Simulations are *not* Models

John Simpson†
University of Alberta

The aim of this paper is to argue that simulation is the activity of inferring future states. I argue that simulations instantiate models and that models are complexes of representations, so the inference in question makes use of the relations between the representations in a simulation's associated model. It follows that simulations should not be properly considered to be models in general, despite it being the case that they are commonly treated, or referred to, as being models, or even models of a special type, namely dynamic models. Further consequences of this position are also discussed.

1. Introduction

The distribution of chemicals by ocean currents and ground water systems, the resilience of buildings to disasters, the behaviour of airplanes in adverse weather conditions, and the spread of infections around the globe are examples of complex systems that are increasingly being investigated by simulations. These are also projects that have important consequences for both specific individuals and the human population as a whole. With human lives depending on the answers that simulations generate, in these and other important areas of research, we should expect to have available a clear account of *why* we are justified in using simulations in the way that we do—why it is that we may use simulations as tools of theory refinement, experimentation, and proof—as a legitimate substitute for, or even an accessory to, other scientific methods. Having this account becomes even more important once it is recognized that simulations are inherently (partially) false characterizations of the world (they are at best only *similar* to actual processes and never identical with them), raising the question of *how*

---

† In preparing this paper I am grateful to: Adam Morton and Vance Schamehorn, for their willingness to provide me with the benefit of their knowledge and experience throughout many drafts; Wes Cooper, for pointing out in the first place that my faith in the value of simulations was overly naïve; and Sarah Dawrant, for suffering through my philosophical tirades and editing them out of my drafts.

we can justify any claim that simulations lead us to (partial) truths whenever they are used.

In addition to simulations, other seemingly uncontroversial pieces of science seem open to question because on investigation their natures are not as innocuous as they might be on first inspection: models parallel simulations in their abstraction and simplifications and theories are known to work only for specific idealizations that do not admit of the complexities and exceptions of real world phenomena (Cartwright, 1984). Still, simulations are new in comparison to such more established pieces of science and so it behooves us to at least establish simulations as having a level of legitimacy on par with elements of scientific practice, such as theories and models, that are already considered legitimate. One plausible way to provide simulations with this level of legitimacy is to connect them to a recognized component of the scientific process that is already seen as legitimate in a way that will allow simulations to inherit this legitimacy.

A cursory glance at writings within the philosophy of science indicates that there is the possibility of such a relationship between simulations and models.[1] Such a glance will also indicate that models are seen as being more important than simulations in general and that the relationship between simulations and models is not symmetric. Such simple comparisons do not reveal the nature of the relation between simulations and models though and any deeper reading through these articles provides only the most general of senses about what the shape of this relation might be.

In what follows, I provide an account of simulations that allows for a legitimacy within scientific practice that is on par with models by linking simulations to a set of, what I call, underlying models. With the hope of being clear and succinct I begin my argument by extracting a general conception of what a model is by examining three general perspectives on what models are. Following this, a set of criteria by which we can determine what candidates will actually qualify as being

---

[1] A search through the Philosopher's Index for "model" and "simulation" returns the following results. "Model" *and* "simulation", 47 hits. "Model", 10613 hits. "Simulation", 385 hits.

simulations for the core arguments in this paper is given. With these tools in hand I then consider, and reject, the possibility that simulations *are* models, despite the ability of such a reduction to grant simulation the required legitimacy outright. The reasons for this rejection are three-fold. First, a simple thought experiment shows that while simulations may require what I refer to as underlying models to take place, these models exist independently of the simulation they underlie and so the two cannot be identical. Second, another thought experiment shows that simulations have a special characteristic of self-contained state-to-state determination which is not shared by those models which we might traditionally refer to as being "dynamic." Third, a look at the ways in which simulations are actually used shows that they are better understood as being a special kind of experiment rather than a special kind of model. With the relationship between models and simulations clarified, the core property of simulations identified, and the role of simulations within scientific practice roughed out, all via the rejection of the possibility that simulation are models, the discussion is quickly turned to the establishment of the legitimacy of simulations within scientific practice from these three results. Finally, with the legitimacy of simulations established, the paper closes with an investigation of two interesting questions on the ontology of simulations: specifically, "How do we distinguish between simulations?" and "What actually counts as a simulation beyond the narrow casting established early in the paper to streamline the discussion?"

2. A General Conception of 'Model'

Exactly what 'model' designates within the philosophy and practice of science remains open for discussion.[2] Fortunately, the details of this debate concern the task at hand only to the extent that by looking at few of the more well established perspectives, we may hope to extract a general notion of what a model is. If successful, this notion will allow us to not only examine the relationship between

---

[2] For a particularly poignant example of a contemporary debate that turns on the interpretation of 'model' see the exchange between Fetzer (1999) and Johnson-Laird and Byrne (1999)

models and simulations, but to have the conclusions drawn from this examination be flexible enough to accommodate any conception of a model that might finally be converged on.

Simplifying the situation a great deal, there are three perspectives that compete for our understanding of 'model.' In all likelihood, no one perspective captures our regular uses of 'model', but what matters for our purposes is the intent to do so. Once identified, these perspectives can be roughly shaken into relief from each other so that the differences can be shaved away from the commonalities. The distinctions chosen are similar to those used by Suppe (1974) and Morgan and Morrison (2001) and in the spirit of divisions made by Morton and Suárez (2001). It is of little consequence that overlaps are possible between these perspectives due to the imprecision with which they are made since the goal of this exercise is to establish a conception of models *in general* and not the production of a theory of models. The three perspectives to be considered are as follows:

