Chapter 7
Computer Simulations: A New Mode of Scientific Inquiry?

Stéphanie Ruphy

Abstract Computer simulations are everywhere in science today, thanks to ever increasing computer power. By discussing similarities and differences with experimentation and theorizing, the two traditional pillars of scientific activities, this paper will investigate what exactly is specific and new about them. From an ontological point of view, where do simulations lie on this traditional theory-experiment map? Do simulations also produce measurements? How are the results of a simulation deem reliable? In light of these epistemological discussions, the paper will offer a requalification of the type of knowledge produced by simulation enterprises, emphasizing its modal character: simulations do produce useful knowledge about our world to the extent that they tell us what could be or could have been the case, if not knowledge about what is or was actually the case. The paper will also investigate to what extent technological progress in computer power, by promoting the building of increasingly detailed simulations of real-world phenomena, shapes the very aims of science.

1 Introduction

In 2013, two projects were selected by the European Commission as "Flagships" projects, receiving each a huge amount of funds (about one billion euros over 10 years). It is telling that one of these two top-priority projects, the Human Brain Project, aims at digitally *simulating* the behaviour of the brain. Computer simulations have not only become ubiquitous in the sciences, both natural and social, they are also more and more becoming ends in themselves, putting theorizing and experimenting, the two traditional pillars of scientific activities, into the background. This major addition to the range of scientific activities is in a straightforward sense directly linked to technological advances: the various epistemic roles fulfilled by computer simulations are inseparable from the technology used to perform it, to wit, the digital computer. Asking to what extent technology (in that case

S. Ruphy (✉)
Laboratoire PPL, Université Grenoble Alpes, F-38000 Grenoble, France
e-mail: stephanie.ruphy@upmf-grenoble.fr

ever increasing computing power) shapes science requires assessing the novelty and the epistemological specificities of this kind of scientific activities. Sure enough, computer simulations are everywhere today in science – there is hardly a phenomenon that has not been simulated mathematically, from the formation of the moon to the emergence of monogamy in the course of evolution of primates, from the folding of proteins to economic growth and the disintegration of the Higgs boson, but what exactly is specific and new about them?

There are two main levels of assertions about their novelty in the current philosophical landscape. A first kind of assertions concerns the extent to which computer simulations constitute a genuine addition to the toolbox of science. On a second level, the discussion is about the consequences of this addition for philosophy of science, the question being whether or not computer simulations call for a new epistemology that would be distinct from traditional considerations centered on theory, models and experiments.

Given the topic of this volume and the direct link between technological progress made in computational power and simulating capacities, I will be mainly interested in this paper in the first kind of assertions, the ones that state the significant novelty of computer simulations as a scientific practice.[1] Here's a sample of those claims, coming both from philosophers of science and scientists. For the philosopher Ronald Giere for instance, the novelty is quite radical: "[. . .] computer simulation is a qualitatively new phenomenon in the practice of science. It is the major methodological advance in at least a generation. I would go so far as saying it is changing and will continue to change the practice not just of experimentation but of science as a whole" (2009, 59). Paul Humphreys, also a philosopher of science, goes even one step further by talking about revolution: "[. . .] computer modelling and simulation [. . .] have introduced a distinctively new, even *revolutionary*, set of methods in science" (2004, 57. My italics). On the scientific side, the tone is no less dramatic as for instance in a report a few years ago to the US National Academy of Sciences: "[But] it is only over the last several years that scientific computation has reached the point where it is on a par with laboratory experiments and mathematical theory as a tool for research in science and engineering. The computer literally is providing a *new window through which we can observe the natural world in exquisite detail*." (J. Langer, as cited in Schweber and Wächter 2000, 586. My italics).

In their efforts to further qualify the novelty of computer simulations and the associated transformative change of scientific activities, philosophers of science have engaged into descriptive enterprises focusing on particular instances of simulation. Given the widespread taste of professional philosophers for accumulation of definitions and distinctions, as well as for fine-grained typologies, efforts have also been made to offer scientifically informed definitions of simulations (distinguishing them in particular from models), as well as typologies ordering the variety of

---

[1]See Humphreys (2009) and Reiss and Frigg (2009) for discussions of the second kind of assertions.

scientific enterprises coming under the banner of computer simulation, by typically classifying them according to the type of algorithm they employ ("Discretization" mathematical techniques, "Monte Carlo" methods, "Cellular automata" approaches, etc.). However interesting and useful these philosophical studies are, I won't talk much here of the various definitions and distinctions they propose, being more concerned by the challenging epistemological and ontological issues *common* to many kinds of simulations. And a widely-discussed first set of issues refers to the relationship between computer simulations and experimenting and theorizing. From an ontological point of view, where do simulations lie on this traditional theory-experiment map? Simulations are often described as "virtual" or "numerical" experiments. But what are the significant similarities or differences between computer simulations and experiments? Do simulations also produce measurements? Do they play similar epistemological roles *vis-à-vis* theory? Another set of challenging issues concerns the sanctioning of a computer simulation. How do computer simulations get their epistemic credentials, given that they do not simply inherit the epistemic credentials of their underlying theories (Winsberg 2013)? Is empirical adequacy a sure guide to the representational adequacy of a simulation, that is, to its capacity to deliver reliable knowledge on the components and processes at work in the real-world phenomenon whose behaviour it purports to mimic? As we shall see, this kind of issues are especially acute for what I will call *composite* computer simulations, developed to integrate as much detail of a given phenomenon as computing power allows. In light of these epistemological discussions, I will offer a requalification of the type of knowledge produced by simulation enterprises, emphasizing its *modal* character. And I will conclude with tentative remarks on the way ever increasing computing power, by promoting the building of fully detailed simulations of real-world phenomena, may progressively transform the very *aims* of science.

