IS SIMULATION AN EPISTEMIC SUBSTITUTE FOR EXPERIMENTATION?[1]

Isabelle Peschard

San Francisco State University

0. Introduction

It is sometimes said that simulation can serve as epistemic substitute for experimentation (e.g. Hartman 1996). Such a claim might be suggested by the fast-spreading use of computer simulation to investigate phenomena not accessible to experimentation (in astrophysics, ecology, economics, climatology, etc.). But what does that mean? As a substitute, computer simulation should be an alternative means to achieve some identical end. In the weakest sense, it may simply mean that simulation produces information. But there is a much stronger and challenging possible interpretation of the notion of epistemic substitute that is the object of this paper. On this interpretation, simulation is construed as 'epistemically on a par' with experimentation. Computer simulation, as epistemic substitute, can be a different means to learning the same thing, that is, what we would have learnt with experimentation had experimentation be possible. Recent comparative analyses of the methodology and epistemology of computer simulation and experimentation offer some support for this view (Winsberg 2009, Morrison 2009, Norton and Suppe 2001). And, even though the strength of this support will be called into question, these studies show two things that really motivate this paper: 1) that this view cannot be simply taken for granted nor easily rejected and 2) that what is at issue is not only the epistemic function of simulation but also that of experimentation and the relation between the two.

A guiding intuition throughout will be that in experimentation, the system under study is interacted with and given the opportunity to somehow express itself via a causal effect on

[1] The author wishes to acknowledge support for this research by NSF grant SES-1026183 and thank Ronald Giere, Eric Winsberg and Frank Varenne for helpful discussions.

the instruments. In simulation, by contrast, what is expressed in the results is the model used in the simulation, assumed here, in the case of computer simulation, to be a mathematical model. Apparently, the epistemic distinction is clear: experimentation produces information about how the system behaves; simulation produces information about a model. Experimental results call for a model to explain the data that are produced. Simulation shows what data could be explained by a given model, the one used in the simulation. As we will see, however, this quick comparison, is, well, too quick. Closer inspection reveals significant methodological and epistemological similarities between simulation and experimentation that preclude hasty conclusions. If there is a basis to distinguish the epistemic function of simulation and experimentation, it is not in the appearances that might feed our intuition.

The paper starts with a clarification of the terms of the issue and then focuses on two powerful arguments for the view that simulation and experimentation are 'epistemically on a par'. One is based on the claim that, in experimentation, no less than in simulation, it is not the system under study that is manipulated but a system that 'stands-in' for it. The other one highlights the pervasive use of models in experimentation.

It will be argued that these arguments, as compelling as they might seem, are each based on a mistaken interpretation of experimentation and that, far from simulation and experimentation being epistemically on a par, they do not have the same epistemic function, do not produce the same kind of epistemic results. If that is so, comparing the epistemic power of these activities yields limited return. Rather we need to get clearer on the differences between their epistemic functions so as to better understand how these functions co-operate in the investigation of a given system. Some suggestions about the main elements of this co-operation will be given and illustrated at the end of the paper.

1. Basics: simulation, experimentation

Let's start with some elementary characterization of experimentation and simulation that can be used as uncontroversial basis for a more precise and, maybe controversial, characterization of the way these activities function as epistemic instrument: instrument to produce information.

Experimental knowledge claims are based on instrumental interactions with a physical system. What sort of claims they are, what they are about, what the relation is between what they are about and what is manipulated, are questions we will come back to later. Most generally and minimally, for now, experimentation on a physical system S can be thought of as a procedure that consists in, at least:

1) Preparing the system S in a certain state, by fixing initial and boundary conditions, and selectively putting under control the parameters that have an effect on the outcomes of measurement, the *active parameters*.[2]

2) Letting the system evolve. The evolution of S is characterized by the evolution of a set of physical quantities characterizing the state of S, the *state variables*.

3) Recording the evolution of S through a sequence of states when the values of some of these parameters are varied; analyzing the results.

"Computer simulation" is used in two main senses[3] often not clearly distinguished:

1) simulation of a mathematical model, that is, a computer-implemented procedure that solves an approximation of a system of equations;
2) simulation of a physical system S, that is, a computer-implemented procedure of imitation, mimicking the evolution of S as a result of the implementation of a mathematical model of S.

In both senses, the simulation involves a process undergone by a computer, a computational process. Hereafter, as is commonly done, I will take this latter sense of simulation as having the former sense built in. When one speaks of the computer simulation of, say, a fluid mechanical system, it is presupposed that the simulation of this physical system is realized by carrying out a computer simulation of a system of equations that represents the evolution of the fluid mechanical system, that is, of a mathematical model of this physical system.

---

[2] Control of the parameters means either neutralizing their effect or manipulating their value.
[3] See Humphreys 1994, Fox Keller 2003, Varenne 2010.

Often attached to the notion of simulation is the idea of “experimenting with a model” (Ören 2011). Unsurprisingly, this aspect of simulation will prove relevant to the comparison of the epistemic functions of simulation and experimentation.

