* just in case h was not already warranted as knowledge for him before he obtained e .⁸ Whether h was already warranted as knowledge for some agent A will depend on the account of knowledge adopted, but I will assume that a sufficient condition for some h to be warranted as knowledge for an agent A is that h is logically or mathematically entailed by what A already knows.⁹ Thus, if an agent already knows that his height is 1.8 meters, that his weight is 75.5 kilograms and that body mass index is defined by one's weight in kilograms divided by the square of one's height in meters, then he gains no genuinely new knowledge when he learns that his body mass index is approximately 23.3.

In the case of simulation then, for an epistemic agent A , knowledge of h gained via simulation will be genuinely new knowledge for A if her background knowledge prior to learning the simulation result did not already warrant h as knowledge. Recall layperson Jane, who consults the highly-reliable weather forecasting system and obtains evidence that (h) [it will be around 22 °C tomorrow]. Her background knowledge before consulting the simulation-based forecast presumably was not sufficient to have already warranted h as knowledge for her. So, if she does come to know that h after consulting the forecast—because she comes to have a justified true belief that h , or comes to have a warranted credence in h above some threshold, etc.—then she will have gained genuinely new knowledge of the world via simulation. Her epistemic situation is similar to that in which, in light of the reliable

⁸ I take it that some proposition p can be *warranted as knowledge* for an agent without the agent *actually knowing* that p , e.g. when the agent does not believe that p or even thinks that p is unlikely, because she does not recognize that her existing knowledge entails p .

⁹ A similar view of novelty is resisted by Lusk (2016) when discussing whether simulations provide novel empirical data, on the grounds that it is too demanding. I aim to show that agents can gain new knowledge from simulation *even with* a demanding standard for novelty.

pathologist's testimony, the medical patient gains genuinely new knowledge that he has a benign tumour.

Of course, knowledge of h that is not genuinely new in this sense might be “new” to an agent in some other significant sense. For example, it might be *psychologically new*, because the agent is unaware that his or her existing knowledge warrants h as knowledge. Knowledge that is “merely” psychologically new can be very valuable—in some cases just as valuable as if it were genuinely new—since the agent still becomes aware of new premises for reasoning and thus new grounds for action. While the focus of this paper is whether computer simulation can provide genuinely new knowledge of the world (because that is what is often denied in debates over the epistemic power of computer simulation), it is worth remembering that, even in cases where an agent gains “merely” psychologically new knowledge via simulation, that knowledge can be extremely valuable.

4 Evidence and Knowledge from Simulation in Science Today

It is not only laypersons like Jane but also scientists and engineers who can obtain evidence and, in some cases, genuinely new knowledge of the world via computer simulation. This is illustrated below by way of two extended examples that, though hypothetical, are inspired by and closely resemble real uses of simulation in the fields of meteorology and astrophysics, respectively. After presenting the examples, I consider how common it is that scientists and engineers today obtain evidence and genuinely new knowledge of the world via simulation.

4.1 Sam: Evidence and Genuinely New Knowledge

Suppose that Sam the meteorologist is involved in a city project that aims to issue public warnings when dangerously high summer temperatures and heat waves are coming. A key ingredient to the process will be forecasts from a particular state-of-the-art weather forecasting system developed at a national meteorological agency (see e.g. Sheridan and Kalkstein 2004 on such warning systems). Sam was not involved in the construction of the weather forecasting system, but he knows quite a bit about how systems like it work: starting from initial conditions that are estimated from recent observations of the atmosphere, they repeatedly estimate solutions to a set of dynamical equations grounded in theoretical knowledge from fluid dynamics, thermodynamics, cloud physics, etc.

Sam himself possesses some of this theoretical knowledge, since it was part of his general training as a meteorologist, but he lacks much of the detailed, specialist knowledge that was used in formulating the model; it has components representing different atmospheric processes (e.g. cloud and rain formation), each developed by scientists who are experts on those particular processes. Nevertheless, after consulting data on the forecast system's performance, Sam knows that, in recent years, it has almost always predicted summer high temperatures for his city to within a few degrees Celsius, a couple of days ahead of time. Let us assume, then, that Sam has

good reason to believe that the forecast system is highly reliable when it comes to that predictive task. Moreover, suppose that the forecast system really is reliable in this way.¹⁰

Then on either of the views of evidence discussed above, the fact that (*e*) [the forecasting system gave result *R* today, indicating that it will be 41 °C in 2 days' time], will count as evidence that (*h*) [it will be around 41 °C in 2 days' time]. *R* is a good indication that *h*, given the forecasting system's reliability, and, since Sam knows of this reliability, its giving result $R=41$ °C warrants an increase in his degree of belief that *h*.

