

Epistemic entitlements and the practice of computer simulation

John Symons and Ramón Alvarado

University of Kansas

Introduction

Arguably, computer simulation has become a third major branch of science, supplementing traditional theoretical and experimental research. Computer simulation has successfully permitted investigation into topics where cognitive, ethical, political, or practical barriers would otherwise prevent traditional experimental inquiry (See Symons and Boschetti 2013, 809). In important and sometimes controversial matters such as nuclear weapons testing, climate science, and studies of the behavior of epidemics, to take just a handful of cases, simulation using computers is often our only viable means of investigation. Because of its increasingly important scientific and policy role, there is an increased focus on how to determine the appropriate level of credence that ought to be assigned to the results of simulations (Boschetti et. al 2012). This is not a concern solely for the scientific and philosophical communities. Understanding the epistemic status of the claims produced using computer simulations presents an increasingly pressing practical challenge for citizens in democratic societies and for commercial decision-makers (See Vallor 2017).

In this paper, we assume that computer simulation has made positive contributions to scientific progress in a variety of fields. Astrophysics, climate science, epidemiology, fluid mechanics, and particle physics are some of the scientific disciplines in which computer simulations have played a crucial role advancing the state of our knowledge (Oreskes et al., 1994; Oreskes, 2004; Norton and Suppe, 2004) What we mean by computer simulation is broadly inclusive. On our view, computer simulation is the scientific or engineering practice of studying the behavior of actual or possible natural or social systems using software implemented in electronic devices. This software often

(but not always) relies on or approximates the behavior of mathematical models of the natural systems of interest. In simulations of dynamic process, which will serve as the focus of this paper, the algorithm coded into the software itself is intended to implement a function that takes as an input a representation of some initial state of the natural system and gives as an output the state of the system at some later time. Typically, the simulation repeats this operation on the outputs of previous iterations of the operation, thereby creating a sequence of representations that can shed light on the unfolding dynamics of the natural system.¹

In this paper, we explore the conditions under which one is justified to treat simulations as authoritative sources of knowledge. The question of when it is rational to trust computer simulations involves a range of distinguishable questions. Prominent among these include the following:

- How should we judge the results of simulation in areas of science where empirical evidence and data are relatively scarce?
- What justifies our credence in areas where we have abundant data but where no settled theory dominates the scientific domain?
- With respect to computational methods themselves, can we be confident that the software generating a simulation is free of error?
- Given the complexity of modern simulations and the increasing importance of machine learning in science, to what extent should we expect human understanding or humanly graspable explanations to play a role in science?

Taken in isolation, these are not completely new questions for philosophers of science. However, the epistemology of computer simulation brings them together in

¹ Some modeling (in geophysics for example) aims to explain static phenomena (modeling magnetic or gravitational fields). In other cases the use of closed-form equations may not need iterative solutions. While many of the points we discuss here also apply to non-dynamical cases, we focus on the dynamical cases here since they are most commonly discussed in the philosophical literature. We are grateful to an anonymous referee for reminding us to emphasize the static cases in which computer simulations are also used.

novel ways (Frigg and Reiss, 2009; Humphreys, 2009).² Moreover, the way that we respond to these questions is likely to have significant practical implications. In addition to shaping the practice of science and engineering, the way we understand the nature and justification of simulation will affect the priorities of funding agencies, as well as the reception of science by policy-makers. Thus, getting the epistemology of computer simulation right is crucially important.

While the questions listed above cannot be answered in detail in the space of a single paper, they all involve the problem of understanding how we are entitled to trust the results of computer simulations.³ In this paper we argue for a conservative position, namely that a high-level of trust in the results of computer simulations is only appropriate when they are grounded in well-curated data, plausible scientific theory, empirical evidence, and good engineering practice. Though these normative criteria may seem obvious, they are not always followed in practice and they are certainly not emphasized in philosophical discussions. Simulations that are not grounded in these ways, we argue, may turn out to be interesting sources of insight or inspiration, but they should not be accepted as guides in decision-making where potentially harmful or expensive risks are involved.⁴

Our general approach to the epistemology of simulation runs counter to two major tendencies in the current philosophical literature. The first attempts to justify our reliance on computer simulation by analogy with our ordinary epistemic practices. Specifically, it is widely held that in ordinary life a condition for the possibility of knowledge is that we must rely on the testimony of other people and on the reliability of our senses. Furthermore, in ordinary experience we must treat our basic cognitive capacities, such as our memory as capable of transmitting information without altering it in epistemically significant ways. Reasoning about the conditions for the possibility of ordinary epistemic judgments is valuable and productive, we do not object to this as an enterprise. However,

² Philosophical treatments of the epistemology of computer simulations are in an early stage of development and work on this topic so far has been covered well by Eric Winsberg in his *Stanford Encyclopedia of Philosophy* article (2015).

³ In this paper we give particular attention to the first two questions. Thorough discussions of the two last questions can be found elsewhere (Symons and Horner, 2014; Symons and Alvarado, 2016; Humphreys and Alvarado, 2017)

⁴ Policy regarding large scale institutional interventions, life-critical systems assessment, pharmaceutical tests and the outcome of otherwise untested medical procedures fall within this category.

one prominent strain in the philosophy of science incorrectly regards the justification of computer simulation as involving epistemic entitlements in a manner equivalent to the way that these entitlements operate in ordinary experience. This use of the idea of epistemic entitlement, largely drawn from Tyler Burge's arguments, has influenced philosophers in the contemporary debate over the epistemology of computer simulation.⁵ We will argue that epistemic entitlements for computer simulations cannot be recast in terms of the principles governing the behavior of individual human agents in everyday epistemic contexts.

Invoking everyday epistemic practices to justify our reliance on computer simulation ignores the qualitatively different standards governing our best scientific inquiry. Insofar as scientific inquiry aims to provide a deeper and more accurate understanding than we can access via unaided commonsense, the epistemic norms and standards guiding science are unlikely to be identical with those guiding ordinary commonsense perceptual judgments. As such, we are also critical of the tendency in the epistemology of computer simulation to take Burge's Acceptance Principle as a starting point for an account of epistemic entitlement with respect to computer simulation in science and engineering. We explain our disagreement with the use of the Acceptance Principle below and emphasize the qualitative differences between the norms governing everyday judgments in ordinary human experience and the norms that we argue should govern scientific practice.

A second major tendency that we object to is a pragmatist strain in discussions of computer simulation and modeling. We argue that accounts of epistemic entitlement for computer simulation simply cannot be grounded in straightforwardly pragmatic terms. There are a variety of ways that philosophers have attempted to ground simulation pragmatically. These range from the crudest forms of atheoretical or commercially-driven pragmatism to more sophisticated efforts by Andreas Kaminski and colleagues to provide fine-grained studies of simulation practice. The overall aim of this work is to reduce the epistemic burden on the use of simulation by emphasizing practical aspects of the craft of modeling and simulation. Kaminski and colleagues focus specifically on the practice of narrowing simulation tasks to tackling engineering problems composed of

⁵ See for example, Barberrouse and Vorms 2014, Beisbart 2017.

discrete, manageable measurement problems (See for example Hubig and Kaminski 2017). While we acknowledge that Kaminski and colleagues are offering an accurate description of current simulation practice, our purpose in this paper is to criticize the relatively impoverished normative role that pragmatism can play in the epistemology of computer simulation. Pragmatism is an unhelpful response to the most challenging philosophical problems arising from reflection on computer simulation for reasons we will explain here.

Our positive proposal is very straightforward. We argue that simulations by themselves are empty. Their use is justified only insofar as they are grounded on established theory, good engineering practice, well-curated data, and empirical evidence. As such, our account differs from pragmatism and is significantly more demanding than the conventional entitlement view that Burge and others advocate.

