

SANCTIONING MODELS: THE EPISTEMOLOGY OF SIMULATION[†]

Eric Winsberg
Department of History and Philosophy of Science
Indiana University
Bloomington IN, 47405, U.S.A.

The Argument: In its reconstruction of scientific practice, philosophy of science has traditionally placed scientific theories in a central role, and has reduced the problem of mediating between theories and the world to formal considerations. Many applications of scientific theories, however, involve complex mathematical models whose constitutive equations are analytically unsolvable. The study of these applications often consists in developing representations of the underlying physics on a computer, and using the techniques of computer simulation in order to learn about the behavior of these systems. In many instances, these computer simulations are not simple number-crunching techniques. They involve a complex chain of inferences that serve to transform theoretical structures into specific concrete knowledge of physical systems. In this paper I argue that this process of transformation has its own epistemology. I also argue that this kind of epistemology is unfamiliar to most philosophy of science, which has traditionally concerned itself with the justification of theories, not in their application. Finally, I urge that the nature of this epistemology suggests that the end results of some simulations do not bear a simple, straightforward relation to the theories from which they stem.

§0. Introduction

Although computers have come to play an increasingly large role in scientific research, a detailed study of their significance for the philosophy of science has yet to emerge. I refer here not to the role of the computer in the study of mind, for which there is a vast philosophical literature, but to the more workaday role of the computer in helping to manage mathematically unsolvable sets of equations. Specifically, this paper will address the scientific practice of computer simulation in the study of complex physical systems. In many instances, these computer simulations are not simple number-crunching techniques. They involve a complex chain of inferences that serve to transform theoretical structures into specific concrete knowledge of physical systems.

In this paper I argue that this process of transformation is also a process of knowledge creation, and that it has its own unique epistemology. It is an epistemology that is unfamiliar to most philosophy of science, which has traditionally concerned itself with the justification of theories, not in their application. I also argue that the complex and motley nature of this epistemology suggests that the end results of simulations often do not bear a simple, straightforward relation to the theories from which they stem. Accordingly, I urge philosophers of science to examine more carefully the process of 'theory articulation', the process by which a general theory is made to conform to a particular application. It is a relatively neglected aspect of scientific practice, but it plays a role that is often as crucial, as complex, and as creative as theorizing and experimenting. Indeed, my conclusion will be that we now need a new philosophy of simulation to complement recent work on the philosophy of experiment.

Why, should simulation — a form of calculation — require an epistemology? We need an epistemology of simulation because simulation modeling is a set of scientific techniques that produces *results*. When science produces results, we would like to have standards for deciding whether or not these results have some degree of reliability. Even though simulation is

fundamentally about replacing analytical solutions with calculation, which at first sight appears to be merely a mathematical transformation, the question of the reliability of the results of simulation modeling goes beyond simple concerns about the reliability of the calculation, and reaches out to the entire simulation process and the conclusions scientists reach using it.

The first step in getting a handle on what an epistemology of simulation might be is to highlight and characterize the different inferential steps that take place during the process of simulation; those that might be subject to epistemic scrutiny. So I will try to make the point that a simulation study embodies a rich inferential process by outlining the essential steps that are involved in the study of complex phenomena using computational techniques. Along the way, though, I want to fulfill a second related task.

The second task is to get clear on the different uses of the word “model”, as they will come up in this discussion. Regrettably, the term “model” is used in far too many ways in both scientific and philosophical parlance. According to Nelson Goodman,

Few terms are used in popular and scientific discourse more promiscuously than 'model'. A model is something to be admired and emulated, a pattern, a case in point, a type, a prototype, a specimen, a mock-up, a mathematical description-- almost anything from a naked blonde to a quadratic equation -- and may bear to what it models almost any relation of symbolization. (Goodman, 1968, p. 171).

In a discussion of simulation modeling, the situation is particularly ambiguous. There are at least five different uses of the notion of model in the context of simulation. All of them are important, and all of them are, I believe, usefully thought of as types of models. So, as I outline the successive steps of a simulation study, I want to develop a small taxonomy, and a lexicon, for the different notions of model that I will be using in my discussion of simulation. The integral part of the practice of simulation is the process of building what I call a hierarchy of these models. This hierarchy includes a mechanical model, a dynamic model, *ad hoc* models, a computational model, and finally, a model of the phenomena.