Moreover, the fact that the forecasting system gave result *R* is aptly characterized as higher-order evidence: it is evidence that there exists first-order meteorological evidence that it will be around 41 °C in two days' time in Sam's city. This first-order evidence consists of facts about recent weather conditions, the observation of which informed the initial conditions of the forecast simulation. Facts about recent weather conditions can be evidence_i regarding soon-to-arrive conditions, if the latter evolve from the former in a way that can be described rather accurately with a deterministic set of equations. (It has long been the standard view in meteorology that the short-term evolution of atmospheric conditions can be so described; the core theoretical equations of dynamic meteorology are deterministic.) Likewise, those facts about recent weather conditions can be evidence_B regarding later conditions for someone who knows enough about the physical processes by which one state of the atmosphere evolves to another (i.e. who can formulate a sufficiently accurate set of dynamical equations).¹¹

Sam himself might not *have* this first-order meteorological evidence; he might be unacquainted with the relevant facts about recent weather conditions that constitute the relevant evidence_i (and that would be evidence_B for an agent with sufficient background knowledge) for the hypothesis that (*h*) [it will be around 41 °C in 2 days' time]. Nevertheless, he knows enough to recognize that the forecast system's producing result *R* is evidence that, *somewhere* among the observations of earlier meteorological conditions that were collected, there is evidence_i for *h*. The forecast system reveals to Sam what the available meteorological evidence indicates about coming temperatures in his city.

If the support for *h* provided by result *R* is sufficient for Sam to come to *know* that *h*—because he comes to have a justified true belief that *h* or a warranted credence in *h* above some threshold, etc.—then, like Jane the layperson, he will have gained *genuinely new knowledge* of the world via simulation, since his background meteorological and other scientific knowledge, while considerable, presumably was not

¹⁰ This is not implausible; see e.g. UK Met Office 2016

¹¹ Evidence for dynamical relationships among meteorological variables, knowledge of which informs the dynamical equations of the forecasting model, are also part of the 'total evidence' that is available. In the analysis given here, knowledge of those dynamical relationships is treated as part of the background knowledge of the meteorology experts who built the model (and of Sam too, insofar as he knows those relationships); facts about the recent state of the atmosphere, observations of which inform the initial conditions for the forecast, are understood as additional, newly-obtained evidence to be taken into account when forming beliefs about future weather. The dynamical relationships will be foregrounded in the second example in this section, involving Sarah.

sufficient to have warranted h as knowledge for him on its own.¹² (And of course even if the support for h provided by R isn't strong enough to warrant h as knowledge for him, it might be enough for him to advise city officials to issue a public warning for high temperatures, given what's at stake.)

4.2 Sarah: Evidence but No Genuinely New Knowledge

Suppose that astrophysicist Sarah wants to explain a surprising feature of objects in the Kuiper belt, a large collection of icy bodies in the region beyond Neptune. In particular, she wants to explain why an orbital parameter, f , tends to take values near zero for these objects. She hypothesizes that the gravitational pull of an undiscovered large, distant planet could be the cause (see Batygin and Brown 2016 for the scientific study on which this example is based). To test the quantitative plausibility of this explanation, she plans to conduct a series of numerical experiments: she will run computer simulations exploring how test particles (putative Kuiper belt objects) behave in the presence of known massive bodies in the solar system as well as a hypothetical additional planet, under different assumptions about the mass, distance, eccentricity of orbit, etc. of the additional planet. She will check whether Kuiper belt objects in these simulations tend to have $f \approx 0$.