1. Why Pragmatism is an Inadequate Account of the Epistemology of Computer Simulation

Our discussion of epistemic entitlements for simulations begins with a thought experiment involving a mechanical oracle. As we shall see below, we imagine an oracle that produces reliable predictions but does not grant its users understanding or explanation concerning the system whose states it predicts or concerning its own inner workings. We should not, and (for the most part) do not treat simulations as oracles. As we shall see, warranted trust in the results of simulations is grounded in much more complex practices than simply deferring to a record of predictive success. The differences we note here between the mechanical oracle and the actual practice of simulation in contemporary science and engineering will help us to explain the limits of pragmatist and Burge-style approaches to the epistemology of computer simulation. We will explain, by contrast with the oracle, where we should look in order to understand how trust in actual simulations is and ought to be warranted.

One implication of our account will be that some prominent uses of “simulation” in the social sciences, for example, much of what is called agent-based modeling, may

provide interesting sources of insight but will not count as trustworthy guidance with respect to policy or other kinds of decision-making.⁶

1.1 The Tale of the Mechanical Oracle

Imagine, that once upon a time, there was a civilization in which people had outsourced the task of predicting and controlling their natural and social environments to black boxes with blinking lights. Let's imagine that the story of how these boxes appeared was lost in the mists of history, but some said that the boxes appeared suddenly, like the obelisks in Arthur C. Clarke's *2001: A Space Odyssey*. Unlike Clarke's obelisks, the boxes with blinking lights were talkative and responded to questions. At first, people might have approached them with trepidation; naturally, they would have been unsure of whether the boxes could be trusted. However, over time, those who trusted the boxes gained an advantage over those who did not. Eventually, the practice developed whereby almost everyone asked the boxes questions with the expectation of receiving predictively accurate answers.

While the mechanical oracles became sources of guidance for their users, the inner workings of these machines remained a mystery. When users asked the boxes for an explanation of their inner workings, the boxes simply refused (or were unable) to provide one. Users *assumed* that the boxes obeyed the usual physical laws and that their inner workings were in accord with mathematical principles of computation, but this assumption was simply part of their shared commonsense metaphysical background theory rather than having anything directly to do with what was known specifically about the artificial intelligences in the blinking boxes.⁷ Neither the details of the software

⁶ Nicole Saam (2017), for example, thinks that given the problem of epistemic opacity which can be a serious obstacle to the reliability of simulation outcomes, only those computer simulations that are generated and treated as *thought experiments* (as opposed to laboratory experiments) should be considered in the social sciences. Simulations treated as experiments have at their core, significant issues concerning access or lack thereof to all the relevant steps of their processes (Humphreys, 2009). Further, conventional error assessment strategies may be inadequate in contexts where software use is prominent (Symons and Horner, 2014). In contrast, simulations that are treated merely as thought experiments, according to Saam, provide "well-founded answers to what-if-things-had-been-different questions" (Saam, 2017 p.81)

⁷ Notice that explicit theoretical or justificatory assumptions from the part of the first people to trust the oracle are not a necessary component of their reliance on the oracle. They could have begun to trust the oracles by default, accident or even superstition (See Skinner's 1947 study called Superstition in the

engineering process that produced the oracles nor the machine learning algorithms that seemed to support their operation were accessible to the users of the artifacts.

The consensus among users was that the boxes rarely made serious errors with respect to important matters like the paths of hurricanes, the likelihood of epidemics, or the behavior of the economy. They were, by all relevant measures, predictively accurate. For the sake of argument, let us suppose that the users of the oracle are correct in their judgment of its predictive power. And thus, users posed questions to the oracle, took past successes as evidence that future answers were reliable, and granted the boxes significant influence in decision-making. Eventually, it was widely agreed that ignoring the oracle in matters involving important decisions would be unethical.

1.2 Are We Oracle-believers?

It is tempting to analogize the situation of the oracle-believers and our own relationship to at least some prominent computational methods. At first glance it seems that in several important contemporary decision-making contexts we already depend heavily on systems that fit the profile of our imagined mechanical oracles. Financial credit scoring algorithms are an unavoidable and more or less accepted presence in the lives of people in the United States. Likewise, risk assessment algorithms which compute the many factors that play a role in an individual's probability of repeated criminal behavior or data analysis techniques used to deploy policing resources in certain neighborhoods are epistemically opaque to those who use them (Symons and Alvarado, 2016; See also O'Neil, 2016).

It is often the case that we do not have sufficient understanding of the workings of these computational methods to evaluate each of the epistemically relevant steps in its operation. Paul Humphreys described systems of this kind as epistemically opaque (2009). Likewise, given the complexity of modern software, we cannot fully survey the operation of computational simulations, nor can we be sure that they operate correctly all the time (Symons and Horner, 2014; Horner and Symons, 2014). Little surprise then that

Pigeon. In it pigeons would continue to perform behaviors they equivocally associated with food rewards solely because the behavior and the distribution coincided successfully in the past).

reliance on software intensive systems in science has struck some thinkers as questionable and perhaps unscientific (Newman, 2015). If we are ignorant of the method by which the simulation generated its answer, then indeed it seems that we do not have rational basis to believe the answers. But this intuitive reaction is too hasty.

In scientific and commercial contexts it certainly seems as though we have already come to rely heavily on software intensive systems. Given that science and commerce are, for the most part, successful, it looks as though our faith in simulation is not unreasonable. However, we are not in the same position as the oracle-believers of our thought experiment. We would argue, for instance that the oracle-believers are not engaged in *science* in any standard sense of the word. Part of the task of the epistemology of computer simulation is to explain the difference between the contemporary scientist's position in relation to epistemically opaque computer simulations and the believers' relation to their mechanical oracles in the story above.

Scientists and engineers are often motivated by pragmatic considerations. However, if we restrict ourselves to pragmatic considerations alone we will not be in a position to explain the differences between the oracle-believers and us. As we will see below, pragmatic reasons should be distinguished from properly epistemic reasons.

Given something like the pessimistic meta-induction with respect to the truthfulness of science over the course of history, if our criterion for trust is the history of past successes or failures, then the oracle-believers are certainly more entitled to believe the blinking boxes than we are to believe our latest, best science. The pragmatist must defer to the oracle, since his criterion for deciding between theories is their relative predictive accuracy. While it is certainly true that predictive power is an important desideratum in choosing between theories or models, it is not the only dimension of scientific inquiry (See Symons and Boschetti 2013). Predictions can be accidentally true, predictions can be trivially true (Everyone alive today will be dead in 300 years), or predictions could be true by magic or clairvoyance. In none of these cases would we ordinarily call the sources of these predictions scientific.⁸

⁸ A position of this kind leads Quine to concede that he would have to count genuine clairvoyance as science. (1974) While genuine clairvoyance would be extremely interesting and useful, it would be a phenomenon that is largely orthogonal to science.

In cases like the oracle, in which the system generates reliably good predictions, most of us would wonder *why* the oracle works so well. Indeed, it is often the case, in the history of science that natural regularities and correlations become the *target* of explanation. Thus, for the scientifically-minded, the mysterious predictive power of the oracle would become the explanandum. For the pragmatist, since predictive power is the preeminent criterion of success, explaining predictive power could only be justified in terms of the pursuit of some additional, enhanced, predictive power.⁹ By stipulation, the oracle in our example has maximal predictive power. Therefore, in this case, there is nothing more for a pragmatist to do or to ask for. The inability for the pragmatist to explain the difference between the oracle and ordinary scientific practice is indicative of the weakness of pragmatism as a philosophical position.