Each of these steps in the simulation modeling process is a step on the way toward an

inference. Figure 1 illustrates the different layers of models involved in simulation and the resources used in each inferential step. Via these steps, the simulationist hopes to infer, from *existing* theoretical knowledge, *new* knowledge about the system being simulated. While it is typical in philosophy of science to talk about deducing results from theories, the inferential moves described above are patently *not* deductive. They have neither the inevitability nor the epistemic certainty associated with deduction.

If one of the central tasks of a philosophy of science is applied epistemology, then any philosophical look at simulation modeling needs to ask when, and why, these nondeductive inferential steps are acceptable—that is, under what sorts of conditions do they produce reliable knowledge, and why. But another task of philosophy of science is interpretation. So we will also need to ask what the nature of the inferred knowledge is, and in what relation it stands to the theoretical structure from which it was inferred.

The structure of this paper will be as follows. In section one I identify and describe roughly the kind of simulation research on which I will be focussing. In section two I will outline the methodological structure of this kind of simulation research. I will do this with the help of a case study involving the simulation of a severe super-cell thunderstorm. I will also introduce, in this section, a ‘hierarchy of models’, which presents the various kinds of models that play important roles as intermediate steps in the process of simulation. Section three will focus on the final model of the hierarchy, what I call the “model of the phenomena,” which represents the end result of simulation research. In sections four and five I argue for the unique character of the epistemology of simulation, and I outline how this epistemology fits (or fails to fit) into the philosophical literature. In section six I explore the relationship that models of phenomena bear to theoretical structures, arguing that the relationship is more complex than one might suppose. Finally, in section seven, I offer some concluding remarks.

§1. Simulation Techniques.

The use of computational techniques in the sciences has become more and more widespread, and the range of techniques and applications is enormous. I will focus on one type of application of these techniques—the practice of modeling very complex physical phenomena for which there already exist good, well-understood theories of the processes underlying the phenomena in question. Though the underlying theories are well understood, the phenomena themselves are not well understood because of the complexity of interactions involved in generating the phenomena. These computational techniques, involving the solution of intractable differential equations, via simulation modeling, have as their aim an understanding of complex phenomena they model. In contrast, my discussion will not necessarily be relevant to simulations that do not draw on a base of accepted theory, such as traffic pattern simulations, or the simulation of a flock of birds in flight.

Let me begin by giving a rough sketch of these kinds of techniques. As an illustration, let us suppose that we are confronted with a physical system of which we would like to gain a better understanding: a severe storm, a gas jet, or the turbulent flow of water in a basin¹. The system in question is made up of underlying components such as solid particles or parcels of fluid that behave according to a strict set of physical laws, and we can assume that we know what these components are and that we know the laws that govern them.

The assumptions we have made so far allow us write down a set of partial differential equations. These differential equations represent an exact determination of how the system will evolve through time, as given by the physical model. In the types of systems that the simulation modeler is concerned with, though, the equations are non-linear, and it is mathematically impossible to find an analytic solution to these equations – the model is said to be nonintegrable. That is, it is impossible to write down closed form equations, equations given in terms of known mathematical functions, which represent an exact solution to the set of differential equations and will thereby tell us what the system will do over time.

This problem is not new to the computer age and nor are attempts to it. In the past, attempts focused on analytic techniques for finding approximate solutions to the differential equations in question. These techniques have succeeded in generating closed form functions that are approximately valid. That is, for problem situations such as the three-body problem, functions can be found which can be shown to have the same qualitative character as the unknown solution to the equations to be solved.

But there are vast regions of possible solutions to interesting equations that are qualitatively different from any known closed form function. The approach that the simulationist takes to this problem is to discretize the equations and "solve" them by brute force.

Discretization turns differential equations, which relate continuous rates of change over infinitesimal intervals, into difference equations, which relate rates of change over finite, or discrete, intervals. The values that these difference equations give can then be calculated by a digital computer, from one discrete moment of time to the next. This technique of simulation is often called "finite differencing."

Of course the transformation of the differential equations into difference equations constitutes an approximation. But by choosing an appropriately "fine grid," that is by using discrete intervals of space and time that are sufficiently small, the simulationist can reduce the "damage" done by the approximation as much as he or she wants. In principle, unlike for analytic techniques, no assumptions such as symmetry or time independence need necessarily be imposed.