Sarah knows that gravitational forces are what matter for her question, and she has good reason to believe that, for the level of accuracy she requires, the dynamics of interest can adequately be represented using Newtonian physics (i.e. relativistic effects and other forces can be ignored). She formulates a set of dynamical equations K accordingly. Since these equations cannot be solved analytically, Sarah implements an approximate version of K —call it K^* —into a computer simulation model and will rely on the computer to carry out the many calculations needed to estimate solutions using numerical methods; for each hypothetical additional planet considered, a few dozen simulations will be run, assuming different initial conditions for the Kuiper belt object in each case. The goal is for the results of interest from each simulation—results related to the value of f —to approximate the corresponding entailments of K (see Fig. 1), and thus to reveal whether, for some of these planetary scenarios, Kuiper belt objects would tend to have orbits with $f \approx 0$.

Suppose Sarah finds that, when planetary parameters are assigned values Φ in her simulations, many of the simulations do give results $\{R\}$ indicating $f \approx 0$. Let us abbreviate this finding as $K_\Phi \vdash (f \approx 0)$. If her simulation procedure is a highly reliable means of estimating (to within an acceptable margin of error) the values that

¹² I am assuming that Sam did not consult weather charts depicting conditions in the region for the preceding days and formulate a forecast of his own; he instead has some prior for the hypothesis that (h) *it will be around 41 °C in 2 days' time* that reflects his background knowledge of the climate of his region (e.g. it's rarely above 40 °C, even in the summer). Even if Sam had consulted weather charts and formulated his own forecast, whether the hypothesis that (h) *it will be around 41 °C in 2 days' time* would be warranted as knowledge for him would depend on the details of his background knowledge, on the extent of information consulted, etc. It is more plausible that, as a result of such independent forecasting efforts, at best a coarser-grained hypothesis like (h_c) *it will get hot in the coming days* would be warranted as knowledge for him.

Fig. 1 *Learning about entailment, via simulation.* K is a set of propositions about a real-world target, in the form of a set of mathematical equations. K entails an answer to a question of interest about the target, but it is unclear to an agent what that answer is. K is (usually imperfectly) implemented as K^* in a computer simulation model M . Running the model produces solution(s) S , which may need to be processed or otherwise transformed to arrive at result R , providing a putative answer to the question of interest about the target. The goal is for R to approximate the (unknown) answer entailed by K

orbital parameter f would take in planetary scenarios like those investigated, then her results $\{R\}$ are evidence_i that (h) [If there were an extra planet in the solar system with properties Φ , then Kuiper belt objects would tend to have orbits with $f \approx 0$]. The evidence that $\{R\}$ provides can be aptly characterized as higher-order evidence: it is evidence that there exists other evidence for h , some of which has already been collected. This time, the already-collected evidence is that which underwrites the formulation of the dynamical equations for the simulation model. It is difficult to precisely characterize this body of evidence, but it presumably includes, among other things, a wide range of facts about the past motions of planets and other masses in the solar system, which have been found to fit the predictions of Newtonian theory well enough (and which lead Sarah and other physicists to believe that the domain of applicability of Newtonian theory includes the sorts of planetary scenarios under investigation).

From a Bayesian perspective as well, an analysis of the evidence that $\{R\}$ provides to Sarah is more complex. The standard Bayesian view assumes that epistemic agents are logically/mathematically omniscient. Thus, if it is true that $K_\Phi \vdash (f \approx 0)$, then Sarah should immediately know this; she should have no need for computer simulation. Actual epistemic agents, of course, are not mathematically/logically omniscient. Before Sarah runs her simulations, she does not know that $K_\Phi \vdash (f \approx 0)$; she presumably has at most a moderate (prior) degree of belief that it is true.

Suppose, however, that Sarah is confident, and for good reason, that her simulations are revealing, to within an acceptable margin of error, the relevant entailments of K . Her confidence in this might stem from her understanding of the mathematical techniques and approximations that she employed in building her computer

simulation model and from tests she has performed on the algorithm (see e.g. Oberkamp and Roy 2010; Winsberg 2010). In that case, the set of simulation results $\{R\}$, many of which indicate $f \approx 0$, will be evidence_B for her for the mathematical/logical hypothesis that $K_\phi \vdash (f \approx 0)$. But since, as noted above, Sarah is already confident that K adequately represents the relevant physics of the planetary scenarios she is investigating, when she obtains evidence_B from simulation that $K_\phi \vdash (f \approx 0)$, she is *at the same time* obtaining evidence_B that (h) [If there were an extra planet in the solar system with properties Φ , then Kuiper belt objects would tend to have orbits with $f \approx 0$]. She adjusts her confidence in h not by updating via Bayes' theorem but because she recognizes that, if she is confident that what K entails about the situation is accurate, and she has good evidence that K entails h , then as a matter of coherence she should be confident that h too.