2 Hybrid Practice

A good starting point to discuss the similarities and differences between simulations on the one hand, and experiments and theories on the other, might be to ask scientists how they would describe their activities when they build and use simulations. Fortunately, some science studies scholars have done just that and I will draw here on Dowling's (1999) account based on 35 interviews with researchers in various disciplines ranging from physics and chemistry to meteorology, physiology and artificial life. One of the most interesting, if not totally surprising lessons of Dowling's inquiry is that for its practitioners the status of this activity is often hybrid, combining aspects partaking of theoretical research and of experimental research. Simulations are commonly used to explore the behaviour of a set of equations, constituting a mathematical model of a given phenomenon. In that case,

scientists often express the feeling that they are performing an experiment, by pointing out that many stages of their digital study are similar to traditional stages of an experimental work. They first set some initial conditions for the system, vary the values of parameters, and then observe how the system evolves. To that extent, as in a physical experiment, the scientist interacts with a system (the mathematical model), which sometimes may also behave in surprising ways. In other words, in both cases, scientists engage with a system whose behaviour they cannot totally anticipate, and that is precisely the point: to learn more about it by tinkering and interacting with it. In the case of a simulation, the unpredictability of the system is no mystery: it usually comes from the nature of the calculations involved (often dealing with non linear equations that cannot be solved analytically). Scientists sometimes talk about "the remoteness of the computer processes": they cannot fully be grasped by the researcher who "black-boxes" them while performing the simulation run (Dowling 1999, 266).

Mathematical manipulation of a theoretical model is not the only experimental dimension of a computer simulation. Producing data on aspects of a real-world system for which observations are very scarce, inexistent or costly to obtain is another widespread epistemic function of a computer simulation. To the extent that these simulated data are then often used to test various hypotheses, computer simulations share with experiments the role of providing evidence in support or against a piece of theoretical knowledge.

As for the similarities with theories, the point has been clearly, if somewhat simplistically, made by one of the physicists being interviewed: "Of course it's theory! It's not real!" (Dowling 1999, 265). In other words, when the issue of the relationship to reality is considered, that is, when simulations are taken as representations, the manipulation dimension of the simulation gives way to the conjectural nature it inherits from its theoretical building materials. So from the point of view of the practioners, computer simulations combine significant features of both theories and physical experiments, and that might explain why simulation practioners are sometimes less inclined than philosophers to describe computer simulation as a radically new way of finding out about the world. That might also explain why expressions such as "in silicon experiments", "numerical experiments", "virtual experiments" have become so common in the scientific discourse. But from an ontological point of view, to what extent exactly should these expressions be read literally?

Philosophers of science have further explored the similarities between computer simulations and physical experiments by asking three kinds of (related) questions. First, can one still talk of experiment in spite of the lack of physical interactions with a real-world system? Second, do simulations also work as measurement devices? Third, does the sanctioning of a computer simulation share features with the sanctioning of an experiment?

3 The 'Materiality' Debate

The first kind of questions is often referred to as the 'materiality' debate (see e.g. Parker (2009) for a critical overview of it).[2] This debate builds on the claim that the material causes at work in a numerical experiment are (obviously) of a different nature than the material causes at work in the real-world system (the target system) being investigated by the simulation, which is not the case with a physical experiment as explained by Guala:

> The difference lies in the kind of relationship existing between, on the one hand, an experiment and its target system, and on the other, a simulation and its target. In the former case, the correspondence holds at a "deep", "material" level, whereas in the latter, the similarity is admittedly only abstract and formal. [. . .] In a genuine experiment, the same material causes as those in the target system are at works; in a simulation, they are not, and the correspondence relation (of similarity or analogy) is purely formal in character (2005, 214–215).