In practice, though, the amount of time and memory required to do these computations goes up very quickly as the simulationist chooses smaller grids. Frequently, the problem posed requires, for reasonably accurate solutions, a grid too small for any reasonable allocation of computer time and memory. If the simulationist uses the full set of laws in the model, and tries to solve the resulting difference equations with a grid fine enough to assure a reasonable degree of accuracy, he or she will often run up against the limits of the available computational power. In

such cases, the equations are said not only to be analytically unsolvable, but computationally intractable as well.

§2. Models, Models, Models.

The solution is to make modeling assumptions. The idea here is to develop computer algorithms that embody some simplified version of the original set of equations. Depending on what aspect of the solution the simulationist is interested in resolving, it is often advantageous to trade away theoretical rigor in the equations in favor of a finer grid. The deciding factor is not which approach is most true to theory but which approach will produce a solution-set that best resolves the features of the system important for understanding it.

For an adequate discussion of this point, we need to understand more about the simulation process, how complex it is, and the number of diverse factors involved. The heart of the practice of simulation is the construction of a hierarchy of models, (illustrated in figure 1) . The first step in any simulation study is to identify the theory under whose domain the phenomena of interest lie. This will form the basis for the simulation. For example, we might start with mechanical laws, force laws, or equations of state. (Kaufmann and Smarr, 1993)

i) Mechanical Models

By itself, the theory tells us very little about anything but the most idealized systems. To apply them to real world systems requires a *mechanical model*. A mechanical model is a bare bones characterization of a physical system that allows us to use the theoretical structure to assign a family of equations to the system. When we characterize a system *as being like a damped, harmonic oscillator*, we have assigned a mechanical model to the system. The *locus classicus* of philosophical discussion of these kinds of models is Nancy Cartwright's simulacrum account of models in *How the Laws of Physics Lie*. (Cartwright, 1983).

ii) Dynamical Models

Even though a mechanical model provides a foothold for the application of a theory to a

set of real world problems, a mechanical model by itself is still a very general entity, for it is not about any particular system. The next step for a simulationist is to specify a class of parameters, boundary values, and initial conditions that restrict the theoretical model to a specific class of phenomena. This conjunction of the theory with parameters, boundary values, and initial data make a concrete *dynamical model* (or really, a family of dynamical models) for a highly specific class of phenomena. The specification of these values for a simulation is rarely a straightforward process, and is often a delicate balancing act between accuracy and tractability (Smarr, 1985).

In order to lend some substance to these remarks, I will consider a project led by Robert Wilhelmson: a numerical (finite difference) simulation of a severe thunderstorm. The purpose of the simulation, as stated by Wilhelmson, is to provide "improved understanding of severe storm structure and evolution" (Wilhelmson, 1990, p. 20). The simulation generated a four-hour period of solution space for a system of nine partial differential equations. These equations describe the time evolution of the dependent variables of the model. The discrete data which comprise this solution were then subjected to a variety of techniques of data visualization in order to resolve the water and ice structure inside of a storm, to be able to see how air moves and rotates in and around a storm, and to discern various physical processes that influence storm rotation near the ground.

The dynamical model of Wilhelmson's simulation is a system of nine partial differential equations that govern temperature, pressure, three components of velocity, water vapor content, cloud water content, rain water content, and "sub-grid-scale" kinetic energy. For initial conditions, the researchers used observed conditions from one vertical column of air in an actual pre-storm environment. The model was then initialized using horizontally homogeneous values for each of the nine variables of the simulations. A storm, however, will not grow under such homogeneous conditions, so the researchers initiated the storm by introducing a small temperature perturbation at the horizontal center of the storm.

iii) Computational Models

Next comes the construction of a *computational model*. Some dynamical models are analytically tractable; their differential equations can be solved and a mathematical function can be given which provides a good representation of the dynamics of the system. But in systems of interest to simulationist, this approach is not possible. The dynamical model needs to be transformed into a computational model so that computational techniques can be used to overcome the problem of analytical intractability.

