

Eric Winsberg

A Tale of Two Methods¹

Abstract

Simulations (both digital and analog) and experiments share many features. But what essential features distinguish them? I discuss two proposals in the literature. On one proposal, experiments investigate nature directly, while simulations merely investigate models. On another proposal, simulations differ from experiments in that simulationists manipulate objects that bear only a formal, (rather than material), similarity to the targets of their investigations. Both of these proposals are rejected. I argue that simulations fundamentally differ from experiments with regard to the background knowledge that is invoked to argue for the “external validity” of the investigation.

1. Introduction.

Imagine two physicists interested in studying the interaction of a pair of fluids at supersonic speeds. Each of them uses sophisticated technological artifacts to generate images of the flow structures that are generated as a shock wave propagates through a fluid. Each of them manipulates the equipment so as to be able to investigate their phenomenon of interest at a variety of values of basic parameters—different relative

¹ Thanks to Martin Carrier, Johannes Lenhard, Janet Kourany, Alfred Nordmann, and other members of the ZiF working group on “Science in the context of application”, two anonymous referees, and especially to Wendy Parker, for helpful comments and advice. Thanks to the ZiF, University of Bielefeld, for financial support during the preparation of the manuscript.

speeds, different densities of fluid, different geometrical configurations, different boundary conditions, etc. And each of them analyzes the data and images they generate in order to try to discern fundamental patterns, scaling relations, and other features of interest in the flow.

The first physicist's equipment is a laboratory setup consisting of a tank of fluid containing simple spherical and cylindrical shapes, bubbles of gas, and a physical mechanism for causing a shock wave to propagate through the tank. The second physicist's only piece of equipment is a digital computer. Using models from the theory of fluid dynamics as a rough starting point, the second physicist builds an algorithm suitable for "simulating" the relevant class of flow problems, and transforms that algorithm into a computer program that runs on her computer. The computer outputs data, including perhaps graphical output depicting flow patterns.

Is there a fundamental difference—a difference of kind—between these two activities? And if so, how should we characterize it? How can we make precise what distinguishes activities of the first kind, the traditional kind of activity that we call "experiment", from activities of the second kind, what we usually call "simulation." And are there, in particular, fundamental respects in which the nature of the epistemological relationship between the artifact and nature—that is to say regarding our abilities to use the artifact to learn about nature—differ in the two cases?

One obvious difference is the special role of the computer in the second example. But there is reason to think that, at least in one respect, this difference is not entirely fundamental. It is useful here to remember that, at least on one common understanding of the notion of "simulation", not all simulations are computer simulations. There is

also a class of techniques for investigating nature often called “analog simulations”. There are plenty of paradigmatic examples. Take the example of using fluids to simulate the dynamics of black holes. Here is the basic idea: we would like to understand the behavior of a class of black holes, and we have Einstein’s field equations at our disposal, but we are unable to solve them for the relevant boundary conditions. Instead of finding a computer program to simulate the black holes, physicists find a fluid-dynamical setup for which they believe they have a good model, and for which that model has fundamental mathematical similarities to the model of the systems of interest. They observe the behavior of the fluid setup in the laboratory in order to make inference about the black holes.

Common language use—the word “simulation” is often used to describe both these kinds of activities and “computer simulations”—suggests that there is perhaps some fundamental quality that is shared by analog simulations and computer simulation, and which jointly distinguishes them from ordinary “experiments.” To repeat the question, then: is this so? And if so, what is that fundamental quality? If analog simulations count as simulations, then it cannot be the special role of the computer that is fundamental. So what, if anything, is?

2. Competing Intuitions

Intuitions here can tug in opposite directions. On the one hand, we are inclined to think that the first two physicists’ activities could not possibly differ from each other more. The first physicist, so this way of thinking suggests, *is generating novel empirical knowledge about fluids by manipulating an actual fluid*. The other physicist is doing no such thing. She is merely *exploring the consequences of manipulating existing*

knowledge—in this case the Navier-Stokes equations—by using brute-force methods to crank out solutions to those equations that are, merely because of practical difficulties, difficult to generate by more traditional paper-and-pencil means. The following quotation encapsulates this intuition rather succinctly: "The major difference is that while in an experiment, one is controlling the actual object of interest (for example, in a chemistry experiment, the chemicals under investigation), in a simulation one is experimenting with a model rather than the phenomenon itself" (Gilbert and Troitzsch, 13).

The opposite intuition fixates on the experimental qualities of simulations—even and perhaps especially computer simulations—and finds no fundamental difference. Why, its proponents ask, does the second physicist often refer to what she does as conducting “numerical experiments”? Why does she call what she generates “data”? Why does simulation practice resemble experimental practice in so many obvious respects²? Must we dismiss all this as just loose metaphor?

Perhaps more significantly, there are troubling questions we can raise about some of the assumptions that lie behind the first intuition—assumptions that play a crucial role in painting such a stark contrast between the activities of the two physicists. The most troubling assumption, I think, is that experimenters “control the actual object of interest.”

More often than not, this is simply not true. What if we were to find out that both of our original pair of physicists’ primary area of interest is astrophysics? The systems that actually interest them both are the supersonic gas jets that are formed when gasses are drawn into the gravitational well of a black hole. Neither physicist, then, is

² See Dowling (1999) for examples.

actually manipulating his or her actual system of interest. Neither one is even manipulating a system of the same type, on any reasonably narrow sense of the term. Each one is manipulating something that *stands in* for the real class of systems that interest them. In one case, that stand-in is a tank of fluid. In the other, it is a digital computer. In both cases, the actual systems of interest are vastly different from the system being manipulated—in scale, in composition, and in many other respects.

