Are computer simulations experiments? And if not, how are they related to each other?

Claus Beisbart

European Journal for Philosophy of Science

ISSN 1879-4912
Volume 8
Number 2

Euro Jnl Phil Sci (2018) 8:171-204
DOI 10.1007/s13194-017-0181-5

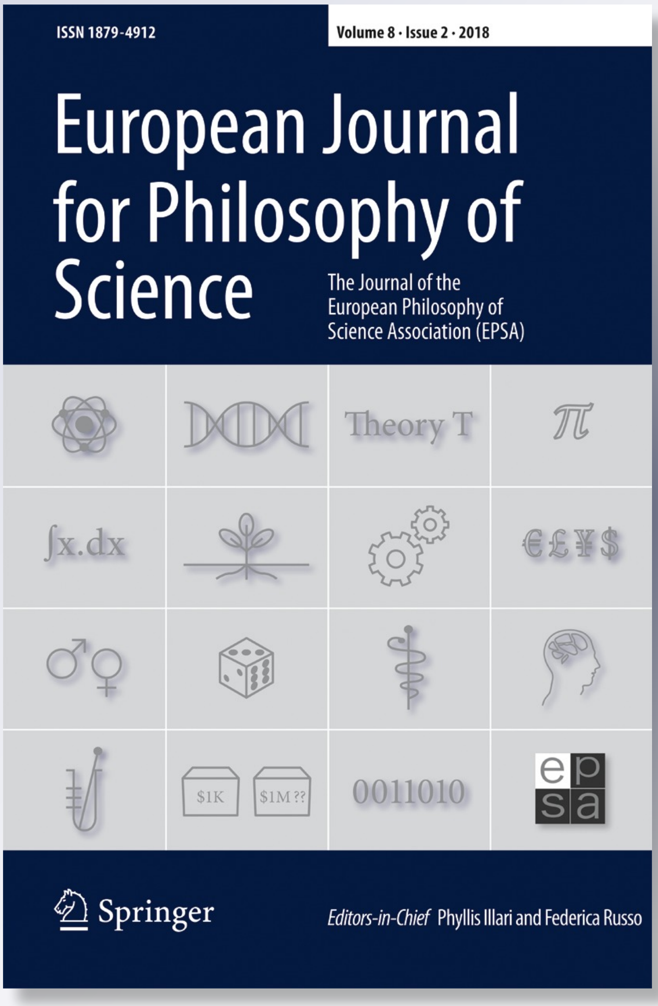

Euro Jnl Phil Sci (2018) 8:171–204
DOI 10.1007/s13194-017-0181-5

ORIGINAL PAPER IN PHILOSOPHY OF SCIENCE

Are computer simulations experiments? And if not, how are they related to each other?

Claus Beisbart[1]

Received: 19 August 2016 / Accepted: 28 June 2017 / Published online: 24 July 2017

Abstract Computer simulations and experiments share many important features. One way of explaining the similarities is to say that computer simulations just are experiments. This claim is quite popular in the literature. The aim of this paper is to argue against the claim and to develop an alternative explanation of why computer simulations resemble experiments. To this purpose, experiment is characterized in terms of an intervention on a system and of the observation of the reaction. Thus, if computer simulations are experiments, either the computer hardware or the target system must be intervened on and observed. I argue against the first option using the non-observation argument, among others. The second option is excluded by e.g. the over-control argument, which stresses epistemological differences between experiments and simulations. To account for the similarities between experiments and computer simulations, I propose to say that computer simulations can model possible experiments and do in fact often do so.

Keywords Methods · Computer simulation · Experiment · Intervention · Experimental control · Observation · Modeling

1 Introduction

A cosmologist runs a computer simulation of structure formation in the Universe. For some very early cosmic time, an initial distribution of matter is assumed. It is almost homogeneous, with tiny fluctuations here and there. The shape of the fluctuations is

✉ Claus Beisbart
Claus.Beisbart@philo.unibe.ch

[1] Institut für Philosophie, Universität Bern, Länggassstr. 49a, CH-3000 Bern 9, Switzerland

constrained using data from the Cosmic Microwave Background, but also depending on theoretical assumptions. Then, the computer is run to follow the evolution of the matter from the tiny fluctuations to the formation of galaxies and up to our time. To this purpose, equations that connect physical characteristics as functions of time are solved to some approximation and evaluated for certain instances of time. The outputs of the simulation comprise tons of numbers that represent the matter distribution at several times. The latter is visualized in images and animations. Since the result does not agree with what the working scientist knows about the Universe, she concludes that a theoretical assumption underlying the initial conditions was wrong. She modifies the initial distribution and runs the simulation again.[1]

This is a rough and ready description of how a computer simulation (CS, for short) might work. No doubt that we are here talking about a computer simulation. But it may seem that the scientist runs an experiment too. After all, she tinkers with the initial conditions, is curious to learn what happens and may be surprised by the outcome. All this is typical of experiment. So are some or, maybe, all CSs experiments? If so, is their being experiments crucial for their epistemic value? And if not, how can we explain the similarities between CSs and experiments?

Much talk about CSs suggests that they are experiments. For instance, CSs are called "computer experiments".[2] Simulation studies in biology are referred to as "experiments in silico".[3] Some people use the expressions "experiments in theory"[4], "numerical experimentation"[5] or "virtual laboratory"[6] when they speak about computer simulations. Further, and more importantly, some authors explicitly claim that CSs are experiments. For instance, (Norton and Suppe 2001, pp. 87 and 92) write:

> Simulations are often alleged to be only heuristic or ersatz substitutes for real experimentation and observation. This will be shown false. [...] Simulation modeling is just another form of experimentation.

But is this correct? Are computer simulations really experiments? This is the research question of this paper. I will deny that CSs are experiments and put forward an alternative thesis that clarifies the relationship between CS and experiment.

There are good reasons to answer the question of this paper. Computer simulation is often described as a new scientific method (e.g.Binder and Heermann 2010, p. 1). It is the task of philosophers of science to understand this method. To this effect, they should explain how the method works and how it achieves the tasks it is supposed to complete, in particular how it can produce new knowledge about the world.[7] Now, if CSs turn out to be experiments, we can explain how CSs do their epistemic work by simply referring to their status as experiments; for we understand quite well how

[1] See Efstathiou et al. (1985) and Bertschinger (1998) and Dolag et al. (2008) for simulations in cosmology.

[2] E.g. Gramelsberger (2010).

[3] E.g. Naumova et al. (2008).

[4] Keller (2003), p. 203, in scare quotes.

[5] Humphreys (1994), p. 103.

[6] Dowling (1999), p. 261, in scare quotes.

[7] Authors as different as Winsberg (1999), p. 277, Stöckler (2000), p. 366 and Barberousse et al. (2009), pp. 558–559 agree that this is an important task for a philosophical treatment of CS.

experiments achieve their tasks. In particular, we can understand how CSs can yield new knowledge about nature because it is agreed that experiments can yield such knowledge.

The recent research literature in philosophy of science has seen an intensive debate on whether CSs are experiments. Important contributions include Hughes (1999), Sec. 5.2.4, Guala (2002), Morgan (2005), Parker (2009), Winsberg (2009b), Barberousse et al. (2009), Morrison (2009), Giere (2009), Arnold (2013), and Durán (2013) and Peschard (forthcoming).[8] Although this literature has produced interesting insights, I do not think that my question has yet been answered in a satisfying way. I will argue this point in detail in this paper. In my view, the challenge is to account for the salient similarities between CSs and experiments even though the former are not experiments properly speaking.

As CSs and experimentation have extensively been discussed in the last few years, there is no need to carry out more case studies. The main philosophical issues on the table are conceptual. Section 2 thus sketches what an experiment is. Section 3 lists a number of pre-theoretical similarities between CSs and experiments. Any account of CSs should elucidate these similarities. Section 4 classifies claims that philosophers have hitherto made about the relationship between experiments and CSs. Two such claims, which take CSs to be experiments, are rejected in Sections 5 and 6. Section 7 shows that, even if CSs are in fact experiments, their epistemic power does not derive from their being experiments. Section 8 introduces my own proposal, viz. that CSs can model possible experiments and do often do so. Using this suggestion, we can account for the similarities between experiments and CSs without unduly assimilating the two methods. I wrap up in Section 9.

This paper is only about computer simulations and does not consider other forms of simulations, e.g. analog simulations (cf. Trenholme 1994). The reason is that computer simulations and other types of simulation differ significantly. If not stated otherwise, "simulation" refers to computer simulation.[9]

2 Intervening and observing: Experiment reconsidered

To answer the question of this paper, it is first useful to briefly characterize experiments. For the purposes of this work, it will suffice to state important necessary conditions on experiments.[10]

[8]Consult (Imbert 2017), Sec. 3.5 for a very useful overview of the recent debate.

[9]But see fn. 68 for a short remark on analog simulations.

[10]I can draw on a rich philosophical literature about experiments. See Hacking (1983), Part B, Janich (1995), Morrison (1998), Heidelberger (2005), Radder (2009), Bogen (2010) and Franklin (2010) for introductory pieces or reviews about scientific experiments. Radder (2003) is a recent collection in the philosophy of experiment. See Falkenburg (2007), particularly Ch. 2, for a recent account of experimentation with applications to particle physics. Biological experiments are philosophically analyzed by Weber (2005), experiments in economics by Sugden (2005). For studies about modern experiments see also Knorr-Cetina (1981) and Rheinberger (1997).

Experimentation is an empirical method; it crucially involves observation. But observation is not sufficient for experimenting. The distinguishing mark of experiments is intervention. Related catchwords are "manipulation", "interference" or "control". In an experiment, a system is set up, prepared or at least somehow causally interfered with before it is observed. Also, while the experiment is running, the experimental set-up is often controlled. Call the system that is interfered with and observed the "object of the experiment" or the "system experimented on".[11]

The crucial idea that underlies experiments may be explained by drawing on Kant. In his „Critique of Pure Reason", he particularly mentions experiment to explain the success of modern natural science, and he characterizes the underlying spirit as follows:

> "Reason, to be taught by nature, must approach nature with its principles in one hand [...] and, in the other hand, the experiments [...] – yet in order to be instructed by nature not like a pupil, who has recited to him whatever the teacher wants to say, but like an appointed judge who compels witnesses to answer the questions he puts to them." (B XIII, Kant 1998, p. 109).

For our purposes, we may understand Kant's metaphor as follows: Experimenters ask nature a well-defined question by preparing the system, by controlling it and by conceptualizing the experiment. They then attend to nature's answer and observe how the system that they have set up reacts.[12]

Intervening on the system that is observed comes with advantages and drawbacks. On the positive side, experimentalists are enabled to investigate systems, conditions and circumstances that do not, or only rarely, occur in nature. They can thus learn about systems that are significantly simpler and easier to understand than systems outside the context of the experiment. Experimentalists can also vary the set-up to isolate relevant factors from irrelevant ones. On the negative side, the conditions probed in an experiment often differ from those typical of phenomena that occur outside the context of the experiment. It is thus a non-trivial task to generalize the results from experiments.[13]

Experiments are always run with the superordinate aim to obtain or to consolidate information about the system that is experimented on. The information concerns the way the system reacts to the intervention and is supposed to make a difference to what the experimentalist or other people know about the system. In typical examples of experiments, the reaction of the system is not known in advance by the experimenter. In other cases, the experimenter may feel confident about what the result of her experiment is, but the expected result counts at least as strengthening her hypothesis.

The condition that an experiment has the superordinate aim of providing information is necessary, since we often intervene on systems and observe them, while not

[11] See e.g. Janich (1995) and Parker (2009), p. 487 for the two components of experimentation. Causal interference in experiments is put into a historical perspective by Tiles (1993).