There are two steps to this process. First, the continuous differential equations of the dynamical model need to be converted into discrete algebraic equations for which solutions can be cranked out by the computer. Even though this solves the problem of analytical intractability, the new model must also be *computationally* tractable. Simulationists use *ad hoc* modeling assumptions to help make their computational models more tractable and manageable.

iv) *Ad Hoc* Models

Ad hoc modeling includes such techniques as simplifying assumptions, removal of degrees of freedom, and even substitution of simpler empirical relationships for more complex, but also more theoretically-founded laws. This model making can be eliminative or creative. The modeling can involve eliminating considerations from the dynamical model, or making up new ones. Sometimes the simulationist will ignore important factors or influences from their computational models because limitations of computational power make their inclusion impractical. This is what I refer to as eliminative *ad hoc* modeling. In this case, the simulationist has one of two options: either to determine that the effects of this neglected factor are negligible or to make use of some sort of empirical "fudge factor"—creative *ad hoc* modeling—to make up for the absence of the neglected factor.

Often, the question of whether some particular aspect of a system under study is crucial to the system's dynamics is not even the issue. There are times when the simulationist is acutely

aware of the important influence of one component of the dynamics and yet it is simply impractical to include it in a full-blown simulation. It is in this sort of situation that the simulationist will resort to what is often described in the scientific literature as "using modeling." For the simulationist, "using modeling" in a simulation, means to use some sort of rough and ready, *ad hoc* model inside of the context of the computational model itself. This way of talking may seem a little strange, given that what we would normally associate *any* attempt to apply theoretical equations to some concrete physical situation with the term "modeling." That is why I prefer to use the term "*ad hoc* modeling." The term "*ad hoc*" serves to distinguish this activity from all the other roles models play in simulation as well as to emphasize the fact that the construction of the model relies on insight gained from *outside* of the context of our best *theoretical* understanding of the phenomenon.

Creative *ad hoc* models typically involve relatively simple mathematical relationships designed to approximately capture some physical effect in nature. When "coupled" to the more theoretical equations of a simulation, they allow the simulation to produce results that are more realistic than they could have been without *some* consideration of that physical effect. For example, here are Wilhelmson's own words describing his simulation:

A very simple model is used to account for the development of rain. In many studies such simple models are sufficient for studying storm dynamics. Although simple, they provide the key storm-driving forces, namely, warming due to the release of latent heat as water vapor condenses and cooling due to evaporation of cloud and rain drops in unsaturated regions. (Wilhelmson, 1990, p. 22)

§3. The Model of the Phenomena

Once the computational model is implemented on a computer in the form of a particular algorithm, the algorithm produces results in the form of a data set, often a very large one. This data set requires interpretation. For this, the data can be visualized, subjected to mathematical analysis, and used in conjunction with other sources of knowledge, including observation, in

order to arrive at the final goal of a simulation study; what I call a *model of the phenomena*.

A model of the phenomena is a manifold representation that embodies the relevant knowledge, gathered from all relevant sources, about the phenomena. It can consist of mathematical relations and laws, images (both moving and still), and textual descriptions. The construction of the model of the phenomena is an attempt to summarize the basic robust qualitative features of a whole class of structurally similar phenomena. It might include such features as:

- An emergent high-level mathematical relationship among certain aspects of the system, such as a scaling law.
- A transport mechanism: any effect, such as diffusion, turbulence, an instability, or viscosity, which explains the movement of some entity or quantity such as mass, energy, or angular momentum, through a particular system.
- Threshold values of parameters; for example a Reynolds number at which a system undergoes the transition from soft to hard turbulence.
- Characteristic coherent structures (like the red spot of Jupiter).
- Characteristic geometries of flow.
- Patterns of interaction and competition among coherent structures.

In Wilhelmson's simulation, the solution space data set is composed of values for each of the nine dependent parameters at each of the points on the space time grid of the simulation. This data set was then subjected to a variety of complicated and labor-intensive visualization techniques designed to "reveal the inner dynamics" (Wilhelmson, 1990, p. 20) of the phenomena. The ultimate goal was to produce a visual record of how the basic internal stable structural features of the storm evolve, and to understand the internal mechanisms that are at work in creating and preserving the stability of these structures.

In this case, the researchers generated images corresponding to naked eye observations of the simulated storm as well as images corresponding to those generated by surface reflectivity radar. The visual viewpoint was generated by rendering images of the surfaces that enclose regions of cloud (small water droplets and ice particles) and regions of rain within the storm. This created a time series of images that depict what the storm cloud would look like to the naked

eye from some particular vantage-point. Images traditionally generated from reflectivity radar are two-dimensional cross sections that are color-coded to graph the concentration of raindrops at every point in the cross section of the storm. The simulationists recreated these images from the data set generated by the model.