This is a fairly common feature of experimental work of all kinds. Laboratory setups often differ in substantial respects from the classes of natural systems for which they are intended to speak. Think of Galileo watching the chandelier swing to learn about how all bodies fall, or Mendel manipulating his pea-plants to learn how traits are passed on from parent to offspring throughout the plant and animal kingdom, and in humans.³ So both of our physicists, and indeed almost all scientists, it turns out, rely in the end on *arguments*—either explicit or implicit—that the results that they get from manipulating their respective pieces of equipment are appropriately *probative* concerning the class of systems that interest them. The same, it seems appropriate to say, is true of almost all experimental work. And of course it is never (at least in the sense that Gilbert and Troitzsch intend) quite correct to say that anyone is “experimenting with a model.” The models Gilbert and Troitzsch are speaking of—the models that inspire the computer

³ It might be argued that Mendel’s peas and Galileo’s chandelier are instances of the systems of interest, and the physicist’s tank is not, but this would be somewhat question begging. In some respects, the physicist’s tank is an instance of the system of interest, since it is in fact an instance of a supersonic interaction of a pair of fluids. And few of Galileo’s contemporaries would have thought of his chandelier as a “freely falling object.” Some, conceivably, might have doubted that cultivated plants are an instance of natural heredity. The point is that what all of these examples have in common is that the object being manipulated or observed speaks for more than itself, and it takes an argument (even if that argument turns out to use, as its major premise, that one is an instance of the other) that it can validly do so.

programs in computer simulations— are abstract entities⁴, and we cannot experiment with them. What simulationists manipulate is a physical entity: either a digital computer or some sort of analog device.

Of course, it could be objected that this line of reasoning does not properly distinguish between a computer program, which might be said to be the basis of a simulation, and the underlying hardware, which might be said to be accidental. On the view of such an objector, it is not a computer that is manipulated by a simulationist, but the computer program—an abstract entity.

It is certainly true that computer programs are multiply realizable, and what seems most salient about a computer, when it runs a simulation, is the program that it runs, and not the particular hardware that runs it. But I want to resist this objection for two reasons. The first is that, while it is tempting, and indeed useful, in many contexts, to think of computation in abstract terms, there can be no real-world computation without some physical system to implement the computer. A computer program, corollarily, cannot be manipulated without manipulating the physical system that implements it—not necessarily by changing the hardware connections in the computer, but certainly by effecting some physical changes. The second reason is that if we hope to get clear on the fundamental relationships between experiment and simulation—where simulation is

⁴ Some commenters take the view that the “object” (in the sense I define below) of an experiment *is* a model. So, on this way of thinking, Mendel’s peas were model organisms, and Galileo’s pendulum is a concrete model of the freefalling object. And of course these kinds of models, and many others, are not abstract. I am perfectly sympathetic to this kind of talk, but notice that if we adopt it, Gilbert and Troitzsch’s claim is still false—since it then becomes wrong to say that only the simulationist manipulates a model. And in any case, I do not think this is the kind of model that they had in mind.

explicitly taken to be a kind that includes both analog and digital cases—then we are forced to take seriously the material characteristics of computation.⁵

The other problematic assumption the line of reasoning typified by Gilbert and Troitzsch is the assumption that the physicist using the computer is not generating new knowledge, but merely exploring the consequences of existing knowledge in the form of the Navier-Stokes equation. This is obviously not true in the case of analog simulations, and I would argue that it is not true in the case of computer simulations either. To think it is true is to assume that anything that you learn from a computer simulation that is based on a theory of fluids is somehow already “contained” in that theory. But to hold this is to exaggerate the representation power of unarticulated theory. As I argued in (1999), simulations involve complex inferences as they move from theory to data, and they certainly generate new knowledge. It is a mistake to think of them as simply tools for unlocking hidden empirical content.

So the intuition that seems to lie behind, for example, the view articulated by Gilbert and Troitzsch, appears to be on shaky ground. Let us see if we can make the reasons for that more precise. Following the literature on this issue⁶, let us call the class of systems in which the physicists are interested (in our case gas jets) the “target” of their investigations. And let us call the artifact that they intervene on and observe the “object” of their investigations. What Francesco Guala, Wendy Parker, and others have made

⁵ There is a third reason to take seriously the physical character of the computer in computer simulation. A number of philosophers have argued for continuity between simulation and experiment by explicitly conceiving of the computer in computer simulation as a physical object that is experimentally manipulated. I would include here Humphreys (1995), Hughes (1999), Norton and Suppe (2001), and to a limited extent Parker (forthcoming). To simply assert that the physical characteristics of the computer are incidental would be to beg the question against these claims. I prefer to follow all of these commenters, and take seriously the idea that computer simulationists use computers as physical stand-ins, and to locate the special characteristics of simulation elsewhere.

⁶ Especially Guala (2002) and Parker (forthcoming).

perfectly clear is that both of our physicists have to establish what is sometimes called the “external validity” of the conclusions they draw from their activities.⁷ They each have the non-trivial task of establishing, that is, that what they learn about the behavior of the object of their investigations can be appropriately informative about their targets. This makes it naïve to think that there is an uncomplicated sense in which the first physicist is studying nature directly, while the second one is studying only a model.⁸

One is still bound to sense, however, that there is some kernel of truth to the first intuition. The above arguments, in other words, are unlikely to shake many from at least the suspicion, if not the conviction, that there is still something fundamentally different, something fundamentally epistemologically different, about the two kinds of activities described above. That suspicion (or conviction) is likely to be that, all of the above notwithstanding, the experimenter simply has more direct epistemic access to her target than the simulationist does. How could we try to salvage that kernel of truth?

3. Material vs. Formal Similarity

One suggestion on how to go about this comes originally from Herbert Simon (1969), but has been made more explicit by Francesco Guala (2002). Guala is acutely aware that both experiments and simulations have objects on the one hand and targets on the other, and that, in each case, one has to argue that the object is suitable for studying the target. Despite this similarity, Guala thinks there is still a profound difference. The

⁷ An experimental result is internally valid when the experimenter is genuinely learning about the actual system he or she is manipulating—when, that is, the system is not being unduly disturbed by outside interferences. An experimental result is externally valid when the information learned about the system being manipulated is relevantly probative about the class of systems that are of interest to the experimenters. Parker (unpublished), cites Campbell (1957) as the original source of this conceptual distinction.