But before going further, it is useful to clarify what is at issue when we compare the epistemic functions of simulation and experimentation. We are not comparing simulation and experimentation ‘in general’. What needs to be compared is what can be learned about a given system. Of course, that cannot be done when experimentation is not possible. But it can be done for cases, probably the most current, where simulation and experimentation are used in tandem, that is, where they are directed at the same system S under study. The use of simulation in tandem with experimentation seems to capture less philosophical interest than its use where experimentation is not possible. But this kind of use is the most appropriate for a comparison of epistemic functions because it enables us to put their epistemic results, so to speak, side-by-side. And there is no reason why the conclusions from this configuration about epistemic functions would not be of relevance to the case where experimentation is not possible.

The tandem configuration has an additional virtue. Simulation has a multiplicity of epistemic functions. They are not all relevant to our concern. We are specifically concerned with the use of simulation in the investigation of system that could be, in principle, object of experimental investigation, and with the epistemic function of simulation associated with this sort of use. The tandem configuration enables us to zoom in on this epistemic function.

2. Simulation as experimentation

The idea that simulation is a form of experimentation plays substantial role in the view that simulation and experimentation are ‘epistemically on a par’, so it is worth making it clearer.

Following our above minimal characterization of experimentation, simulation can be analyzed into three similar sub-processes:

1) The preparation of a system: it consists in the transformation of a set of initial

equations into an algorithm suitable for computation given certain constraints of time, computational power, and accuracy (Winsberg 2003), and the implementation on the computer of this algorithm, the simulation model, with an assignment of numerical values to constants and parameters.

2) The evolution of the system: the autonomous transformation over time of the physical system that implements the computation (Humphreys 1994; Norton and Suppe 2001). It may be argued that, in a certain sense, the model that is simulated also undergoes a transformation (Krohs 2008)[4].

3) The third is the recording, organization, and classification of the results in the form of models of the data (Winsberg 2003).

That simulation is a form of experimentation is an important ingredient in the discussion but it is not determining. It is not determining because we need to know how this form of experimentation, by contrast to what is traditionally identified as experimentation, relates to the system under study, the target system. But it is important because how it relates to this system will depend on what sort of activity it is. To take it that we can start with a clear distinction between simulation and 'what is traditionally identified as experimentation' may seem to be begging the question in favor of a distinction between these activities and their epistemic functions. It may turn out that on closer inspection at least some of these distinctions will vanish as illusory. But that we generally use the terms 'simulation' and 'experimentation' to distinguish two kinds of activities would remain (Winsberg 2009). If we were not able, whatever the basis for it, to make such a distinction, the question of the comparison would not even arise. What is questioned is what these activities consist in, in particular regarding experimentation. So I will propose a distinction between simulation and experimentation which will have decisive consequences regarding their epistemic function, but rather than begging the question I take this contrast to be an offer to be 'put to the test'.

[4] There may be something unsettling about speaking in this way of the transformation of the model simulated. But it becomes quite intuitive if one considers that the simulation produces a series of sets of values forming different realizations of the simulation model, in the same way as, say, a series of values of temperature would characterize different states of the same fluid system.

The starting point of the discussion is a system S under investigation and a conception of simulation as experimenting on a model of the system S. The immediate results of the simulation are sets of values that are successive realizations of the model. In that sense, they are about the object manipulated, the model. By contrast, it is proposed that experimentation results come directly from a manipulation of and causal interaction with the system S itself. In this sense of causally mediated relation, experimentation comprises a 'direct access' to S. It is anticipated that a causally mediated relation with S should enable experimentation to produce information about S that the manipulation of a model of it might not be able to produce.

In the next section, I will discuss two arguments against the distinction between simulation and experimentation in terms of direct access to the system under study.

3. False similarities

There are two main objections to the idea that direct access to the target system distinguishes experimentation from simulation:

1) no more than simulation does experimentation consist in the actual manipulation of the system of interest;

2) experimentation no less than simulation involves the use of and dependence on models.

It is important for the discussion to be clear on the different elements that are involved in simulation and experimentation for, as we will see, there are some distinctions that are not always made and are crucial to the issue.

An epistemic activity such as simulation or experimentation is specifically designed to produce information about a certain system: its **target system.** Both in experimentation and simulation, the investigation of a given system involves manipulating a certain object: the **object manipulated**. In the case of a simulation, that the target system, whose behavior is simulated, and the object manipulated in the simulation, the model implemented on the computer, are distinct, is not controversial.

In addition to these notions, common in recent studies on simulation and experimentation, we need two new concepts. In either simulation or experimentation there is certain information about the target system that the procedure aims to obtain: its **epistemic target.** This needs to be taken narrowly. Suppose that the motivation for some research is to find out the effect of a certain drug on humans, while the experiment is conducted on rats. The effects of the drug on rats are reported as the research result: that was the epistemic target. But the **epistemic motivation** for the study was to obtain information concerning effects on humans. At the cost of multiplying terms we might speak of a motivating 'remote' target system and a motivating 'remote' epistemic target. These might be mentioned in the original grant proposal or in an introductory paragraph of the eventual journal article, but will have to be distinguished from the actual reported or hoped for research result.