Once again, the evidence here seems aptly characterized as higher-order evidence: simulation results $\{R\}$ constitute evidence for Sarah that *she already had evidence for h* , insofar as she had evidence that the dynamical equations of K adequately represent planetary systems like those being investigated. When she adjusts her confidence in h in light of $\{R\}$, she is in effect *correcting* it, so that it better reflects the import of the evidence that she already had. Insofar as this correction serves to align Sarah's confidence in h with her confidence that K adequately represents the planetary scenarios under investigation, she presumably will not come to *know* that h in light of the correction unless she *knows* that K adequately represents such planetary systems. Thus, if Sarah comes to know that h via her simulation study, that knowledge might be psychologically new, because she wasn't previously aware of the entailments of her existing knowledge, but it will not be *genuinely new knowledge*.

Note that Sarah's simulation results $\{R\}$ can be higher-order evidence not just for the hypothesis that an as-yet-undiscovered planet *could* be the cause of the surprising fact that Kuiper belt objects tend to have orbits with $f \approx 0$, but also for the bolder hypothesis that (H) [there actually exists an additional planet in the solar system, with properties Φ]. To see why, note that Sarah takes the implications of K_ϕ to be a close approximation to the implications of H . So when her simulations provide evidence_B that $K_\phi \vdash (f \approx 0)$, they in effect provide evidence_B for her that $H \vdash e$, where e is the surprising fact that [Kuiper belt objects tend to have orbits with $f \approx 0$]. Had this entailment relationship been known to Sarah before the discovery of e —as it would have been for someone logically/mathematically omniscient—then her learning of e would have increased the probability that she assigns to H , in accordance with Bayes' theorem. By providing Sarah with evidence_B of the entailment relationship, her simulation results $\{R\}$ seem aptly characterized as *evidence that she already had evidence_B for H* , in the form of e .

In this case too, it seems that Sarah ought to adjust her degree of belief—this time in H . But here we encounter the well-known 'problem of old evidence' for Bayesian confirmation theory: by the time Sarah becomes confident that $H \vdash e$ via simulation, she already knows that e , so she assigns $p(e) \approx 1$; but then H gains little or no boost in probability from e when updating in accordance with Bayes' theorem, even though, intuitively, the fact that H can account for surprising fact e should count in its favor. Various solutions to this problem have been proposed (see e.g. Sprenger 2015), including counterfactual approaches that recommend updating the

probability of H assuming $p(e)$ takes the value that one would have assigned before learning e and approaches that take the evidence for H in such situations to be the fact that that $H \vdash e$, rather than e itself. In effect, these proposals provide an agent like Sarah with guidance on how to correct her degree of belief in H , once she realizes that $H \vdash e$.

4.3 How Common Is It?

In the examples above, the scientists have good reason to be confident that the simulation systems they are consulting are highly reliable in the domains of application at hand. While these examples are inspired by real uses of simulation in science today, it is fair to say that such epistemic situations are not the norm; there is often significant uncertainty about the extent to which simulations provide accurate information about the world. There are various reasons for this. One is that scientists often lack the data they need to directly test the performance of simulation systems in domains of application of interest. This might be because the system or phenomenon of interest is spatially and/or temporally remote, or because the question of interest pertains to the behaviour of a system under boundary conditions quite different from those experienced so far. Conclusions about reliability are made more difficult insofar as these systems and phenomena are also nonlinear, complex, and/or not-well-understood theoretically.