[12] Kant himself is eager to stress the conceptual work necessary to ask nature a well-defined question, but this is not important in what follows; our focus concerning the idea that nature is asked a question is rather on the causal interference with the system that is observed.

[13] See Hüttemann (2000) and Falkenburg (2007), Ch. 2 for a related discussion.

running an experiment. For instance, a potter causally interferes with the clay, while also observing it. In typical cases, this does not count as experimentation though. The reason is that the information obtained by the potter is typically only used to improve upon the process of molding a figure in clay. That is, the information is only fed in the very process again and thus an aim subordinate to the intervention. Only if the potter takes the information obtained to have significance of its own, e.g. because she wants to learn about a new sort of clay, are we talking of an experiment. Thus, in an experiment, the aim of gaining information is superordinate to the intervention: The latter is done for the sake of the former rather than the other way round.[14]

Many experiments are run with an eye to systems different from the system experimented on. Morgan (2002) reports class-room experiments by Chamberlin and Smith that were used to understand the behavior of humans in real markets outside the class-room. If the experiment is run in a laboratory, but supposed to be interesting for a system outside the lab, it is tempting to mark the contrast between both systems by speaking of the lab and the real world. But this terminology does not make sense more generally because experiments need not be carried out in the lab. You can manipulate a natural system to carry out an experiment; for instance, you can pour some chemical substance in a pond to see what happens (this experiment isn't hereby recommended!). Thus, in what follows, no restrictions will be assumed on the system on which an experiment is done, if only it is a real system, i.e., if it exists in space and time.

Since some experiments are run with an interest in systems distinct from the system interfered with, some authors take it that experiments have not just an *object* that is experimented on, but also a *target*, which is typically different from the object of the experiment.[15] Under this view, experiments include an inference from the object of the experiment to the target. Call this claim the *inferential view of experiments*.

It is certainly true that experimentalists often run inferences to other systems when they think about their results. For one thing, experiments are stipulated to be repeatable, and this can only be so if certain low-level generalizations hold (e.g. to the effect that time and the exact place of the experiment do not matter). The results of an experiment are thus typically inferred to hold in different realizations of the same experimental set-up. For another thing, the system experimented on may be causally related to another system that is of prime interest to the working scientist. The scientist may then run a causal inference to the latter system. But the inferences of both sorts just mentioned, i.e., the lower-level generalizations and the causal inferences, do not lead scientists very far, while proponents of the inferential view refer to inferences that are much stronger because they refer to different types of systems that are causally unrelated to the system experimented on. For instance, (Winsberg 2009b, p. 577) mentions the example of an experiment on a tank that is supposed to allow for inferences to an astrophysical application.

[14]In the example of the potter, I cannot exclude that the potter runs an experiment, if additional conditions are fulfilled. But even if she does, the respective experiment would not count as *scientific*. Scientific experiments are embedded in a broader scientific practice. As a consequence, the epistemic difference an experiment is supposed to make is more pronounced.

[15]See Guala (2002) and Winsberg (2009b), pp. 52–53.

I do not see any compelling reason to say that such stronger inferences form a proper part of an experiment. Even if an experimentalist is interested in such an inference, there is no reason to treat the inference as proper part of her experiment. It seems more natural to say that the experiment was made with the further purpose to run an inference to some system different from the one manipulated. Note also that many experiments such as the Michelson-Morley experiment do not comprise an inference to other systems. I thus reject the inferential conception of experiments as an over-generalization (cf. Parker 2009, p. 487 and Peschard (forthcoming), Sec. 3.1). This is not to deny though that the inferentialist view of experiments holds true of some significant subclass of experiments.

To obtain information on how the object of the experiment behaves, experimentalists carry out observations. What is observed about the system can be summarized using *data*, which is either qualitative or quantitative.[16] Today, many experiments use *measurement devices* or *instruments* such as detectors to obtain data, and I shall include the use of measurement devices under observation as well.[17] Since the functionings of instruments often depend on theories, natural descriptions of the observations are frequently theory-laden. Suppose, for instance, a scientist tells us that she has observed an electric current of a magnitude of 1 ampere. What the scientist has seen more strictly speaking is only the position of a pointer in a measurement apparatus. Taken as a result about electric current, her observation thus relies upon the assumption that the position of the pointer reflects electric current in an appropriate way.[18]

We can summarize the results of this section by formulating the following necessary conditions on experiments:

Conditions on Experiment 1 A person A is experimenting on system X only if

1. A intervenes on X;
2. A takes observations about X; and
3. A does 1 and 2 with the superordinate aim to obtain information about the way X behaves and reacts to the intervention.

We will say that A runs an experiment if and only if (iff, for short) she runs an experiment on some X.[19]

It is important that the first two conditions refer to the same object X. It is also possible to observe X and causally interfere with some Y≠X. This is typical of *instrument-based observation*, where Y may be a telescope and X the star observed through the telescope.

The conditions given here are best thought of as part of an explication in the sense of Carnap (1962), Ch. I. They try to reflect the way in which the term "experiment" is used in talk by scientists and by laypeople, but they may regiment this talk slightly.

[16] See e.g. Balzer (1997), p. 139.

[17] See Shapere (1982) and Hacking (1983), Chs. 10–11, Falkenburg (2007), pp. 65–71 and Humphreys (2004), Ch. 1 for broad accounts of observation.

[18] A recent monograph about theory-ladenness is Adam (2002).

[19] See Peschard (forthcoming), Sec. 1 for a similar account of experiment.

For instance, the conditions exclude so-called natural experiments and thought experiments from being proper experiments, because both types of methods do no include an intervention.[20] If natural experiments or thought experiments are indeed experiments, this is not reflected in my conditions.[21]

Consequently, not everybody may agree with the proposed conditions on experiment. Fortunately, this does not matter too much for the arguments to follow. The reason is that most of my arguments draw on a single condition, respectively, and different arguments draw on different conditions. In consequence, my case against CSs being experiments is robust because it does not depend on how exactly experiment is defined.[22]

3 Pre-theoretical similarities between computer simulations and experiments

CSs bear a number of similarities to real experiments. Here is a list of some such similarities, which is not meant to be exhaustive.

1. Something like intervention is crucial for experiments and for many simulations. This point has already been argued concerning experiments in the previous section. As regards CSs, consider the traffic simulation package TRANSIMS (Casti 1997, pp. 131–142). It can be run with many different inflows of the traffic. The traffic inflow is thus subject to what might be called intervention by the simulation scientist. Note also that, in a CS, the hardware of the computer is causally interfered with when the computer is programmed.
2. CSs produce outputs that take a similar form as do data from experiments (or from observations, more generally), at least if the outputs and data are suitably processed. For instance, in a recent study, Lim et al. (2015) present what they call a "comparison between simulated and experimental data" about neuron activity (their Fig. 5 on p. 1808). They show the probability density as a function of firing rate for two different types of neurons, one time for experiments, the other time for simulations. Here, the probability density is not directly measured, but rather constructed from the raw data/simulation outputs. The example shows that outputs of CSs and experimental data can be of exactly the same type, at least if they are suitably processed.[23] Also, data from experiments and outputs of CSs are often analyzed in the same way (cf. Winsberg 2010, p. 33).
3. Both experiments and CSs have results in the sense that they can show something. Scientists can obtain information about a system by experimenting on it or by simulating it. The point should be obvious for experiments and is no less true

[20]For instance, Zimmerman (2003) calls his study a natural experiment. See Brown and Fehige (2017) for a recent overview of thought experiments

[21]One may of course argue that natural and thought experiments are not really experiments, but this is not the place to do so.

[22]I'm grateful to an anonymous referee for pointing me to this fact.

[23]See Barberousse et al. (2009) and Humphreys (2013) for useful discussions of the notion of data in the context of CSs.

for simulations; for instance, simulations of the type first proposed by Gillespie (1976) provide us with knowledge about the kinetics of chemical reactions under various conditions (see Hasty et al. 2001, p. 270 for testimony).

4. The results of both experiments and CSs can often not be foreseen (cf. Humphreys 1990 and Rohrlich 1990). Experiments and CSs can thus have results that are unexpected and surprising (e.g. Morgan 2005, pp. 323–326).
5. The epistemologies of experiments and CSs share important similarities in that scientists use similar strategies when they use the outputs of experiments and simulations to justify certain claims. For instance, both simulation scientists and experimentalists use comparable techniques to ensure that their results are genuine and not mere artifacts (Winsberg 2010, p. 21 and pp. 43–44; see also Parker 2008). This epistemological similarity is likely a consequence of other similarities noted so far (and, maybe, some background knowledge). It is natural to expect that scientists use similar strategies to avoid artifacts in experiments and CSs if both methods produce outputs that are similar (cf. similarity 2) and that arise from a similar sort of variation of the initial conditions (cf. similarity 1).

A noteworthy fact related to these similarities is that simulations can often *replace* experiments. To learn something about a system or to better understand it, scientists simulate the system using the computer instead of experimenting on it (cf. Humphreys 2004, p. vii; see also Skaf and Imbert 2013 for an illuminating case study). Thus, in the terms of Skaf and Imbert (2013), experiments sometimes are functionally substitutable by CSs. As the authors correctly point out, functional substitubility of an experiment by a CS does not imply that both perform the function in exactly the same way.

The similarities listed refer at least to typical examples of experiments and simulations. This does not imply though that they hold universally. I will elaborate on this point later in this paper.

To sum up, there are a number of affinities between experiments and CSs. At least some of the features shared by experiments and CSs, viz. nos. 1, 2 and 5 are epistemologically significant because they relate to the way in which both methods are used to obtain knowledge and to achieve related epistemic aims.

4 A map of positions and further directions

The similarities are easily explained if we say that computer simulations are experiments. But what exactly does it mean to say this? Well, CS and experiment constitute methods and as such define types of actions: We classify what scientists do by saying that they run a computer simulation or an experiment. The task of this paper then is to relate two types of action to each other in a sensible way.

One option is type identity, i.e., we say that the two types of action are the same. Under this option, whenever you perform the first type of action, you perform the second, and the other way round.[24] Type identity is not plausible for CSs and

[24]Cf. the "identity thesis" mentioned by Winsberg (2009a), p. 840.

experiments simply because running a computer simulation involves working on a computer, and for ages there have been experiments that do not involve computers. Also, type identity would imply that CSs can do everything that experiment can do, but this is not plausible, as will become clear in Section 6 below.

Another option is to say that performing one type of action always *includes* performing the other one. Since a lot of experiments have been run before the age of computers, we can of course not say that experimenting always includes simulating. But it may well be the other way round: Computer simulation always includes an experiment. This then is a sensible way in which computer simulations may be said to be experiments:

CE: Each CS includes a real experiment.

Note that the experiments claimed here to be included in CSs need not be typical, but may rather be of a peculiar sort.

In what follows, we allow that a CS includes an experiment if it *is* an experiment. Under this interpretation, the claim that CSs are a subset or subtype of experiments is a limiting case included in CE. Thus, if CE is wrong, so is the view that CS is a subtype of experiment. Therefore, for the purposes of this paper, it will suffice to refute CE.

Now when it comes to explaining what CS can achieve, CE on itself is not very interesting. Even if we accept CE, the experimental part of a CS may not be crucial for epistemic matters. Let us thus formulate a second claim, which presupposes CE and strengthens it:

CE+: The epistemic power of a CS exclusively derives from the epistemic power of the experiment included in the CS.[25]

Here the epistemic power is meant to be the disposition to produce knowledge, understanding and related achievements. Of course, knowledge and understanding are very broad notions that may be interpreted in various ways and that come in several varieties. For instance, Baumberger (2011) distinguishes between understanding-why and objectual understanding of some subject matter and argues that each cannot be reduced to the other. At this point, however, it does not make much sense to restrict epistemic power by using such distinctions. We are here interested in an attempt to explain the work that CSs do by pointing to their alleged status as experiments. The idea of CE+ must thus be that, whatever epistemic tasks CSs achieve, they do so due to their being experiments. But note that the restriction of CE+ to some achievements would not matter for most of my arguments.[26]

If each CS is to include an experiment, the simulationalist has to experiment on some object X in space and time according to my conditions on p. 6. For simulations, there are only two plausible candidates for X, which lead to the following two claims.