3.1 Experimentation doesn't manipulate the system of interest either?

As we said, the object manipulated in a simulation and the target system of the simulation are clearly distinct. The question is whether there is such a distinction between object manipulated and target system in the case of experimentation. According to Francesco Guala (2008), there is a " 'representative analogy' between experiments and models: both stand for some other system, or set of systems, that is the ultimate target of investigation". Similarly, Eric Winsberg (2009) argues that both in simulation and in experimentation, results about the target system will be inferred from the results about the object that was manipulated. Guala locates the difference between simulation and experimentation in "the kind of relationship existing between, on the one hand, an experimental and its target system, and, on the other, a simulating and its target system": "the mapping in a simulation is purely formal, whereas experiments are materially analogous to their target system". For Winsberg, the relevant distinction is strictly epistemological, in the kinds of justification supporting the inference from the result of the manipulation to claims about the target system, but with no reason to believe that experimentation is "epistemically more powerful than simulation" (Winsberg 2009:591).

Winsberg makes a compelling argument to the effect that if in experimentation, just as in simulation, what is manipulated is a system standing in for the target system, there is no basis for drawing a principled distinction between the epistemic functions of simulation and experimentation. But the premise of the argument, that in both cases what is manipulated is a system standing in for the target system, is questionable. First, it does not seem to be necessarily the case in experimentation, by contrast with simulation, that the system manipulated is different from the target system. And secondly, when the two are distinct in an experiment, the relation between them is different from what it is in the case of a simulation.

Regarding the first point, the distinction between system manipulated and target system finds its intuitive motivation in the observation that what we want to learn about is the world but the world is too complex to be manipulated. We want to learn about wakes around islands or wings or poles, but we manipulate tame wakes around polished cylinders in shielding laboratories. We want to learn about human reactions to drugs, but we manipulate rats. The system manipulated, it then seems, is not the one we really want to learn about. On the basis of this distinction, Guala (2003) distinguishes the problem of internal validity from the problem of external validity. The former is related to the validity of inferences about the system manipulated whereas the latter concerns the validity of inferences about the target system on the basis of the results about the system manipulated.

Interestingly, however, Guala (2008) notes that "experimental physicists do not recognize external validity as a separate problem of inference", and more generally, that "[i]t should be stressed that experimenters are often concerned with proving the existence of certain mechanisms or phenomena in the lab only, and leave it to policy-makers or applied economists to apply such knowledge in the field". But if scientists are not concerned with drawing inferences about the system 'in the field', then it seems inappropriate to take this system to be the target system. The idea here is not to regiment the use of a term. Rather, it is that if we are to compare the epistemic functions of simulation and experimentation, we should better make sure that the term 'target system' picks out the same system for simulation and experimentation when they are used in tandem. The target system of the simulation is the system represented by the model manipulated: it is the system which the simulation is designed to produce information about. It is this system that needs to be identified as target system of the experimentation as well. And if that is so, the system in the field is not the target system of experimentation, not in the sense in which the target system of the simulation is the system that is represented by this model used in the simulation.

The confusion at the source of this misidentification of the target system of experimentation is the lack of distinction between the **target system**, identified clearly through the reported research results, and the **epistemic motivation**. The latter certainly motivates the procedure and the epistemic function of this motivation is important. Still, going back to the experimental study on rats, the motivation for such a study is, we said, human reaction to the drug. The same experimentation might have had a different motivation. And the same experimental results might even turn out to be unexpectedly informative about systems different from the original motivation. On the other hand, the same epistemic motivation may motivate different sorts of experimental study, manipulating different systems.

To illustrate this distinction, let's introduce briefly a particular experimental study in fluid mechanics, which will be used again in the next section to compare the epistemic functions of experimentation and simulation. A classic object of study in fluid mechanics is the wake that forms behind a blunt obstacle. Wakes are found everywhere in nature, behind islands, rocks, poles, wings. In a laboratory, it is typically created by placing a cylinder perpendicular to the upstream direction of a flow and increasing the velocity of the flow. Beyond a certain threshold of the control parameter vortices are emitted periodically behind the cylinder and form a wake flowing downstream (see e.g. Williamson 1989).

Some of these experiments aimed to specify the evolution of the shedding frequency of the vortices with the control parameter. That is what the results published as conclusion of the study were about, and they were about the system that was manipulated. This system is the one under study, the target system. Of course, these results may be used for new studies interested, for instance, in what happens behind coupled wakes in a lab, or in what happens behind an island, in the field. The system envisaged for further application might have been an epistemic motivation for the experimentation that was conducted. But those uses of the results would pertain to a different experimental enterprise.

More will be said later about what exactly we learn from these results—the reported research results do not include merely incidental features of the very particular system that was manipulated. If there is no artifact in the results, we learn something about any system that is relevantly similar to the one that is manipulated. This is what reproducibility and reliability are all about. But the point remains that in these experiments the object manipulated and target system coincide.

What is then the target system of the experimentation? Just as in simulation, it is what the procedure of manipulation is specifically designed to learn about. It is what the conclusions of the experimental study are about, which is different from the speculations they may encourage. It is the system that is manipulated; at least it is so in a large number of cases, and especially in physics.

That may not, however, always be the case. There are cases of experimentation where object manipulated and target system do not coincide: typically, when the system manipulated is a sample from a population about which the manipulation is designed to make conclusions.