⁸ The quotation from Gilbert and Troitzsch comes from Parker (dissertation) and there she develops a sustained argument against this view that is similar to the one I offer here. Other related arguments can be found in Guala (2002) and Winsberg (1999).

difference, Guala argues, is that there are two fundamentally different kinds of relationships that can exist between an object being investigated on the one hand, and the target of that investigation on the other.

The difference, according to Simon and Guala, is this: In an experiment, the relationship between object and target is that they share a “deep, material” similarity. In a simulation, the similarity between the object and the target systems is only abstract and formal. In the first case, “the same material causes” are at work in the object as in the target, but not in the second case. To make Guala’s proposal somewhat more precise, “experiment” and “simulation” are really two-place predicates: we should count an investigation as an “experiment” just in case the object of the investigation bears a deep material similarity to its intended target and if the same material causes are at work, and call it a “simulation,” if the object bears only an abstract, formal similarity to its intended target.

Mary Morgan (2002, 2003) has argued for a similar view. She, in fact, urges that this difference between what she calls “material” experiments and simulations is precisely what makes experiments epistemically privileged compared to simulations. The fact that the object of a simulation bears only a formal similarity to its target, according to her, makes the task of establishing a simulation’s external validity—of establishing that the object is a suitable sort of entity for studying the target—that much more difficult than in the case of an experiment.

These suggestions are fairly compelling, and our example seems to support them. After all, the first physicist’s apparatus, despite bearing some significant dissimilarities to an intergalactic gas jet, is still, after all, primarily composed of fluids. These fluids really

have different densities, and they really flow past each other at supersonic speeds. The second physicist's apparatus is made out of silicon and wire. It has none of the significant material properties of a gas jet. There is indeed an appealing sense in which the first pair shares a "material similarity" that the second pair lacks; that the same material causes (conservation of momentum, viscous forces, advection, etc.) are at work in one case, but not the other. Furthermore, we have the general impression that the material similarity between object and target—the fact that the same material causes are at work; the fact that both are fluids—in the first case will guarantee that there will automatically be at least *some respects* in which the results will be informative about the target. In the second case, the impression urges us, the computer can only be informative about a gas jet in virtue of being suitably programmed; the reliability of the results depends *entirely* on having chosen the right model, and the right algorithm. Thinking of analog simulations, a fluid can only be informative about black holes *if* a fairly substantial assumption turns out to be correct: if the researchers are correct that there is indeed relevant formal similarity between a good model of black holes and a good model of fluids. One cannot help but be struck by this difference and find it significant, and perhaps even conclude, with Morgan, that this difference results in a significant disparity between the epistemic power of the simulation and that of the experiment. But despite the appeal of these suggestions, there are some obstacles that they need to overcome.

In particular, I think there are two problems with this account, which, following Parker in her (forthcoming), we ought to worry about. Take first the claim that simulations and experiment can be distinguished by the type of similarity that obtains between the object and the target of the investigation—whether it is deep and material, or

merely abstract and formal. The notion of material similarity here is too weak, and the notion of mere formal similarity too vague, to do the required work. Consider, for example, the fact that it is not uncommon, in the engineering sciences, to use simulation methods to study the behavior of systems fabricated out of silicon.⁹ The engineer wants to learn about the properties of different design possibilities for a silicon device (often a computing or a communications device), so she develops a computational model of the device and runs a simulation of its behavior on a digital computer. Naturally, there are deep material similarities between, and some of the same material causes are going to be at work in, the central processor of the computer, and the silicon device being studied. Should we therefore conclude that the nature of this investigation is more like that of our first physicist than our second? Probably not. One problem is that, in this case, it seems quite clear that the *relevant* similarities are not material. This is easy to tell in this case because we know that the simulation would run equally well if the computer were made out of gallium arsenide.

The peculiarities of this example illustrate the problem rather starkly, but the problem is in fact quite general: Any two systems bear some material similarities to each other and some differences. The clear lesson of the Gallium Arsenide processor is that, what Guala, Morgan, and Simon must have had in mind was that that the *relevant* similarity between the two systems be either a material or a formal one. But this idea might be difficult to spell out in detail in a way that works. Indeed, once we put it in its proper context, the whole idea of two material entities having formal similarities becomes rather obscure. We will return to this point shortly.

⁹ It is common, for example, in the design of nanoelectromechanical systems. The locus classicus for methods of simulating solid-state silicon is (Stillinger and Webber, 1985)

The second thing that we need to recognize is that on the Simon/Guala definitions of simulation and experiment, they are both success terms. An investigation will count as an experiment only if it is successful in the sense that the relevant material similarity between object and target actually obtain, and a simulation will be successful only if the relevant formal similarity between object and target actually obtains. But this seems wrong. Surely there can be failed experiments and failed simulations. That is, surely there can be examples of experiment and simulations that fail, in the end, to be externally valid. But on the kinds of accounts offered by Guala, Morgan and Simon, there cannot be.

There is a related worry: if experiment and simulation are success terms, then investigators may never be in a position to know if they are conducting a simulation or an experiment, since they may not know if the relevant similarity they have established a material, or merely a formal one. Think about Galileo's famous example of dropping a mass from the mast of a "moving" ship. What Galileo wanted to show, of course, was that the mass would fall at the bottom of the mast and that, by extension, a mass dropped from the top of a tower would fall at the base of a tower on a rotating earth. But a critic of Galileo's argument could presumably have doubted whether the extension was legitimate. He could have doubted, as I assume some did, whether the same causes are at work when a ship is in motion as when the entire world rotates. And so according to the material similarity criterion, Galileo and his critics would have disagreed about whether the ship study was an experiment or a simulation. But this seems troubling. True, not all semantic categories need to be epistemically accessible. It does not seem to be the case, however, that we need a God's eye perspective to know whether something

is an experiment or a simulation. It would be especially peculiar for Morgan if this were so, since she thinks that experiments are more epistemically powerful than simulation. But what good is knowing that if we can never be sure if something is an experiment or a simulation?