a) *Models as abstract entities*. From this perspective models are principally relational entities given existence insofar as they are contained in someone's thoughts and ideas. These relations have a representational component to them, but this may be indirect. The syntactic and semantic conceptions of scientific theories most properly belong under this perspective since both require no more of models than they remain abstract entities held in the mind (Morgan and Morrison, 1999; Giere, 1988; Suppe 1974). On the syntactic view it is theories first and models second. Here, theories are defined as axiom sets within a specific language—typically first-order mathematical logic with equality—and models are defined as interpretations on which the axiom set is true. As *interpretations* models do not require physical existence in the world and it is this characteristic that places the syntactic view of theories within the models as abstract entities perspective. It may be the case that our interpretation connects things that exist in the world to the theory, but the interpretation remains abstract in the sense that it has no existence outside the head of an interpreter of the theory. Even claims that the model can be written down establish, at most, that physical entities can be triggers for bringing the model to the forefront of our minds and does not constitute a rejection of the claim that a mind is needed to do the interpretation precisely because the interpretation happens in the mind. The semantic view of theories turns this hierarchy on its head, defining the theory in terms of a set of models and ultimately removing the special status given to logic by the syntactic view because a wide range of languages could be used to generate sets of axioms that are true from the models provided. Swapping the precedence of theories and models does not make models any less abstract though, since the closest that models get to being physical entities is their empirical substructures (Giere, 1988), which can be understood as being the form that our experiences in the world impart on models as abstract entities. The value of the semantic and syntactic towards an adequate conception of models is not their individual definitions of 'model', but the position implicit between them that without minds to hold the similarity relations that compose models any claim that some thing or another is a model would be nonsensical.

b) *Models as physical instruments*. The conception of a model as a principally abstract entity is not one held by most people. This is not an indication of ignorance or misinformation about models but a reflection of how models are predominantly used in practice. Kits purchased from a hobby shop for building planes, the coloured balls and little sticks common to chemistry classes and computer programming code all have a physicality that contributes strongly to the conception that those unfamiliar with philosophy of science, logic, and higher order mathematics will attach to 'model.' Each of these models can be seen, touched, or directly manipulated, in virtue of the common reality that they exist in, in ways that abstract models cannot. The key idea here is that models are objects in the world that have an imperfect isomorphic relation to some particular thing or other (Suppe, 1974). Despite the imperfections of this relationship, these models serve as excellent instruments of instruction and as platforms for further investigation. An example of the explanatory power of physical models can be easily had by considering how an orrery or film can quickly convey the movements of the planets around the sun in a way that a lecture alone could never achieve.[3] There is something immediately accessible in the physical representation of a system normally beyond our observational capacities that makes any attempt to explain the concept of planetary motion without the model clumsy in comparison. The gaps that these models leave in their simplifications may be looked at as cause for concern, but in practice they are invitations for further exploration.

c) *Models as mediators*. This perspective is due to a recent compilation by Mary Morgan and Margaret Morrison (1999) in which it is argued that models should be properly understood as autonomous instruments of investigation that mediate between theory and world. It is the their view that models are: a) *autonomous* from theory and world because they are not wholly determined by either, but are the products of attempting to square the two together and; b) that models are *instruments,* in that through their representative nature they function as tools which aid in fashioning our understanding of world and theory. The models as mediators perspective is best understood as an attempt to extract the elements from the models as abstract entities and the models as physical instruments perspectives that reflect how scientists actually use models and synthesize them into a single perspective. The ideas in *Models as Mediators* have roots in earlier works, notably Cartwright (1983), Nagel (1961), and Hesse (1966). Morgan and Morrison point out that their view of models within science is not yet a theory of models that would help differentiate between model and theory, but rather a platform of information about

---

[3] An orrery is a clockwork device for representing the movement of heavenly bodies.

the place of models in practice that will “help to yield that information” (Morgan and Morrison, 1999, p. 12). It is the spirit of their project that is most important for our purposes here since it will be in this same spirit that the similarities between all three perspectives will be used to form a general conception of 'model' with which to compare simulation against.

While there are clear differences in these three perspectives on what a model is, an important commonality remains: each perspective requires that its models bear a set of relations such that it can be recognized as being a representation of the thing that it is claimed to be a model of. These relations are the source of the representative capacity of models because they *are* representational relations that the object being modeled bears, either within itself or to other objects. In the case of the syntactic theory of models for something to be considered a model it must make the theory it is to be a model of true. To do this the prospective model must not only specify entities that can give specificity to the general variables provided by the axioms, but these specified entities must maintain the relations generalized by the axioms. On the semantic theory a model is a structure, providing both a set of entities and a specification of how these entities relate to each other. It could be argued that some of those things that are called models on the semantic view are basic and unconnected to anything else, but this would not refute the claim being made here. Typically, these models are inspired by states of affairs observed in the world and thus are representations of these states of affairs. Of all the perspectives, the general need for models to be composed of representational relations is easiest to see in physical models where both the relevant components of the model and the thing being modeled can be pointed to, and the success of the representations judged on how well the relations in the actuality are maintained. Finally, the models-as-mediators perspective has models taking their representative cues from both theory and the world in ways similar to those of the other perspectives given here. While there may be other commonalities across these perspectives, it is the notion that models are collections of representative relations that will form the basis of our investigation into the relationship between models an simulations.

Before moving on and selecting a working conception of what counts as a simulation, it is important to handle a possible line of objections to this generalization of what counts as a model. These objections arise because taking models as basically being representative relations means that they are really only mental entities and thus be a claim that many physical things that we call models are no such thing. In particular this objection might be raised by those holding the models-as-physical-instruments view (and possibly some of those who believe that models may be instruments that mediate between theory and world) and phrased something like this: "But look at this tiny plastic car! Surely it is a model regardless of what anybody thinks. It even came from a model shop." The problem is only an apparent one though, resting on a confusion between the semantic content of the word and the conditions necessary for that content to be there in the first place. I am not asserting that models are models because we happen to call them "models" instead of "kumquats" or "rigmaroles". Rather, that without *someone* to use the things we call "models" as loci of representational relations, and thus impart them with meaning, these objects would never be able to operate as models do, in the case of models with physical instantiations, or exist at all, in the case of purely abstract models.

The principally mental nature of models could be challenged further though on the ground that the need for a mind to make the kinds of representational relations necessary for there to be a model is an endorsement of relativism that is ultimately self-defeating. Before any thoughts in this direction go too far—say to the point of claiming that the conception being painted here is self-defeating because if it were true then everyone would have a different conception of what a model is and thus make any attempt to talk about what a model is, including this one, nonsensical—we need to take stock of what we actually observe in our regular uses of "model." There is clearly variation in what different people call and use as models: lines drawn on a napkin can function as a model of the layout for a city's streets for one person and be meaningless for another; the models that a logician might use to capture a particular theory may be symbolized into obscurity for most everyone else; and a piece of abstract art can serve as a model of the human form to some and be nothing more than a blob of steel for others. Despite this variation there is a great overlap in the things that we consider to be and use as models: toy planes, instructive diagrams written on napkins, and plastic cadavers are all commonly agreed to be models. This overlap counts as evidence that there are boundaries on what what we conceive of as a model. The source of such boundaries is unclear and beyond the scope of this paper, but they can likely be found in an appeal to a combination of physiology and social environment.