One reason in favor of taking a pragmatic stance towards computer simulation is that, given the novelty of their use in scientific inquiry, current simulation practices are more like a craft than a science- in which simulation skills include an interpretative artistry to “reflect and make decisions based on the outcome of computer simulations” (Resch, Kaminski and Gehring, 2017 p.3). While the craft element of scientific practice is undeniable, its ultimate role in epistemic justification of the results of computer simulation to the consumers of scientific claims will be minimal. Typically, there will be situations in which practitioners in many areas of scientific research will develop skills that are not easily explained in explicit terms. Clearly, even in highly cognitive projects like teaching or scientific research, one develops skills that are difficult to articulate or justify verbally. For example, in some areas of applied mathematics, practitioners use heuristics and trial and error in a way that would be unacceptable in pure mathematics. One might argue for example, that the justification of techniques such as stochastic gradient descent relies on judgment and accumulated expertise.¹⁰ One could generate many examples from experimental science where the use and calibration of instruments

⁹As helpfully noted by a reviewer of this paper, in the context of computer simulation, other secondary criteria are also at play. These include Simplicity of the algorithm and greater unification of algorithmic calculations. However, these criteria will generally be subordinate to predictive power as a mark of success for the pragmatist.

¹⁰We thank one of our anonymous referees for drawing our attention to this issue.

involves skills that are not even explicitly teachable, but require something like a period of apprenticeship.

The cultivation of skill and judgment over the course of a career is part of what it means to be an expert and indeed in some areas it makes sense to rely on the judgments of experts simply in virtue of the fact that they are experts. Thus, for example in cases where an applied mathematician uses techniques that require tacit knowledge or judgment developed through practice to reach some conclusion, we could recognize and rely on this fact in the same way we might recognize that someone who successfully rode a bicycle to work arrived to work without their being able to articulate explicitly, the details of how one rides a bike. The rider's ability to ride a bicycle will not be entirely explainable by the rider, but the fact that he rode to work successfully can be regarded as undeniable and we can believe his claim to have made the journey by bicycle in virtue of knowing that he has the ability to ride a bicycle. We can grant all this and can acknowledge that expertise is often difficult to articulate. However, as we shall see later in the paper, the fact that experts are trustworthy in some contexts is not sufficient for justifying our faith in computer simulations, since simulations themselves are not credentialed as experts in the same way as human experts.¹¹ We can trust human experts in the use of techniques like stochastic gradient descent in applied mathematics because of their accumulated judgment and experience. However, as we will argue in more detail below, in the case of the computer simulation itself, consumers of scientific claims would be mistaken if they confused the expertise of modelers with the reliability of simulations – the simulation itself does not have the credentials that we would grant an experienced applied mathematician for example.

In response, the defender of expertise in these contexts might argue that expertise is relevant to the consumer of claims generated via simulation insofar as simulations can be certified by experts and experts are reliable sources of judgments about the reliability of computer simulation. In Section Three we will argue that endorsements from expert

¹¹ As we will see, even human expert credentials are not given out the same way we give out credentials to qualify as minimally reasonable interlocutors. I can trust my son, as Burge suggests, when he says that there is no more milk in the fridge based on a minimum intelligibility requirement and an assumption of healthy perceptual capacities. I wouldn't use the same criteria to choose a scientific expert.

practitioners are relevant to the evaluation of simulations, but they are by no means the source of epistemic justification for computer simulation.

Another broadly pragmatist defense of computer simulation sees it as a set of engineering tasks tackling one measurement problem at a time (Hubig and Kaminski, 2017). While this is an accurate characterization of some simulation practices and their deployment, it does not by itself constitute a substantive reason to adopt a pragmatist approach to the epistemology of computer simulation. While it is important to account for actual uses and practices of scientific inquiry, reporting on the practice does not account for the normative judgment concerning whether we should continue to use or expand the use of that practice. The pragmatic approach seems to presume that science using computational methods is so fluid that the methods and instruments themselves can only be assessed in an ad hoc manner. The focus on reading simulations as sets of isolated measurement tasks misses, for example, the role of carefully curated data and principled theoretical background knowledge in the judgment of the expert interpreter.

However, the problem here is not merely pragmatism's inability to provide satisfactory normative principles for the epistemology of computer simulation. By failing to capture the actual factors by which we rationally come to trust simulations, we risk neglecting the practical processes by which rational trust in science that depends on computational simulation ought to be cultivated in practice. Defenders of Big Data, for example, argued that the way that practitioners were successfully deploying data analysis tools already constituted proper scientific practice and that in many domains, theory and theoreticians were now useless (Steadman, 2013; Symons and Alvarado 2016). The dangers of big data hubris are now appreciated. An example of the pitfalls of atheoretical Big Data studies, is discussed by Lazer et al. (2014). Ultimately, advocates of Big Data were misguided in spite of being able to point to well-established practices, successful predictive patterns and expertise within the data science community. Modern computer simulation is in a far better situation than Big Data. It is a broad and diverse enterprise involving a range of distinct techniques and technologies with sensitivity to empirical and theoretical constraints. It therefore has many constraining and correcting influences from the scientists that engage with it (Roush, 2017) and as a result is not subject to some of the excesses of the Big Data hype of recent decades. Nevertheless, defending computer

simulation by reference to the well-established practices of or expertise of practitioners would be just as questionable as similar defenses of Big Data during its brief period of prominence.

2. The Reasonableness of Acceptance in Everyday Judgments

One way that philosophers of science have suggested explaining *why* we trust computer simulations appeals to the ways that epistemic entitlements work in other less controversial forms of inquiry (See for example Barberousse and Vorms, 2014). For example, it might make sense to trust simulations in the same way that we trust our perceptual faculties and our memories, or in the same way that we trust expert testimony. We are entitled to base beliefs on the evidence of our senses or the testimony of experts in spite of not having full access to either the underlying workings of the senses or the education and abilities of the expert. It seems that, for all ordinary purposes, one is warranted in believing that the universe is over one billion years old without knowing more than a few paragraphs worth of cosmology. Similarly, we can trust that we are correctly identifying the smell of freshly baked bread even though we might have no understanding of the operation of the olfactory system.

Perhaps something similar is the case for computer simulations. Let's call this the analogical approach to entitlements. Computer simulations are enormously complicated and in general they operate in ways that are not surveyable by ordinary human minds. Computer simulations are also epistemically opaque in some cases. Under these circumstances, some argue, it might still be reasonable to trust computer simulations given the assumption that everything is working smoothly (Beisbart, 2017). When this line of reasoning appears in the philosophical literature, it frequently draws on insights from Tyler Burge's work. In particular it takes as a starting point Burge's view of how human beings *ordinarily* maintain a posture of acceptance in epistemic matters. Just as we ordinarily trust our senses, we also ordinarily tend to accept what other people say unless we have reason to disbelieve them.

“Justification in acquiring beliefs from others may be glossed, to a first approximation, by this principle: *A person is entitled to accept as true something that is presented as true and that is intelligible to him, unless there are stronger reasons not to do so.* Call this the *Acceptance Principle*. As children and often as adults, we lack reasons not to accept what we are told. We are entitled to acquire information according to the principle-without using it as justification-accepting the information instinctively.”(1993 p. 467)

Burge points out that this ordinary disposition to accept the testimony of others is a necessary condition for acquiring language, and for a range of other social phenomena.¹² As far as science is concerned, failure to act in accordance with the Acceptance Principle makes rational collaborative projects of inquiry impossible. On this view, testimonial sources such as published research, standard pieces of scientific equipment, and expert opinion should be trusted by default as a precondition for the possibility of ongoing inquiry.