4. Simulation as activity and as representation.

What should we conclude? One possibility is to give up on drawing a conceptual distinction between computer simulation and analog simulations on the one hand, and experiments on the other. Parker, for example, can be read as being skeptical of attempts to distinguish the kinds of activities our first two physicists are engaged in terms of the one doing an experiment, and the other a simulation. Instead, she argues that the distinction between the terms “simulation” and “experiment” should not be drawn in anything like the way that we have so far been inclined to draw it. In fact, I would argue that on Parker’s view, the two terms refer, roughly speaking, to two ontological sides of the same coin. Here is what I mean by that:

Parker defines a simulation as “a time-ordered sequence of states that serves as a representation of some other time-ordered sequence of states; at each point in the former sequence, the simulating system’s having certain properties represents the target system’s having certain properties.” An experiment, for Parker, is “an investigative activity that involves intervening on a system and observing how properties of interest of the system change, if at all, in light of that intervention.” (Parker, forthcoming, pp. 4-5).

I say that, on these definitions, the terms refer to two ontological sides of the same coin because the distinction is roughly analogous to the distinction between a car and driving. For example, according to Parker’s definitions, both of our physicists at the

beginning of the paper are engaged in activities that involve both simulation and experiment, with each term merely emphasizing different aspects of their activities. For the first physicist, the tank can *serve as a simulation* of the intergalactic gas jet, insofar as the tank undergoes a time-ordered sequence of states, and the physicist believes that the tank in each of these states is representative of what state the intergalactic gas jet would be in. What the physicist is *doing with the tank is an experiment* in so far as he intervenes on the tank in order to investigate its properties in light of that intervention. Symmetrically, for the second physicist, the computer *serves as a simulation* of the gas jet as it moves through a series of computational states, and what the physicist is *doing with the computer is an experiment*, in so far as she intervenes on the computer by putting it in a particular initial state and observes its subsequent states to learn about its properties in light of that intervention.

It is clear that on Parker's definitions there are examples of experiment that do not involve simulation (such as when I intervene on an object in order to learn only about that very object in particular). And there are also examples of simulations that are not used for experimenting (I might build an orary that simulates the solar system just to display the motions of the planets, or program a computer simulation for educational purposes). But these examples are exceptional, just as are examples of driving without a car, or using a car without driving it.

In principle, I have no problem with using these two terms in this way. The definitions are clear and useful. They are in fact no doubt actually used that way in

ordinary scientific parlance in some contexts.¹⁰ Indeed, on the view of experiment that Parker and I share, this use of the term simulation is particularly useful. The view we share, I take it, is that, in a large class of experiments, there is some object involved that stands in for the target of interest. And Parker's definitions of simulation is useful for describing, for many of those experiments, precisely what kind of standing in is involved.¹¹

But I am reluctant to give up trying to draw a clear conceptual distinction between the kinds of activities that are exemplified by our two physicists. And I also think—though points of language are not that important to me—that the two sorts of activities can be usefully distinguished using the pair of terms “experiment” and “simulation”. I think, indeed, that in most ordinary contexts, the terms are used in opposition to each other precisely to distinguish the two kinds of activities we have been discussing. And using the terms in this way helps to make sense of the fact that the word “simulation” is used to talk about computer simulations, and analog simulations.

So I think it is worth distinguishing two rather different uses of the term “simulation.” In one sense of the word, a simulation_R is a kind of representational entity. This sense of the word is covered well by Parker's definition. But in the other sense of the word, simulation_A is a kind of activity on a methodological par with, but different from, ordinary experimentation. This second sense of the word unifies computer

¹⁰ It is perhaps interesting to note, moreover, that there are examples of computer simulation studies that do not involve “simulation” on Parker's definition at all. Many so called “Monte Carlo simulations”, for example, produce results without doing anything like going through a sequence of states that represents the sequence of states that the target system goes through. Stop such a simulation half way through its evolution, and the state it is in does not correspond in any way to a state of the target system. Perhaps the correct response is to deny that these are genuine simulations. But that seems far from common practice.

¹¹ One could even refer to this view of experiment, using Parker's definition of simulation, as the “simulation account of experiment.”

simulation and analog simulation. And while simulation_R is close to co-extensive with experiment (many simulations_R are experimented on, and many experiments involve simulation_R), simulation_A is meant to *distinguish* certain kinds of activities from ordinary experiments (though there may, of course, be some rare borderline or hybrid cases.) The contrast class for simulation_A is ordinary experiment; there are ordinary experiments¹² on the one hand, and there are computer simulations and analog simulations on the other.¹³

Indeed, I think it is fairly important to clarify the distinction between what we traditionally call “experiment”—the kind of activity exemplified by the first physicist—and “simulation”—the kind exemplified by the second. Let me say first why: One of my principal interests is in the epistemology of simulation. And I think that this enterprise depends crucially on our ability to sort out the epistemological respects in which simulations and experiments resemble each other, and the respects in which they differ. Recall Morgan’s claim that “traditional experiments” have greater epistemic power than simulations, in part because they “have greater potential to make strong inferences back to the world”. Parker, quite correctly I think, disputes *this* claim. She points out that there are some circumstances in which it will be substantially easier to establish the validity of a traditional experiment, but there are others in which, for example, a computer simulation would provide arguably more reliable results.