This ontological difference emphasized by Guala has epistemic consequences. For Morgan (2005) for instance, an inference about a target system drawn from a simulation is less justified than an inference drawn from a physical experiment because in the former case, and not in the latter case, the two systems (the simulation/experimental system and the target system) are not made of the "same stuff". In other words, as Morgan puts it, "ontological equivalence provides epistemological power" (2005, 326). And computer simulations, if conceived as experiments, must be conceived as *non-material* experiments, on mathematical models rather than on real-world systems. This lack of materiality is precisely what Parker wants to challenge. Parker makes first a distinction between a computer simulation and a computer simulation study (2009, 488). A computer simulation is a "sequence of states undergone by a digital computer, with that sequence representing the sequence of states that some real or imagined system did, will or might undergo" (2009, 488). A computer simulation *study* is defined as "the broader activity that includes setting the state of the digital computer from which a simulation will evolve, triggering that evolution by starting the computer program that generates the simulation, and then collecting information regarding how various properties of the computing system [. . .] evolve in light of the earlier information" (2009, 488). Having defined an experiment as an "investigative activity involving intervention" (2009, 487), Parker then claims that computer simulation *studies* (and not computer *simulations)* do qualify as experiments: when performing a computer simulation study, the scientist does intervene on a *material* system, to wit, a programmed digital computer. So in this particular sense, concludes Parker, computer simulation studies are *material* experiments: 'materiality' is not an exclusive feature of traditional experiments that would distinguish them from computer studies.

For all that, acknowledging this kind of materiality for computer simulation studies does not directly bear on Morgan's epistemological contention. Recall that

---

[2]See also Barberousse et al. (2009) and Norton and Suppe (2001).

the epistemic advantage granted to traditional experiments follows from the fact that the system intervened on and the target system are made of "the same stuff" and not only from the fact that both systems are material systems. In the case of a computer simulation study, the system intervened on and the target system are both material but obviously they are not made of the same kind of "stuff". Parker does not deny this distinction but contends that its epistemological significance is overestimated. What is significant is not so much that the two systems (the experimental and the target systems) are made of "the same stuff", it is rather that there exist "relevant similarities" between the two. And being made of "the same stuff", in itself, does not always guaranty more relevant similarities between the experimental system and the target system. In the case of a traditional experiment, scientists must also justify making inferences from the experimental system to the target system.

4   Measurements

Another well-discussed kind of similarities between experiments and simulations concern the status of their outputs, and that leads us to our second issue – can the output of a computer simulation count as a *measurement*? Philosophers of science provide various and sometimes conflicting answers, depending on how they characterize measurements.

Morrison (2009) offers an interesting take on the issue by focussing on the role of models in a measurement process. Models do not only play a role when it comes to the interpretation of the outputs of an experiment; the measurement process itself involves a combination of various kinds of models (models of the measuring apparatus, correction models, models of data, etc.). This close connection between models and experiment is commonly acknowledged by philosophers of science. But Morrison (2009) goes one step further by adding that models themselves can function as "measuring instruments". To ground her claim, Morrison gives the example of the use of the physical pendulum to measure the local gravitational acceleration at the surface of the Earth. In that case (as in many other experiments), a precise measuring of the parameter under study (here the local gravitational acceleration) requires the application of many corrections (taking the air resistance into account for instance). So that many other, sometimes complex models are used, in addition to the simple model of a pendulum, to represent the measuring apparatus in an appropriate way. In other words, says Morrison, "the ability of the physical pendulum to function as a measuring instrument is completely dependent on the presence of these models." And she concludes: "That is the sense in which models themselves also play the role of measuring instruments" (2009, 35). To reinforce her point, Morrison gives another, more intricate example of measurement, where the role of models is even more central. In particle physics, when measuring a microscopic property such as the spin of an electron or the polarization of a photon, Morrison stresses that the microscopic object being measured is not directly observed. On the one hand, there is a model of the microscopic properties of the

target system (the electron or the photon). On the other hand, an extremely complex instrument is used together with a theoretical model describing its behaviour by a few degrees of freedom interacting with those of the target system. And this is the comparison between these models that constitutes the measurement (Morrison 2009, 43). So for Morrison, in this kind of experimental settings, models function as a "primary source of knowledge", hence their status of "measuring devices". This extension of the notion of measuring goes hand in hand with the downplaying of the epistemological significance of material interaction with some real-world system: "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" (Morrison 2009, 53).

Dropping the traditional emphasis on material interaction as characterizing experiment allows Morrison to contend that a computer simulation can also be considered as a measurement device. For once you have acknowledged the central role played by models in experimental measurement, striking similarities, claims Morrison, appear between the practice of computer simulation and experimental practice. In a computer simulation, you also start with a mathematical model of the real-world target system you want to investigate. Various mathematical operations of discretization and approximation of the differential equations involved in the mathematical model then give you a discrete simulation model that can be translated into a computer programme. Here too, as in a physical experiment, tests must be performed, in that case on the computer programme, to manage uncertainties and errors and make sure that the programme behaves correctly. In that respect, the programme functions like an apparatus in a traditional physical experiment (Morrison 2009, 53). And those various similarities, according to Morrison, put computer simulations epistemologically on a par with traditional experiment: their outputs can also count as measurements.