Nevertheless, the examples involving Sam and Sarah are intended to make plausible that at least *some* computer simulation results today can be evidence for hypotheses about the world, and can be recognized as such; sometimes scientists *do* have trustworthy theoretical resources for building simulation models and *are* able to subject their simulation models to a range of tests on relevant data. Additional examples undoubtedly could be given, especially from engineering practice, where computer simulation is now commonly used to explore designs for products and for a range of other purposes. Moreover, in the examples involving Sam and Sarah, the simulation results constitute evidence even on a relatively demanding view of what counts as evidence (the good indication view). On less demanding views, simulation results will qualify as evidence more often. Under a standard Bayesian view, for instance, an increase in an agent's degree of belief in some hypothesis $h = R(\pm \epsilon)$ can be warranted by some result R even when it is thought to be only slightly more likely than not that the source will have given a result that is within ϵ of the truth. So results from today's computer simulations might relatively often constitute at least weak evidence_B for hypotheses about the world.¹³

¹³ Some philosophers have argued that, even when the reliability of a simulation system in a domain of application is largely unknown, its results might nevertheless constitute evidence that some state of affairs remains a *real possibility*, i.e. something that we don't have good reason to think couldn't occur (Katzav 2014; see Betz 2015 for complications). Detailed examination of such situations must be left for another time. Likewise, there is not space to consider here whether very simple models that are not intended to represent any particular real-world system—such as Schelling's checkerboard model—can provide evidence for hypotheses about the world and, if so, how this works (see e.g. Grüne-Yanoff 2013 for a related discussion). Both sorts of cases merit further investigation, however.

Similarly, how common it is that today's scientists gain genuinely new knowledge from simulation depends on the view of knowledge that one adopts, on what individual epistemic agents already know and on facts about today's computer simulation systems (e.g. their reliability). Given that the reliability of today's simulation systems is often significantly uncertain, it may be a relatively rare occurrence that today's scientists gain knowledge of the world via simulation and, ipso facto, a relatively rare occurrence that they gain genuinely new knowledge. Nevertheless, it is important not to overlook that there *can be* situations like Sam's—indeed, it is plausible that actual meteorologists and various other users of weather forecasts today often *do* gain genuinely new knowledge of the world via simulation.¹⁴

5 What's Special About Computer Simulation?

The perspicuous reader might have noticed that an analysis similar to that given in the last two sections can also be given for results derived from models or theories without the help of computers; these results too can be evidence for hypotheses about the world. For instance, if a procedure used to derive a result R from theoretical equations using pencil and paper is highly reliable (to within ϵ) for the domain of application at hand, then R can be evidence_i for a hypothesis $h = R(\pm \epsilon)$ about the world. Once again, the evidence seems aptly characterized as higher-order evidence—it is evidence that other evidence for h has been collected (which informed the choice of premises for the derivation). The same will be true of results derived with a handheld calculator or even in someone's head.¹⁵ It does not follow, however, that there is nothing distinctive or special about computer simulation as a means of obtaining evidence for hypotheses about the world.

Most notably, computer simulation is helpful when the body of first-order evidence for a hypothesis, and/or the background knowledge needed to interpret that body of first-order evidence, is *large, complex and distributed across multiple agents*. For some questions about the world that scientists would like to answer with the help of existing knowledge—such as questions about earth's climate system or about the results of particle collisions in the Large Hadron Collider—the body of knowledge relevant to the question is enormous and complex, spanning multiple fields of expertise. Even if an individual scientist possessed that body of knowledge, learning its implications in particular cases via paper-and-pencil calculation or

¹⁴ These simulation systems can be used to learn about past weather events too (e.g. diagnosing causal relationships). Thus even if one thinks that knowledge of the future is in principle impossible, there is almost certainly some knowledge of the world gained from these simulation systems.

¹⁵ For creatures like us, the application of mathematics or logic is always at the same time an empirical process; calculations or derivations are carried out using pencil and paper, or in a brain, or with a computer, or with a handheld calculator, etc. Facts about those empirical processes—that they culminated in particular inscriptions or brain states or computer states—can constitute evidence_i for hypotheses about other natural or social systems insofar as they are good indications of the truth/falsity of those hypotheses. Usually, when such facts about the results of derivation/calculation are evidence in this way, it is because the derivations/calculations started from propositions informed by knowledge of those real-world systems.

unaided reasoning would be practically infeasible. In many cases, however, no single scientist does possess the full body of existing, relevant knowledge; it is distributed in various ways among a broader scientific community. In such cases, computer simulation models not only serve to augment the computational powers of human beings (Humphreys 2004) but also serve as sites of community knowledge integration (Galison 1997; Gramelsberger 2011), where the implications of (one distillation of) collective knowledge can be explored. In situations like these, it may be that *every* individual scientist in the community can obtain genuinely new knowledge that h via the simulation, because none of them individually knows enough for h to be already warranted as knowledge.¹⁶