¹² Whether we want to contrast simulations with “experiments” or with “ordinary experiments”—that is to say whether or not we want to think of simulations as a particular kind of experimental activity—seems to be to an issue of whether or not to award them an honorific title. And that motivation, it seems to me, is grounded in the misguided intuition that “experiments” are intrinsically epistemologically superior that Parker is so keen to overthrow. Whenever it is convenient, I try to remember to contrast simulation with “ordinary experiment” so as not to pre-judge this question. I avoid the term “material experiment” because, as this paper should make clear, I do not think “materiality” is fundamental to what separates simulation of ordinary experiment.

¹³ For the remainder of the paper, unless I specify otherwise, I will mean simulation_A when I use the word.

The simple point is that the details matter. Consider once again our two physicists. If we want to know which physicist has greater potential to make strong inferences about intergalactic gas jets, we will need to know a great deal about the details of their work. How closely do the conditions in the tank mirror the conditions in intergalactic space? How much “noise” does this kind of apparatus generate? How much has its credibility been established with past performance? Similar questions need to be asked about the computer model: How credible is the underlying model? How crude are the approximations used in the computing scheme? How fine is the discretization grid? How many factors (viscosity, compressibility, electromagnetic forces, etc.) have been included or omitted?

So it is true that experiments are not *intrinsically* more epistemically powerful than simulations. But there may still be important epistemological differences between experiments and simulations. Indeed I think there are. Just because it is the case that experimentalists often face epistemological challenges that are just great as those faced by simulationists, it does not follow that the *kinds* of challenges they face do not have fundamental differences. And I think they are worth spelling out.

5. Arguments and Background Knowledge

So let us review where we stand. It was overly simplistic to say that experiments differed from simulations in that the first investigates nature directly, while the second merely investigates a model. Both experiments and simulations involve an object and a target. And in both cases, the task of establishing the validity of using the object to make inferences about the target can be substantial and non-trivial. And in both experiments and simulations, the object of investigation is a material entity.

The material similarity proposal—the idea that the object and target of a simulation lack the kind of deep material similarity that one finds in experiments—came in response to our recognition of this state of affairs, but it was unable to overcome two obstacles. The first obstacle was that the distinction between a deep material similarity and a mere formal similarity was too vague. We considered refining the proposal to focus on the *relevant* similarities. But then the second obstacle was that it seemed wrong to define simulation and experiment in such a way that they are both success terms.

What can we do about these problems? Parker offers the helpful suggestion that the following amendments to Guala’s proposal might do the trick: that simulation studies are characterized by the fact that the investigators *aim* for their objects to have *relevant* formal similarities to their targets and that ordinary experiments are characterized by the fact that the investigators *aim* for their objects to have *relevant* material similarities to their targets (forthcoming, pp.4-5). Adding the word “relevant” is supposed to take care of the first obstacle, and saying that the investigators *aim* for the (material or formal) similarity instead of saying that the similarity *actually obtains* is supposed to take care of the second obstacle.

I do not think this works. I think the whole idea of formal vs. material similarity is confused, no matter how much it is tempered by "relevant", "aimed for", or whatever. First, I am puzzled by the idea of two concrete entities having objective formal similarities. Give me any two sufficiently complex entities and I can think of ways in which they are formally identical, let alone similar. And I can think of ways in which

they are formally completely different. This fact is at the heart of one of John Searle's basic complaints about computational theories of mind.¹⁴

Now, we can speak *loosely*, and say that two things bear a formal similarity, but what we really mean is that our best formal representations of the two entities have formal similarities. Take the case of the use of fluids as analog simulators of black holes. What would it mean to say that a black hole is formally similar to a vat of fluid? This cannot be a statement about an objective relation between two entities. The only thing it could mean is: we believe that there are formal similarities between the mathematical structure of our best models of fluids on the one hand and our best models of highly-curved space-time manifolds on the other—and that it is this fact, rather than any material similarity between the two entities, that is being exploited by the researchers.

Properly speaking therefore, when one claims that the researchers “aim for relevant formal similarities” between two concrete entities, what one must really mean is: they have a way in mind of modeling their target, and they have a way in mind of modeling their object, and what they hoping is that on that way of thinking about the two entities, formal similarities will exist between those two models. But when you phrase it like that, what you have, in a nutshell, is the claim that simulationists aim for their objects to properly stand in for their targets—to be simulations_R of their targets. But *that* is precisely what is *aimed* for in *both* cases of simulation_A and ordinary experiment. At least so says the “simulation_R account of experiment” to which Parker and Guala presumably subscribe.

¹⁴ I happen to think that Searle's specific worries about CTM have been answered by Chalmers (1996), but not in a way that is of any help here.

What distinguishes simulation from ordinary experiment is what forms the *basis* for that hope. We will have to spell out the details with more care in what follows, but roughly, it is this: In one case, we base that hope on the fact that we know how to build good models of our target systems, and in the other case, we (in many cases) base that hope on the fact that the object and the target belong to the same *kind* of system—or in some cases, if you prefer, that they are materially similar. It is wrong to say that experimenters aim for their objects and targets to have material similarities. They *aim* for the one to stand in for the other, and (in many cases) they *rely* on the fact that the two belong to the same kind—and hence perhaps have material similarities—to try to argue that they are likely to achieve that aim.

It might be obvious what is coming next: if we want to characterize the difference between an experiment and a simulation, we should not focus on what objective relationship actually exists between the object of an investigation and its target, nor even on what objective relationship is being aimed for. We should focus instead on epistemological features—on how researchers *justify* their beliefs that the object can stand in for the target. When we do, here is what we find: what distinguishes simulations from experiments is *the character of the argument given* for the legitimacy of the inference from object to target and the *character of the background knowledge* that grounds that argument. Simulations, in particular, are legitimated by a very special kind of argument and background knowledge. In simulation, the argument that the object can be used to stand in for the target—that their behaviors can be counted on to be relevantly similar—is supported by, or grounded in, certain aspects of model building practice. We will now need to spell out what those are.