For Burge, acceptance is the epistemic default position for human beings and it grounds what he calls an *apriori entitlement* whereby a person is entitled to accept a proposition that is “presented as true and that is intelligible to him unless there are stronger reasons not to do so, because it is *prima facie* preserved (received) from a rational source, or resource for reason; reliance on rational sources- or resources for reason-is, other things equal, necessary to the function of reason” (1993 p.469). Given that the Acceptance Principle serves as a necessary precondition for the function of reason and therefore of science, it has been natural for philosophers influenced by Burge to imagine that a priori warrants can ground computer simulations. This is particularly the case if one sees computational methods as mere extensions to existing mathematical methods.

Burge’s reflections on computer- assisted mathematical proofs in the context of computer simulations emphasize the role of what he calls transparent conveyers of warrants. We will discuss the idea of transparent conveyers in more detail below, but for

¹² Burge’s account of the Acceptance Principle (as he acknowledges) is very similar in spirit to the Principle of Charity, as it figures in Quine (1960) and Davidson (1973). The principal difference between these principles is the role that Burge’s notion of preservation of content plays in his account.

now it suffices to note that a transparent conveyer is one that does not modify content in any epistemically relevant way. One immediate problem for the application of the idea of transparent transmission of warrants in simulation practices is that we cannot, in fact, be confident that computer simulations do not introduce epistemically relevant changes to the content being manipulated. As Eric Winsberg notes, simulation involves what he calls *motley* sets of practices and technologies (Winsberg, 2010). When we compare simulation to the process of computer-assisted proof, the latter clearly involves a more homogenous epistemic context (Barberousse and Vorms, 2014). Computer aided mathematical proof is a venue where one is more likely to be able to ensure the transparent transmission of warrants and the preservation of a priori content (Arkoudas and Bringsjord, 2007; McEvoy, 2008; 2013).

The more fundamental problem with adopting Burge's strategy as the basis for an account of computer simulation lies in its reliance on everyday epistemic stances with respect to the testimony of other people. As we shall see in the following sections, computer simulations and computational methods in general- as instruments deployed in scientific inquiry- are neither transparent conveyers, nor can they be regarded as equivalent to expert sources of testimony.

Clearly there are occasions in which one is entitled to a belief independently of one's subjective grasp of the epistemic rights or warrants supporting that belief. We can be entitled to believe, for example, what the weather forecaster has told us independently of the state of our knowledge of meteorology (Dretske, 2000; Williams, 2000; Adler, 2015). Given the epistemic entitlements described by Burge, we are frequently exempted from such justificatory practices as citing evidential support or of giving reasons more generally (Wright, 1985; Burge, 1993; Lackey, 1999; Dretske, 2000; Davies, 2004; McGlynn, 2014). In Sections Four and Five we will discuss when we are legitimately exempted from justificatory practice in more detail.

3. Computer Simulations are not Transparent Conveyers

Following Burge's application of the Acceptance Principle to the epistemology of computer-assisted mathematical proofs, Anouk Barberousse and Marion Vorms (2014)

argued that the epistemic warrants supporting our reliance on computer simulations are not intrinsically constituted or enhanced by appeal to empirical evidence. They argue that the warrants behind our trust in computational methods can and do take the form of Burge-style entitlements. Just like other devices whose role, like that of memory or perception, is to transmit and preserve intact information, computational methods that manipulate mathematical content appropriately are called *transparent conveyers* (Burge, 1993; Barberousse and Vorms, 2014). In a transparent conveyer if the propositions being conveyed are justified a priori, for example, then these justifications will not change from a priori to empirical.¹³ Following Burge, they call this *content preservation*.¹⁴ From this perspective, a warrant for belief can be a priori even if the only way one can attain the warrant depends on some contingent fact about the world, such as the fact of having a brain or some particular perceptual capacity. The fact that a human being needs a brain to do arithmetic is irrelevant to the justification of an arithmetical proposition. One implication of this position, is that a cognitive capacity such as memory “is no more intrinsically an empirical faculty than it is a rational faculty. Its function in deductive reasoning is preservative” (Burge, 1999)¹⁵. In other words, when memory does what it is supposed to do it conveys information without altering it.

However, as Burge acknowledges, when it comes to the transmission of information outside one’s own cognitive processes things are more complicated. Although Burge believes that there are similarities between the memory example and the way we can gather knowledge from other sources (humans or otherwise), he accepts that it is “only in special cases that a priori knowledge can be preserved through interlocution” (Burge, 1999 p.5).

¹³ This is sharp contrast to the view, called the materiality thesis (Morgan, 2005; Parker, 2009), which stipulates that the devices by which information is manipulated bear significant epistemic import.

¹⁴ Though they acknowledge that there is a substantial difference between computer assisted mathematical proofs, such as the ones Burge focused on and complex computer simulations in terms of content preservation, they justify the apriority of a scientist’s entitlement to trust a simulation in virtue of a second strategy, which we will inspect in detail in section three: trusting computer simulations, they argue, is like trusting expert testimony. Though expert testimony may be fallible, for this view, casting a general doubt on the practice absent specific reasonable doubt can be seen as irrational.

¹⁵ Similarly, when we rely on our senses we grant that when they are working the way they are supposed to they transmit information without altering it. That is, as explained above, they are transparent conveyers. Thus, though one can acknowledge their fallibility, in the absence of a plausible reason to doubt their well-functioning, it is rational to rely on the senses (Burge, 1993)

It is very important to note, for example, that Burge's account of content preservation and transparent conveying requires that the recipient already has reason not to doubt the source.¹⁶ Much of Burge's argument relies on the intelligibility of other minds, and the inescapable need for trust in testimony as a condition for the possibility of the most basic epistemic practices. However, it is important to recognize that scientific inquiry *is not a basic epistemic practice* but rather a very special cultural practice that is designed, in part, to overcome the evident limitations of our ordinary epistemic conditions.

Let us begin by pointing out how content preservation fails in the context of instruments deployed in the aid of inquiry. That is, how the warrants that underwrite scientific instruments are inadequate to warrant trust in the instruments themselves. Unlike perception and memory, or even the human tendency to trust others, the introduction of computational methods into scientific inquiry has been the product of artifacts and practices whose reliability has involved a gradual process including laborious efforts on the part of scientists and engineers (Winsberg, 2010). As Evelyn Fox-Keller points out, contemporary uses of computational methods in science are the product of massive collaborative efforts since the Second World War involving trial and error approaches to practical problems (Fox-Keller, 2003).

As Eric Winsberg argues, this historical development includes an array of distinguishable evidential benchmarks and other features such as considerations of fit, calibration, and even extra-theoretical and extra-mathematical engineering practices (Winsberg, 2010). These practices involve both hardware and software innovations that are well-known in practice but are seldom discussed in the epistemology of science. From architectural considerations for optimal processing to mathematical discretization, any trust assigned to computer simulations involves facts about the history of their successful deployment but also the appropriate management and assessment of uncertainties, errors and calibration procedures (Ruphy, 2015). These sanctioning processes, and therefore the computer simulation itself, should not be regarded as transparent in the sense required by Burge as we will explain below. The kind of homogeneity that transparent conveyers of

¹⁶ As we will see below, whether one has reason not to doubt, no reason to doubt, or reason to trust represent significantly distinct challenges for this prerequisite.

epistemic warrants require is simply not available. For example, given the role of engineering constraints, implementation, discretization, questions of fit, calibration, and countless other non-explicit features of computer simulation it would be a mistake to think that the warrants underwriting our trust in these artifacts are simply derived from their code.