Not everybody agrees though. Giere (2009) for instance readily acknowledges the central role played by models in traditional experiment but rejects Morrison's extension of the notion of measuring on the ground that the various correcting models involved in the measurement process remain abstract objects that do not interact causally with the physical quantity under study (in Morrison's pendulum example, the Earth's gravitational field). And that suffices to disqualify them as measuring device. The disagreement thus seems to boil down to divergent views on the necessity of having some causal interaction with a physical quantity to qualify as a measurement. Consider a computer simulation of the solar system (another example discussed by Morrison and Giere). Do the outputs of this simulation (say, the positions of Saturn over the past 5,000 years) count as measurements? The laws of motion of the planets being very well established, the values provided by the simulation are no doubt more precise and accurate than the actual measurements performed by astronomers. Are they nevertheless only calculations and not measurements? It seems that legitimate answers and arguments can be given on both sides, depending on what you think is central to the notion of measurement. If you give priority to the epistemic function of a measurement, that is, providing reliable information of the values of some parameters of a physical system, then

Morrison's proposed extension seems appealing (provided that the reliability of the information can be correctly assessed in the case of a simulation – I will come back on this important issue later). But if you give priority to the ontological criteria stipulating that a measurement must involve some kind of causal interaction with a real-world physical quantity, then Morrison's proposition will seem too far-fetched. In any case, the very existence of this debate indicates a first way in which computer technology shapes scientific practice: the growing use of computer simulations directly bears on what it means to perform a measurement in science.

Another well-debated issue concerns the similarities – or lack thereof – between the way the output of an experiment is deemed reliable and the way the output of a computer simulation is. And that will lead us to the general issue of assessing the reliability of a simulation.

5   Internal Validity (Verification)

As briefly mentioned earlier, management of uncertainties and errors and calibration are essential components of a simulation enterprise. Simulationists must control for instance errors that might result from the various transformations the initial equations must go through to become computationally tractable (e.g. discretization), or errors resulting from the fact that a computer can store numbers only to a fixed number of digits, etc. And, as Winsberg (2003, 120) puts it: "developing an appreciation for what sorts of errors are likely to emerge under what circumstances is as much an important part of the craft of the simulationist as it is of the experimenter". Drawing on Alan Franklin' work (1986) on the epistemology of experiment, Winsberg adds that several of the techniques actually used by experimenters to manage errors and uncertainties apply directly, or have direct equivalents, in the process of sanctioning the outputs of a simulation. For instance, simulationists apply their numerical techniques on equations whose analytical solutions are known to check that they produce the expected results, just as experimenters use a new piece of experimental apparatus on well-known real-world systems to make sure that the apparatus behaves as expected. Also, simulationists may build different algorithms independently and check that they produce similar results when applied on the same mathematical model, just as experimenters use different instrumental techniques on a same target (say, optical microscopes and electronic microscopes) to establish the reliability of the techniques.

These various strategies aim at increasing our confidence in what is often called the *internal* validity or the *internal* reliability of a computer simulation. The point is to ensure that the solutions to the equations provided by the computer are close "enough" (given the limits put by computing power) to the solutions of the original equations. When refering to these checking procedures, scientists usually talk of *verification*. But verification is only (the first) half of the story when one wants to assess the reliability of a computer simulation. The other half, usually called *validation*, has to do with the relationship between the simulation and the real-world

target system whose behavior it purports to investigate. And assessing this *external* validity will depend on what kind of knowledge about the target system you expect from the simulation.

6 External Validity (Validation)

A distinction similar to the traditional distinction between an instrumentalist view of the aims of a scientific theory and a realist one can be made about simulations. By instrumental aims, I mean here the production of outputs relative to the past (retrodictions) or future (predictions) *observable* behaviour of a real-world system. Retrodictions are very common in "historical" natural sciences such as astrophysics, cosmology, geology, or climatology, where computer simulations are build to produce data about past states of the simulated system (the spatial distribution of galaxies one billion years after the Big Bang, the position of the continents two billion years ago, the variation of the average temperature at the surface of the Earth during the Pliocene period, etc.). A very familiar example of predictions made by computer simulation is of course weather forecast. Realist aims are – it is no surprise – epistemically more ambitious. The point is not only to get empirically adequate outputs; it is also to get them for the right reasons. In other words, the point is not only to save the phenomena (past or future) at hand, it is also to provide reliable knowledge on the underlying constituents and mechanisms at work in the system under study. And this realist *explanatory* purpose faces, as we shall see, specific challenges. These challenges are more or less dire depending on how the simulations relate to well-established theoretical knowledge, and especially their degree of 'compositionality', that is the degree to which they are built from various theories and bits of empirical knowledge.