Noting the examples used so far, models are pervasive entities in our existence and we might even go so far to suggest that there is an important reason for this in that physical models, whether miniaturizations or symbolizations, capture complex relationships that we might well not be able to fully hold in our heads. It is for this reason that some theorists might consider physical models to be extensions of our minds into the world around us (Wilson, 2004). While interesting, the exploration of such controversial ideas is beyond our purposes though. We can thus set such concerns aside and, taking our broad understanding of models as representational relations with us, turn our attention towards determining what the exemplars of simulation upon which our investigation of the relationship between simulations and models should be based.

3. A set of simulation exemplars

There are two ways that defining what is meant by 'simulation' can be done: we can either make an explicit definition that ensures that we will have the legitimacy we are seeking, or we can look at the things scientists are using as simulations and try to extract something definition like from this set of exemplars. Since simulations are already used so widely and since the project we are concerned with is to *determine* what legitimacy simulations have within current scientific practice, rather than simply *making* them legitimate, it is the latter method of determining what counts as a simulation that will be pursued. To this end, the processes typically referred to as "computer simulations" will form the basis of these paradigm examples, for two reasons. First, the frequency with which "simulation" and "computer simulation" are taken as being perfect substitutes suggests that they are the most easily recognizable sources of simulations available. Second, it is rise of computer simulations as a method within science that has brought about the need for an account of the legitimacy of simulations within scientific practice in the first place, and so it is fitting that they be the source from which our set of exemplars is drawn. Exactly what is meant by "computer simulation" though? Without further qualification "computer simulation" is rather ambiguous—it could be bent to capture any of the myriad of processes that a computer could undertake—and so further refinement must be provided before we can move on and examine the relationship between models and simulations.

Fortunately, such refinement is easily had by noting that we seem to have an intuitive sense of what processes carried out by computers as simulations and what do not—the programs that compose a spreadsheet or word processor seem to be importantly different in nature from the investigation into star thermodynamics discussed by Winsberg (2001) and the multi-agent systems examined by Sawyer (2004). The source of this intuitive division seems to be the that those processes considered to be computer simulations have acquired their status as such through our recognition that they are intended to stand-in for some actual process, whereas other processes do not have this character. This further refinement is still too rough to give us a class of exemplars though since there are many things that computers are used for that would seem to be both simulations and not simulations at the same time. Consider a flight-simulator as an example of a problematic case for our current casting of computer simulation.

A stereotypical flight simulation involves a human pilot who interacts with a computer through a set of controls that resemble those found in the cockpit of an actual plane. Based on the inputs provided through the controls the computer produces a set of outputs meant to be taken by the pilot as representations of what would actually happen in a real plane. This case is problematic for two reasons. First, the whole process is superficially no different from a computer program which we would intuitively classify as not being a simulation because the behaviour of the pilot and the computer are so intertwined that it is not clear we have a simulation at all. In both the cases of a spreadsheet and flight simulator, a human provides a set of inputs to the computer which then provides a set of outputs that are then used by a human as a basis for a new set of inputs, starting the whole process again. Second, the presence of the human pilot makes it difficult to call the *entire* system a simulation because the human pilot stands to be considered an actuality. Since we are seeking a set of examples that can be uncontroversially seen as being simulations, uses of computers that complicate matters by regularly involving outside intervention, such as flight simulators, need to be excluded until we have developed the conceptual tools to deal with their status. We can bring about this exclusion by distinguishing between a closed computer simulation and an open one.

A closed simulation is one that is entirely self-contained once the initial inputs have been provided, it contains all the necessary information to process the initial inputs or the data that it created as a result of processing this input. In contrast, an open simulation accepts new input on various occasions throughout running its program because it is meant only to be a component of a larger system which has not been represented within the program. Constraining our set of simulations exemplars to contain only those that qualify as closed computer simulations will allow us to set aside many problematic cases until later, when we have the insight necessary to accommodate them.

One final qualification needs to be put on our criteria for selecting a set of simulation examples before we move on: it must be understood that by “simulation” we are concerned with what is conventionally referred to as the *running* of a simulation and not the *construction* of whatever elements are required in order for the simulation to be run. I raise this distinction so that Winsberg's argument that “simulation is a rich inferential process, and not simply a “number crunching” technique” (Winsberg 2001, p. S442), will not be seen to be at odds with the narrower conception of simulation being drawn here. Reflection on Winsberg's discussion makes it clear that his reason for including the construction of a simulation under the name “simulation” is to highlight the entire simulation process—including: construction, use, and refinement—as being worthy of its own epistemology within the philosophy of science, an argument that I agree needs to be made.

Any further refinements beyond simply insisting that any candidate computer simulation be closed will run the risk of stripping the general notion of computer simulation of its accessible quality. So, with closed computer simulations as our paradigm referent for 'simulation' let us turn our attention towards establishing their legitimacy.

4. Simulations are *not* Models

With the general notion of the essence of what counts as a model and the criteria for determining what can count as a member of set of simulation exemplars available, respectively sets of representational relations and closed computer simulations, we can now turn to our attention towards trying to establish the legitimacy of using simulations in scientific practice via an exploration of the relationship between models and simulations. This section carries out this investigation in three parts, ultimately providing an argument in three parts for the demarcation of simulations from models in a way that passes the legitimacy of models onto simulations. Parts one and two are thought experiments that juxtapose our conception of what a simulation is with how they are actually used, with part one showing that simulations are not models *in general* and part two showing that simulations are not dynamic models *in general*. The third part takes the conception of a simulation refined by these thought experiments and argues that simulations should not be considered as simply special kinds of dynamic models because such a designation is not reflected in their use. Of equal importance to the distance that these arguments place between models and simulations is the connection that remains between them and which bestows the legitimacy of models on simulations. While this connection will be made apparent throughout this section, specific attention will not be drawn to it until the next.