Consider discretization techniques. Many computer simulations are the result of a process that includes the transformation of differential equations in a mathematical model into manageable step-by-step algebraic expressions that represent approximate values for specific finite spatiotemporal states of a system.¹⁷ Discretization techniques involve determining practical ways of implementing calculations in manageable chunks. This transformation/translation procedure introduces epistemically relevant decisions on the part of the modeler that are distinct from the original mathematical model. That is, the reason a specific set of equations is chosen in the *mathematical* model has nothing to do with the challenge of approximating the model via a computer simulation. That is, computational constraints such as whether a machine can handle the mathematical dynamics observed in the world need not be a constraint on mathematical models. By contrast, the choice of specific discretization techniques in a computer simulation will be responsive to the practical necessity of implementing a model in a digital device. In fact, when it comes to discretization techniques, the decision to select one technique over another is often a matter of engineering trade-offs (Winsberg, 2010 p.12, 23).

Discretization techniques, a fundamental aspect of computer simulations, undermine the possibility that computer models can serve as transparent conveyers insofar as they alter the nature of the epistemic justification of the content being manipulated. And this is so even when one grants that they are working as intended, as Beisbart (2017) suggests. Further, even if the reasons to trust the results of the equations in the original mathematical model are grounded in well-established theory, or are supported a priori on purely mathematical grounds, the introduction of discretization techniques involves an engineering element that alters the justificatory considerations

¹⁷ One can think of discretization as trying to approximate a circle by drawing one regular polygon after another with more sides each time starting from a square. Of course a square is a terrible circle, but a polygon with millions of sides may be visually indistinguishable for practical purposes. Nevertheless at each point, one is not drawing a continuous curve but rather a series of straight lines at an angle from each other.

involved in a computer simulation of the mathematical model. Thus, in a very basic sense, computer simulations cannot count as transparent conveyers given Burge's characterization because justificatory elements distinct from the ones warranting the original content do in fact enhance, decrease or constitute the epistemic warrants of the manipulated content (McEvoy, 2008).¹⁸ Furthermore, as we noted above, computer simulation involves independent epistemic warrants at different stages. Consider the following. Even if the mathematical model is fully warranted and works as intended and even if the discretized version of the model also works as intended there is no reason to think that we have reason to trust the latter *because* of the former.

When a discretized model is ultimately implemented in a device, for example, it requires yet another epistemically relevant transformation in the process of coding. In coding, considerations of fit, trust and/or reliability of a given algorithm will depend on independent factors from those involved in the discretization process. Unlike discretization techniques that involve established techniques and theories, code is often the result of highly idiosyncratic problem-solving approaches. Coding practices are error-prone. Consider for example the way many significant software bugs are 'patched'. That is, they are not 'fixed' per se, as in actively engaging with their malfunctioning components or erasing an erroneous block of code. Rather, code is added that 'patches' the original code by superseding previous functionality. These patches almost always introduce their own new bugs, making the process of assessing the reliability of the software even harder than it already was. In the process of patching, as is often observed, software, which is at the core of any and all processes of computer simulations, tends to grow over time merely because it can (Holzmann, 2015). This is especially the case now that memory has become so inexpensive in modern computing. The 'true' command in Unix, for example, which originally consisted of an empty file with nothing to execute grew to nearly 23,000 bytes from 1979 to 2012 (2015). The inevitability of errors in the practice of coding is an empirical reason to decrease trust in any content manipulated via

¹⁸ McEvoy in his response to Tymoczko and Kitcher on the apriority of computer assisted mathematical proofs concedes as much by saying "What determines whether a proof is a priori is the type of inferential processes used to establish the conclusion of that proof. If the method of inference for any of the steps in the proof is a posteriori, it is a posteriori" (2008 p.380).

computational methods. Thus, even if one grants that a computer simulation seems to work as intended, it cannot be regarded as a transparent conveyer in Burge's sense.

If, as Burge suggests (1999 p.4), one of the main issues to address is whether an entitlement or justification has any independent justification apart from empirical evidence, or as Barberousse and Vorms (2014) put it, whether the justification of a warrant is *in any way constituted or enhanced* by empirical means, then the answer concerning the warrants underlying our trust in computer simulations is clear. Typically, computer simulations are altered by and dependent upon empirical considerations. In fact, empirical and engineering constraints are often the reason that software for simulations is written the way it is.

In a similar vein as the views illustrated above and in contrast to our own position, Beisbart (2017) argues that warrant transmission takes place from one stage in the simulation process to the next. On his view, the results of computational methods can be said to provide knowledge in virtue of a few inferential transfers from one level of propositional content to another. That is, at each step, there is a result, a proposition to be considered: from target phenomenon to mathematical model, from mathematical model to computerized model, from computerized model to computational implementation, and from single implementation to iterations (Beisbart 2017 p.160-162). At any given step our trust is justified if we assume that our methods are the adequate ones and we assume they work as intended. In other words, we are warranted in inferring from one stage to the next if we grant that each level is somewhat/somehow warranted to begin with. For Beisbart an agent is "inferentially justified in believing a propositional result constructed from a computer simulation if she is justified in believing the dynamic equations used to feature the system under scrutiny and if she is justified to think that the simulation works as intended" (2017 p.171). However, for the agent to be justified in believing that the simulation works as intended, all that is required is that the "epistemic agent has sufficient reason to believe that it does so" (p.169). This "sufficient reason" in turn implies, the *assumption* that at every inference step, from mathematical model to final display, certain types of errors (round-off errors, modeling errors, hardware failures) are "excluded or are sufficiently small" (p.161-163). At every step however, trust is warranted with the proviso that there are no significant epistemic challenges (2017,162).

On Beisbart's account, a scientist is ultimately justified in believing a proposition derived from a computer simulation because what the computer itself has done is drawn inferences from the propositions of a discretized conceptual model (2017, p.165). The scientist draws the final inference that ties the results back to the target phenomenon by assuming that at any given step the methods involved *worked as intended*. Note that Beisbart recognizes the distinction between warrants supporting the content of simulations and warrants supporting our reliance on their results. His analysis highlights the empirical and historical considerations that ground our trust in computational methods.

However, as we have seen, the practices involved in simulation processes are such that they are not the product of transparent conveyers *even when they seem to work as intended*. The reasons why each of the steps can be said to be reliable and/or trustworthy are independent from the step before or after it, offering no single kind of warrant transmission that spans the entire process of computer simulation (Beebe, 2001; Davies, 2004; Pryor, 2012; Moretti and Piazza, 2013). As we will explain below, whatever reason we have for trusting the results of one stage in a simulation as we move to the next should generally not be ad hoc; should not be based solely on the factors that are unique to particular instances of a simulation practice itself but should be the product of established theoretical principles and engineering standards.

It is important to note that the methods and technologies used in practice are not simply deployed in an unchanged form across different applications and platforms. They are modified, often in non-trivial ways, to suit the task at hand. Throughout these modifications there are ways in which one can establish benchmarks that help to sanction the results of a simulation. These might involve, for example, repeated runs and internal comparison, comparison to outputs of other simulations, or most importantly, comparison to real-world data (Symons, 2008; Winsberg, 2010; Gramelsberger, 2011). Thus, when computer simulations can be trusted it is because of their adherence to theoretical principles, empirical evidence, or engineering best practices and not because of their output alone. Like other instruments deployed in scientific inquiry, computer simulations do not simply inherit their warrants in virtue of their manipulating formal syntax in a rule-governed manner. They also don't directly inherit them from the fact that theoretical

principles themselves, or the data analysis methods themselves, are warranted. Rather, just like instruments, they must invoke their own set of warrants (which in turn must include considerations of coherence with independently justified theoretical principles and empirical observations).

It must be noted however that unlike other instruments in science whose implementation and reliability has been a slow process, digital computers and computer simulations have had a relatively short period of research and development. And while some other derivative scientific instruments have had quick success in certain areas of inquiry (e.g. the electron microscope) computational methods are applied so ubiquitously that the assessment of their success is not as straightforward as it is for many more specifically targeted scientific instruments.