However, features of computer simulation also bring some distinctive epistemic challenges. Because of the complexity of many simulation algorithms, and because of the sheer number of calculations performed in a simulation, the process of obtaining evidence via simulation often has an *opacity* (Humphreys 2004) from the human agent's perspective that is not present when the agent works out the implications herself. Likewise, trust in the results of the simulation can be built in various ways, but not usually by having human agents verify that each individual step of the simulation was carried out as the model-builders intended. In addition, because simulation models are often the product of a range of ingredients, including simplified implementations of various scientific theories, numerical solution methods, compilers—and sometimes ad hoc assumptions and mathematical tricks as well—the process of determining whether a simulation system has succeeded in revealing (something close to) the implications of an existing body of knowledge, or provides evidence for some hypothesis h about the world, can be a much more *motley* affair (Winsberg 2001) than the process of determining the same for a pencil-and-paper calculation. It is precisely because of these and other distinctive features of computer simulation that its epistemology—especially how and when trust in simulation results can be justified and how uncertainties associated with results should be probed and characterized—has been of particular interest to philosophers of science (see e.g. Winsberg 1999, 2010; Parker 2008).

6 Conclusion

The foregoing goes some way toward reconciling the competing intuitions about computer simulation introduced at the outset: computer simulation is a methodology capable of providing evidence for hypotheses about real-world systems and phenomena, but evidence from simulation is of a different sort than the (first-order) evidence that is typically obtained from observation and experiment. In particular, evidence from simulation is more aptly characterized as a type of higher-order evidence: it is evidence that other evidence for a hypothesis h about the world has been collected.

¹⁶ The closer a situation is to this extreme, the more misleading it seems to say that no genuinely new knowledge of the world was gained via simulation, even though it is true that simulation models do not provide information about the world beyond that which is already implicit in their assumptions.

This already-collected evidence—a kind of ‘old evidence’—informs the construction of the simulation model, whether its dynamical equations, its initial conditions, or both. Put differently, running a simulation can help epistemic agents to see that some available information or data is *evidence for h*, though they might not have recognized it before.

If the evidence for *h* that a simulation provides is strong enough that an agent comes to *know* that *h*, then this will be genuinely new knowledge for the agent if *h* was not previously warranted as knowledge for them. This might happen, for instance, if the agent is unacquainted with the already-collected evidence whose existence is indicated by simulation, or if the agent lacks background knowledge in light of which that already-collected data or information can be interpreted *as* evidence for *h*. Such epistemic situations are common not just for laypersons but also for many scientists and engineers today, since the latter often use complex simulation models that they did not construct themselves and that reflect a range of highly-specialized knowledge. If it is relatively rare that such scientists and engineers do obtain genuinely new knowledge of real-world systems and phenomena from simulation, this is because of uncertainty about the reliability of particular simulation systems—and thus about the strength of the evidence they provide—rather than some inherent limitation of simulation.

This is not to deny that simulation models can only provide information about the world that is already implicit in their assumptions. For a hypothetical agent who somehow possessed in a coherent way the knowledge of a scientific community (or communities) as a whole, and who employed that knowledge to build reliable simulation systems, there would be no genuinely new knowledge of the world from simulation. In practice, however, many epistemic agents, including many individual scientists and engineers, are simply not in anything close to that epistemic situation. For them, it is possible to gain not only evidence but also genuinely new knowledge of the world from simulation. And it is plausible that, in some cases, they actually do so.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Achinstein, P. (2014). Evidence. In S. Psillos & M. Curd (Eds.), *The Routledge companion to philosophy of science* (pp. 381–392). New York: Routledge.
- Barberousse, A., Franceschelli, S., & Imbert, C. (2009). Computer simulations as experiments. *Synthese*, 169, 557–574.