6.1 Duhemian Problem

At one end of the compositionality spectrum, you find computer simulations built from one piece of well-established theoretical knowledge (for instance computer simulations of airflows around wings built from the Navier-Stoke equations). In most cases, the models that are directly "read-off" a theory need to be transformed to be computationally tractable. And, depending on the available computer resources in terms of speed and memory, that involves idealizations, simplifications and, often, the deliberate introduction of false assumptions. In the end, as Winsberg (2003, 108) puts it, "the model that is used to run the simulation is an offspring of the theory, but it is a mongrel offspring". Consequently, the computer simulation does not simply inherit the epistemic credentials of its underlying theory and establishing its reliability requires comparison with experimental results. The problem is that when the simulated data do not fit with the experimental data, it is not always clear what

part of the transformation process should be blamed. Are the numerical techniques the source of the problem or the various modelling assumptions made to get a computationally tractable model? As noticed by Frigg and Reiss (2009, 602–603), simulationists face here a variant of the classical Duhemian problem: something is wrong in their package of the model and the calculation techniques, but they might not know where to put the blame. This difficulty, specific to computational models as opposed to analytically solvable models, is often rephrased in terms of the inseparability of verification and validation: the sanctioning of a computer simulation involves both checking that the solutions obtained are close "enough" to the solutions of the original equations (verification) and that the computationally tractable model obtained after idealization and simplification remains an adequate (in the relevant, epistemic purpose-relative aspects) representation of the target system (validation), but these two operations cannot always, in practice, be separated.

6.2   The Perils of Accidental Empirical Adequacy[3]

At the other end of the compositionality spectrum lie highly *composite* computer simulations. By contrast with the kind of simulations just discussed, yielded by a single piece of theoretical knowledge, composite computer simulations are built by putting together various submodels of particular components and physical processes, often based on various theories and bits of empirical knowledge. Composite computer simulations are typically built to mimic the behavior of real-world "complex" phenomena such as the formation of galaxies, the propagation of forest fires or, of course, the evolution of the Earth climate. Typically, this kind of simulations combines instrumental and realist aims. Their purpose is minimally to mimic the observable behaviour of the system, but often, it is also to learn about the various underlying physical components and processes that give rise to this observable behaviour.

Composite computer simulations face specific difficulties when it comes to assess their reliability, in addition to the verification issues common to all kinds of computational models. The main problem, I will contend, is that the empirical adequacy of a composite simulation is a poor guide to its representational adequacy, that is, to the accuracy of its representations of the components and processes actually at work in the target system. Let me explain why by considering how they are elaborated throughout time. Building a simulation of a real-world system such as a galaxy or the Earth climate involves putting together submodels of particular components and physical processes that constitute the system. This is usually done progressively, starting from a minimal number of components and processes, and then adding features so that more and more aspects of the system are taken into account. When simulating our Galaxy for instance, astrophysicists

---

[3]This section (and the following) directly draws on Ruphy (2011).

started by putting together submodels of a stellar disc and of a stellar halo, then added submodels of a central bulge and of spiral arms in order to make the simulations more realistic. The problem is that the more the simulation is made realistic, the more it incorporates various submodels and the more it will run into a holist limitation of its testability. The reason is straightforward: a composite simulation may integrate several inaccurate submodels, whose combined effects lead to predictions conformed to the observations at hand. In other words, it is not unlikely that simulationists get the right outcomes (i.e. in agreement with the observations at hand), but not for the right reasons (i.e. not because the simulation incorporates accurate submodels of the actual components of the target system). And when simulationists cannot test the submodels independently against data (because to make contact with data, a submodel often needs to be interlocked with other submodels), there is unfortunately no way to find out if empirical adequacy is accidental. Therefore, given this pitfall of accidental empirical conformity, the empirical success of a composite computer simulation is a poor guide to the representational accuracy of the various submodels involved.

6.3 Plasticity and Path Dependency

Looking at simulation building processes reveals other, heretofore underappreciated features of composite computer simulations that also directly bear on the issue of their validation, to wit, what I have called their path-dependency and their plasticity. Let me (briefly) illustrate these notions with the example of a simulation of the evolution of our universe.[4] As is well known, cosmology starts by assuming that the large-scale evolution of space-time can be determined by applying Einstein's field equations of gravitation everywhere. And that plus the simplifying hypothesis of spatial homogeneity, gives the family of standard models of modern cosmology the "Friedmann-Lemaître" universes. In itself, a Friedmann-Lemaître model cannot account for the formation of the cosmic structures observed today, in particular the galaxies: The "cold dark matter" model is doing this job. To get off the ground, the cold dark matter model requires initial conditions of early density fluctuations. Those are provided by the inflation model. This first stratum of interlocked submodels allows the simulation to mimic the clustering evolution of dark matter. Other stratums of submodels, linking the dark matter distribution to the distribution of the visible matter must then be added to make contact with observations.

The question that interests us now is the following: at each step of the simulation-building process, are alternative submodels with similar empirical support and

---

explanatory power available? And the (short) answer is yes (see Ruphy (2011) for a more detailed answer based on what cosmologists themselves have to say). Moreover, at each step, the choice of one particular submodel among other possibilities constrains the next step. In our example, inflation, for instance, is appealing only once a Friedmann-Lemaître universe is adopted (which requires buying a philosophical principle, to wit, the Copernican principle). When starting, alternatively, from a spherically symmetric inhomogeneous model, inflation is not needed anymore to account for the anisotropies observed in the cosmic microwave background. So that the final composition of the simulation (in terms of submodels) turns out to depend on a series of choices made at various stages of the simulation building process.