In addition to the importance of recognizing different kinds of epistemic practices in different stages of the construction of the simulation, it is also the case that even when the foundational components of the simulation are sound, the processes of combining them into a working application can fail in challenging ways. To put this point simply, we can say that the inference from warrants at the level of “parts” to warrants at the level of the “the whole” computer simulation is not sound. The reason that it is a mistake to regard an epistemic warrant as being successfully conveyed from relatively simple, “component-level” inferences to the behavior of the simulation as a whole, even in cases where the components of a software system are well understood, is because of the role of software engineering itself. Software, which is an essential part of any general purpose computing system, is a source of error insofar as it is created via human engineering and is built to serve human purposes. Human engineered software contains errors at the level of coding and vagueness at the level of specification (Tymoczko, 1979 p.74). These errors mean that normal notions of epistemic-warrant-transmission from elegant and well-understood foundations to the behavior of aggregates of these components simply do not hold (Winsberg, 2010). In other words, as Stéphanie Ruphy notes, “the computer simulation does not simply inherit” its epistemic credentials (2015 p.139) Rather, its reliability is rather the product of a range of diverse sanctioning processes.

In practice, our trust in computer simulation is fundamentally different from the epistemic relation of the oracle believers to their mechanical oracles in our thought

experiment. In themselves, computer simulations do not constitute good sources of testimony as we will argue in greater detail in the following section. As we have seen in this section, our trust in computer simulation takes place within a motley tradition of engineering and scientific practice. The actual process of building a simulation involves distinct stages, each of which has its own epistemic standards and warrants. This heterogeneous context makes the idea of transparent conveyers inapplicable to the epistemology of computer simulation.

We have argued that there is a non-trivial epistemic difference between formal systems and computer simulations. The epistemic status of computer simulations is more akin to scientific instruments, than to the results of a lengthy mathematical proof.

4. Computer Simulations are not Sources of Expert Testimony

This section distinguishes between ordinary epistemic thresholds involved in trusting others and those that ought to be deployed in trusting computer simulations as expert sources in scientific inquiry. According to Burge-style views of epistemic entitlement, when we are dealing with a given piece of testimony that we have *no plausible reason to doubt*, the default rational position is to accept it as truthful (Burge 1993). If, for example, I hear that weather models predict that a hurricane is likely to hit my city, it would be rational for me to heed the meteorologist's warning without rigorously investigating the methods and evidence supporting the prediction. In this scenario, it would seem that I am trusting the model in the same way that I would trust an expert, or perhaps even in the same way that the oracle believer in our thought experiment trusted the boxes with the blinking lights. This, however, is not the right way to understand the non-scientist's attitude towards the prediction. First, we would argue that in the weather forecasting case I do not trust the models as experts. Instead, I trust the judgment of the human meteorologists with respect to the models in question.

Second, while it is reasonable for me to regard the meteorologist as a reliable source of testimony (Borge, 1993; Jenkins, 2007), it would not be reasonable for the meteorologist him- or herself to maintain this epistemic attitude towards the computer simulation. It is reasonable for us, in ordinary life, to trust the weather forecaster insofar

as they have expert standing. Neither we, nor the weather forecaster should treat the simulation itself as an expert. Again, the simulation in itself is not self-validating. When things go well, *the community of experts and the tradition of using simulations in successful scientific practice grounds* the confidence of experts in the use of simulations.

A scientist may be justified in relying on the computer simulation's pronouncements once it is integrated into her practice or into the practice of other experts (Audi, 1997). As explained above, this is a long and gradual process. It is never the case that a scientist would give as a reason to trust a simulation the claim that the simulation is an expert. Treating our trust in simulations by analogy with our trust in human experts, would be to miss the actual warrants that ground the trust granted by human experts to simulations. When non-experts take the testimony of experts seriously, they are not granting that same level of credence to the tools used by the experts they trust. For example, they might not even know of the existence of those tools. In our case, most consumers of the weather forecast simply trust the experts to interpret computer simulations and other tools for them.

Justification to trust testimonial evidence comes from observing a "general conformity between facts and reports [by which] with the aid of memory and reason, we inductively infer that certain speakers are reliable sources of knowledge." (Lackey, 1999 p.474). It is the conformity between facts and reports that does the heavy epistemic lifting when we trust experts. This may be the case for those humans that we consider experts, but it is not the case for computer simulation. For simulations, the reasons that experts would count them as reliable sources of knowledge is their relationship to theoretical principles and engineering best practices in addition to their predictive successes. For example, a reason to trust a given computer simulation of a system is not merely because it is able to predict a given state of the system, but rather *because it does so in the appropriate way*- in conformity with the laws of physics, in conformity with engineering practices that are amenable to error assessment, etc. (Resnik, 1997; Symons and Alvarado, 2016) The task for experts in the evaluation of computer simulations is fundamentally different from the consumption of computer simulations by non-experts. Non-experts depend on experts to have certified those simulations as worthy of attention. The testimony that non-experts are trusting is the testimony of expert communities, not

directly the output of their simulations.

5. Scientific Inquiry and Everyday Epistemic Practices

We have argued that Burge-style approaches to epistemic entitlement are inappropriate in the context of computer simulations when it comes to scientific or other high-stakes applications. In areas of life that involve great risk or expense, appropriate epistemic diligence will be called for. However, we contend that science *should be*, and often is also organized around stringent epistemic norms for the evaluation of computer simulations. In fact, it might even be the case that since time constraints are less pressing in much of science than they would be in, for example, military or policy contexts, *scientists have the luxury of being even more rigorous and demanding with respect to the epistemic standards governing their devices.*¹⁹

Clearly there are aspects of scientific practice that are just like the epistemic practices that we would find in ordinary life. Scientists are people too, and in this sense, Burge's reflections on the Acceptance Principle will have an important role. We are not denying the fact that in order for scientists to engage in research, they need to trust one another, they cannot maintain an attitude of radical skepticism, and they need to make reasonable trade-offs between time spent in various epistemically relevant tasks.

However, there are a few distinctions that ought to be drawn in this discussion. First, philosophical treatment of computer simulations should distinguish between the kinds of entitlement being appealed to in distinct contexts. Is the entitlement in question, for example, an entitlement to trust, or an entitlement to not doubt?²⁰ Second, there are epistemic and non-epistemic entitlements that can be invoked as justifications to trust

¹⁹ In saying this we are not discounting the serious time constraints faced by individual scientists in their careers e.g. the need to publish novel findings, the pressure of funding agencies, or the pursuit of tenure requirements. Rather, we are referring to scientific inquiry as a series of methods aimed at furnishing the best available understanding of our world. As such, scientist (from astrophysics to geology to biology- and some social sciences) can and ought to ensure a rigor in their methods that the nature of industry and war seldom afford.

²⁰ It is unclear, for example, the extent to which Barberousse and Vorms assume a negative (entitlement to not doubt) or positive (entitlement to trust) role for entitlement in their arguments. However, this particular paragraph focuses on the completely distinct categories between kinds of entitlements (a kind of warrant that gives someone *a reason* to believe (x)) beyond their epistemic variety.

something. For instance, an epistemic entitlement might be playing a role in our decision to trust a medical device given our belief that it may be beneficial to a patient. A non-epistemic reason to believe might involve our commitments to our affiliation group. There might be good social reasons to believe what a particular person says in cases where the matter under consideration is trivial or involves no harms or significant expense. If the cost of error is low, then arguably it is reasonable to simply accept the other person's position, in cases where the social (or other) benefits are important. An individual might accept or reject some scientific hypothesis for social reasons. This can be regarded as reasonable insofar as that person's individual commitments have little practical import. Her views for example on climate change, evolution, or other socially controversial matters have little real significance for either the practice of science or for policy decisions based on the scientific hypothesis in question. Insofar as their position on these matters causes no significant harm, it is at least understandable why consumers of science would choose to practice bad epistemic hygiene with respect to the science in favor of strengthening their position within an affiliation group. They strengthen their place within their social network by assenting to the consensus position of that group, perhaps even assenting to an exaggerated version of that consensus position. In those cases one may indeed have reasons to believe/trust, but they are not strictly speaking *epistemic reasons* (See for example Dretske, 2000 p.594).