Pause-and-Play

Imagine that you have before you a computer carrying out a closed simulation—it does not matter what is being simulated so long as the process as a whole is in the spirit of our set of simulation exemplars. Now, given that you have a simulation taking place, consider what remains if you were to halt the simulation in such a way that the computer ceased processing without loss or degradation of any of the states involved in the simulation. Is this suspended process a simulation? No: we speak of simulations as processes and since the process has ended (you stopped it after all) any claim that there was a simulation would lack the necessary existence of a process to even get started. What we do have is a computer with a collection of states within the hardware of the computer that represent some set of data: data that is both a description of a state of affairs and the instructions of what to do with these descriptions. The instructional data can be traced back to the high-level programming code in which the program was designed. Providing that this code was written to capture some set of relations observed elsewhere then we can consider this program to be representational of those relations and thus a model. The data that remains on stopping the simulation will retain this set of relations and, through its lineage, the representative character of these relations qualifying it as a model. While modeling with code potentially has a host of functional differences from other types of modeling that could be linked to the same process (e.g. diagrams, written descriptions of process, physical replicas, etc), program code meets all the necessary criteria for being models by acting as representations that bear the appropriate similarity relations. Despite the continued existence of the appropriate model within the data we do not have a simulation on halting, the existence of this model within the data does not constitute a process in anything like the sense in which we demand that simulations are processes, indicating that the two are not the same.

Imagine now that you let the process originally being carried out by the computer resume. Is there now a simulation? From the perspective of the data within the computer it is as if the program was never halted at all so it seems that we should conclude that there is. This situation of stopping and starting a simulation is no different than the situation would be for us if time stopped and then restarted; we would have no idea time had stopped at all because everything in our frame of reference would have been affected along with us (For all we know, this happens regularly!), and the same is true for any possible observations to be made from within the states of the computer. Further, since we were allowed to *start* a program afresh and claim that it the result is a simulation then *resuming* a paused simulation should be uncontroversially considered to be a simulation as well since restarting a paused simulation is technically no different from starting a new simulation that shares the same set of initial states.

The model contained in the data that resides within the computer is unaffected by the halting and restarting procedure that we have been imagining. We know the model remains throughout these interruptions because we could, at least in principle, extract it from the computer at anytime throughout the pause and play process because it is held within actual physical states of the *computer*. It is the resilience of this model to change under these conditions that allows the simulation to be resumed as if it was both started for the first time from that position and never halted at all. It is not merely coincidence that the model exists whether the simulation exists or not, it is indicative of an important relationship between the model and the simulation—the model is necessary for the simulation to exist as it does because it is the model that directs the behaviour of the simulation in the processes that it carries out. The models that stand give simulations their existence and instruction can be seen to be the basis for the simulation in the first place and can thus be referred to as *underlying models* to distinguish them from models that do not have a connection with a simulation.

The disappearance of a simulation on the pausing of the computer's operation, the return of the simulation once the computer is allowed to resume, and the continued *existence* of an underlying model through all this indicates that simulations are not identical with their underlying models. The identity relation, after all, demands that the things to which it can be applied must behave the same way under the same conditions and we do not see this property exhibited by simulations and their underlying models: when we pause the computer we see the simulation cease to exist while the model remains. Simulations may not be identical with the models that they make use of but it remains possible that they are models in their own right, models that subsume the underlying models that they make use of in the way that a Matryoshka doll stacks smaller dolls within larger dolls.[4] We see this property displayed in many cases of actual models—a model ship includes a model rudder, a model of the economy includes a model of financial markets, and a model of star thermodynamics will include a models of various layers of the star—and while it is true that not every model ship needs to have a model rudder, the kind of inclusion being suggested here is so commonplace within models that is almost trivial. If it is the case that simulations are models in their own right then just what kind of models would they be? Even if the pause-and-play argument does not prompt at least the possibility that models and simulations are not the same thing, it does indicate that simulations have a dynamic component—they are active processes that lose some important criteria for their existence when stopped—and so, if we are going to consider that simulations might be models unto themselves, positing that they may be

---

[4] Matryoshka dolls are also known as Russian dolls, nested dolls, and stacking dolls.

dynamic models is a good place to start.

The Flip-Book

In order to consider the possibility that simulations might be dynamic models, we will need a conception of what a dynamic model is and this conception needs to be broad enough to capture any overlaps that might exist between 'model' and 'simulation.' In the interests of brevity, let the working characterization of dynamic models be a model that represents one, or more, processes, i.e. they represent changes between states and the mechanisms that bring about such *changes*. Given this general characterization should simulations be seen as dynamic models? No: while this general casting of what it means to be a dynamic model can be seen as encompassing what composes a simulation at least insofar as simulations also represent change and the mechanisms that bring about such change, the essential nature of simulations extends beyond representations of change. In short, a simulation is a step-by-step process actively driven forward by the logic of an underlying model—the future states in the machine are a direct result of the computer making use of the current states represented within itself under the direction of the relational logic also modeled within these or other states—and this is *not* true of dynamic models in general.

To see that simulations are driven by the logic of their underlying models, consider the difference between a computer simulation that depicts the movements of the nine standard planets within our solar system through a set of images on a computer screen and these same images presented in a child's flip-book. In the case of the flip-book it is clear that there is a representation of the solar *system* and the major bodies within it (we can point to them on each page). It is also clear that by flipping the pages we can represent the movement of theses planets around the sun (each page captures a state successive to the one on the previous page and the difference between these states indicates change). Given the appropriate perspective and temporal coverage/reference we may even be able to predict certain heavenly events of interest, such as conjunctions. In comparison, the simulation seems, at least on a cursory initial consideration, to be no different from the flip-book: by showing the same movement of the planets around the sun, the simulation has the same representational qualities of the flip-book and stands in the same relation to the actual solar system. Such a cursory initial examination does not go deep enough though because it does not look at the process by which each produces the images, but only at the outcome of these processes.