Take, for instance, the investigation of conditional reasoning in cognitive psychology. A typical experimental set up for this investigation is the Wason task. Subjects are shown four double-sided cards, showing the symbols A, K, 2, 7. They are told that each card has a number on one side and a letter on the other. The subjects are asked which cards need to be turned over in order to test or falsify the hypothesis that "if there a vowel on one side, there is an even number on the other" (Wason 1968). It is clear that on the basis of the experimental results, scientists draw conclusions that go well beyond the individuals who performed the task. The procedure is rather similar to polls probing some features of a population. Under the assumption that the sample is representative of the whole, the results are about the whole population. In this case, the former seems to qualify as much as the latter for the status of target system[5]. One may well say that being representative amounts to standing in, in the way that in a simulation the model stands in

[5] Note that in the case of the wake also one could speak of the particular system that was manipulated as representative of something else: a class. But a class doesn't fit well the role of target system whereas in the Wason experiment, the system manipulated is representative of another *system*: the whole population from which the sample is taken.

for the system it is a model of. But there is an important difference though: the model is not *representative* of the system it represents, it is meant to be a *representation* of it.

Mary Morgan (2003) already pointed out this distinction between being a representation and being representative. The distinction is between two different ways in which a system may stand in for another: by contrast to being a representation, the system that is representative of another is only different from it in the way that a part is different from the whole. Morgan speaks of the representative as being 'of the same stuff' as what it is representative of. 'Being of the same stuff' should be taken literally: 'being a part of'. Wendy Parker (2009) argues that material similarities between the system manipulated and the target system are not necessarily more informative than the formal similarities relied upon in simulation, between the target system and a model of it. This is true. But what is important in the idea of 'being of same stuff' is not the idea of a different type of similarities; it is the idea that what is manipulated is not a different sort of system. It is rather a sub-system.

Winsberg (2009)'s analysis of the epistemological consequences of the difference between manipulating a model and manipulating an experimental system still applies: the justifications for the results will be of different sorts. But more importantly for our concern, there will be consequences of the difference between manipulating a *representation* and manipulating a *representative* of the target system. And these consequences are epistemic. They have to do with what we can learn about the target system.

There is finally another case of experimentation where the target system and object manipulated may be said to be distinct: the case where the manipulation of a system X is designed to obtain information about a different system Y, via the information obtained about X. Experimental results about Y are then not the results of a manipulation of Y, no more than the results from a simulation using a model of Y. True. But when, for instance with the investigation of an astrophysical or climate system, simulation is said to be epistemic substitute for or epistemically on a par with experimentation, it is not substitute for experimentation on a physical system different from the astrophysical or climate system. It is substitute for experimentation on the system that is simulated and that cannot be manipulated.

So when the simulation of the system X is claimed to be 'epistemically on a par', it is with experimentation on and manipulation of X or a representative part of X. But in that case, experimentation has, contrary to simulation, a direct access to the target system, in the sense that the target system is the system that is manipulated in experimentation.

3.2 Both simulation and experimentation rely on models?

Experimentation, we conclude from the previous section, manipulates the target system of the investigation and in this sense has, by contrast with simulation, direct access to it. But that does not mean a direct access to the information about the system. Getting experimental information about the system that is manipulated is a complicated, delicate matter that involves preparation, control, calibration, interpretation of the results of the manipulation, not to speak of the selection and arrangement of these results. All these different aspects of experimentation involve theoretical background and assumptions about different elements of the experimental process, and thus rely on models. So just like simulation, experimentation involves reliance on models, that much is clear.

Focusing on one particular use of models in experimentation, their use in measurement, Margaret Morrison argues that the similarity of the role of models leaves "little basis on which to epistemically differentiate the activities involved in some types of simulation and experimentation" (2009: 40). Taking epistemic differences to be differences in epistemic results, the claim looks like an objection against the idea of a difference in the epistemic functions of simulation and experimentation.

Morrison's strategy is not to deny that, in simulation, all that is manipulated is a putative model of the target system whereas performing measurements on the system involves an interaction with the system: "a measurement… bears some type of *causal connection to the property being measured*"(ibid p.52, italics added). Rather, the role models play in measurements seems to make this interactive dimension of experimentation epistemologically irrelevant:

> Experimental measurement is a highly complex affair where appeals to materiality as a method of validation are outstripped by an intricate network of models and inferences. (ibid p. 53)

That models play a crucial role in measurement should be uncontroversial. That, however, leaves open the possibility that this role be epistemically different from the role it plays in simulation. But Morrison goes further in the function she assigns to models, speaking literally of models as measuring instruments.[6] Her claim seems to be that the physical interaction between the target system and the instrument that takes place in experimentation makes no epistemological difference because experimentation, no less than simulation, consists in manipulating models. In what sense does experimentation consist in manipulating models?

As a simple illustration let's consider the wake again: suppose we are interested in how the shedding frequency, F, of the vortices behind the cylinder, depends on the velocity, V, of the flow before it reaches the cylinder, the upstream velocity. We are thus interested in F = f(V). Suppose the instrument used to measure the shedding frequency is a hot-wire anemometer. To put it simply, the sensible part of this instrument is a thin wire that undergoes a brief change in temperature as the local velocity of the flow is increasing. The state of the instrument is characterized by the temperature of the wire, T. We place the anemometer behind the cylinder, on the way of the vortices traveling downstream with the flow. The periodic change in local velocity, v, due to the emission of vortices will then cause an alternating variation in the temperature of the wire. The use of the instrument relies on a relation between v and T: $v = g(T)$ so as to produce, as an outcome of the measurement, the local velocity v. As the value of the upstream velocity V is changed, the values of T and v change, and the relation $v = g(T)$ goes from one realization to another, in the same way as happens when a model is run in a simulation. And, for the simulation, we did talk of the manipulation of the model described by the equation implemented. So experimentation could be seen as well as the manipulation of a model, the model of the instrument(s). The question now is whether this manipulation of the model is really all what matters to the epistemic function of experimentation.