CE_H: Each CS includes a real experiment on the hardware of the computer.

[25] Similar claims as CE and CE+ figure in J. Norton's reduction of thought experiments to arguments, see Norton (1996); cf. Beisbart (2012).

[26] So far, the focus of the philosophy of computer simulations was on knowledge.

CE_T: Each CS includes a real experiment on the target of the simulation.

Both claims have had their proponents or at least their sympathizers. Parker (2009) accepts CE_H. She thinks that

> "a computer simulation study *does* qualify as an experiment – an experiment in which the system intervened on is a programmed digital computer." (p. 488).

Morrison (2009), by contrast, seems to support CE_T, at least she is very sympathetic to it. The general aim of her paper is to call into question the distinction between experiment and simulation (e.g. ibid., pp. 33 and 55–56). She tries to realize this aim by arguing that the outputs of some CSs should be seen as measurements (e.g. p. 36). There is no evidence that she is talking about a measurement of characteristics about the computer; rather she stresses that the "the material [in a CS] is, strictly speaking, what the experimenter can represent in the simulation models" (p. 45). She also suggests that there is a causal link (albeit a very indirect one) between the target system and the outputs produced by simulation, as we expect it for experimental measurements (p. 55). But if CSs are seen as experiments because they are claimed to measure the characteristics of the target system, this is very close to CE_T.[27]

CE_T cannot apply to every CS because some simulations do not have a real-world system as their target and because one cannot run a real experiment on a system that does not exist. Other simulations are known to rely on heavy idealizations and approximations and it is thus counterintuitive to say that they are experiments on the target system (Winsberg 2009a, p. 841). Consequently, CE_T has to be restricted to some kinds of simulations. In what follows, I will discuss CE_T in this restricted sense.

There are also other claims that may explain why CS and experiments share many features. Some authors propose to say that simulations are experiments *on mathematical models* (e.g. Morgan 2003, p. 225; cf. also Fritzson 2004, p. 7) or *on theories* (e.g. Dowling 1999). But it is difficult to understand what an experiment and thus an intervention on a mathematical model or a theory should be, literally speaking. I thus take it that the term "experiment" is merely meant in a metaphorical way here and that the authors do not subscribe to CE.

Some authors, e.g. Guala (2002), Morgan (2002), and Morgan (2003) and Winsberg (2009b) seem to reject the claim that CSs are experiments. At least they wish to maintain a distinction between CSs and experiments, as does (Winsberg 2009b). Note though that we need to be very careful when we relate their claims to CE etc. For one thing, a mere distinction between CSs and experiments is compatible with CE because, even if both methods are different, CSs may include experiments. For another thing, other authors hold different views of experiments than do I, e.g. if they adopt the inferential view of experiment. Winsberg is a point in case; he rejects the idea that simulations are experiments because he thinks that the inference to the

[27] To be fair, I should mention that the paper by Morrison provides also indications that she does not fully support CE_T. For one thing, the wordings of her central claims are very cautious; she never says that CSs *are* experiments, but rather e.g. that there are no reasons to maintain the distinction (e.g. p. 55). This claim seems also to be restricted to some "contexts" (p. 33). She further admits that computer simulations do not involve the manipulation of the target system (fn. 16 on p. 55). This concession does not seem to matter much for her argument; so, maybe, she does not think that intervention is crucial for experiment.

target is different between an experiment and a simulation. This is not incompatible with the claim that a simulation includes an experiment on the hardware under the non-inferential view of experiments.

On a natural interpretation, CE etc. are *conceptual claims*. They then are supposed to follow from the definitions of experiment and CS and are assumed to hold necessarily and universally. For the purposes of this paper, however, I need not adopt this interpretation. What matters for my argument is only that the claims are supposed to hold universally.

My own position is that both CE and CE+ are false. This does of course not exclude that *some* CSs happen to include experiments. To account for the affinities between experiments and computer simulations I will later argue for the following claim:

CME_1: CSs allow scientists to model possible experiments.

I also claim that

CME_2: Many CSs do model possible experiments.

While CME_1 may be understood as a conceptual claim in the same way as CE etc., CME_2 only applies to some simulations, so it is not a conceptual truth.

As mentioned above, CE is often advocated because people think that it can explain the epistemic power of CSs. Likewise, CME_1 and CME_2 can to some extent explain the epistemic power of those computer simulations to which they apply.

My position is close to what Morgan (2003) says about models (see also Morgan 2005). I am also indebted to Barberousse et al. (2009).[28]

5 Why computer simulations do not in general include experiments on the hardware

To proceed, I will first argue against CE. Proponents of CE have to accept either CE_H or CE_T, and I will take both theses in turn. In this section, I will argue against CE_H, i.e., the claim that CSs include experiments on the hardware.

Suppose that person P runs a CS and assume for the sake of the hypothesis that the simulation includes an experiment on the hardware of the computer. According to my conditions on experimentation (see p. 6), this implies that P intervenes on the hardware and observes it to learn about it.

[28] CE, CE+ and CME try to clarify the relationship between experiments and simulations. But what exactly do they mean by CSs? There are broadly two ways of conceiving CSs depending on whether a CS is supposed to be one run with a simulation program or whether it is what Parker (2009), p. 488 calls a "computer simulation study", which also includes writing the program, testing it, etc. (e.g. Frigg and Reiss 2009, p. 596; Parker 2009, p. 488). Analogous questions can be raised about experiments too, e.g. is the construction of the detector used in an experiment part of the latter or not? In what follows, I will not rely upon any specific proposal as to what is included in an experiment or a CS. Rather, I will assume that experiments and CSs are identified in a similar way such that the claims under consideration have a chance of being true. For instance, when we discuss CE, it would be too uncharitable to assume that experiments include detector building and similar activities, while a computer simulation is simply one run of a simulation program. What is important though for my argument is that every experiment includes an intervention on the object of the experiment.

A first problem with this is that a simulationalist is not normally interested in the hardware and that she does not want to learn about the hardware.[29] If this is true, then the simulation does not include an experiment on the hardware, because my third condition on experiments is violated.

By itself, this point does not carry much conviction because it invites the following objection: It is certainly true that P's *ultimate* aim is not to learn about the computer hardware. But we may still say that P does have the aim to learn about the computer hardware, if this aim is supposed to serve the superordinate aim to learn about the target system through an inference from the hardware to the target. If this is so, the simulation may include an experiment on the hardware.

To rule this out and to make a successful case against CE_H, a stronger argument is needed. To present such an argument, I'd like to show that the person P from the example does not in general make observations about the hardware. This is to say that the second condition on experimentation from p. 6 is violated.

Why will P not in general make observations about the hardware to learn how it reacts? To make such observations, P would have to observe that *p*, where *p* is a proposition about the hardware. The reason is that the observations must have some content, and this content can be stated using propositions.[30]

The problem now is that there is in general no suitable proposition *p* with the following properties:

- *p* is about the hardware;
- P observes that *p* during the simulations;

To show this, let us for definiteness assume that the output of the simulation is printed on paper. This was indeed the case in many simulations. The print-out consists of arrays of signs that denote numbers. Person P is looking at the paper. What observations is she making? Well, she is observing that, e.g. '1.2' etc. is written on the paper. The problem is that this does not tell her anything about the computer hardware (e.g. the processor) because the numbers refer to the system that is simulated on the computer. For instance, 1.2 may be the value of the simulated surface temperature of a star in some units (cf. Barberousse et al. 2009, p. 561 for this point).

It may be objected that one can re-describe her observations such that they are more telling concerning the computer hardware. Compare this to the following example. A thermometer takes the temperature in a certain room and makes a print-out on a sheet of paper every ten minutes. If a working scientist reads and understands the signs on the paper, we can say that she observes that the temperature at noon was 10 degrees Celsius and so on. The reason is that there is a causal connection between the temperature in the room and the signs on the paper and that the scientist knows how to interpret the signs in terms of their relevant causes.

It may seem tempting to construct a similar story for the scientist who runs CSs. Here is a suggestion. Suppose, the sign '1.2' was caused by a hardware state in which

[29] Cf. Hughes (1999), p. 137.

[30] Note that we are here not talking about observing in the sense of looking at. Observation on this interpretation does not suffice for experimenting.

a certain register was in a state with such and such potentials.[31] So can't we say the following: By looking at the sheet of paper, the scientist observes that, at a certain time, a certain part of the hardware is in this and this state? Well, we cannot say this because the working scientist does not come to know this. She will in general have no clue about what happens in the hardware. Most people working with computers do not know how the hardware is functioning (cf. Barberousse et al. 2009, p. 573). Thus, the working scientist will in general not be able to infer anything substantial about the hardware from the numbers on the sheet of paper.[32]

So far, my argument assumes that the hardware includes the processor of the computer, but not the print-out. The problem I'm pointing out then is that the print-out is not used to infer anything about the hardware. But it may be objected that the paper with the print-out is indeed part of the hardware. This is very natural in view of the fact that the Turing machine is often said to include a tape (e.g. Barker-Plummer 2016, Sec. 1).[33] Now if the print-out is part of the hardware, then the scientist's observation that '1.2' is printed is an observation of the hardware, and there does not seem to be a problem with CE_H any more.

One possible strategy to reject this objection is to argue that the paper is not part of the hardware, properly speaking. It is in fact not entirely clear whether the tape is part of the Turing machine proper. For instance, (Turing 1937, p. 232) defines the machine in such a way that it does not include the tape. Nevertheless, this strategy is not promising, because it may rightly be rejected as being based on an ad hoc individuation of hardware.

But we can slightly modify our example to make it even less plausible to say that the computer scientist observes the computer hardware. There are cases in which the simulation scientist does not even know what the hardware is. For instance, there are computing networks that allow for the following service: You submit a job with your simulation program. The job is randomly assigned to a computer from the network and executed there. The output is temporarily stored on the computer and then transferred to your desktop PC, where you can look at it. For the sake of the example, we may stipulate that the computer scientist doesn't come to know which computer has executed the simulation program. But how then can we say that she has learned about the hardware of a computer? A possible opponent of my argument may insist that, in this example, the hardware is really the computer network. But now the opponent's argument seems to be based upon an ad hoc way of individuating hardware.

My modified example of running a computer simulation is rather peculiar, and it may be doubted that it can carry much weight. However, we are here taking issue with a general claim to the effect that every simulation includes an experiment on

[31] See Rechenberg (2000), Chs. 2–3 for a brief description of the hardware of computers; the details do not matter for my purposes.

[32] Barberousse et al. (2009) seem to agree with my claim that the working scientist does not observe the hardware, for they write that

> "the running computer is not observed qua physical system, but qua computer, i.e., qua *symbolic* system" (p. 564).

[33] I'm grateful to an anonymous referee for raising this objection.

the hardware, and one counterexample is enough to refute it. Further, my rejection of CE_H seems also plausible in the view of other examples. Suppose, for instance, that a computer scientist runs a simulation, the results of which are presented as an animation on screen. What the scientist sees is most naturally described in terms of the computer model. When asked what she sees, the scientist may e.g. say: "I see that two galaxies are merging in my model," and statements of this type are the basis for her further work. It seems artificial to say that the scientist first observes something about the computer and then runs an inference to the model.