In the case of the flip-book planetary movement is represented by the succession of images being presented to us. The order of this succession was determined by the understanding of how the planets and sun interact (that is to say, by theories they had governing the movements of large astrological objects, the most important idea of which is gravity) held by the creators of the book. It follows that once the book is constructed we have at our finger tips a collection of images representing the solar system that we can flip through in rapid succession to get a sense of how the positions *of* these bodies changes over time. On the other hand, while the simulation provides the same set of images, it is not the case that the computer should be seen as carrying out a process reducible to simply flipping an ordered set of images across the screen (although computers are certainly capable of doing this). Before each image is presented the computer must actually *determine* what the next image will be and make this determination based on information it holds, describing the previous or initial state in conjunction with the logic of the model contained in the program. The flip-book, in contrast, makes no such determinations in and of itself. Yes, its pages are ordered in accord with the same theories that helped shape the models that underlie the simulation and yes there is a sense in which a combination of the ordering of the pages and the hands of the person holding the flip-book contribute to what is being presented to us, but these are outside the flip-book itself. The flip-book is unable to answer the question, “What do I do next?” from within—it requires the intervention and interpretation of something beyond itself, namely us and our fingers—whereas the simulation is moved from state to state without the need of outside guidance.

The existence of state-to-state determination internal to simulation allows us to demarcate them from only *some* dynamic models thought, specifically those ones that do not also have this property. It does not, unfortunately, allow us to demarcate simulations from an interpretation of 'dynamic model' that includes internal state-to-state determination processes. Further, it *seems* unlikely that such a demarcation could ever be made because once internal state-to-state determination has been granted to dynamic models they will share both the outputs and the production mechanisms of any comparable simulation. In short, dynamic models with internal state-to-state determination will be identical to simulations on all observable accounts meaning that simulations are a special case of dynamic models. Before accepting that this somehow closes the need to continue this investigation because simulations clearly are dynamic models, we need to consider that the possibility that simulations just are dynamic models can be taken in two ways. The first (and seemingly sensible) option is to accept that simulations are just special dynamic models and conclude that simulations are legitimate parts of the scientific process because models are legitimate parts of the scientific process. The alternative is to block the assimilation of simulations by models by claiming that those things that considered to be dynamic models with internal-state-to-state determination are not models at all. Despite the initial appeal of the first option, it is the second option that should be taken because, while it is the case that simulations do have characteristics in common with models, the ways in which they are used make a stronger case for considering them to be a special kind of experiment.

Simulations as experiments

Since we are allowing that simulations and those things captured by the phrase "dynamic models with internal state-to-state determination" amount to the same thing, I will take it for granted that what is true for one will be true for the other. With this proviso in place I will show that simulations, and thus many dynamic models, should not be understood as models at all, but rather special kinds of experiments: experiments in our understandings. I begin by arguing that despite some blurring of the division between models and experiments brought about by the models-as-mediators perspective championed by Morgan and Morrison that a division can still be sensibly made. From this division I then examine the purposes for which simulations are used and argue that, although simulations should be understood as an interweaving of models and experiments, they fall most clearly on the experiment side of this contrast.

Traditionally, models and experiments have been taken to be distinct entities, but recently this clear division has been blurred by arguments that models can and regularly do contribute heavily to the shaping and running of experiments in ways that suggest the two are more closely related than is usually thought. The ideal model of a pendulum discussed by both Giere (1988) and Morrison (1999) is a case in point for the overlapping of these concepts. On Morrison's account, the model of the pendulum provides a context wherein the theory of motion set out by Newtonian mechanics can be applied regardless of whether the model is entirely abstract (Morgan and Morrison call such abstract models “notional objects”), entirely physical, or some hybridization between these extremes. By providing a context upon which theory can be examined and refined in this way the model acts as an exemplar that we can test our ideas on to see if they fit with our observations of the world. This testing makes the model pendulum the locus of a kind of experiment in our understandings—we make some predictions and see to what degree they are borne out by the model at which point we then decide whether it is our model or our theory that is in need of adjustment (Morgan and Morisson, 1999, p. 19). While it is true that models have been a part of experiments for a long time, the models as mediators account highlights the importance of this relationship for scientific practice and the multiple ways that models contribute to it.

Despite the intertwining of model and experiment through the mediating role that models can play as instruments of investigation, it remains the case that we can, and should, conceptually separate the models and experiments, that is, we should not think of experiments as trials of models nor of models as experiments in our idealizations. Models, as has been argued above, are *representations* related to other objects, abstract or physical, that constitute our understandings of the world and its operations. In this role models are abstract entities, even though they may be so closely with actual physical objects that we might be tempted to refer to these objects as "models." It seems that models can thus be, at best, apparati with which experiments can be performed. Experiments, on the other hand, are *activities* by which we seek to acquire observational data and in so doing either support or challenge our perspectives on the world. The simplest of all possible experiments are those whereby we observe the world repeatedly and with only as much interference as is necessary to prompt the required situation and make our observations. Through the observations we make of the world we infer to the best of our abilities what the underlying structure of the world is most likely to be, producing a set of understandings about the world that have all the makings of a set of abstract models. We treat these in-head models as more than mere descriptors of past performances, putting them to use as predictors as well. Through the process of predictive successes and failures, roughly amounting to a stream of small experiments, we further refine our conceptual models and our theories.[5] Despite the partnership that models and experiments have in furthering our understandings of the world they remain markedly distinct. With this division in hand, we are in a position to argue that simulations are experiments and not models *if* it can be shown that simulations can be understood as experiments. The evidence for simulations being experiments rests in a comparison of the reasons that simulations and

---

[5] We may well have a genetic disposition to see the world in certain ways. The process that I have suggested here is framed in terms of an individuals learning capacity, but is compatible with natural selection as well.

experiments are employed and noting that they are the same.