By contrast, according to Dretske, for example, if what we are appealing to in accepting some proposition as true is to be regarded as an epistemic reason then the only grounds are its truth or possible truth (p.593-4). The truth or possible truth of the values produced by an instrument deployed in the aid of science can only be a product of a reliability assessment, which is in turn distinct from personal justification (Williams, 2000). Epistemic entitlement for computer simulations, should not be reduced to, or confused with the effort to provide what Dretske calls "a pragmatic vindication of the practices we use to fix beliefs" in place of efforts to "validat[e] the beliefs themselves" (2000, 598). Merely showing that a "practical purpose or need is served by accepting certain practice does not address the problem of our *epistemic* right to accept the individual beliefs that occur in this practice." (ibid)

Scientists are held to a higher epistemic standard than ordinary epistemic agents (Nola and Sankey, 2014). This is one of the reasons why they can serve the role of providing expert testimony. Consider Burge's own use of the example of Newtonian calculus (1999 p.8). Burge's claim is that non-demonstrative reasoning can underwrite sound mathematical beliefs since Newton's "knowledge" of elementary truths of calculus was available to him before formal explanations were available. And yet, without such formal explanations few would have considered Newton's knowledge to have had any epistemic value. Similarly, Burge provides an example in which people accept the Pythagorean theorem solely on the basis of a diagram or the word of another without being able themselves to produce its proof. This acceptance only takes place in circumstances where no meaningful inquiry is being conducted on the theorem itself, or in which the theorem does not play a central role. In fact, to accept the Pythagorean theorem in this way is to treat it on a par with, for example, one's neighbor explaining that recycling is collected with the trash every second week. The significance of knowing when to put out the recycling with the trash is (under normal circumstances) low enough, that believing one's neighbor is a reliable strategy. If it were very important, for some reason to know the recycling and trash collection schedule with greater certainty - imagine, for example, that we were the person in charge of the agency collecting the trash-, , then one might make additional efforts to determine the truth.

Burge's example of how one might come to accept the Pythagorean Theorem without really understanding the proof is convincing only insofar as we have relatively low standards for what we take to be true in mathematics. In ordinary life, it is true that our standards will often be low. There are cases where knowledge of the Pythagorean Theorem will have the same priority and role in our everyday lives as knowledge of the trash/recycling schedule. However, in the practice of science, our standards are higher and in the judgment of scientific evidence, a critical educated public will also hold important matters to a higher epistemic standard. Thus, expert scientific testimony is not on a par with ordinary testimony and ordinary epistemic norms are not adequate to formulate an epistemology concerning scientific practice or the public understanding of science.

While no individual scientist is an epistemic saint, the scientific enterprise collectively seeks to ensure truth-aptness and reliability independently of the shortcomings of each individual practitioner. In relation to modeling and simulation, this has practical import. Humphreys, for example thinks that practitioners should always possess working knowledge of a given instrument and he argues that the background theory of its principles should always be at hand to the practitioner. This is not only because some instruments, for example, an MRI scanner, require this knowledge in order to be used effectively (Humphreys, 2004 p. 38), but also because, the theory of how the instrument operates is “routinely needed to protect the user against skeptical objections resulting from ways in which the instrument can produce misleading outputs.” (p.38)

Computer simulations are often used in contexts where highly demanding scientific standards are in place and where even ordinary well-grounded scientific practice falls below the threshold for acceptability (Ruphy, 2011; Morrison, 2015). The famous five sigma standard governing evidence at CERN in the discovery of the Higgs Boson (where computer simulation played a central role) is an extreme, but important instance of the kinds of high epistemic standards governing scientific research.²¹

The fact that scientific standards would be regarded as unreasonably strict if we were to adopt the perspective of everyday epistemic practices such as normal decision-making, trust in perception, and credence with respect to the testimony of others should make it clear that the norms governing ordinary epistemic practices are not appropriate for the evaluation of computer simulation.

6. Conclusion

The epistemology of computer simulation is a highly complex domain where assumptions concerning the transmission of warrants and the characterization of epistemic entitlements must be understood in relation to the diversity of scientific practices. The actual epistemic situation is fundamentally dissimilar to the situation

²¹ The five-sigma standard corresponds to a p-value, or probability, of 3×10^{-7} , or about 1 in 3.5 million. In this case it is the probability that *if* the particle does *not* exist, the CERN team would find what they observed. It is extremely unlikely that they could have generated the data by accident. How unlikely? 1 in 3.5 million unlikely!

depicted in our thought experiment. Neither pragmatism, nor an analogy with ordinary epistemic practices will prove useful as guides to the epistemology of computer simulations.

Instead, we have argued that trust in simulations should be grounded in empirical evidence, good engineering practice, and established theoretical principles. Without these constraints, computer simulation risks becoming little more than unmoored speculation.

ACKNOWLEDGMENTS ETC.

Bibliography

Adler, Jonathan, "Epistemological Problems of Testimony", *The Stanford Encyclopedia of Philosophy* (Summer 2015 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2015/entries/testimony-episprob/>.

Arkoudas, Konstantine, and Selmer Bringsjord. "Computers, justification, and mathematical knowledge." *Minds and Machines* 17, no. 2 (2007): 185-202.

Audi, Robert. "The place of testimony in the fabric of knowledge and justification." *American Philosophical Quarterly* 34, no. 4 (1997): 405-422.

Barberousse, Anouk, Sara Franceschelli, and Cyrille Imbert. "Computer simulations as experiments." *Synthese* 169.3 (2009): 557-574.

Barberousse, Anouk, and Marion Vorms. "About the warrants of computer-based empirical knowledge." *Synthese* 191.15 (2014): 3595-3620.

Beebe, Helen. "Transfer of warrant, begging the question and semantic externalism." *The Philosophical Quarterly* 51.204 (2001): 356-374.

- Beisbart, Claus. "Advancing Knowledge Through Computer Simulations? A Socratic Exercise." In *The Science and Art of Simulation I*, pp. 153-174. Springer, Cham, 2017.
- Bonjour, Laurence. "A priori Knowledge" in Bernecker, Sven, and Duncan Pritchard, eds. *The Routledge companion to epistemology*. Routledge, (2011): 283-293
- Boschetti, F., E. Fulton, R. Bradbury, J. Symons "What is a model, why people don't trust them and why they should" in *Negotiating Our Future: Living Scenarios for Australia to 2050*. Australian Academy of Science (2012) 107-118
- Borge, Steffen. "The word of others." *Journal of Applied Logic* 1, no. 1-2 (2003): 107-118.
- Burge, Tyler. "Content preservation." *The Philosophical Review* 102.4 (1993): 457-488.
- Casullo, Albert. "What is entitlement?." *Acta Analytica* 22, no. 4 (2007): 267-279.
- Davies, Martin. "II—Martin Davies: Epistemic Entitlement, Warrant Transmission and Easy Knowledge." *Aristotelian Society Supplementary Volume*. Vol. 78. No. 1. The Oxford University Press, 2004.
- Dretske, Fred. "Entitlement: Epistemic rights without epistemic duties?." *Philosophy and Phenomenological Research* 60, no. 3 (2000): 591-606.
- Dummett, Michael. "Testimony and memory." In *Knowing from words*, pp. 251-272. Springer Netherlands, 1994.
- Floridi, Luciano, Nir Fresco, and Giuseppe Primiero. "On malfunctioning software." *Synthese* 192.4 (2015): 1199-1220.