A straightforward consequence of this path-dependency is the *contingency* of a composite simulation. Had the simulationists chosen different options at some stages of the simulation building process, they would have come up with a simulation made up of different submodels, that is, with a different picture of the components and mechanisms at work in the evolution of the target system. And the point is that those alternative pictures would be equally plausible in the sense that they would also be consistent both with the observations at hand and with our current theoretical knowledge. To deny this would clearly partake of an article of faith. Path-dependency puts therefore a serious limit to the possibility of *representational* validation, that is, to the possibility of establishing that the computer simulation integrates the right components and processes.

Plasticity is another (related) source of limitation. Plasticity refers to the possibility of adjusting the ingredients of a simulation so that it remains successful when new data come in. Note, though, that plasticity does not boil down to some ad hoc fine-tuning of the submodels involved in the simulation. Very often, the values of the free-parameters of the submodels are constrained independently by experiment and observation or by theoretical knowledge, so that the submodels and the simulation itself are progressively "rigidified". Nevertheless, some leeway always remains and it is precisely an essential part of the craft of the simulationist to choose which way to go to adjust the simulation when new data come in.[5] It is therefore not possible to give a general analysis of how these adjustments are achieved (they depend on the particular details specific to each simulation building process). Analysis of actual cases suggests, however, that the way a composite simulation is further developed in response to new data usually does not alter previously chosen key ingredients of the simulation.[6] Hence the *stability* of the simulation. In other words, there is some kind of inertial effect: one just keeps going along the same modelling path (i.e. with the same basic ingredients incorporated at early stages), rather than starting from scratch along a different modelling path. This inertial effect should come as

---

[5]See for instance Epstein and Forber (2013) for an interesting analysis of the perils of using macrodata to set parameters in a microfoundational simulation.

[6]This is the case for instance for the astrophysical and cosmological simulations discussed in Ruphy (2011) and for the Earth climate simulations analyzed in Lenhard and Winsberg (2010).

no surprise, given the pragmatic constraints on this kind of simulation enterprise. When a simulation is built over many years, incorporating knowledge from various domains of expertise, newcomers do not usually have the time nor the competences to fully investigate alternative modelling paths.

The overall lesson is thus the following: because of its plasticity and its path dependency, the stability and empirical success of a composite computer simulation when new data come in cannot be taken as a reliable sign it has achieved its realist goal of representational adequacy, i.e. that it has provided accurate knowledge on the underlying components and processes of the target system (as opposed to the more modest instrumental aim of empirical adequacy).

Let us take stoke here of the main conclusions of the previous epistemological discussions.

Computer simulations may fail or succeed in various ways, depending on their nature and on our epistemic expectations. We have seen that sanctioning a computer simulation involves minimally verification issues. Those issues might be deemed more of a mathematical nature than of an epistemological nature.[7] In any case, they are clearly directly dependent on the evolution of computing power and technology. Then come the validation issues, that is, sanctioning the relationship between the computer simulation and the real-world system whose behaviour it purports to mimic. A first level of validation is empirical: do the outputs of the simulation fit with the data at hand? In most cases, however, simulationists are not merely seeking empirical adequacy, they also aim at representational adequacy. The two are of course interdependent (at least if you are not a die-hard instrumentalist): empirical adequacy is taken as a reliable sign of representational adequacy, and representational adequacy justifies in its turn trusting the outputs of a simulation when the simulation is used to produce data on aspects of a real-world system for which observations or measurements are impossible (say, the radial variation of temperature at the centre of the Earth). When assessing empirical adequacy, simulationists may face a variant of the Duhemian problem: they might not be able to find out where to put the blame (on the calculation side or on the representational side) when there is a discrepancy between real data and simulated data. Sanctioning the representational adequacy of an empirically successful simulation may be even thornier, especially for composite computer simulations. For we have seen that, because of the path-dependency and the plasticity that characterize this kind of simulations, the more composite a simulation gets to be more realistic (i.e. to take into account more aspects and features of the system), the more you loose control of its representational validation. In other words, there seems to be a trade-off between the realistic ambition of a simulation and the reliability of the knowledge it actually delivers about the real components and processes at work in the target system.

For all that, taking the measure of these validation issues should not lead to a dismissal of the scientific enterprise consisting of developing purportedly realistic simulations of real-world complex phenomena. Rather, it invites to reconsider the

---

[7]See the exchange on this topic between Frigg and Reiss (2009) and Humphreys (2009).

epistemic goals actually achieved by these simulations. My main claim is that (empirically successful) composite computer simulations deliver *plausible* realistic stories or pictures of a given phenomenon, rather than reliable insights on what is *actually* the case.

7 Modal Knowledge

Scientists (at least epistemologically inclined ones) often warn (and rightly so) against, as the well-known cosmologist George Ellis puts it, "confusing computer simulations of reality with reality itself, when they can in fact represent only a highly simplified and stylized version of what *actually* is" (Ellis 2006, 35, My italics). My point is, to paraphrase Ellis, that computer simulations can in fact represent only a highly simplified and stylized version of what *possibly* is. That models and simulations tell white lies has been widely emphasized in the philosophical literature: phenomena must be simplified and idealized to be mathematically modelled, and for heuristic purpose, models can also knowingly depart from established knowledge. But the problem with composite computer simulations is that they may also tell non-deliberate lies that do not translate into empirical failure.