When asked why simulations are being used as part of any particular investigation the response will likely make reference to their cost effectiveness and ability to overcome impracticalities and ethical boundaries. These are reasons for choosing simulations *over alternatives*, though, and not reasons for choosing simulations *in the first place*—casting runes is cheap, possible, and, *ceteris paribus,* free from ethical constraints, but scientists do not cast runes because they do not believe it capable of doing the tasks they are involved in. When we ask for reasons why simulations are used, what we are really asking about is what the roles are that simulations play for those who make use of them, of which there seem to be three, as follows:[6]

> *Proof.* The mathematical precision with which computers can complete tedious calculations is well established. With the speed of modern processors we are now able to direct the combinatorial force of computers towards proving conjectures that would take many lifetimes if a human were to attempt it on their own or even, given the historical truth in Moore's Law, with the aid of a computer from a decade ago. A well known example is the use of a computer to prove the four colour theorem—the conjecture that there is at least one colouring scheme for any two dimensional map such that with only four colours are needed to ensure that no adjacent cells are the same colour (Staff, 2005)—through generating and dismissing thousands of possibilities. Such proofs are still controversial though because they run counter to unspoken components of the standard conception of what counts as a proof. Proofs by simulation can be almost impossible to check by hand, possibly requiring even more time than it would take someone to actually complete the proof themselves, and that any would be proof-tester/referee have the requisite level of skill in the programming language used, and computer science and

---

[6] The separation of the roles made here is for clarity and should not be taken to imply that they can easily be disentangled from each other in practice as each of the examples cited could be easily bent to exhibit the other roles.

philosophy in general, to judge the success of the implementation of the ideals. Even granting that the use of computers in the proving of conjectures like the four-colour theorem does not mean that the computer is making the proof via simulation though since it is not clear exactly what system the process is a representation of. This is a fair challenge and one that the paper deals with later, when the question of what counts as a simulation is considered, by noting that simulations, like models, do not need to be representative of actual processes but might also be representative of hypothetical or ideal processes. In this specific case the computer is simulating the process that a mathematician would go through *if* they were willing to devote a lifetime (or two) to following the same procedures carried out by the computer.

Tuning. When the complexity of our theories and models is minimal we are typically able to judge their value by quickly comparing what they say about the world with how the world actually is and thus begin a looping process of refinement. Many of our theories and models are too complex for us to accurately determine their value through a comparison between prediction and results though, typically because the relations they describe are complicated or the number of objects to which they must be applied is large, making it unlikely that we could carry-out the process without error. Simulations take on the burden of such computationally tedious and exacting tasks, telling us what our theories and models imply for many interactions beyond any consequences we could ever hope to foresee on our own. This is particularly true during the construction and the return of initial results of any particular simulation. During this time a scientist will be especially keen to "get it right" by thinking deeply and critically about the model by which they will characterize the situation in question and that will ultimately underly the simulation. It is during this time that Winsberg's (2001) claim that simulations are not simply acts of number crunching but rich inferential processes is most applicable, making us aware of the possible consequences of our pictures of the world and in doing so, suggesting alternatives and refinements.

Inference. Once satisfied with the underlying model, the simulationist is then free to explore the entailments of the underlying model given any set of states (possible or impossible, probable or improbable). It is in this role that simulations act as artificial worlds sufficiently similar to our own such that we can draw useful inferences and parallels to our own that is foremost in the minds of most people when the topic of simulations comes up. Robert Axelrod's (1984) prisoner's dilemma tournament is probably the most widely know computer simulation to provide such inferences—its champion, tit-for-tat, is often cited as evidence for the validity of moral claims such as initially trusting strangers and being forgiving.[7] Other prominent examples include the historical use of simulations in the development of the atomic bomb and the contemporary use of them to make predictions regarding climate change.

Experiments have a long history of helping scientists to prove, tune, and infer. Given this overlap in roles asking whether we should consider simulations as models will be quite similar to asking whether we are willing to consider experiments as models. In our earlier comparison of models and experiments we saw that the character of each is quite different: models *are* our understandings and experiments are tests by which we can gain or refine these understandings. Without going into all the possible variations that collapsing 'experiment' and 'model' together could take, we should be strongly opposed to any such maneuver because it would make it difficult for us to claim the same “arms length” objectivity that experiments are seen as having and which is in turn used to lend credibility to our models (and ultimately our theories, beliefs, etc.). It is this separation that allows scientists to regularly make statements like “We'll have to wait and see what the data says,” “Experimental evidence

---

[7] Tit-for-tat was only victorious in the first two tournaments that Axelrod ran. In a third, anniversary, tournament the victor was a set of programs that “conspired” to elevate one of their members to victory.

confirms...," and "Based on what we now know we have revised our position..." without having to defend themselves from charges of circular reasoning. Indeed, the whole scientific endeavor is premised on the ability of experiments to contribute to the shaping of our models in a way that would be much more open to question if the tests were somehow seen to be the same things as what was being tested.

Since we should not be comfortable treating experiments as models then we should not be comfortable treating simulations as models either—were we to do so we would put in jeopardy the ability of simulations to play the useful roles that they do. This said, we should not simply accept simulations as being on par with what might be called "traditional experiments" either, because while simulations fill the same roles as experiments, they do so through importantly different ways. Traditional experiments are performed in the world, meaning that, as much as possible, the actual entities in question are submitted to investigative processes, and it is the actual entities in question that provide the responses. In contrast, we have seen that simulations can be used to carry out investigations into the nature of the same entities, but via representations of the entities and their surrounding environment—a set of representations that is synonymous with our understandings of the world and how it works. It is on this difference that we may say that traditional experiments are experiments in the world and simulations are experiments in our understandings.

Even though simulations more properly qualify as experiments than models, they do retain close ties to models through their underlying models that traditional experiments lack. In a traditional experiment the world responds to our prodding in the ways it does without any influence from the models that may have inspired our investigations. In a simulation though the underlying model is all important since it takes the place of the structure of the world. In this way, even though simulations are experiments in our understandings they may also be seen as being models of traditional experiments, even if these traditional experiments never existed. It is at this point that reopening the claim that all dynamic models with internal state-to-state determination *are* really just simulations must be returned to for we now have a reason to differentiate between the two. Morgan and Morisson (1999, p.11) have suggested that a model is like a hammer by being built for a specific task (driving nails), but suitable for many tasks (paperweight, measuring stick, digging claw, etc.). Simulations are also like hammers in this way, but with the additional observation that depending on the task which the hammer is being used for it may at times be advantageous to refer to it by a different name: “paper weight,” “measuring stick,” or “digging claw.” When simulations are typically used in science they are used as special experiments and calling them “simulations” can likely be seen as designating this use. Simulations can be used as models of experiments though and when this is the case referring to them as “dynamic models” can be useful in declaring that this is how they are being used. Even with changes in the descriptive label the underlying structure remains the same and so a “simulation” remains a 'simulation' even if we chose to call it a “dynamic model,” just as the “hammer” remains a 'hammer' even if we call it a “digging claw.” With this observation made, we can draw on the common convention of having the name that is commonly used to refer to an object as the one associated with the commonly recognized use of that object to lend one final, albeit weak, piece of support to the claim that simulations are most properly considered experiments: if they were commonly considered models then not only would we expect to see them used as models more often than they are, but we would also expect to see them predominantly referred to with a name that captured their nature as models.