Note also that observation is usually much professionalized in the sciences. Observation is most often made in terms of quantitative measurements and the raw data are stored to allow for a re-analysis. But in a CS, no quantitative measurements about the hardware are taken, since, even if a scientist reads numbers from a printout, this is not a quantitative measurement. Also, no information is stored that may allow the scientist to recover the states realized by the hardware. Thus, if a CS included an experiment on the hardware, the observation part of the experiment would be rather unprofessional.

Here then is the catch: The most natural interpretation of the outputs of CSs is in terms of the target of the simulations. Indeed, scientists themselves commonly describe their simulations in terms of assumptions about the target. It needs substantial knowledge about the hardware to squeeze results about the latter out of the simulations. Working scientists need not have this knowledge for running a simulation on the hardware.

To be fair, I should admit that the scope of my argument against CE_H is limited. It does not exclude that quite a lot of computer simulations include experiments. The argument does not work for CSs in which the working scientist has sufficient knowledge about the hardware such that she can be said to make observations about the hardware. Also, it is sometimes recommended that scientists run certain programs to test the accuracy of the calculations carried out by the hardware[34], and these programs may be simulation programs. In this sense, simulations may be used to learn about the hardware, and they may thus be said to be experiments. But this is not the general case of a simulation, and it is arguable that related runs of a simulation program do not really qualify as simulations because they are not carried out to learn about the target of the simulations.

It may also be possible to slightly stretch the meaning of "observe" to make conceptual space for the claim that, in a CS, the hardware is observed, and thus for something like CE_H. There is a more general problem here because the notions that we are discussing, e.g., those of experiment, computer simulation or observation do not refer to natural kinds and thus leave space for borderline cases, in which it is debatable whether they apply or not. Even granted all this, I still regard the proposal that all CSs include experiments as unnatural.

So far, my argument against CE_H is entirely based upon the idea that CSs need not always involve observation, for which reason the second condition on experiment is violated. This is significant. Even if somebody rejects the first condition on

[34] See Press et al. (2007), p. 9 for an example of a suitable program.

experiment, e.g. because she takes natural experiments to be experiments proper, then this does not matter for the no-observation argument.

But the discussion so far also casts doubt on the idea that CSs involve manipulation of the hardware. For instance, in the example with the network of computers, we cannot say that a specific hardware was intentionally intervened on by the scientist. This is at least so if "intervention" (or "manipulation") is meant to be intentional: I only intervene on an object in this sense, when I know what I'm doing on which object. Such a notion of intervention seems crucial for experiment, since, according to our third condition on experiment, the experimentalists wants to learn how the system reacts to an intervention, and she does not do so if she does not know what the intervention is. Now intervention in this thick, intentional sense does not take place in a computer simulation if a simulation scientist does not know what the hardware is. Nor does it take place when the scientist cannot describe what she did to start the computer as an intervention on the computer. There thus is a no-intervention argument that parallels the no-observation argument.[35]

So we have two arguments against CE_H, one showing that it violates the condition of observation; the other showing that it violates the condition of manipulation.[36] Both arguments are complementary. If somebody isn't swept away by the no-observation argument, she may still be impressed by the no-manipulation argument, and the other way round. Admittedly though, both arguments raise similar issues. So I'll refrain from discussing possible objections against the no-manipulation argument.

If somebody isn't swept away by either of the arguments, we may try to combine them in the following way. As mentioned above, the first two conditions on experiments are connected because they refer to the same system (which is here called the system experimented on). We may thus argue that the conjunction of the conditions is violated in some computer simulations because there is not one object that is intervened on and then observed. For instance, in some simulations, the simulation program is parallelized and run on several machines.[37] Whereas one computer may take the initial values, another may be responsible for the procession of the data. It then is natural to say that, if a computer is manipulated, it is the first one, and if a computer is observed, it is the second one. Both are not the same, which is to say that the first and the second conditions are not fulfilled for the very same object.[38]

Parker (2009) is a defender of CE_H, so it is worthwhile to compare to her position. The notions of experiment that Parker and I assume, respectively, are quite similar; in particular, we both stress the importance of intervention and observation (see her p. 487). So we have more than a merely verbal disagreement about CE_H. We mainly differ because Parker thinks that a simulation scientist observes the computer, while I have cast doubt on this claim.

[35] Such an argument is suggested by Imbert (2017), Sec. 3.5.

[36] This is so if intervention in the second condition is meant to be intentional. If this is not so, then the argument needs the third condition which makes it very likely that the intervention is intentional.

[37] The example of a simulation that is parallelized is also used by Barberousse et al. (2009, p. 565), albeit in a different argument.

[38] My arguments against CE_H from this section can easily be generalized to show that *computations* carried out on computers (and not just simulations) do not include an experiment on the hardware.

But why is Parker so eager to stress the experimental status of CSs? Parker's ultimate concern seems about the epistemology of CS (ibid., p. 491):

> the epistemology of computer simulation must attend to this materiality [i.e., to the fact that "computer simulation studies are *material* modeling exercises and *material* experiments", ibid.], or it will be impoverished, if not fundamentally confused.

To substantiate her point, Parker mentions the possibility of hardware failures and the fact that simulationalists are often worrying about modifications of the hardware of computers because the simulations have to be changed accordingly (p. 491). Her thought seems to be that the significance that the material structure of a computer has for the epistemic value of the CSs can only be explained by saying that CSs are experiments.

Hardware failures are certainly an important issue in the epistemology of CSs. But this fact can be accommodated without saying that CSs are experiments. We can think of computers as *instruments* or tools that may or may not function as required. Whether or not they do so, can be a matter of legitimate concern. All this does not imply that simulationalists do experiments on computers. We can illustrate this by using a simple analogy. Suppose that we store data on a computer. Later, we look up the data. This is certainly not an experiment. There may nevertheless be hardware failures that prevent the data from re-emerging on the screen in an undisturbed fashion, and we may worry about this possibility.[39]

6 Why computer simulations do not include experiments on the target

Even if computer simulations do not include experiments on the hardware of the computer, they may include experiments on the target. This is at least what CE_T claims, and Morrison (2009) seems sympathetic to this view.[40] The aim of this section is to argue against CE_T. Since CE_T has only a fighting chance for CSs that have real-world systems under fairly realistic conditions as their targets, I will concentrate on such simulations.

Even under this restriction, the proposal that a simulation includes an experiment on the target is very dubious. First, experimentation involves intervention, but CSs do not normally include an intervention on the target. The target is only intervened on in the sense that the scientist is free to set the initial conditions of the simulations. This is not a real intervention on the target system, but rather something like an intervention in the simulated world. I will come back to this point shortly.

[39] Some measurement apparatuses used in experiments function in a similar fashion. We do not observe certain measurement apparatuses, but rather use them to observe something else (cf. fn. 2 on p. 7 above). To do so we have to trust the instruments. There is thus a close parallel between computers as instruments and instruments in experimentation. Cf. Humphreys (2004), Ch. 1 and Parker (2010).

[40] For the purposes of this paper, we need not engage with Morrison's argument in detail because Giere (2009) has made a convincing case against it.

Second, the target system is in general not observed in a simulation, as would be required for an experiment according to my second condition (e.g. Guillemot 2010, p. 244). True, some simulations are initialized or calibrated using data from the target, and we may say that the outputs of such simulations contain empirical information about the target. But even then it does not seem that we are observing the target properly speaking (Parker 2010, pp. 42–43).[41]

Although CE_T looks very implausible in view of this criticism, let me add a third argument against CE_T. The main idea is that experiments are subject to a condition called "no over-control" that simulations do not fulfill. I thus call the argument the over-control argument. It can be seen as an extension of an argument by Morgan (2003) and Morgan (2005). Morgan compares experiments and *models* and claims that both can surprise, but that only the latter can confound (Morgan 2005, p. 324). I extend the argument to CSs and flesh it out a bit more. I also elaborate an argument by Arnold (2013, pp. 59–60). Note that the over-control argument is of particular interest since it is significant for the epistemology of CSs (see below).

The gist of my argument is nicely introduced with the idea that experimenters ask nature a question in the way judges ask witnesses (see p. 4 above) and then register her answer. Now nature does not reply to a question if her answer is completely determined by the question or the way the question is asked. Accordingly, person P is only running an experiment if the result is not over-controlled. However, as I shall argue below, the results of CSs are in fact over-controlled. So the results of CSs are not experimental, and CSs do not include an experiment on the target.[42]

This, at least, is the argument in a nutshell. To elaborate it, I will first explain why over-control is a problem for experiments. I will then argue that computer-simulations violate the requirement that there be no over-control.[43]

According to our third condition on experiment, an experimenter wants to learn about how a system reacts to an intervention. Now she can only have this aim if she thinks that it can be realized. So she must think that she can learn about the reaction of the system, and that the system does in fact react to her intervention. But if there is over-control, then there is no reaction of the system properly speaking. The behavior of the system is completely determined by the set-up and thus nature doesn't have anything to say.

Thus, granted that the scientist infers an immediate consequence from her belief that she can learn about the reaction of the system experimented on, she will conclude that there is no over-control. So over-control is excluded in a subjective sense. Further, granted that the scientist is right in thinking that she can learn about the reaction

[41] My first two arguments against CE_T may be summarized by the claim by Arnold (2013, pp. 58–59) that CSs do not operate on the target itself.

[42] Here, the last inferential step excludes the possibility that CSs include an experiment on the target even though the results of CSs are not experimental. This possibility can indeed be dismissed as being far-fetched and useless. Even if it were realized, we could not directly appeal to experimentation to explain the epistemic power of CSs.

[43] Properly speaking, it is the results of CSs (or experiments) that are (not) over-controlled; but for convenience, I will sometimes say that CSs (experiments) are (not) over-controlled. See fn. 42 for a justification.

of the system experimented on, over-control is excluded in an objective sense: There is no over-control in the experiment.[44]

All this is supposed to show that the absence of over-control in experiments is a consequence of necessary conditions on experiments and of some natural assumptions. This is important since it would not be fair to argue from a new requirement on experiment.[45]

The requirement that there be no over-control in experimentation fits nicely with M. Morgan's views about experiments from economics. Here is Morgan (2005), p. 324:

> This is an important consideration in the design of [economic] experiments: experiments need to be set up with a certain degree of freedom on the part of participants so that their behavior in the experiment is not totally determined by the theory involved, nor by the rules of the experiment. If the experimental behavior is totally predetermined, there is no potential for unexpected patterns to emerge.

But what exactly is over-control or determination? The idea cannot be that the outcomes of experiments should not be nomologically dependent on the setting. For, in a deterministic world, the set-up of the experiment and its environment must lead to a certain outcome, given the laws of nature, and we cannot exclude a deterministic world at this point. The idea of over-control is rather that a logical relationship between the assumptions on which an experiment is based and its result be avoided. The assumptions that are needed to describe the set-up and the understanding of the experiment do not imply what the outcome of the experiment will be.

It may be objected that the condition of no over-control is too strong. There are many experiments the results of which are expected by the scientific community. There are further experiments on systems the dynamics of which is in principle completely known. The point of the experiment may then be to learn about the behavior of the system under specific conditions, because this very behavior cannot be predicted since the dynamical equations cannot be solved. The objection then is that experiments of these types are indeed over-controlled because the results are either part

[44] To elaborate a bit: The argument starts from what the working scientists wants. The reason is that the third condition on experiment is cast in terms of what the experimenter wants; it does not require that the experimenter be successful. The argument further assumes that the scientist is rational in that she only wants something that she takes to be possible and that she draws immediate consequences of her beliefs. For the objective absence of over-control, it is further assumed that the scientist *correctly* believes that she can learn about a reaction of the system. We can grant the additional assumptions because proponents of CE_T will not want to save their claim by claiming that scientists are not rational or that they do not know basic things about the concrete setting. If we don't find the additional assumptions convincing, we may restrict ourselves to successful experiments in which the goal mentioned in the third condition is in fact fulfilled. Proponents of CE_T will not want to exclude such experiments. We can then argue that over-control would prevent success.