---

6 For a view of models as measuring instruments see also Boumans 2005.

With a simulation, the simulationist intervenes on the model directly by choosing the input, which are the initial values for the constants and parameters. In experimentation the input of the model, e.g. the temperature of the wire in the anemometer, is the state of the target system. This state itself results from the initial conditions that were fixed by the experimenter when she chose a value for V. The experimenter does not intervene directly on the model; she intervenes on the system. And via the effect the system has on the instrument, and the model representing the relation between this effect and its cause, she finds out about this cause, which is the state of the system resulting from the intervention on the system. In experimentation, the model is then intervened on by the system and the output of the model, the outcome of measurement, tells us how it was intervened on, e.g. what was the value of *v,* the local velocity, resulting from the experimenter's intervention on V, the upstream velocity. So it is certainly true that the experimenter makes crucial use of the effect of an intervention on/manipulation of the model. But how the model is intervened on/manipulated seems to make all the difference with regards to what we can learn from the use of the model. The experimenter using a model of an instrument is not interested in learning about the behavior of the model given a certain input. It is assumed that she knows how the instrument works, how it behaves under given conditions. She wants to learn about the input to which the instrument, or its model, 'reacts'. Because in experimentation, this input, the source of the intervention on the model, is the state of the target system, the manipulation of the model is a means to learn about the state of the target system. And if this state of the target system is solely the effect of the intervention on the target system, then the manipulation of the model is also a means to learn about the relation between this intervention on and the resulting state of the target system.

Morrison's examples of experimentation are different in that the model used to make the measurement is the model of the whole target system, e.g. a model of the pendulum, rather than a model of a probing instrument. But the same reasoning applies. The case considered is the experimental measurement of the value of some constant that figures in the model of the physical target system, like the value of the gravitational constant. As Ron Giere (2009) notes, however, the measurement procedure supposes that the pendulum interacts *causally* with the Earth's gravitational field. We learn something about the environment by having the model intervened on by the environment and by having this intervention satisfy specific conditions. So in this case or in the previous one, in spite of the use of a model, the material interaction in experimentation does seem epistemically relevant.

Admittedly, we 'know' of the features of the system that affect the instrument only in so far as we 'know' of the relation between these features and the state of the instrument, that is, only in so far as we have and are justified in using a given model of the instrument. But to say that this mediating role of model makes causal interaction in experimentation epistemically irrelevant looks like saying that the role of language in expressing our sensory experience makes the sensory character of this experience epistemically irrelevant.

It is also true that materiality is not unique to experimentation. The reliability of experimentation results depends on having a proper functioning of the instrument and proper conditions of its use; the reliability of the simulation results depends on the proper functioning of the computer in which the simulation is implemented. A discussion of the role of materiality in simulation is beyond the scope of this paper but it is worth noting some differences between the contribution of materiality to the epistemic functions of simulation and experimentation. In experimentation, the material conditions are essential to the interpretation of what we learn. Depending on where we position the anemometer the outcome produced might have to be interpreted differently, be the measurement of V or the measurement of v. How we measure is essential to what we measure. In a simulation, when everything goes well it seems that materiality becomes, in principle, epistemically transparent[7]. This is why we can talk of manipulating the model, rather than the computer. This is also why, as we will see in the next section, in experimentation, material 'errors' may be instructive *about the system under study*: from inconsistent results, we may learn, for instance, about the existence of a new factor, a new feature of the system. From a material flaw in simulation, we do not learn *about the target system*.

The two arguments that were just examined called into question the idea of experimentation as direct access to the system under study. They fall short however of

---

[7] This is, of course, not to say that what computer is used for the simulation will not make a difference to what results can be produced.

showing that there is no basis for some general epistemic difference between simulation and experimentation. To get a better understanding of what these epistemic differences are we need first a better understanding of what the respective epistemic targets are.

4. Real (epistemic) difference

This section will focus on the epistemic targets of simulation and experimentation, that is, the epistemic results that experimentation and simulation are respectively designed to produce about a given target system, and the difference between these epistemic targets.

4.1 Epistemic target in experimentation

Part of the epistemic result of experimentation is the measurement of the values of quantities characterizing the target system. But the characterization of the epistemic target of *scientific* experimentation must also include the objectives of 1) reliability and 2) formal generalization.

One condition for the reliability of measurements is that experimental outcomes be reproducible. But not any reproducibility will do. Results on gravitational waves were claimed to be reproducible, but the conditions in which they were reproducible were not the 'right' conditions. What the right conditions are is neither always clear from the outset nor written in stone. It is rather typical of 'exploratory experimentation'[8] that it aims at specifying what they are.