5. Consequences

Accepting that simulations stand apart from models as a type of experiment in rigorously hypothetical realms leads to a number of interesting consequences that are relevant for tracking the legitimacy of simulations. I will discuss four of these that follow quite directly.

Always a model underlying a simulation

As implied by the preceding discussions, every simulation has an underlying model. These underlying models may not be formalized or immediately apparent, but there must be some underlying set of representations and associated logic for a process to possibly exist as a simulation. If there is no underlying model, then there is no simulation and whatever is being claimed to be a simulation is either a) a whirring thing, detached from any legitimate claim to parallel the world, or b) an actuality, in the way that a model ship identical to the original would be no model at all.

The attachment of a simulation to an underlying model is an important point of division between simulations and other computer programs. If only the processes/behaviour of a computer program and a simulation are compared then there will be no discernible difference between them as they will simply be collections of data carrying out operations that shift the insides of a computer between states. This distinction does not rest in process, but in the character that simulations gain from having a connection to an underlying models. This connection is observer dependent and therefore simultaneously (slightly) arbitrary and ethereal, but this does not detract from its importance. Recognizing that things act as models and simulations because of how we see and use them and not because they have some special property inherent in themselves is important in understanding that simulations exist because of the relations they bear to their underlying model.

Simply knowing that every simulation must have an underlying model does not make it the case that answering the question of *what* these models are, especially when the original programming code is unavailable or complicated into obscurity, is going to be an easy, or even possible, task. Still, the task can be made easier by clarifying exactly what the terms "model" and "simulation" are meant to pick out in the question.

There are two reasonable interpretations of "model" that we must distinguish between: the *specific* model that was used to program the simulation which created the results and what I will call the *ideal* model, being the intersection of all the models capable of producing the simulation. By being the smallest possible model that could produce the simulation, ideal models stand for all models that might have produced the data because, to return to the Matryoshka doll analogy, they will be contained in all the models that may have produced the data on simulation. It follows that the ideal model will thus be contained in the specific model on which the simulation in question was run.

We must also be concerned with what simulation we are actually being asked to determine the underlying model for because if we are only concerned with a small set of data then there are many more possible underlying models than if we are concerned with the underlying model based on data that produced by what we might refer to as a *complete* simulation—namely a simulation run through every possibly variance in input data and which consequently has a returned every data set it is capable of returning.[8] Of course, a complete simulation may never be realized due to the infinite nature of possible values that some variables might take and the possibility that some simulations may treat two different inputs as being the same, but the concept remains accessible.

By making the distinction between the specific and ideal models that might underlie a simulation and the distinction between an instance of a simulation and a complete simulation determining what model underlies a given situation is made easier because it can be broken into four separate questions, not all of which may necessarily need to be answered, namely:

---

[8] We might further refine this notion by referring to bounded complete simulations—those complete simulations where a significant portion of the input data has upper and lower limits—and unbounded complete simulations—those complete simulations where a significant portion of the input data does not have upper or lower limits. Exactly what constitutes a "significant portion" is left to the community to determine as a matter of convention, but as a guide I suggest that bounding random inputs not be grounds for calling a complete simulation bounded.

		“Model” refers to:	
		Specific model	Ideal model
“Simulation” refers to:	Instance(s) of a simulation	What specific model underlies these instances of simulation?	What ideal model could underlie these instances of simulation?
	Complete simulation	What specific model underlies this complete simulation?	What ideal model could underlie this complete simulation?

It may prove to be the case that which of these four questions is being asked will be apparent from context, but this will not always be the case. The most useful of these questions is “What *ideal* model underlies this *complete* simulation?” because once we have an answer to this question, we will have all the information necessary to answer the other three questions that might be asked. It is in the interests of clarity that I suggest the question “What model underlies this simulation?” be taken to be asking “What *ideal* model underlies this *complete* simulation?”, unless there is further qualification.

Simulations are as legitimate as their ideal model

The litmus test of a simulation is the degree to which it is successful in paralleling actual events in the world. We have seen that what determines this parallel is the ideal model that underlies the simulation. It follows directly that a simulation is only as legitimate as the ideal model that underlies it and this is true for both any given instance of a simulation and a complete simulation. Note that it may not be possible to determine the legitimacy of a model until it underlies a simulation though due to complexities within the model beyond our capacities. In this way, simulations can be used as microscopes or growth chambers by which we can examine our models in more detail. This is only one component of the constant feedback loop between models, simulations, and observations that makes it problematic to talk about *the* model or *the* simulation, because these are both entities that when used may lead to changes in the other which in turn changes them. The rate of this change is fastest during tuning of a model or the initial construction of the simulation as preliminary data is brought in and the model that the simulation instantiates is brought in-line with the model that the simulation is intended to instantiate. Once the gap between the model intended to be expressed by the simulation and the model actually expressed becomes sufficiently small the focus on changes becomes more about adjusting the intended model and less about making sure that the intended model is being properly expressed if the results of running the simulation data continues to diverge from observational data. In some cases, the simulation may suggest that additional observational data be sought and this again has potential for changing the model and ultimately the simulation. To alleviate the confusion as to exactly what 'simulation' and 'model' are being talked about a distinction should be made wherever possible between the complete simulation and any particular set of instances of the complete simulation.

Identity of Models and Simulations

It may well be of interest at times to know whether two models amount to the same thing and where, for reasons of complexity, it is difficult to determine whether they are identical and, if they are not identical, than in what respects they differ? In such cases, simulations are a possible method of comparison. Because simulations are instantiations of the models that underly them, two simulations will be congruent in their outputs to the degree that the underlying models are congruent. Here is a case in point where a precise terminology introduced earlier is particularly useful. If this question of identity is turned towards two complete simulations then we may say that they are identical if it is the case that the ideal model underlying each one are the same. Checking this is relatively straightforward, assuming that both use the same conventions for presentation of outputs, since a single comparison of outputs may be run to determine if this is the case.[9] In cases where the size of the data sets is an impediment to a timely response an analysis of the underlying code may be performed. Depending on

---

[9] Straightforward does not necessarily mean short or possible. Some complete simulations may have infinite data sets making an output by output comparison intractable. Worse, determining whether two computers behave the same for all inputs is undecidable. Even determining whether a computer rejects all inputs is undecidable.

the length and complexity of the programs under review, this may prove to be even less desirable than a full comparison of outputs. Fortunately, what we are concerned with is not an absolute proof of the identity between two underlying models, but some reasonable suggestion that they are roughly identical which can possibly be had by simulating each for some critical set of inputs and taking into account the degree of divergence in the outputs between them.