- Batygin, K., & Brown, M. E. (2016). Evidence for a distant giant planet in the solar system. *The Astronomical Journal*, 151(2), 22.
- Beisbart, C. (2012). How can computer simulations produce new knowledge? *European Journal for Philosophy of Science*, 2(3), 395–434.
- Beisbart, C. (2018). Are computer simulations experiments? And if not, how are they related to each other? *European Journal for Philosophy of Science*, 8(2), 171–204.
- Betz, G. (2015). Are climate models credible worlds? Prospects and limitations of possibilistic climate prediction. *European Journal for Philosophy of Science*, 5(2), 191–215.
- Christensen, D. (2010). Higher-order evidence. *Philosophy and Phenomenological Research*, 81(1), 185–215.
- Conee, E., & Feldman, R. (1998). The generality problem for reliabilism. *Philosophical Studies*, 89(1), 1–29.
- Feldman, R. (2009). Evidentialism, higher-order evidence and disagreement. *Episteme*, 6(3), 294–312.
- Galison, P. (1997). *Image and Logic: A Material Culture of Microphysics*. Chicago: University of Chicago Press.
- Giere, R. (2009). Is computer simulation changing the face of experimentation? *Philosophical Studies*, 143, 59–62.
- Gramelsberger, G. (2011). Generation of evidence in simulation runs: Interlinking with models for predicting weather and climate change. *Simulation and Gaming*, 42(2), 212–224.
- Grüne-Yanoff, T. (2013). Appraising models non-representationally. *Philosophy of Science*, 80(5), 850–861.
- Hájek, A. (2007). The reference class problem is your problem too. *Synthese*, 156(3), 563–585.
- Howson, C., & Urbach, P. (2007). *Scientific reasoning: The Bayesian approach* (3rd ed.). Chicago: Open Court Press.
- Humphreys, P. (2004). *Extending ourselves: Computational science, empiricism and scientific method*. New York: Oxford University Press.
- Humphreys, P. (2013). What Are Data About? In E. Arnold & J. Duran (Eds.), *Computer simulations and the changing face of experimentation* (pp. 12–28). Cambridge: Cambridge Scholars Publishing.
- Ichikawa, J. J., & Steup, M. (2018). The analysis of knowledge. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Summer 2018 Edition). Retrieved February 20, 2020, from <https://plato.stanford.edu/archives/sum2018/entries/knowledge-analysis/>.
- Katzav, J. (2014). The epistemology of climate models and some of its implications for climate science and the philosophy of science. *Studies in History and Philosophy of Modern Physics*, 46(May), 228–238.
- Kelly, T. (2016). Evidence. In: E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Winter 2016 Edition). Retrieved February 20, 2020, from <https://plato.stanford.edu/archives/win2016/entries/evidence/>.
- Lusk, G. (2016). Computer simulation and the features of novel empirical data. *Studies in History and Philosophy of Science Part A*, 56, 145–152.
- Morrison, M. (2009). Computer simulation and the changing face of experimentation. *Philosophical Studies*, 143, 33–57.
- Oberkampff, W. L., & Roy, C. J. (2010). *Verification and validation in scientific computing*. Cambridge: Cambridge University Press.
- Parker, W. S. (2008). Franklin, Holmes and the epistemology of computer simulation. *International Studies in the Philosophy of Science*, 22(2), 165–183.
- Parker, W. S. (2009). Does matter really matter? Computer simulations, experiments and materiality. *Synthese*, 169, 483–496.
- Parker, W. S. (2017). Computer simulation, measurement and data assimilation. *British Journal for the Philosophy of Science*, 68(1), 273–304.
- Sheridan, S. C., & Kalkstein, L. S. (2004). Progress in heat watch-warning system technology. *Bulletin of the American Meteorological Society*, 85(12), 1931–1941.
- Sprenger, J. (2015). A new solution to the problem of old evidence. *Philosophy of Science*, 82(3), 383–401.
- Tal, E., & Comesaña, J. (2017). Is evidence of evidence evidence? *Nous*, 51(1), 95–112.
- UK Met Office. (2016). How our forecasts measure up. Retrieved February 20, 2020, from <https://blog.metoffice.gov.uk/2016/07/10/how-our-forecasts-measure-up/>.
- Winsberg, E. (1999). Sanctioning models: The epistemology of simulation. *Science in Context*, 12(2), 275–292.

Winsberg, E. (2001). Simulations, models, and theories: Complex physical systems and their representations. *Philosophy of Science*, 68, S442–S454.

Winsberg, E. (2010). *Science in the age of computer Simulation*. Chicago: Chicago University Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.