The confusion with reality feeds on the very realistic images and videos that are often produced from simulated data, thanks to very sophisticated visualization techniques. These images and videos "look" as they had been obtained from observational or experimental data. Striking examples are abundant in fields such as cosmology and astrophysics, where the outputs of the simulations are transformed into movies showing the evolution of the structures of the universe over billions of years or the collision of galaxies. In certain respects, the ontological status of this kind of computer simulations is akin to the status of richly realistic novels, which are described by Godfrey-Smith (2009, 107) as talks about "sets of fully-specific possibilities that are compatible with a given description".

The stories or pictures delivered by computer simulations are *plausible* in the sense that they are compatible both with the data at hand and with the current state of theoretical knowledge. And they are *realistic* in two senses: first because their ambition is to include as many features and aspects of the system as possible, second because of the transformation of their outputs into images that "look" like images built from observational or experimental data. I contend that computer simulations do produce useful knowledge about our world to the extent that they allow us to learn about what could be or could have been the case in *our* world, if not knowledge about what is or was actually the case in our world. Note that this *modal* nature of the knowledge produced by simulations raises resistance not only among philosophers committed to the idea that scientific knowledge is about *actual* courses of events or states of affairs, but also among scientists, as expressed for instance by the well-known evolutionary biologist John Maynard Smith "[. . .] I have a general feeling of

unease when contemplating complex systems dynamics. Its devotees are practicing fact-free science. A fact for them is, at best, the output of a computer simulation: it is rarely a fact about the world" (1995, 30).[8]

This increasing modal nature of the knowledge delivered by science via the use of computer simulations is not the only noticeable general transformation prompted by the development of computing power. Also on a quite general note, it is worth investigating to what extent ever increasing computing power, by stimulating the building of increasingly detailed simulations of real-world phenomena, shape the very *aims* of science.

8   Shaping the Aims of Science: Tentative Concluding Remarks

Explanation is often considered as a central epistemic aim: science is supposed to provide us with explanatory accounts of natural (and social) phenomena. But do the growing trend of building detailed simulations mean more and better explanations? There is no straightforward answer to that question, if only because philosophers disagree on what may count as a good scientific explanation and what it means for us to understand a phenomenon. Some indicative remarks may nevertheless be made here. Reporting a personal communication with a colleague, the geologist Chris Paola wrote recently in *Nature*: "... the danger in creating fully detailed models of complex systems is ending up with two things you don't understand – the system you started with, and your model of it" (2011, 38). This quip nicely sums up two kinds of loss that may come with the increasing "richness" of computer simulations.

A much-discussed factor contributing to the loss of understanding of a simulation is "epistemic opacity". Epistemic opacity refers to the idea that the computations involved in many simulations are so fast and so complex that no human or group of humans can grasp and follow them (Humphreys 2009, 619). Epistemic opacity also manifests itself at another level, at least in composite computer simulations. We have seen that these simulations are often built over several years, incorporating knowledge and contributions from different fields and different people. When using a simulation to produce new data or to test new hypotheses, the practitioner is unlikely to fully grasp not only the calculation processes but also the various submodels integrated in the simulation, which are then treated as black boxes.

As regards the loss of understanding of the target system, at least two reasons may be put forward to account for it. Recall first one of the conclusions of the previous epistemological discussion about validation: the more detailed (realistic) a simulation is, the more you loose control of its representational validation. So if the explanatory virtue of a simulation is taken as based on its ability to

---

[8]I borrow this quotation from Grim et al. (2013), which, in another framework, also discusses the modal character of the knowledge produced by simulation.

deliver reliable knowledge about the real components and mechanisms at work in the system, then indeed, fully detailed computer simulations do not score very high. But even though the reliability of the representation of the mechanisms at work provided by the simulation could be established, there would be another reason to favour very simplified simulations over simulations that include more detail. This is the belief that attention to what is truly essential should prevail on the integration of more details, when the simulation is built for explanatory purpose (rather than instrumental predictive purposes). Simulations of intricate processes of sedimentary geology, as analysed in Paola (2011), is a case in point. Paola (2011, 38) notes that "simplified representations of the complex small-scale mechanics of flow and/or sediment motion capture the self-organization processes that create apparently complex patterns." She explains that for many purposes, long-profile evolution can be represented by relatively simple diffusion models, and important aspects of large-scale downstream variability in depositional systems, including grain size and channel architecture, can be understood in terms of first order sediment mass balance. Beyond the technicalities, the general lesson is that "simplification is essential if the goal is insight. Models with fewer moving parts are easier to grasp, more clearly connect cause and effect, and are harder to fiddle to match observations" (Paola 2011, 38). Thus there seems to be a trade-off between explanatory purpose and integration of more details to make a simulation more realistic. If fewer and fewer scientists resist the temptation to build these ever more detailed simulations, feed on the technological evolution of computing power, explanation might become a less central goal of science. Predictive (or retrodictive) power may become more and more valued, since increasingly complex computer simulations will allow to make increasingly detailed and precise predictions, on ever finer scales, on more and more various aspects of a phenomenon.