Both the experimental investigations introduced earlier, on the wake formed behind a cylinder and the Wason task, illustrate this aim. In the former, experimentation aimed to determine how the shedding frequency of the vortices forming the wake evolves when the upstream velocity of the flow increases. Two different forms of evolution were found, both being reproducible. It took a 30-year long controversy to identify which, if either, was the right one, the reason being that it was not clear what factors were responsible for the difference between these two forms of evolution. Depending on which factor is

[8] On the distinction between exploratory experiment and theory-testing experiment, see Steinle 2002.

responsible, the effect will count as genuine, intrinsic, or as interference. The controversy ended with the identification of the factor responsible for the effect (Williamson 1989).[9]

In the Wason task experiments, the object of investigation is the ability to use conditional reasoning. Given the rule: "if there is a vowel on one side then there is a even number on the other side", and the task of testing it, it is expected that a proper understanding of the conditional will lead to the choice of turning over the card with a vowel and the card with an odd number. These are the two choices that could falsify the rule: if there is a vowel on one side, there must be an even number on the other face. If there is an odd number on one face, there should not be a vowel on the other face.

Experimental results always show this response as only one of those displayed by the subjects and not the most frequent one. A common explanation of the results is in terms of correct vs. incorrect understanding of the material conditional. A recent experimental study argues, however, that some differences in the exact terms of the formulation (e.g. 'testing' vs. 'falsifying') can have an influence on the reaction to the task that has been overlooked by previous studies (Stenning and van Lambalgen 2001). If that is correct, the conditions of reproducibility of the results are not those in which most studies were conducted since in these conditions, this semantic factor was ignored and had a non-measured effect on the result. The conditions of reproducibility must have all the factors having an influence on the effect of interest be under control: either fixed, if they are part of the background conditions, or systematically varied and measured if their influence is relevant to the investigation of the effect. Reproducibility must be in the right conditions: conditions that properly distinguish the background from the factors that are causally relevant.

When a complex phenomenon is investigated, the aim of the experimentation is then, first, given a certain variable of interest, to identify the factors causally relevant to its evolution, the relevant parameters, and measure their effect. This phase aims at a measurement of the behavior of the system of interest. But the aim of experimentation is also to go from the data that are collected to *relations* between variables of interest and

[9] For a detailed discussion of this case study see Peschard 2011.

relevant factors that these data instantiate. It is to go from the data to the phenomenon.[10] The information that experimentation aims to produce about this system, the target system, is the structural pattern, the phenomenon, that its behavior instantiates. If the aim only were the behavior of the system, reproducibility of the measurement would not be such an issue. More precisely, the distinction between the background and the relevant factors would not be an issue. Reproducibility *in the right conditions* is required because what is at stake in experimentation is neither the behavior of this or that target system, nor the pattern that only this or that target system instantiates. It is the structural pattern that any system that is relevantly similar to the target system instantiates.

The behavior of the system manipulated, as particular instantiation of a structural pattern is, in general, only an intermediary or preliminary epistemic target for experimentation. I will hereafter use the term *epistemic target* to refer to the structural pattern.

When experimentation is used merely as a testing procedure, it may seem that the epistemic target is actually the behavior itself of the system manipulated, the instantiation of a pattern rather than the pattern instantiated. But one should remember that what is tested is a (theoretical) structural pattern, which would be instantiated not only by the system that is the object of measurement but by any system relevantly similar to it. And the measurement outcomes are epistemically significant in that they can be seen as the instantiation of a pattern that does or does not match the one that is under test.

What about the epistemic target of the simulation?

4.2 Epistemic target in simulation

A simulation is the manipulation of a putative model of the target system S, that is, of a set of relations between some variable(s) of interest and parameters. Such a set of relations is what we just called a structural pattern. The results of the simulation are sets of values that satisfy these relations, that is, an instantiation of this pattern, corresponding to some initial conditions. Note the contrast: experimentation ultimately aims to elucidate the structural pattern that is instantiated by the evolution of the system manipulated;

[10] About the distinction between data and phenomena see e. g. Bogen and Woodward 1988, McAllister 1997, Feest 2009.

simulation aims to produce an instantiation (numerical instantiation) of the pattern that is manipulated.

For experimentation, we made a distinction between the aim of measuring, the acquisition of data, and the ultimate aim of elucidating phenomena, structural patterns. We can make a similar distinction between two levels of aims for simulation. To produce an instantiation of the pattern that is manipulated is only a preliminary result. The *ultimate* epistemic target of the simulation is the evolution of some physical system S, the target system, in the physical conditions represented in the simulation. By learning how the structural pattern that is simulated is numerically instantiated, the simulationist aims to learn how it *would be* physically instantiated by a system correctly represented by the model.

Granted, there are exploratory simulations, where no physical system, actual or even possible, is represented by the model run on the computer. In fact, in this case, just as in experimentation, there is no distinction between system manipulated and target system. There is no such distinction because the model plays both roles. Obviously, if no physical system is simulated, if there is no physical system in the role of target system, then the ultimate epistemic target of simulation cannot be the physical instantiation of the model manipulated. But it also means that we are outside the domain where the comparison between the epistemic functions of experimentation and simulation makes sense.

This objection, however, suggests another one, more interesting. If we put aside for a moment the view of simulation as imitation, we can see that, like the experimenter, the simulationist is interested not just in the instantiation of a pattern but also in the pattern instantiated. We need to remember the two senses of simulation that we have distinguished earlier: in one sense, it is an imitation, but, in the other sense, it is the calculation of the solution of the algorithm implemented on the computer. The form of this solution is generally opaque to the simulationist (Lenhard 2006, Humphreys 2009). But this is what is instantiated by the data produced by the simulation. And beyond the data, the simulationist is interested in identifying the form of this solution, the pattern that the data instantiate. In addition, there is a compelling argument to be made that this pattern is not already somehow contained in the model used for the simulation. And that enables one to say, that just like for experimentation, that simulation produces ‘new knowledge’ (Winsberg 2010).