[45] Note though that the notion of a reaction does a lot of work for the argument since we understand reaction in a way that excludes objective over-control. It may be objected that, in the conditions on experiment, "reaction" may instead be taken to be any consequence of the intervention. The assumption that experiments exclude over-control would then need additional justification. I think that the discussion provided in this section does provide this justification.

of the background knowledge of the scientists or follow from the known dynamical equations.[46]

In response, I deny that the types of experiments just mentioned are in fact over-controlled. As defined above, over-control is about the assumptions that are *necessary* to run an experiment. There is over-control if these assumptions imply the results. If these assumptions only imply the results if they are combined with other assumptions, there is no over-control. Now, assumptions that concern the dynamics of the system are usually not necessary to run the experiment. Suppose for instance, that I'm interested in the dynamics of a fluid under very specific conditions. I know that the system is completely described in terms of the Navier-Stokes equations with certain boundary conditions and parameters. But I can't solve the equations and want to know what the dynamics of the system is. Then I can run the experiment on the system to learn this. My knowledge of the Navier-Stokes equations does not matter for what this experiment is. I would run the same experiment if I didn't know the equations. So my claim that experiments are not over-controlled stands the objection.[47]

The problem for CE_T now is that CSs are over-controlled and even taken to be over-controlled. This is to say that there is over-control in the objective and the subjective senses, as indicated above. It follows that the simulation cannot include an experiment on the target. To see this, suppose that a scientist runs a CS with a deterministic algorithm. There are thus no random elements in the algorithm. Such an algorithm or a program implementing it defines a unique mapping from inputs to outputs. Provided that the hardware functions as intended, the results of the simulations are completely determined by the computer program and the input. As Beisbart (2012) shows, CSs can be reconstructed as deductive arguments, where some premises reflect the input and the algorithm and the conclusions the output. The premises describe the initial conditions and the dynamics, while the conclusions characterize the system for later times and thus deliver the result of a CS. According to Beisbart (2012), the argument is deductive. CSs thus are over-controlled in the sense that their results are entailed by the assumptions that define the set-up of the simulations. There is no space left for a possible answer by nature. This is even known to the scientist. Thus, an important condition on experimentation is violated and there is an epistemologically significant difference between experiments and CSs.

For an illustration of the argument, consider first the Michelson-Morley experiment (Michelson 1881; Michelson and Morley 1887). This experiment is thought to have refuted the assumption that the earth has a non-zero velocity with respect to the ether postulated by many physicists at the time around 1890. The experiment has a complicated set-up, and a number of assumptions are needed to interpret its data as

[46] I'm grateful to an anonymous referee for raising this objection.

[47] There may be exceptions. For instance, I may use a system that is known to follow certain dynamical equations to identify a solution to these equations. Here I need to know the equations, otherwise I can't interpret the experiment in the way I wish to. To save my claim that experiments are not over-controlled, I may either deny that we are really talking about an experiment here. Alternatively, I may claim that knowledge of the equations is not needed for the experiment proper, but only for an inference that is based upon it.

having implications about the ether. But all this does not imply what the result of the experiment is.

Suppose now that we would like to replace the experiment with a computer simulation. To this end, we would have to make an assumption as to whether the earth has a relative velocity with respect to the ether or not, at least in an implicit way. But the results of the simulation will then follow from this assumption. This is what I have called over-control. And my claim is that simulations are over-controlled and even taken to be over-controlled.

A possible objection to my argument is that the results of a CS follow only *in principle* from the assumptions that enter the simulation, while, *in practice*, there is no way to see how they follow. If this was false, the simulation would not be run. As noted above (cf. our fourth pre-theoretical similarity on p. 8), the results of CSs are often unforeseeable, as are the results of experiments. But then the results of the simulations are not really taken to follow from the set-up of the simulations from the viewpoint of the practicing simulation scientists, or so it may be suggested. As a consequence, the over-control argument would not work, at least if it is restricted to the subjective level. This objection may be supported with the general claim that a philosophical appraisal of practices such as experimentation or CS should stick to what is possible in practice. Some philosophers (see e.g. Winsberg 2010; Humphreys 2004, Sec. 5.5.) have claimed that lofty in-principle claims do not provide a fruitful perspective on CSs.

This objection is not very convincing though. Working scientists often think that their CS successfully traces its target system. This much must also be admitted by proponents of CE_T, because CSs can only qualify as an experiment on the target if they are thought to faithfully reflect the target. But at least in most cases, the simulation scientist will only think this because she trusts the assumptions upon which the simulations are built. The idea must be that trust in the assumptions or the set-up of the simulations can be transferred to the outputs of the simulations. But this transfer is only possible in a simulation if the outputs are entailed by the assumptions of the simulations. The working scientist has thus to think that the outputs are entailed in the set-up of the simulations. Conversely, a simulation scientist should be very wary of the results of her simulation if she did not think that the latter could in principle be derived from the assumptions upon which the simulation is built. All in all, in a lot of cases of CSs, the thought that the results of a CS are implied by the assumptions necessary to run the simulation is endorsed by the working simulationalist herself, at least implicitly. It is not attractive to appeal to examples of simulations for which this thought is not implicit in the understanding of the scientist, because these simulations are likely to be found problematic by herself.

The proposition that the assumptions of a CS determine the outputs, and thus the results, is only true of deterministic CSs. There are other simulations, viz. Monte Carlo simulations, which include a random element.[48] Such simulations may feel more like an experiment. It is telling in this respect that Humphreys (1994) uses Monte Carlo simulations to argue that simulations form a method between

[48] See e.g. Beisbart and Norton (2012).

experiment and theory. It may thus seem doubtful whether the over-control argument developed in this section can be generalized to Monte Carlo simulations.

But as Beisbart and Norton (2012) argue, Monte Carlo simulations are not really experiments. One reason is that the outputs of simulations do not interest the scientists in every detail, but only in some relevant respects. Random fluctuations that arise due to the use of random elements are usually hoped or even controlled to be too small as to make any difference of interest. If this is so, then at least the relevant aspects of the results are implied by the assumptions upon which the simulations are built, which means that there is over-control.

The over-control argument thus seems successful and broadly applicable. As such, it has important consequences for the epistemic purposes to which computer simulations may be used: Computer simulations alone do not suffice to refute a combination of assumptions in the way an experiment can.[49] The reason is, roughly, that simulations do not leave anything to say for nature, while experiments do. Simulations implement a model that entails the results of the simulation. The simulations can thus not refute these very assumptions. Nor can they refute other assumptions, unless the model of the simulations has independent support.

The results of a simulation can of course be incompatible with other assumptions, and this may provide a decisive case either (a) against the simulations or (b) against the other assumptions. But in neither case can we say that something is refuted only because of the simulations. In case (a), the scientists have to draw on prior knowledge to make a case against the simulations; for instance, they have to have data that are incompatible with the results of the simulations. In case (b), a contradiction between the results of the CSs and the other assumption can only be used to argue against the latter, if the simulations, and thus the model that they implement, have independent support.[50]

The fact that CSs cannot be used to empirically refute a set of assumptions in the way experiments can, would suffice to show that CSs are not experiments. This point is easily overlooked because CSs can take the role of experiments, given that certain conditions are fulfilled. This does not imply though that CSs can *always* or *generally* take the role of experiments.[51]

To sum up this section, there are three reasons to deny CE_T: In a CS, the target system is not intervened on. Nor is it observed. Finally, the results of CSs are over-controlled by the assumptions necessary to run the simulation, while experiments

[49] This is also what Hughes (1999), p. 142 claims. I am here speaking of *combinations* of assumptions because Duhem (1954, Ch. 10, §§2–3) has taught us that many hypotheses cannot be tested or refuted in isolation.

[50] Here, simulations are distinguished from arguments to the effect that the simulations get it right. When arguments of this type are part and parcel of simulations, the latter can of course empirically refute a set of assumptions, as can do experiments. But it is trivial that CSs in this extremely thick sense can do this. This trivial point cannot be the rationale for CE_T. – Note also that purely mathematical theories may be falsified using computer simulations. But mathematical theories are not my concern here.

[51] What then are the conditions under which a CS can replace an experiment? Well, it can do so if the CS is known to reflect the intervention and the reaction of the system experimented on in a sufficiently faithful way. What sufficiently faithful representation means depends on the aspects that the working scientists are interested in and on the desired level of accuracy.

are not. This last argument is certainly most interesting because it is directly connected to epistemological questions. At the same time, it seems most vulnerable to objections. It may e.g. turn out that the notion of assumptions necessary to run an experiment/simulation is not stable enough to do the work it is supposed to. But fortunately, the argument is not needed for my case against CE_T.

7 What if simulations happen to be experiments nevertheless?

CE_H and CE_T are the only sensible specifications of CE I can think of. Since both fail, so does CE. Let me nevertheless assume for a moment that all CSs include experiments. The aim of this section is to show that even then it is not decisive for the epistemic power of CSs. This is to reject CE+ (defined on p. 10).

To do so, we run the following thought experiment.[52] Suppose that there is a mathematical genius who can do every calculation we wish her to do in a few milliseconds. We can give her an algorithm and initial values of the variables that figure in the algorithm, and she carries out the instructions of the algorithm and tells us about the outcomes. Let us also stipulate that the genius is subject to the same types of errors a computer may make. Suppose now, we give the genius an algorithm that evaluates how certain physical characteristics behave as functions of time according to a model. The genius would return results in a few seconds. We could use these results in exactly the same way in which we could use results from a CS that follows the same algorithm. We would have similar worries about both types of results, and we could use both of them for exactly the same purposes. This shows that the epistemic power of the CS is the same as that of the calculation of the genius. But if the genius carries out the algorithm, this does not constitute an experiment. It is obviously not an experiment on the hardware of a computer (or an experiment on the brain). Nor do we have an experiment on the target. Consequently, an alleged experimental status of CSs is immaterial for their epistemic power.

8 Computer simulations as modeled experiments

Since CE and CE+ have both failed, it is now time to clarify the relationship between CS and experiment with a new idea. My proposal is as follows

CME_1: CSs allow scientists to model possible experiments.
CME_2: Many CSs do model possible experiments.

The idea here is always that a CS with target X can model a possible experiment on X and does often do so. I'm talking of possible experiments since simulations may model experiments that are never run in the real world.

[52] Barberousse et al. (2009), pp. 562 and 565 make a similar point concerning the hardware of a computer. Their target is not so much an alleged experimental status of simulations rather than the idea that the physicality of the simulations is crucial for the epistemic power of simulations.