What separates ordinary experiments from simulations are the answers to these questions: Why do the researchers believe the object can serve as a good stand in for the target? What kind of background knowledge do they invoke, implicitly or explicitly, and does their audience need to accept, in order to be convinced that one can learn about the target by studying the object?

In an experiment, the argument that the object can stand in for the target can be based on a variety of possible considerations¹⁵. It might, for example, be based on something like the belief that the object and target are members are presumed to be of the same kind, or have the same material composition—this is the kernel of truth behind Simon and Guala’s proposal.¹⁶ What then, is the nature of the background knowledge that grounds belief in the external validity of a simulation?

The first pass at an answer would be this: In a simulation, the background knowledge that is required to argue for the external validity of the study is *trust in a model of the target systems*. But this will not work for a variety of reasons. The requirement is too weak and too strong. It is too weak, because in experiments, we also need to trust models of the target system. For one thing, we need to have some kind of model of the target system in order to decide, for example whether the target and object are of a kind, or have material similarities, or whatever. For another, in an experiment, when we use an object to study some target class of systems, we are holding out the

¹⁵ If there even is such an argument. Recall that in some experiments, the object is not distinct from the target, and hence no such argument is required.

¹⁶ Note that the difference between what I am suggesting here, and the Guala and Simon proposal is somewhat subtle. The difference, certainly, does not hinge whether the requirement is that object and target be “of the same kind” or “materially similar”. These are similar requirements. The difference is that Guala says that an experiment is characterized by a material similarity between object and target. I say it is characterized by a belief that there is a material similarity, or a common kind, or something of that nature, and that this belief plays the role of background knowledge in support of the argument that the object resembles the target in a relevant, formal way—that it simulates_R it.

object as a model of the target¹⁷. So on one construal, experiments require us to have trust in model of the target systems. But the requirement is also too strong. It is too strong because simulationists often do not begin, by way of background knowledge, with trust in a particular model of their target systems. The hard work of simulation involved the construction of such a model. This construction makes use of other kinds of background knowledge to sanction the trustworthiness of that model. Our job here, therefore, is to clarify what kind of background knowledge that is.

So let us return to our second physicist, and ask ourselves what kinds of background knowledge simulationists bring to the construction of models in computational fluid dynamics. I would argue, as I have elsewhere, that there are three kinds of background knowledge that they can bring to bear. The first is knowledge that comes from the theory of fluids. Simulationists will argue for the reliability of their simulation, in part, on the basis of the fact that the construction of their models has been guided by sound theoretical principles from the theory of fluids. Second, they will also rely on the soundness of their physical intuitions about the fluids they study. If you are simulating the flow of a river, for example, you might rely on the quasi-equilibrium assumption, which says that the flow of the river does not deviate too much from the steady state. How reliable this assumption will be depends entirely on how right you turn out to be about your physical intuition –that the steady state assumption is a good one. Finally, simulationists will rely, by way of background knowledge, on the soundness of computational tricks they employ. Simulations in computation fluid dynamics, for example, often rely on techniques like artificial viscosity, eddy viscosity,

¹⁷ See fn. 3

vorticity confinement, and others to increase the accuracy of their results. Trust, which presumably comes from a history of past successes, in these kinds of computational tricks is a third kind of background knowledge that grounds the trust in the external validity of a simulation.

Let us call these three kinds of background knowledge, (and there may very well be other, similar ones), “principles for model building”. More precisely, let us say that simulationists argue for the external validity of their simulation on the basis of their belief that they possess reliable such principles for building models of the targets of their investigations.¹⁸

We can now, I believe, properly distinguish simulation from experiment. It is the nature of a simulation that the argument for the suitability of using the object to stand in for the target depends, by way of background knowledge, on the researchers’ belief that they have reliable principles for building models (in the sense articulated above) of the very features of the target systems they are interested in learning about. Since simulations are generally used to study the dynamics of target systems, we should say that, simulations are investigations in which the choice and configuration of the object of the investigation is guided and constrained by principles taken to be reliable for building *dynamical* models (abstract models that depict temporal evolution) of the target systems.

¹⁸ The details of the example that I use to illustrate this account may lead some to think this account is a bit physics-centric. Perhaps it is. The extent to which basic theory plays a role is almost certainly higher in the physical sciences than in other disciplines. But I do think that the construction of most simulation models is guided by some mixture of theory, intuitive or speculative acquaintance with the system of interest (what one might, in the case of physical examples, call physical intuition), and tried and true computational methods. In any case, I trust that even if this is not so, the expression “model building principles” can be appropriately fleshed out in a like manner in any discipline one is inclined to study.

And it is constitutive of simulation that it is the purported reliability of those principles that provides the background for belief in the external validity of the investigation.¹⁹

It is true that, in a traditional experiment too, we make use of various modeling principles in selecting/constructing the experimental system (e.g. in helping us determine what sorts of things might be confounding factors). And those modeling principles can be part of the background knowledge that sanctions the results of the experiment. But I think there are two features of the background knowledge of simulation that make it distinctive. First, the relevant model building principles are specifically principles for building models of the *target* of the investigation. An ordinary experimentalist worries whether he can suitably control his object. For that, he needs to know how to model the object, not the target—and he is more interested in how the object is coupled to outside interferences than in its internal dynamics. Second, and perhaps more important, in simulation, the reliability of the model-building principles are invoked in arguing for the *external* validity of the study—when modeling principles are invoked in sanctioning an ordinary experiment, they are invoked on behalf of the *internal* validity of the study (for example, in arguing that the object has been adequately shielded, etc.).