That both simulation and experimentation not only produce the instantiation of a pattern, the data, but aim at identifying the pattern instantiated in the data, certainly makes their epistemic functions and targets look very similar to one another. But if simulation is, as was mentioned early on, a form of experimentation, that should not be surprising. On the other hand, it underscores an essential difference between these results: the pattern that experimentation aims to uncover is the one that the evolution of the target system, and any system relevantly similar to it, in similar conditions, instantiates. The pattern that the simulationist aims to uncover is the one instantiated by the data produced during the simulation. The epistemic force of the distinction relies heavily on the argument made previously that in experimentation there is a direct access to the target system.

“Models” as Robert Sugden (2002) writes, “are suggestions about how to set about explaining some phenomena in the real world. […] [T]hey are sketches of processes which, according to their creators, *might* explain phenomena we can observe in the real world”. Of course, the simulationist hopes, as Winsberg nicely puts it, that the simulation does imitate the evolution of the target system and that the pattern instantiated by the simulation data is the one instantiated by the evolution of the target system. And she will offer some reasons to make his cases. But all the simulation can show is how the system would evolve if it were the case that the only factors that make a relevant difference to this behavior are those represented in the model, assuming they interact in the way represented in the model.

The results of the simulation are, putatively, actual claims about the model and counterfactual claims about the target system. By contrast, the results of the experimentation are, putatively, actual claims about the target system, which can be interpreted, as we will see, as counterfactual claims about the simulation. It is this difference in the epistemic targets of simulation and experimentation that makes the use of simulation so productive even in a tandem configuration, where experimentation is possible. This difference makes possible a real epistemic co-operation between simulation and experimentation.

4.3 Articulation of simulation and experimentation

Simulation and experimentation have different epistemic targets, i.e., they are designed to produce different kinds of epistemic results about a given target system. This difference should incite us to investigate more closely the way in which these activities co-operate. On the other hand, to look closer at this co-operation helps to make more precise the difference in their respective epistemic targets.

If simulation and experimentation have different epistemic targets, simulation cannot serve, in the strong sense of the term, as an epistemic substitute for experimentation. The activities are not epistemically on a par. But the co-operation shows another way in which they are on a par: they are both for each other a source of information and of constraint. Again, an episode in the investigation of the wake will serve as brief illustration.

The object of the investigation was the relation between the frequency with which the vortices forming the wake are emitted behind the cylinder and the control parameter (velocity of the flow before it reaches the cylinder). Experimental measurements produced conflicting results. Some showed a continuously linear evolution whereas others showed a discontinuously linear evolution. A simulation was finally conducted that showed a continuously linear evolution.[11]

Granted some confidence in the basic structure of the model, the experimenter is *informed* about what the behavior of the target system would be, if the only factors that contribute to it were those represented in the model. Where the results from experimentation and simulation disagree, the experimenter is *constrained* to account for the difference in terms of factors not represented in the model used for the simulation. But the simulation does not produce information or constraints about what to do with these factors once they are identified: they might come to be regarded as part of the background, in which case they will have to be experimentally neutralized. But they might instead be deemed relevant to the understanding of the evolution of the target system, in which case, their effect needs to be systematically recorded. New experiments

[11] See Williamson 1989 for a historical overview of the investigation.

on the wake identified a new factor that was not represented in the model, and it was successfully argued by the experimenter that it was a relevant factor. In the case of the Wason experiment, experimenters also argue that they have identified a factor that has an effect on the system and was not represented in the model used for the simulation (Stenning and van Lambalgen 2001).

On the other hand, when a new relevant parameter is identified, the *constraint* for the simulationist is to integrate this new information about the target system into the model used for the simulation. The simulationist is *informed* about what the results of the simulation would be, were all the relevant factors taken into account by the model. In the process, the simulationist may discover that some assumptions made in the representation of the system were mistaken. In the case of the wake, the geometry built into the simulation was that of a flow around a disc, justified by the assumption that the physical system was ideally equivalent to a flow around an infinite cylinder. The new relevant factor identified in experimentation contradicted this assumption. Similarly in the Wason case, the simulation[12] assumes a uniform interpretation of the terms in which the task presented to the subject is formulated. Experimenters claim that there are strong differences in the interpretation of these terms which influence the response to the task.

Conclusion

If simulation is to qualify as substitute for experimentation, in the strong sense that was examined, the two should be able to produce the same epistemic results: they should have the same epistemic targets. If simulation can serve as a substitute when experimentation is not possible, it should be able to serve as a substitute as well in a tandem configuration, where experimentation and simulation have the same target system, S, and experimentation on this system is possible. But in this situation, it was argued, experimentation, by contrast with simulation, provides a direct access to S via causal interaction and, thereby, provide information about the state of S. And it was anticipated that this difference would result in a difference in the epistemic results these activities are able to produce.

[12] The simulation uses a Bayesian model. See Oaksford & Chater 2007.