Winsberg (2003) makes a similar suggestion when he titles a paper "Simulated experiments," but he doesn't elaborate the idea in this paper, while rather critically examining other views about simulation and experiment. In a revised version of the article (Winsberg 2010, Ch. 3) the expression "simulated experiments" has been dropped.[53]

What does it mean to say that a simulation models an experiment? I assume that modeling is a distinct mode of scientific inquiry (Weisberg 2007). A modeler uses a substitute of her target system to learn something about this target, e.g. its dynamics. The substitute may be a real physical system or a purely imagined one. Typically, the substitute is easier to study than the target itself e.g. because it is more accessible, because it abstracts from certain details of the target or because it is in one or the other aspect idealized. The substitute (often called "model system") is analyzed and the results of this analysis are transferred to the target system. This transfer is possible because the target and its substitute are similar to each other in relevant aspects. This view of modeling draws on Suárez (2003), Scholz (2004), Giere (2004), Suárez (2004) and Weisberg (2007). It combines insights of the so-called use and resemblance theories of representation. Thus, two conditions are required for modeling: First, there has to be a model system that resembles the target; second, the model system has to be used as a surrogate for the target.[54]

Why now can we say that a CS can model an experiment? As an example, recall the simulations of cosmic structure formation mentioned at the beginning of this paper. The simulations evaluate solutions to some equations. The equations specify the values of physical characteristics (e.g. the positions of particles) depending on some initial state and we can thus say that the solutions obtained in the simulations describe the dynamic behavior of a certain system. This system is defined in such a way that its physical characteristics take exactly those values that the simulation returns as solutions of the equations. The system is often characterized informally when simulations are described by saying e.g. "In the simulations, the particles attract each other ...". Ideally, the system described in this way coincides with the target system. But we cannot in general assume this coincidence because the simulations may fail to reflect the target system in every respect. For instance, cosmological N-body simulations refer to artificial particles that are not assumed to exist (see Bertschinger 1998 for details). We are thus on the safe side when we say that the simulations describe a substitute for the target system. The substitute is fictional, but its behavior is similar to that of the target in one or the other way. It is therefore used to learn about the target. Simulations thus define a fictional system, a substitute the behavior of which is used to represent the behavior of the target.

[53]That CSs model experiments is sometimes assumed in the sciences, too, see e.g. Haasl and Payseur (2011), p. 161. The claim also makes sense of the following remark by Metropolis and Ulam (1949) about Monte Carlo simulations:

> "These experiments will of course be performed not with any physical apparatus, but theoretically." (p. 337).

[54]As I shall show below, my argument also goes through for an alternative account of modeling that does not assume similarity.

Does this suffice to say that a simulation models a possible *experiment* on the target? To be sure, the model system displays a certain behavior, and we may think of it as the reaction to a possible intervention part of an experiment. We may then say that this behavior models the reaction of the target to a possible intervention and thus represents an important part of a possible experiment. The problem though is that, so far, the model does not have any intervention and observation in it. There is no experimenter in the model who intervenes on the system modeled and observes it.

But we can go further and say that intervention and observation by the experimenter can indeed be modeled in a CS *study* (which also includes setting up the simulation etc.; cf. footnote 28), if only in a different way. For there is more to a CS study than the model system that stands for the target. In a CS study, the simulationalist can set the initial conditions and the values of important parameters, and this is in fact what is often done. This is similar to manipulation and activities of control on the part of the experimenter in an experiment. The similarity is not merely accidental; intervention in an experiment and quasi-intervention in a simulation (as I would like to call it) are done for similar reasons, for instance, to learn about circumstances that do not naturally occur or to isolate causally relevant factors. In some cases, the simulationalist may even consciously imitate the activities typical of an experimenter. We can thus say that simulation scientists can make quasi-interventions that reflect possible interventions in experiments.

To be fair, I should mention that some CSs are supposed to trace how a real-world target system behaves if it is not at all manipulated. In such CS, quasi-intervention does not take place. But the simulation program or a slight modification of it *may* be used to study the effects of a quasi-observation, and thus, as a method, CS allows for models of possible experiments. So as a modal claim, CM_1 does hold true.

As far as observation is concerned, the simulationalist can take notice of the outputs from her simulation, e.g. she can monitor movies produced by the simulation and analyze the outputs in the same way in which data from experiments are analyzed, and this is in fact what is very often done. There is a close parallel to what an experimenter does to observe the system experimented on. The similarity is again not merely accidental; observation in an experiment and quasi-observation in a simulation (as I would like to call it) are done for similar reasons, viz. to learn about the system under scrutiny. In some cases, the simulationalist may even consciously imitate the activities typical of an experimenter. We can thus say that quasi-observation done in simulations reflects observation in experiments.

Admittedly, quasi-intervention and quasi-observation do not represent intervention and quasi-observation in the same was as the dynamics of the computational model represents the reaction of the system experimented on. The computer scientist and her activities are not an imagined system that is described using equations, as is the target system of the simulation. Rather, the activities of the simulation scientist mimic intervention and observation in a way that is familiar from physical or material models: The behavior of one real-world system imitates that of another. In this sense, a computer simulation study is a hybrid model of an experiment because different types of models (viz. models qua imagined systems that are described with equations vs. models qua physical systems) are combined. I do not take this to be a problem for my claim. But if a coherent model is required, i.e., a model in which all parts of the experiments are modeled via the same type of model, we may say that there is in fact such a model, viz. an imagined experiment. This imagined experiment in turn is partly described using mathematical equations, and in other parts modeled by real quasi-intervention and quasi-observation. We would thus be talking about two layers of models, properly speaking, some of them being higher-order models (i.e., models of models). This suggestion is not ad hoc, since Küppers and Lenhard (2005) argue that computer simulations provide second-order models.[55]

In sum, then, in computer simulation studies, simulation scientists can carry out activities that closely resemble the activities constitutive of experiment: They can engage in quasi-intervention and quasi-observation, and do in fact often do so. Indeed, a study based upon interviews with working simulation scientists, Dowling (1999, pp. 268–270) stresses "tinkering" and "noticing" as important activities related to simulations. Both quasi-intervention and quasi-observation are plausibly thought to be used as surrogates for real intervention and real observation, respectively: Tinkering with the initial conditions in a simulation serves as a surrogate for a real intervention, where such an intervention is not possible. And clearly, quasi-observation is a surrogate for real observation. The simulation scientist wants to learn about the model (and then about the target), but she cannot observe the model (nor the target). What she can do though is to look at movies and other outputs that the simulations provide in order to learn about the model (and then the target).[56]

Note also that, if a CS includes quasi-intervention and quasi-observation, the latter are done with the superordinate aim of learning about the system on which an experiment is modeled. This is parallel to experiments, in which intervention and observation are done with the superordinate aim of learning about the object of the experiment (recall our third condition on experiment).

All this establishes CME_1: CSs can model possible experiments. To repeat, the idea is that CSs can not only model the reaction of a system to what may be thought of as possible intervention; also, quasi-intervention and quasi-observation can serve as surrogates for intervention on, and observation of, a real system, respectively. CSs thus allow scientists to learn by imitating the way in which experimenters learn from their experiments.[57]

CME_1 is meant to be conceptual; it follows from the fact that CSs evaluate a dynamical model of a (real or merely imagined) target system. If the model is evaluated, we may always say that the dynamics followed reflects the way in which the target systems reacts to a possible intervention. Further, almost any such model will allow for a variation in initial conditions or parameters, thus quasi-intervention is

[55] See also Beisbart (2014) for the various ways in which computer simulations are related to models.

[56] This analysis is focused on deterministic simulations, but can be generalized to Monte-Carlo simulations. The latter produce many sample trajectories of the computer model. Each sample trajectory arises from initial conditions subject to quasi-intervention and produces outputs that may be quasi-observed.

[57] I have here concentrated on the conditions on experimentation introduced in Section 2 above. These conditions have not been shown to be sufficient. This is not a problem because we here not interested in the claim that an experiment is indeed run, but only in the proposition that an experiment is modeled. Now a model need not fully reflect its target. Thus, not every condition on experiment needs to be reflected in the model, if only crucial aspects of experiments are represented. This clearly seems to be the case.

possible.[58] Also, whenever a CS is run, one can obtain outputs that reflect the states in which the target is according to the model. Thus, simulation scientists can carry out quasi-observation. They can do all this to learn about the system in a similar way as experimenters learn about their object.[59]

Apart from CME_1, my proposal has a second part, viz. CME_2: CSs do often model possible experiments. This is to add to CME_1 that simulations do often engage in quasi-intervention, that a run of the model does in fact reflect a reaction of the target system to a possible intervention and that quasi-observation is in fact done in order to learn in a way that closely reflects the way in which experimenters learn from their experiments. This is not a conceptual claim any more, but rather an empirical one. For the most part, CME_2 should be uncontroversial: It is almost always true that a CS reflects a reaction of its target to what may be thought of as result of a possible intervention and that there is quasi-observation with the aim of learning about the target. What is less trivial is quasi-intervention. There are some simulations that purport to represent the time evolution that their target actually takes (at least in some respect), and not the evolution it *would* take *if* a certain intervention took place. In such cases, we cannot say that setting the initial conditions of the simulations serves to model an intervention on the system. The initial conditions are rather meant to reflect the actual state of the target, and the simulationalist is not free to choose them. This means that such simulations are at most modeled experiments in a very thin sense. This is why CME_2 is restricted to many simulations and does not apply to all simulations. But we can safely assume that there are a lot of CSs in which quasi-intervention is done in order to learn in a way similar to experiments. They provide support for CME_2.[60]

The more restricted scope of CME_2 is not a problem for our purposes though because such simulations are not naturally called experiments or experiment-like even pre-theoretically. Since there are less similarities between these simulations and experiments, there is less to be accounted for, and we do not need to say that these simulations are modeled experiments.[61]

Some simulations refer to target systems that cannot be causally interfered with by humans. Simulations of cosmic structure formation are cases in point. Other simulations assume that the laws of nature are different from those in this world.[62] Setting the initial conditions and tinkering around with several parameters can nevertheless be conceptualized as a surrogate for an intervention, if only one that is not physically

[58] A model in which all initial conditions and parameter values are fixed seems highly artificial, and even in this case, one may vary some of the model assumptions, which would suffice for quasi-intervention.

[59] My claim that a CS can be or even is a modeled experiment is not meant to imply that this CS can be or is an experiment. A modeled experiment is not an experiment as fake snow is not snow.

[60] If a CS does not actually model a possible experiment because it is supposed to reflect the way the target system does behave as a matter of fact, we may still say that it models a *natural* experiment. In such an experiment, no intervention is needed simply because the system that is observed happens to fulfill the conditions that are at the focus of the inquiry.

[61] Some similarities on the list from Section 3 too apply to CSs of which CME_2 does not hold true. We can explain such similarities by saying that the simulations model *only the behavior of a system* as a reaction to certain conditions rather than also an intervention in which the system is subjected to the conditions.

[62] See e.g. Mainzer (1995), p. 467.

possible to us, but which would be of interest. For instance, we may say that the simulations are surrogates of experiments that only God can carry out. An analogous point holds about observation. Many systems that are modeled in simulations cannot be observed at all or are traced using unobservable characteristics. Nevertheless, quasi-observation can be thought of as a surrogate of an observation that is not possible to us, but that would be highly desirable.

So far, my argument for CME_1 and CME_2 is based upon a specific account of modeling that requires some similarity between the model system and the target. But some deny that such a similarity is needed for modeling. This is not a problem for my argument though. Note first that, for the purposes of my argument, the conditions on modeling need to be jointly sufficient. Whether all are necessary doesn't matter for the argument. Thus, even if similarity is not necessary for all models, it may be important for some models (this is the position of e.g. Hughes (1997)). If my argument so far goes through, CSs model experiments on the basis of similarity, even though other models may represent their targets without being similar to them.

It is second plausible to assume that my view may also be established under a different account of modeling. According to Goodman (1968), for instance, denotation is crucial for representation (see Ch. 1, particularly p. 5; a Goodmanian account of scientific modeling is defended by Hughes (1997)). In particular, in scientific models, parts of the model denote (i.e., stand for) parts of the target. Now if quasi-intervention denotes intervention, if the simulated dynamics denotes the reaction of the system and if quasi-observation denotes observation, then we can say that a CS study models an experiment. To show that e.g. quasi-intervention denotes intervention would lead us too far afield for the purposes of this paper.