The conceptual distinction between experiment and simulation is now clear: When an investigation fundamentally requires, by way of relevant background knowledge, possession of principles deemed reliable for building models of the target systems, and the purported reliability of those principles, such as it is, is used to justify using the object to stand in for the target, when a belief in the adequacy of those

¹⁹ It is worth noting that this definition of simulation would rule out activities in which there is no real world target of interest, such as a computational study designed to probe how the world would look like, if, e.g., gravity was proportional to $1/r^{2.5}$. These sorts of activities simply fall outside of the scope of my interests here.

principles is used to sanction the external validity of the study, then the activity in question is a simulation. Otherwise, it is an experiment.²⁰ Sometimes, especially in the physical sciences, some of the model building principles involved are guided by the theory under whose domain the behavior of the target system falls.²¹

George Platzman, speaking at a famous Meteorological conference in the 1960's, made a comment that nicely reinforces this view:

I may add to this another point mentioned by Dr. Charney, a somewhat philosophical comment concerning [ordinary] experiments. I think that I agree with Dr. Charney's suggestion that machines are suitable for replacing [ordinary] experiments. But I think it is also necessary to remember that there are in general two types of physical systems which one can think of modeling. In one type of system one has a fairly good understanding of the dynamical workings of the system involved. Under those conditions the machine modeling is not only practical but probably is more economical in a long run... But there is another class of problem where we are still far from a good understanding of the dynamical properties of the system. In that case [ordinary experiments], I think, are very effective and have a very important place in the scheme of things. (Siono, 1962)²²

In fact, Guala makes similar remarks in his (2002) regarding the nature of the methodologies of experiment and simulation. He notes that “the knowledge needed to run a good simulation is not quite the same as the one needed to run a good experiment.”

(p.70) This is exactly right. But I think of this not as symptomatic of the difference

²⁰ Obviously, there is a small problem here. Since I have characterized “experiment” negatively—experiments get picked out by what they do not require—one might worry that I am letting too much in. Perhaps washing my car, and whistling Dixie, count as experiments. But I am just assuming, from the point of view of this paper, that we have a pretty good pre-analytic notion of what kinds of activities fall under the union of the concepts experiment and simulation, and I am only trying to characterize the difference between the two. Presumably, Parker is right that what these activities have in common is that some object is carefully set up, intervened on, and then observed in order to learn about some target.

²¹ These two features of computer simulations in the physical science—that there is an object that stands in for a target on the one hand, and that the relevant model-building principles are close to theory on the other—are responsible for motivating the intuition that computer simulation “lies somewhere between experiment and theory” that one finds so often in the literature.

²² Thanks to Wendy Parker for pointing me to this quotation. I would not go so far as to argue that Platzman is advocating precisely the same view as I am. But I do think it resonates nicely with my view.

between (ordinary) experiment and simulation, but constitutive of it. Guala also, I think, fails to get just right what the knowledge is that one needs. In the case of simulation, he says that one needs to know “the relationships describing the behavior” of the target systems. This is not a very precise claim. But on an ordinary reading of it, it is clearly too much—if one had at our disposal all the relationships that described the behavior of a system, one would not need to conduct an investigation of it. As I argued above, it is too strong to say we need to have a model of the target system’s behaviors. What one needs, I urge, are reliable principles for building models of those behaviors. This is a very different requirement, and confusing the two is at the heart of a lot of misunderstanding about simulation.²³

The confusion is understandable, since many computer simulations in physics, for example, begin with differential equations. And it is tempting to think that differential equations perfectly describe the behavior of a system. But an unsolved set of coupled partial differential equations do no such thing. They describe how portions of the system would behave under counterfactual conditions. But until one has closed form solutions to a set of equations, one has no description of actual behavior.

I agree with Platzman. What we need is “a fairly good understanding of the dynamical workings of the system”. I would cash that out by saying that the background knowledge that a simulationist needs, in sum, are reliable principles for building dynamical models. In the case of an analog simulation, one needs reliable principles for constructing an abstract model of both the object and the target, and an argument—based

²³ See, for example my discussion of Norton and Suppe's account of simulation in my (2003).

in part on those principles—that the object of the investigation has been configured in such a way that the two models of these systems will have relevant similarities.

In the case of computer simulation, the object being so configured is a stored-program digital computer, and so it is configured by programming. The simulationist uses the principles deemed suitable for building models of the target to guide the construction of a computational model, and uses this model to write the simulation's computer code. When a digital computer is programmed, the computer program—an abstract entity—becomes a model of the behavior of the computer qua physical system. Since the simulationist has an argument from the fact that the computer program's writing has been guided and constrained by reliable principles for building models of his target systems, he has an argument that it is also a good model of the target systems. Hence there is an argument for the external validity of the simulation that is just like the one offered for analog simulations—that both object and target have models with relevant similarities.

We can now review some of the examples from above with these criteria in mind. We can start with the more obvious ones. Take the two physicists discussed at the beginning of the paper. The first physicist studies tanks of fluid to learn about astrophysical gas-jets. What makes this an experiment? What is important is the argument, and its background knowledge, legitimating the study. The first physicist does not need a toolkit for building dynamical models of her target to sanction the external validity of her study. She believes the inferences she will make are legitimate because she is prepared to argue that the two systems are, in relevant respects, the same kind of system, made out of the same material, and can be expected to exhibit relevantly similar

behavior. The opposite is the case with the second physicist. He has no commitments what-so-ever to the object and target being of a kind, but he must be willing to express the non-negligible hope, and to argue that the hope is well founded, that the theory of fluids, his physical intuitions about the situation of interest, and the model-building methods that have worked well in computational fluid dynamics in the past, provide him with a reliable means of constructing a model of his target system. He will want to argue, in other words, that the programming of his computer has been sufficiently guided and constrained by good principles for building models of fluids such that the computational model of his computer is relevantly similar to a good model of the behavior of the gas jets that interest him. And his knowledge of the theory of fluids, along with other model-building principles, plays a central role in underwriting that argument.