- Frigg, Roman, and Julian Reiss. "The philosophy of simulation: hot new issues or same old stew?." *Synthese* 169.3 (2009): 593-613.
- Gramelsberger, Gabriele. "Generation of evidence in simulation runs: Interlinking with models for predicting weather and climate change." *Simulation & Gaming* 42, no. 2 (2011): 212-224.
- Horner, Jack, and John Symons. "Reply to Angius and Primiero on software intensive science." *Philosophy & Technology* 27, no. 3 (2014): 491-494.
- Humphreys, Paul. *Extending ourselves: Computational science, empiricism, and scientific method*. Oxford University Press, 2004.
- Humphreys, Paul. "The philosophical novelty of computer simulation methods." *Synthese* 169, no. 3 (2009): 615-626.
- Jenkins, Carrie S. "Entitlement and rationality." *Synthese* 157, no. 1 (2007): 25-45.
- Keller, Evelyn Fox. *Models, simulation, and "computer experiments"*. na, 2003.
- Lackey, Jennifer. "Testimonial knowledge and transmission." *The Philosophical Quarterly* 49, no. 197 (1999): 471-490.
- Lackey, Jennifer. "Testimonial Knowledge" in Bernecker, Sven, and Duncan Pritchard, eds. *The Routledge companion to epistemology*. Routledge, (2011): 316-325.
- Lazer D, Kennedy R, King G, et al. (2014) The parable of Google Flu: Traps in big data analysis. *Science* 434: 343

- McEvoy, Mark. "The epistemological status of computer-assisted proofs." *Philosophia Mathematica* 16, no. 3 (2008): 374-387.
- McEvoy, Mark. "Experimental mathematics, computers and the a priori." *Synthese* 190, no. 3 (2013): 397-412.
- Moretti, Luca, and Tommaso Piazza. "When warrant transmits and when it doesn't: Towards a general framework." *Synthese* 190.13 (2013): 2481-2503.
- Morgan, Mary S., and Margaret Morrison. *Models as mediators: Perspectives on natural and social science*. Vol. 52. Cambridge University Press, 1999.
- Morrison, Margaret. "Models, measurement and computer simulation: the changing face of experimentation." *Philosophical Studies* 143.1 (2009): 33-57.
- Morrison, Margaret. *Reconstructing Reality*. Oxford: Oxford University Press (2015)
- Naylor, Thomas H., J. L. Balintfy, Donald S. Burdick, and Kong Chu. "Computer Simulation Techniques, John Wiley & Sons." Inc., New York, New York (1966).
- Newman, Julian. "Epistemic opacity, confirmation holism and technical debt: Computer simulation in the light of empirical software engineering." In *International Conference on History and Philosophy of Computing*, pp. 256-272. Springer, , 2015.
- Norton, Stephen, and Frederick Suppe. "Why atmospheric modeling is good science." *Changing the atmosphere: Expert knowledge and environmental governance* (2001): 67-105.
- Nola, Robert, and Howard Sankey. *Theories of scientific method: an introduction*. Routledge, 2014.

Oreskes, Naomi. "The scientific consensus on climate change." *Science* 306, no. 5702 (2004): 1686-1686.

Oreskes, N., Shrader-Frechette, K. and Belitz, K., 1994. Verification, validation, and confirmation of numerical models in the earth sciences. *Science*, 263(5147), pp.641-646.

Pryor, James. "When warrant transmits." *Mind, meaning, and knowledge: Themes from the philosophy of Crispin Wright* (2012): 269-303.

Quine, Willard V. "The roots of reference." (1974).

Resnil, Michael. *Mathematics as a Science of Patterns* (1997) New York, Oxford University Press.

Ross, James F. "Testimonial evidence." In *Analysis and Metaphysics*, pp. 35-55. Springer Netherlands, 1975.

Roush, Sherrilyn. "The epistemic superiority of experiment to simulation." *Synthese* (2015): 1-24.

Ruphy, Stéphanie. "Limits to modeling: balancing ambition and outcome in astrophysics and cosmology." *Simulation & Gaming* 42, no. 2 (2011): 177-194.

Ruphy, Stéphanie. "Computer Simulations: A New Mode of Scientific Inquiry?." In *The Role of Technology in Science: Philosophical Perspectives*, pp. 131-148. Springer Netherlands, 2015.

Steadman, Ian. "Big data and the death of the theorist." *Wired Online*, January 25 (2013): 2013.

Symons, John. "Computational models of emergent properties." *Minds and Machines* 18, no. 4 (2008): 475-491.

Symons, John and Jack Horner. "Software intensive science." *Philosophy & Technology* 27, no. 3 (2014): 461-477.

Symons, John and Jack Horner "On Some Limits to Model-Based Proof of Software Correctness." In Powers, T. (ed.) *Philosophy and Computing: Essays in Epistemology, Philosophy of Mind, Logic, and Ethics*. Springer pp. 2017

Symons, John, and Ramón Alvarado. "Can we trust Big Data? Applying philosophy of science to software." *Big Data & Society* 3, no. 2 (2016): 2053951716664747.

Symons, John, and Fabio Boschetti. "How computational models predict the behavior of complex systems." *Foundations of Science* 18, no. 4 (2013): 809-821.

Tymoczko, Thomas. "The Four-Color Problem and its Philosophical Significance." *Journal of Philosophy* vol. 76 (p57-82)

Vallor, S. (2017). AI and the Automation of Wisdom. In: Thomas Powers (Ed.) *Philosophy and Computing: Essays in Epistemology, Philosophy of Mind, Logic, and Ethics*. Philosophical Studies Series, vol 128. (pp. 161-178). Springer

Williams, Michael. "Dretske on epistemic entitlement." *Philosophy and Phenomenological Research* 60, no. 3 (2000): 607-612.

- Wright, Crispin, and Martin Davies. "On Epistemic Entitlement." *Proceedings of the Aristotelian Society, Supplementary Volumes*, vol. 78, 2004, pp. 167–245.
www.jstor.org/stable/4106950.
- Weisberg, Michael. *Simulation and similarity: Using models to understand the world*. Oxford University Press, 2012.
- Wiener, Norbert. *The human use of human beings: Cybernetics and society*. No. 320. Da capo press, 1954.
- Wiener, Norbert. "Cybernétique et société (1950)." *Paris: Union Générale d'Éditions* (1971).
- Wikipedia contributors, "Patch (computing)," *Wikipedia, The Free Encyclopedia*, [https://en.wikipedia.org/w/index.php?title=Patch_\(computing\)&oldid=784124212](https://en.wikipedia.org/w/index.php?title=Patch_(computing)&oldid=784124212) (accessed June 23, 2017).
- Winsberg, Eric. "Simulated experiments: Methodology for a virtual world." *Philosophy of science* 70.1 (2003): 105-125.
- Winsberg, Eric. *Science in the age of computer simulation*. University of Chicago Press, 2010.
- Winsberg, Eric, (2015) "Computer Simulations in Science", *The Stanford Encyclopedia of Philosophy* (Summer 2015 Edition), Edward N. Zalta (ed.), URL = [<http://plato.stanford.edu/archives/sum2015/entries/simulations-science/>](http://plato.stanford.edu/archives/sum2015/entries/simulations-science/).