This proposal, i.e. CME_1 and CME_2, can account for the similarities and the affinities between experiments and simulations noted in Section 3. For the similarities on the list on p. 7, we can show how they follow from our proposal or at least how this is plausible.[63]

1. *Something like intervention is crucial for experiments and for many simulations.* Intervention is crucial for experiments according to one of our conditions on experiments. It is useful because it enables scientists to study a system under various conditions that do not occur naturally (there may also be other reasons to intervene on a system, e.g. to discover causal relationships). Now, according to CME_2, many CSs do in fact model experiments. They thus contain a surrogate of intervention in that the scientist sets the initial conditions and the values of parameters according to her purposes. Very often, this is to model circumstances that do not occur naturally.
2. *CSs produce outputs that take a similar form as do data from experiments, at least if the outputs and data are suitably processed.* This follows from CME_1 because simulations allow for quasi-observation, which provides a surrogate for

[63]It would be too much to claim that my proposal provides an independent explanation of the similarities listed. The reason is that some similarities are built into the proposal. What the proposal does though is to re-organize the similarities in a useful way. Of course, CE (plus, maybe, CE+) potentially reorganizes the similarities too, but CE and CE+ have been rejected on independent grounds.

observation in experiments. It thus is also natural that the outputs of simulations and the data from experiments are often analyzed in the same way.

3. *Both experiments and CSs have results in the sense that they can show something. Scientists can obtain information about a system by experimenting on it or by simulating it.* It is uncontroversial that experiments produce results that show something about the system experimented on. CSs can have results too because they can model possible experiments (CME_1); in particular, they can model the reaction of their target to a potential intervention. The immediate result of an experiment is that the system experimented on behaves such and such under these and these conditions. Likewise, the result of a simulation is that the system that is modeled behaves such and such under these and these conditions. Simulationalists can thus learn about the behavior of their model system. This is of course only possible if the simulations provide a faithful model of their target. This needs to be shown using validation.
4. *The results of both experiments and CSs can often not be foreseen.* This point should be uncontroversial concerning experiments. Now, if we cannot predict the results of an experiment and if we model this experiment faithfully in a CS (CME_1), we cannot expect that we can predict the results regarding the model before analyzing the model, and this analysis is most often done by running the simulation. So we cannot expect to know the results before running the CS. Note though that the unpredictability of simulation results is not intentionally built into CSs in order to represent real experiments faithfully. It is not only a consequence of the fact that the result of the experiment cannot be foreseen, but also from the fact that having defined a simulation model does not yet yield a prediction for the outcome of the result. This is so because simulation models are often extremely complex such that their consequences cannot be foreseen.
5. *The epistemologies of experiments and CSs share important similarities in that scientists use similar strategies when they use the outputs of experiments and simulations to justify certain claims.* This is a consequence of the similarities 1 and 2 given above.

We can now comment on the status of the similarities listed. While the second and the third follow from essential characteristics of experiments and the conceptual truth CME_1, the first draws on an essential feature of experiments, but is contingent concerning simulations. The fourth is more contingent because it holds only contingently already as far as experiments are concerned. If the first is important for establishing the fifth, the latter is contingent. But the truth seems that some similarities in the epistemologies of CS and experiment only depend on the second similarity, which would make the similarities less contingent.

Another noteworthy fact about simulations and experiments was that the latter can sometimes be replaced using the former. This is clearly entailed by my proposal, which takes at least some simulations to be surrogates for possible experiments.

Even though my proposal accounts for the affinities between experiments and simulations, it does not force us to weep disanalogies under the carpet. Nor does it imply dubious tenets about simulations. We are not committed to say that the targets of CSs are intervened on. The target as well as an intervention on it are at most modeled.

Nor have we to say that the target is observed. The behavior of the target is modeled, and there is a surrogate for observation. Nor is lack of over-control implied by the proposal. The results of simulations are over-controlled because they are implied by the related model, while there is no over-control in experiments.

My proposal is also relevant for the epistemology of computer simulation (cf. the claim CE+ above). For it allows us to explain how CSs can produce new knowledge about the world. For, if CSs can model possible experiments and if experiments can produce new knowledge about the world, then CSs can produce such new knowledge. An important condition though is that the possible experiment be properly modeled in the CS. This is as it should be since computer simulations can fail to faithfully reflect the outcome of a possible experiment. Note also that this explanation of how CSs *can* produce new knowledge does not imply that every and each CS produces new knowledge in the way just described. As we have seen, some CSs do not model experiments. As a consequence, such CSs have to produce new knowledge in a different way.[64]

A seemingly different way of accounting for the similarities between CSs and experiments is to say that both instantiate a more general type or pattern. In this vein, Skaf and Imbert (2013) argue that at least some computer simulations and experiments share the same function, viz. what they call unfolding and what they take to be a very important task of the sciences.[65] All this is compatible with what I claim in this paper. If my account succeeds, I take it to have the advantage that it does not introduce more general terms, but rather sticks with notions that we need anyway to account for the sciences.

A possible problem for my account is that there seems to be the following tension: On the one hand, I stress that experiments and CSs differ in that the former are free of over-control, while the latter are not. But on the other hand, I claim that CSs can model possible experiments and even replace them under suitable conditions. How does this fit together?[66]

In my view, the tension can to some extent be argued away. For one thing, note that a model need not be similar to its target in every respect. So it is not a problem per se if CS are subject to over-control, while experiments are not. For another thing, it is plausible to think that the modeled experiment is indeed *not* over-controlled. There would be over-control if the assumptions needed to carry out the modeled experiment determined the result. Now, the simulation scientist needs to have assumptions that determine the result of the modeled experiment. But this does not imply that these assumptions are needed to carry out the modeled experiment. The modeled experiment may be thought of as part of an imagined world. Knowledge that is needed to sustain the imagination need not part of the imagined world.[67]

[64]They may do so by imitating *natural* experiments, which are not experiments according to our partial explication.

[65]See also Imbert (2017), Sec. 5.2 for a similar strategy.

[66]All this was very helpfully pointed out by an anonymous referee who also noted that the problem is not just restricted to my account, but rather affects any view that embraces the following two claims: i. Experiments are not over-controlled, while CSs are. ii. CSs can replace experiments.

[67]In a similar way, Nagel (1986), p. 93 criticizes Berkeley's so-called master argument because Berkeley confuses something that is needed to create an image with the content of the image.

To wrap up, my proposal that simulations can model possible experiments and that many do in fact do so has a number of merits. It accounts for the similarities between simulations and experiments while not concealing the distinction between simulations and experiments. The proposal is also compatible with the view that simulations function in the same way as thought experiments do (see Beisbart 2012).[68]

9 Conclusions

Computer simulations bear a number of pre-theoretical similarities to experiments. The similarities may be accounted for by saying that CSs are or include experiments and one may go further and say that the epistemic power of CSs derives from their being experiments. In this paper, I have argued against such views. CSs do not include experiments *on the hardware* of a computer because the hardware of the computer is not observed to learn about it. Nor is the hardware always interfered with, properly speaking. They do not include experiments *on the target*, e.g. because experiments require that nature gives an answer to a question posed by an experimenter. Nature does not give such an answer because simulations are over-controlled by modeling assumptions. There are thus epistemologically significant differences between experiments and simulations. The epistemic power of simulations cannot be explained by saying that they are experiments.

The existing similarities between many simulations and experiments should rather be re-organized by saying that CSs can model possible experiments at two levels and that they do in fact often do so. First, simulation scientists can carry out activities that serve as surrogates for intervention and observation in real experiments with an superordinate aim that parallels that of an experiment. Second, the reaction of the system is modeled using the model that the simulations define. In this way, we can explain how a lot of computer simulations produce new knowledge: They do so because they represent possible experiments, which can in turn produce new knowledge. Importantly, my account stresses that computer simulation is a type of modeling. It draws our attention to the modeling assumptions on which simulations rely.

Acknowledgments Thanks to Christoph Baumberger and Trude Hirsch Hadorn for extremely useful comments on an earlier version of this manuscript. I'm also very grateful for detailed and helpful comments and criticisms by two anonymous referees. One of them provided extensive, constructive and extremely helpful comments even about a revised version of this paper – thanks a lot for this!

[68] Can my main thesis, viz. that CSs can, and do often do, model possible experiments be generalized to what is called analog simulation? I here take it that, in an analog simulation, the target and a physical model of it may be described using the same type of dynamical equations. The dynamics of the model then is investigated to learn about the dynamics of the target (Trenholme 1994). For instance, electric circuits may be used to study a fluid in this way (see Kroes (1989) for this example). Now, my claim that CSs can model possible experiments and often do so seems to apply to such analog simulations, too. For instance, if I set up an electric circuit to model a particular fluid, I'm modeling an experiment on the fluid. Note, however, the following difference between analog and computer simulations: If a possible experiment on the target is modeled in an analog simulation, this is a real experiment on the model system (the analogue). As I have argued above in Section 5, this is not so in computer simulations.

References

Adam, M. (2002). *Theoriebeladenheit und Objektivität. Zur Rolle von Beobachtungen in den Naturwissenschaften*. Frankfurt am Main und London: Ontos.

Arnold, E. (2013). Experiments and simulations: Do they fuse? In Durán, J.M., & Arnold, E. (Eds.) *Computer simulations and the changing face of scientific experimentation* (pp. 46–75). Newcastle upon Tyne: Cambridge Scholars Publishing.

Balzer, W. (1997). *Die Wissenschaft und ihre Methoden*. Freiburg und München: Karl Alber.

Barberousse, A., Franceschelli, S., & Imbert, C. (2009). Computer simulations as experiments. *Synthese*, *169*, 557–574.

Barker-Plummer, D. (2016). Turing machines. In Zalta, E.N. (Ed.) *The Stanford encyclopedia of philosophy.* Winter 2016 edn, Metaphysics Research Lab, Stanford University.

Baumberger, C. (2011). Understanding and its relation to knowledge. In Löffler, C.J.W. (Ed.) *Epistemology: contexts, values, disagreement. Papers of the 34th international Wittgenstein symposium* (pp. 16–18). Austrian Ludwig Wittgenstein Society.

Beisbart, C. (2012). How can computer simulations produce new knowledge? *European Journal for Philosophy of Science*, 2(2012), 395–434.

Beisbart, C. (2014). Are we Sims? How computer simulations represent and what this means for the simulation argument. *The Monist*, *97/3*, 399–417.

Beisbart, C., & Norton, J.D. (2012). Why Monte Carlo simulations are inferences and not experiments. *International Studies in the Philosophy of Science*, *26*, 403–422.

Bertschinger, E. (1998). Simulations of structure formation in the Universe. *Annual Review of Astronomy and Astrophysics*, *36*, 599–654.

Binder, K., & Heermann, D. (2010). *Monte Carlo simulation in statistical physics: An introduction, graduate texts in physics*. Berlin: Springer Verlag.

Bogen, J. (2010). Theory and observation in science. In Zalta, E.N. (Ed.) *The stanford encyclopedia of philosophy.* Spring 2010 edn. http://plato.stanford.edu/archives/spr2010/entries/science-theory-observation/.

Brown, J.R., & Fehige, Y. (2017). Thought experiments. In Zalta, E.N. (Ed.) *The stanford encyclopedia of philosophy.* Summer 2017 edn.

Carnap, R. (1962). *Logical foundations of probability*, 2nd edn. Chicago: University of Chicago Press.

Casti, J.L. (1997). *Would-be worlds. How simulation is changing the frontiers of science*. New York: Wiley.

Dolag, K., Borgani, S., Schindler, S., Diaferio, A., & Bykov, A.M. (2008). Simulation techniques for cosmological simulations. *Space Science Reviews*, *134*, 229–268. arXiv:http://arxiv.org/abs/0801.1023v1.

Dowling, D. (1999). Experimenting on theories. *Science in Context*, *12/2*, 261–273.