We can make some similar remarks about analog simulations—such as the black hole example discussed above. Here, the physicists must believe that they have good principles and methods for modeling black holes, good principles and methods for modeling fluids, and that these methods allow them to argue that the setup of the fluids they study has been guided and constrained by reliable principles for modeling black holes. They can then argue that a relevant similarity exists between a good dynamical model of the fluid, and a good dynamical model of the black holes that interest them. It is not that there is a (merely) formal similarity between black hole and fluid that makes this a simulation rather than an experiment. The relevant consideration itself is the need, by way of background knowledge, for a commitment to basic principles that guide and constrain our reasoning about these models and their similarity in the way spelled out above.

6. Conclusion: Epistemic Power.

What about the claim that experiments are epistemically privileged relative to simulations—the claim that they “have greater potential to make strong inferences back to the world”?

I think it is easy to see, following Parker, that this claim is false. A good computer simulation of the solar system—one that calculates orbits carefully from Newton’s laws—will provide me with better grounds to make inferences back to the world of the planets than almost any experimental setup I can imagine; because in such a case the relevant background knowledge—our ability to build good, reliable models—is virtually unassailable. How trustworthy or reliable an experiment or simulation is depends on the *quality* of the background knowledge, and the skill with which it is put to use, and not on which *kind* it belongs to.

But there are significant epistemological differences between simulations and experiments, and some of them might help to explain the appeal of the claim that experiments are intrinsically more epistemologically powerful. Take for example, the role that experiments can play in what Hempel would call hypothesis testing. Certainly, early 20th century philosophy of science, both the positivists and the Popperians, overstated the importance of hypothetico-deductivism and related activities in their foundational accounts of the role of experiment in science²⁴. But experiments do often play the role of providing crucial tests for theories, hypotheses, or models. And this is a

²⁴ I take this is in part what the “new experimentalists” like Ian Hacking taught us with such slogans as “experiments have a life of their own”.

role that cannot ordinarily be played by simulation²⁵, since simulations, as we have noted, assume as background knowledge that we already know a great deal about how to build good models of the very features of the target system that we are interested in learning about.

What this highlights is an important epistemological facet of the difference between simulation and experiment: for epistemic agents like us, experiments are epistemologically prior to simulations. In both the cases of simulation and experiments, you need to know something to learn something. But what you need to know in a simulation is quite abstract and sophisticated knowledge, and it usually depends on things you learned from a long history of experiment and observation. That is because we do not commit ourselves to the reliability of model building principles unless they have been tested against experiments and observations.

One might be tempted to think that the related claim—that experiments are more epistemically powerful than simulations—follows from what I call the epistemological priority of experiments. But I do not think this is correct. There may have been a time in the history of science, perhaps before Newton, perhaps even earlier, when we did not have sufficient systematic knowledge of nature—enough of a toolkit of trustworthy model building principles—for a simulation to ever be as reliable a source of knowledge as even the crudest experiments, but that time has long passed.

²⁵ To be more precise, there can be a role for simulation in the testing of models; but not in the same sense that I intend above. That is, computer simulation can be used to calculate what the model predicts about a particular situation, and that prediction can be compared with data from experiments and observations. But that is not the same as the role that experiments can play as the thing against which the prediction of a model, theory or hypothesis is compared. In fact, I am perhaps inclined to say that when computational methods are used to calculate what a model predicts in order to test the model against experimental results, we should refrain from calling this genuine simulation. R.I.G. Hughes in fact takes something like this view in his (1999).

References

- Campbell, D. (1957), "Factors Relevant to the Validity of Experiments in Social Settings", *Psychological Bulletin* 54: 297-312.
- Chalmers, David (1996) "Does a Rock Implement Every Finite-State Automaton?", *Synthese* 108:309-333.
- Dowling, Deborah (1999), "Experimenting on Theories", *Science in Context* 12(2): 261-273.
- Guala, Francesco (2002) "Models, Simulations, and Experiments" in Lorenzo Magnani and Nancy Nersessian (eds.), *Model-Based Reasoning: Science, Technology, Values*. New York: Kluwer, 59-74.
- Gilbert, N. and K. Troitzsch (1999), *Simulation for the Social Scientist*. Philadelphia: Open University Press.
- Hacking, I. (1983) *Representing and Intervening*. Cambridge: Cambridge University Press.
- Hughes, R. (1999), "The Ising Model, Computer Simulation, and Universal Physics", in Mary Morgan and Margaret Morrison (eds.), *Models as Mediators*. Cambridge: Cambridge University Press.
- Morgan, Mary (2002) "Model Experiments and Models in Experiments" in Lorenzo Magnani and Nancy Nersessian (eds.), *Model-Based Reasoning: Science, Technology, Values*. New York: Kluwer, 41-58.
- , (2003) "Experiments Without Material Intervention: Model Experiments, Virtual Experiments and Virtually Experiments" in Hans Radder (ed.), *The Philosophy of Scientific Experimentation*. Pittsburgh: University of Pittsburgh Press, 216-235.
- Norton, Stephen, and Frederick Suppe (2001) "Why Atmospheric Modeling is Good Science" in Clark Miller and Paul N. Edwards (eds.), *Changing the Atmosphere: Expert Knowledge and Environmental Governance*. Cambridge: MIT Press, 67-105.
- Parker, Wendy (this volume) "Does Matter really Matter: Computer simulations, experiments, and materiality."
- Simon, Herbert (1969) *The Sciences of the Artificial*. Boston: MIT Press.

Sigekata Siono, ed. (1962). *Proceedings International Symposium on Numerical Weather Prediction in Tokyo, November 1960*, Tokyo: Meteorological Society of Japan.

Stillinger, F. H. and Weber, T. A., (1985) "Computer simulation of local order in condensed phases of silicon." *Physical Review B* 31, pp. 5262–5271.

Winsberg, Eric (1999) "Sanctioning Models: the Epistemology of simulation" *Science in Context* 12(2), 275-292.

---, (2003) "Simulated Experiments: Methodology for a Virtual World" *Philosophy of Science* 70: 105-125.