Another general impact calling for philosophical attention is of a methodological nature: very powerful computer means make bottom-up approaches more and more feasible in the study of a phenomenon. In these approaches, the general idea is to simulate the behaviour of a system by simulating the behaviour of its parts. Examples of these microfoundational simulations can be found in many disciplines. In biology, simulations of the folding of proteins are built from simulations of amino-acid interactions; in ecology, simulations of the dynamics of eco-system are based on the simulations of preys-predators interactions, etc. This shaping of general scientific methodology by technology sparks sometimes vivid discussions within scientific communities: bottom-up approaches are charged with reductionist biases by proponents of more theoretical, top-down approaches. This is especially the case for instance in the field of brain studies. There has been a lot of hostility between bottom-up strategies starting from the simulation of detailed mechanisms at molecular level and studies of emergent cognitive capabilities typical of cognitive neurosciences. But discussions may end in the future in a more oecumenical spirit, given the increasing ambitious epistemic aim of brain simulations. Or at least it is what is suggested by our opening example, the European top-priority HBP project (Human Brain Project), whose aim of building multiscale simulations of neuromechanisms explicitly needs general theoretical principles to move between

different levels of description.[9] The HBP project is also representative of the increasingly *interactive* character of the relationship between epistemic aims and computer technology. Simulationists do not only "tune" their epistemic ambitions to the computing power available: technological evolutions are anticipated and integrated into the epistemic project itself. The HBP project for instance includes different stages of multiscale simulations, depending on computing power progress. Big simulation projects such as the HBP or cosmological simulations also generate their own technological specific needs, such as supercomputers that can support dynamic reconfigurations of memory and communications when changing scale of simulation, or new technological solutions to be able to perform computing, visualization and analysis simultaneously on a single machine (given the amount of data generated by the simulations, it will become too costly to move the generated data to other machines to perform visualization and analysis).[10] That epistemic progress is directly linked to technological progress is of course nothing new, but the ever increasing role of computer simulations in science makes the two consubstantial to an unprecedented degree.

References

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

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

Ellis, G. (2006). Issues in the philosophy of cosmology. http://arxiv.org/abs/astro-ph/0602280. (Reprinted in the *Handbook in Philosophy of Physics*, pp. 1183–1286, by J. Butterfield & J. Earman, Ed., 2007, Amsterdam: Elsevier)

Epstein, B., & Forber, P. (2013). The perils of tweaking: How to use macrodata to set parameters in complex simulation models. *Synthese, 190*, 203–218.

Franklin, A. (1986). *The neglect of experiment*. Cambridge: Cambridge University Press.

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

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

Godfrey-Smith, P. (2009). Models and fictions in science. *Philosophical Studies, 143*, 101–126.

Grim, P., Rosenberger, R., Rosenfeld, A., Anderson, B., & Eason, R. E. (2013). How simulations fail. *Synthese, 190*, 2367–2390.

Guala, F. (2005). *The methodology of experimental economics*. Cambridge: Cambridge University Press.

Humphreys, P. (2004). *Extending ourselves*. New York: Oxford University Press.

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

---

[9]As attested by the fact that the HBP project will dedicate some funds to the creation of a European Institute for Theoretical Neuroscience.

[10]I draw here on documents provided by the Human Brain Project at www.humanbrainproject.eu. Accessed 25 June 2013.

Lenhard, J., & Winsberg, E. (2010). Holism, entrenchment, and the future of climate model pluralism. *Studies in History and Philosophy of Modern Physics, 41*, 253–262.

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

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

Norton, S., & Suppe, F. (2001). Why atmospheric modeling is good science. In C. Miller & P. N. Edward (Eds.), *Changing the atmosphere: Expert knowledge and environmental governance* (pp. 67–105). Cambridge: MIT Press.

Paola, C. (2011). Simplicity versus complexity. *Nature, 469*, 38.

Parker, W. (2009). Does matter really matter? Computer simulations, experiments, and materiality. *Synthese, 169*, 483–496.

Ruphy, S. (2011). Limits to modeling: Balancing ambition and outcome in astrophysics and cosmology. *Simulation and Gaming, 42*, 177–194.

Schweber, S., & Wächter, M. (2000). Complex systems, modelling and simulation. *Studies in History and Philosophy of Science Part B: History and Philosophy of Modern Physics, 31*, 583–609.

Smith, J. M. (1995). Life at the edge of chaos? *New York Review of Books, 42*(4), 28–30.

Springel, V., et al. (2005). Simulations of the formation, evolution and clustering of galaxies and quasars. *Nature, 435*, 629–636.

Winsberg, E. (2013). Simulated experiments: Methodology from a virtual world. *Philosophy of Science, 70*, 105–125.