Duhem, P.M.M. (1954). *The aim and structure of physical theory, Princeton science library*. Princeton, NJ: Princeton University Press.

Durán, J.M. (2013). The use of the materiality argument in the literature on computer simulations. In Durán, J.M., & Arnold, E. (Eds.) *Computer simulations and the changing face of scientific experimentation* (pp. 76–98). Newcastle upon Tyne: Cambridge Scholars Publishing.

Efstathiou, G., Davis, M., White, S.D.M., & Frenk, C.S. (1985). Numerical techniques for large cosmological N-body simulations. *Ap J Suppl*, *57*, 241–260.

Falkenburg, B. (2007). *Particle metaphysics. A critical account of subatomic reality*. Heidelberg: Springer.

Franklin, A. (2010). Experiment in physics. In Zalta, E.N. (Ed.) *The Stanford encyclopedia of philosophy.* Spring 2010 edn.

Frigg, R.P., & Reiss, J. (2009). The philosophy of simulation: Hot mew issues or same old stew? *Synthese, 169*, 593–613.

Fritzson, P. (2004). *Principles of object-oriented modeling and simulation with Modelica 2.1*. IEEE Press.

Giere, R.N. (2004). How models are used to represent. *Philosophy of Science*, *71*, 742–752.

Giere, R.N. (2009). Is computer simulation changing the face of experimentation? *Philosophical Studies, 143*(1), 59–62.

Gillespie, D.T. (1976). A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, *22*, 403–434.

Goodman, N. (1968). *Languages of art: An approach to a theory of symbols*. Indianapolis: Bobbs-Merrill.

Gramelsberger, G. (2010). *Computerexperimente. Zum Wandel der Wissenschaft im Zeitalter des Computers*. Transcript, Bielefeld.

Guala, F. (2002). Models, simulations, and experiments. In Magnani, L., & Nersessian, N. (Eds.) *Model-based reasoning: science, technology, values* (pp. 59–74). New York: Kluwer.

Guillemot, H. (2010). Connections between simulations and observation in climate computer modeling. scientist's practices and bottom-up epistemology lessons. *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics*, *41*, 242–252. Special Issue: Modelling and simulation in the atmospheric and climate sciences.

Haasl, R.J., & Payseur, B.A. (2011). Multi-locus inference of population structure: a comparison between single nucleotide polymorphisms and microsatellites. *Heredity*, *106*, 158–171.

Hacking, I. (1983). *Representing and intervening*. Cambridge: Cambridge University Press.

Hasty, J., McMillen, D., Isaacs, F., & Collins, J.J. (2001). Computational studies of gene regulatory networks: In numero molecular biology. *Nature Reviews Genetics*, *2*, 268–279.

Heidelberger, M. (2005). Experimentation and instrumentation. In Borchert, D. (Ed.) *Encyclopedia of philosophy. Appendix* (pp. 12–20). New York: Macmillan.

Hughes, R.I.G. (1997). Models and representation. *Philosophy of Science (Proceedings)*, *64*, S325–S336.

Hughes, R.I.G. (1999). The Ising model, computer simulation, and universal physics. In Morgan, M.S., & Morrison, M. (Eds.) *Models as mediators. Perspectives on natural and social sciences* (pp. 97–145). Cambridge: Cambridge University Press.

Humphreys, P. (1990). Computer simulations. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, *1990*, 497–506.

Humphreys, P. (1994). Numerical experimentation. In Humphreys, P. (Ed.), *Patrick Suppes. Scientific philosopher* (Vol. 2, pp. 103–118). Dordrecht: Kluwer.

Humphreys, P. (2004). *Extending ourselves: Computational science, empiricism, and scientific method*. New York: Oxford University Press.

Humphreys, P.W. (2013). What are data about? In Durán, J.M., & Arnold, E. (Eds.) *Computer simulations and the changing face of scientific experimentation* (pp. 12–28). Newcastle upon Tyne: Cambridge Scholars Publishing.

Hüttemann, A. (2000). Natur und Labor. Über die Grenzen der Gültigkeit von Naturgesetzen. *Philosophia Naturalis*, *37*, 269–285.

Imbert, C. (2017). Computer simulations and computational models in science. In Magnani, L. & Bertolotti, T. (Eds.) *Springer handbook of model-based science* (Vol. 34, pp. 733–779), Cham, chapter: Springer.

Janich, P. (1995). Experiment. In Mittelstraß, J. (Ed.) *Enzyklopädie Philosophie und Wissenschaftstheorie. Band 1, Metzler, Stuttgart* (pp. 621–622).

Kant, I. (1998). *Critique of pure reason*. Cambridge: Cambridge University Press. translated by P. Guyer and A. W. Wood; Cambridge Edition of the Works of Kant.

Keller, E.F. (2003). Models, simulation, and computer experiments. In Radder, H. (Ed.) *The philosophy of scientific experimentation* (pp. 198–215). Pittsburgh: University of Pittsburgh Press.

Knorr-Cetina, K. (1981). *The manufacture of knowledge: An essay on the constructivist and contextual nature of science*. Pergamon international library of science, technology, engineering, and social studies, Pergamon Press.

Kroes, P. (1989). Structural analogies between physical systems. *British Journal for the Philosophy of Science*, *40*, 145–154.

Küppers, G., & Lenhard, J. (2005). Computersimulationen: Modellierungen 2. Ordnung. *Journal for General Philosophy of Science*, *36*(2), 305–329.

Lim, S., McKee, J.L., Woloszyn, L., Amit, Y., Feedman, D.J., Sheinberg, D.L., & Brunel, N. (2015). Inferring learning rules from distributions of firing rates in cortical neurons. *Nature Neuroscience*, *18*, 1804–1810.

Mainzer, K. (1995). *Computer – neue Flügel des Geistes? Die Evolution computergestützter Technik, Wissenschaft, Kultur und Philosophie*, 2nd edn. Berlin, New York: de Gruyter Verlag.

Metropolis, N., & Ulam, S. (1949). The Monte Carlo method. *Journal of the American Statistical Association*, *44*(247), 335–341.

Michelson, A.A. (1881). The relative motion of the earth and the luminiferous ether. *American Journal of Science*, *22*, 120–129.

Michelson, A.A., & Morley, E.W. (1887). On the relative motion of the earth and the luminiferous ether. *American Journal of Science*, *34*, 333–345.

Morgan, M.S. (2002). Model experiments and models in experiments. In Magnani, L., & Nersessian, N. (Eds.) *Model-based reasoning: science, technology, values* (pp. 41–58). New York: Kluwer.

Morgan, M.S. (2003). Experimentation without material intervention: Model experiments, virtual experiments, and virtually experiments. In Radder, H. (Ed.) *The philosophy of scientific experimentation* (pp. 216–235). Pittsburgh: University of Pittsburgh Press.

Morgan, M.S. (2005). Experiments versus models: New phenomena, inference and surprise. *Journal of Economic Methodology*, *12*(2), 317–329.

Morrison, M. (1998). Experiment. In Craig, E. (Ed.) *Routledge encyclopedia of philosophy* (Vol. III, pp. 514–518). London: Routledge and Kegan.

Morrison, M. (2009). Models, measurement and computer simulation: The changing face of experimentation. *Philosophical Studies*, *143*, 33–57.

Nagel, T. (1986). *The view from nowhere*. Oxford: Oxford University Press.

Naumova, E.N., Gorski, J., & Naumov, Y.N. (2008). Simulation studies for a multistage dynamic process of immune memory response to influenza: Experiment in silico. *Annales Zoologici Fennici*, *45*, 369–384.

Norton, J.D. (1996). Are thought experiments just what you thought? *Canadian Journal of Philosophy, 26*, 333–366.

Norton, S.D., & Suppe, F. (2001). Why atmospheric modeling is good science. In Edwards, P., & Miller, C. (Eds.) *Changing the atmosphere* (pp. 67–106). Cambridge, MA: MIT 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*(3), 483–496.

Parker, W.S. (2010). An instrument for what? Digital computers, simulation and scientific practice. *Spontaneous Generations*, *4*(1), 39–44.

Peschard, I. (forthcoming). Is simulation a substitute for experimentation? In Vaienti, S., & Livet, P. (Eds.) *Simulations and networks*. Aix-Marseille: Presses Universitaires d'Aix-Marseille. Here quoted after the preprint http://d30056166.purehost.com/Is_simulation_an_epistemic%20_substitute.pdf.

Press, W.H., Teukolsky, S.A., Vetterling, W.T., & Flannery, B.P. (2007). *Numerical recipes. The art of scientific computing*, 3rd edn. New York: Cambridge University Press.

Radder, H. (2009). The philosophy of scientific experimentation: A review. *Automatic Experimentation* 1. open access; http://www.aejournal.net/content/1/1/2.

Radder, H. (Ed.) (2003). *The philosophy of scientific experimentation*. Pittsburgh: University of Pittsburgh Press.

Rechenberg, P. (2000). *Was ist Informatik? Eine allgemeinverständliche Einführung*, 3rd edn. München: Hanser.

Rheinberger, H.J. (1997). *Toward a history of epistemic things: Synthesizing proteins in the test tube*. Writing science, Stanford University Press.

Rohrlich, F. (1990). Computer simulation in the physical sciences. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, *1990*, 507–518.

Scholz, O.R. (2004). *Bild, Darstellung, Zeichen. Philosophische Theorien bildlicher Darstellung*, 2nd edn. Frankfurt am Main: Vittorio Klostermann.

Shapere, D. (1982). The concept of observation in science and philosophy. *Philosophy of Science*, *49*(4), 485–525.

Skaf, R.E., & Imbert, C. (2013). Unfolding in the empirical sciences: experiments, thought experiments and computer simulations. *Synthese*, *190*(16), 3451–3474.

Stöckler, M. (2000). On modeling and simulations as instruments for the study of complex systems. In Carrier, M., Massey, G.J., & Ruetsche, L. (Eds.) *Science at the century's end: Philosophical questions on the progress and limits of science* (pp. 355–373). Pittsburgh, PA: University of Pittsburgh Press.

Suárez, M. (2003). Scientific representation: Against similarity and isomorphism. *International Studies in the Philosophy of Science*, *17*, 225–244.

Suárez, M. (2004). An inferential conception of scientific representation. *Philosophy of Science*, *71*, 767–779.

Sugden, R. (Ed.) (2005). *Experiment, theory, world: A symposium on the role of experiments in economics*, Vol. 12/2. London: Routledge. Special issue of Journal of Economic Methodology.

Tiles, J.E. (1993). Experiment as intervention. *British Journal for the Philosophy of Science*, *44*(3), 463–475.

Trenholme, R. (1994). Analog simulation. *Philosophy of Science*, *61*(1), 115–131.

Turing, A. (1937). On computable numbers, with an application to the entscheidungsproblem. In *Proceedings of the London mathematical society* (Vol. s2–42, no. 1).

Weber, M. (2005). *Philosophy of experimental biology*. Cambridge: Cambridge University Press.

Weisberg, M. (2007). Who is a modeler? *British Journal for Philosophy of Science, 58*, 207–233.

Winsberg, E. (1999). Sanctioning models. The epistemology of simulation. *Science in Context*, *12*, 275–292.

Winsberg, E. (2003). Simulated experiments: Methodology for a virtual world. *Philosophy of Science*, *70*, 105–125.

Winsberg, E. (2009a). Computer simulation and the philosophy of science. *Philosophy Compass*, *4/5*, 835–845.

Winsberg, E. (2009b). A tale of two methods. here quoted from Winsberg (2010) Ch. 4 pp. 49–71.

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

Zimmerman, D.J. (2003). Peer effects in academic outcomes: Evidence from a natural experiment. *The Review of Economics and Statistics*, *85*(1), 9–23.