What counts as a simulation

So far, the discussion of what constitutes a simulation has been primarily abstract. Where attention has been turned to practical examples each simulation discussed was a closed computer simulation. This was done to keep as much clutter away from the initial explication as possible, but the question of what else should properly be considered a simulation remains. The answer is rather a messy one, not because there are no criteria to which we can appeal to judge whether a process is a simulation or not, but rather that so much of this criteria admits of degrees and that many processes are what might be termed "partial simulations." What is partial about the simulation activity going on within these simulation candidates is that the portion of the activity driven forward by the underlying logic of the model only constitutes a portion of the overall activity taking place, the rest is being filled in by exogenous factors. Consider again the flight simulator that was barred from being a paradigm example of simulation because it was unclear just how it to uncontroversially classify it as a simulation. This union of human being, hardware and software is difficult to classify as a simulation because in this instance the computer does not contain *every* piece of data relevant to the running of the simulation. The complicating factor is of course, the human pilot, who interacts with the hardware and software throughout the candidate simulation process. This human is not a representation, but an actuality. As such, the system composed of pilot and computer should not be considered as a proper simulation. We can still sensibly refer to the flight simulator as being a simulation though insofar as the *between input* *moments* that take place within the computer constitute proper simulations. Some value is also given to using "simulation" to describe the flight simulator because from the perspective of the pilot the *experience* is a simulated one because a plane is not actually being flown. It could be objected here that allowing the status of a process like a simulation to be determined depending on the perspective taken introduces relativism that could threaten to bury us in confusion. As with earlier potential charges of relativism, I believe that such attacks fail because they are not borne out in our observations: in this particular case the same objection can be made against fringe cases of any word, but we are able to navigate these cases as they come and so we should not be troubled by this reality in this instance anymore than we are generally.

More important than using our newly forged conception of a simulation to determine what components of partial simulations are actually simulations and what are questionable is to recognize that on this conception a great deal of what is commonly considered to qualify as "thinking" counts as simulation. To see that this is so, imagine that you have just learned to play chess and are now occupied in a game. In order to decide what to do on your turn, you would ideally consider all the moves you can possibly make, the counter moves that would then be possible from your opponent, your possible counter moves to these counter moves, and so on to some to the limit of our care or mental capacity. This process of consideration requires that we have a representation of the current state of the board in our head as well as whatever we can remember of the rules of play; that is to say that we have a model of the game in our heads. With this model at our mental disposal we proceed to apply the rules to the represented state of the board and in so doing generate a set of moves that might be available to us and which we can then use to determine the and counter moves. The outcome of this process is driven by the relational logic contained within our model of the game and so constitutes a simulation of what would happen were we to make various moves. This capacity for simulation is not confined to toy games, but can be seen in a myriad of our cerebral activities. Wondering what would happen if you won the lottery, deciding what your friend might like for their birthday, and figuring out what is entailed by Newton's laws are all cases of simulation in virtue of the underlying representational relations they contain and their being used to determine what follows from them. The irony of this conclusion must not be missed: if thinking, the very process by which the questions that drive the scientific process as a whole are raised and answered, is quite often really just an instance of simulation, then asking whether simulations are legitimate pieces of the scientific process is rather strange since without simulations this question—and indeed science and philosophy as we typically conceive them—would not be possible in the first place.

Works Cited

Staff writer(s), 2005, “Proof and Beauty”, *The Economist*, http://www.economist.com/science/PrinterFriendly.cfm?Story_ID=3809661

Axelrod, Robert, 1984, *The Evolution of Cooperation*, Basic Books, New York, NY.

Cartwright, Nancy, 1983, *How the Laws of Physics Lie*, Oxford University Press, New York, NY.

Fernández, J., 2003, “Explanation by Computer Simulation in Cognitive Science”, *Minds and Machines*, 13:269-284.

Fetzer, James H., 1999, “The Role Of Models In Computer Science”, Monist, 82:20-37.

Fetzer, James H., 1999, “Mental Models: Reasoning without Rules”, *Minds and Machines*, 9:119-125.

Giere, R. N., 1988, *Explaining Science: A Cognitive Approach*, University of Chicago Press, Chicago, IL.

Hesse, Mary B., 1966, *Models and Analogies in Science*, University of Notre Dame Press, Notre Dame, IN.

Johnson-Laird, P. N., and Byrne, R. M. J., 1999, “Models Rule, OK? A Reply to Fetzer”, *Minds and Machines*, 9: 111-117.

Morgan, M. S., and Morrison, M. (eds.), 1999, *Models as Mediators: Perspectives on Natural and Social Sciences*, Cambridge University Press, Cambridge, MA.

Morton, A, and Suárez, M., 2001, “Kinds of Models”, *Model Validation: Perspectives in Hydrological Science*, M. G. Anderson and P. D. Bates (eds.), John Wiley and Sons, Ltd., New York, NY.

Morton, A., 2003, *The Importance of Being Understood*, Routledge, New York.

Sawyer, R. K., 2004, “Social Explanation and Computer Simulation”, *Philosophical Explorations*, 7:219-231.

Skyrms, B., 2004, *The Stag Hunt and the Evolution of the Social Contract*, Cambridge University Press, New York, NY.

Suppe, F., 1974, *The Structure of Scientific Theories*, University of Illinois Press, Chicago, IL.

Wilson, R., 2004, *Boundaries of the Mind: The Individual and the Fragile Sciences: Cognition*, Cambridge University Press, New York, NY.

Winsberg, E., 2001, “Simulations, Models, and Theories: Complex Physical Systems and Their Representations”, *Philosophy of Science*, 68:S442-454.