Two objections to the idea that experimentation provides a direct access to S were examined. One is that in experimentation just as in simulation, what is manipulated is only a system that stands in for the target system. This objection was answered by stressing two distinctions: one between the target system of the experimentation (system manipulated) and the system that might be its 'epistemic motivation' (e.g. a system in the field); the other one is between being a representation (as a model aims to be) and being representative (as a sample from a population aims to be). The other objection is that experimentation no less than simulation relies on models (models of instruments). The objection was answered by stressing the difference in the way in which the models are intervened on in each case. In experimentation, but not in simulation, the intervention comes from the target system itself via the physical relation between the state of the system and the state of the instrument.

A closer inspection of the epistemic targets of simulation and experimentation brought forth a similarity between these two targets: in both cases, the target is not only the instantiation of a pattern (data) but the pattern instantiated by the data. That might be one additional reason why the two activities are sometimes said to be epistemically on a par. But in the light of the distinctions just pointed out this similarity is only a similarity. In experimentation, the pattern in question is the one that the evolution of the target system instantiates; in the case of simulation, it is the one instantiated by the data produced by the implementation of the model. So these activities are not epistemically on a par. But there is also an important way in which they are actually on a par. The difference in epistemic targets makes possible an epistemic co-operation and in this co-operation, both simulation and experimentation are, through their respective results, informing and constraining each other.

REFERENCES

Bogen, J. and Woodward, J. (1988) Saving the Phenomena. *The Philosophical Review* 97 (3): 303-352.

Boumans, M. (2005). *How Economists Model the World into Numbers*. London and New York: Routledge.

Feest, U. (2009) What exactly is stabilized when phenomena are stabilized? *Synthese*, Online First, DOI 10.1007/s11229-009-9616-7.

Fox Keller, E. (2003) Models, Simulation and ‘Computer Experiments. In Radder 2003: 198-215.

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

Guala, F. (2003) Experimental localism and external validity. *Philosophy of Science* 70: 1195-1205.

Guala, F. (2008) Experimentation in economics. In U. Mäki (ed.), *Elsevier Handbook of the Philosophy of Science, Volume 13: Philosophy of Economics*. Amsterdam: Elsevier.

Hartman, S. (1996) The World as a Process: Simulations in the Natural and Social Sciences. In R. Hegselmann et al. (eds.), *Simulation and Modeling in the Social Sciences from the Philosophy of Science Point of View*. Theory and Decision Library. Dordrecht: Kluwer Academic Publishers, pp. 77-100.

Humphreys, P. (1994) Numerical Experimentation. In P. Humphreys (ed.) *Patrick Suppes: Scientific Philosopher*, *Volume 2.Philosophy of Physics, Theory Structure and Measurement Theory*. Dordrecht: Kluwer Academic Publishers, pp. 103-118.

Humphreys, P. (2009) The philosophical novelty of computer simulation methods. *Synthese* 169 (3): 615-626.

Krohs U. (2008) How digital computer simulations explain real-world processes. *International Studies in the Philosophy of Science* 22 (3): 277-292.

Lenhard, J. (2006) Surprised by a nanowire: Simulation, Control and Understanding. *Philosophy of Science*, 73(5): 605-616.

McAllister, J. (1997) Phenomena and Patterns in Data Sets. *Erkenntnis* 47: 217–228.

Morgan, M. (2003) Experiments without material intervention: Model experiments, virtual experiments and virtually experiments. In Radder 2003: 216-235.

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

Norton, S., and F. Suppe, (2001) Why atmospheric modeling is good science. In C. Miller & P. N. Edwards (eds.), *Changing the atmosphere: Expert Knowledge and Environmental Governance*. Cambridge, MA: MIT Press, pp. 67–105.

Oaksford, M. & Chater, N. (2007) *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford: Oxford University Press

T. Ören, T. (2011) The many facets of simulation through a collection of about 100 definitions. *SCS Modeling and Simulation Magazine* 2: 82-92.

Parker, W. (2009*)*. Does Matter Really Matter? Computer Simulations, Experiments, and Materiality. *Synthese* 169 (3): 483 - 496

Peschard, I. (2011) Modeling and Experimenting. In *Models, Simulations, and Representations*, New York: Routledge: 42-61

Radder, H. (ed.) (2003) *The Philosophy of Scientific Experimentation*. Pittsburgh: Pittsburgh University Press.

Steinle, F. (2002). "Experiments in History and Philosophy of Science." *Perspectives on Science* 10 (4): 408–432.

Stenning, K. and M. van Lambalgen, (2001) Semantics as a foundation for psychology: a case study of Wason's selection task. *Logic, Language and Information*, 10(3): 273-317.

R. Sugden, (2000) Credible worlds: the status of theoretical models in economics. *Journal of Economic Methodology* 7: 1-31.

Varenne, F. (2010) Les simulations computationnelles dans les sciences sociales. *Nouvelles perspectives en sciences sociales : revue internationale de systémique complexe et d'études relationnelles* 5(2): 17-49.

Wason, P. C. (1968) Reasoning about a rule. *Quarterly Journal of Experimental Psychology* 20: 273-281.

Williamson, C.H.K. (1989) Oblique and parallel modes of vortex shedding in the wake of a circular cylinder at low Reynolds number. *Journal of Fluid Mechanics* 206: 579-627.

Winsberg, E. (1999) Sanctioning models: The epistemology of simulation. *Science in Context* 12 (20: 275-292.

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

Winsberg, E. (2009) A tale of two methods. *Synthese* 169 (3): 575-592.

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