Next, imaging techniques were used to study the patterns and mechanisms of airflow inside the storm. Wilhelmson's team used the computed velocity field data in order to simulate the trajectories of imaginary "weightless tracer particles" through the storm environment. The team also used long streamers to display the trajectories of selected air particles inside the storm. These streamers allowed the researchers to view the major stable, long-lived, air currents. Another important aspect of the flow is the vorticity. In particular, researchers are especially interested in depicting the patterns of streamwise vorticity, the rotation of air around an axis parallel to the direction of flow. For this purpose they used differently colored ribbons whose degree of twist is in proportion to the quantity of streamwise vorticity in that region of the flow line. All of these visualizations were preserved as both still images and full motion video.

Once the researchers succeeded in visualizing these aspects of the flow, they were able to make use of these visual representations to identify some of the key structures and trajectories in the inner dynamics of the flow of air, ice, and water through the storm system. They were able to use this knowledge to construct "a model of storm evolution and persistence"—a model of the phenomena for severe storms. Meteorologists researching storm dynamics are particularly interested in the question of how severe storms maintain their longevity and develop and maintain their rotational character. The researchers seek an answer to these questions by analyzing how the basic geometry of the main flow features works to create the features of the storm which are known to be important in preserving its basic structure.

A good example of this kind of explanation involves the updrafts in the storm and vorticity of this flow. Wilhelmson and his colleagues have shown that an updraft with a high

degree of streamwise vorticity will become helical in character, and have argued that this type of flow is essential for reducing the energy dissipation in a severe storm, thus prolonging its life. Because of the importance of streamwise vorticity, Wilhelmson and his team have used the visualization techniques mentioned above in order to do "an in-depth analysis of the processes that govern the development of the vorticity." The researchers identified four processes in the storm that contribute to vertical vorticity: advection (horizontal transport of air due to temperature variation), convergence of air, tilting of horizontal vortex lines into the vertical, and dissipation.

This example illustrates well the type of model I am attempting to characterize. That is, what we have here is an attempt to understand a crucial feature of the storm's dynamics—streamwise vorticity in the updraft—by uncovering the geometrical patterns of the dynamical processes in terms of large-scale durable structures.

§4 Epistemology.

In trying to apply the resources of philosophy of science to an epistemology of simulation, there are two obstacles that we face. The first is that there is very little philosophical discussion of epistemological issues in the practice of building novel applications for existing theories. The typical focus of philosophy of science is on the most theoretical aspects of science, and much less attention is paid to the subsequent *application* of this theoretical knowledge in the modeling of higher level, more complex phenomena.

The second obstacle hearkens back to the old philosophy of science tradition of the "layer-cake" image of science. The layer cake provides a picture wherein science operates by gathering the most basic low-level facts and then uses these to build up higher-level generalizations, then laws, and finally theories. The layer cake is foundational. Epistemology, for the philosophy of science, is about deciding when these moves *up* the layer cake, are justified.

While the more modern, hypothetico-deductive, account of epistemology remains agnostic about where theoretical claims come from, it retains a fundamental feature of the layer-

cake model. It retains the idea that the epistemology of scientific knowledge is fundamentally about justifying the top part of the layer cake by appealing to the low-level facts of the bottom of the cake. Our epistemological concerns, however, focus themselves in the opposite direction. We are concerned with the autonomous sanctioning of conclusions that we draw *from* our scientific theories. But despite the lack of philosophical literature on applied science, we do have one fortuitous inroad.

One of the most influential accounts of epistemology in the philosophy of science is the hypothetico-deductive model, and any discussion of the hypothetico-deductive model requires a discussion of the deduction of empirical consequences from theoretical statements. Even Carl Hempel realized early on (in a preliminary way) that the 'deductive' half of the hypothetico-deductive model is not as deductive as the name implies. (Hempel, 1966, p.31)

But Hempel's concerns steer clear of such issues as approximations, simplifications, and other forms of transformations designed to tease new applications out of theories. His approach foreshadows a persistent limitation of the philosophical literature regarding the derivation of empirical conclusions from theoretical structures. The literature follows Hempel in focusing on the problem of *testing* theoretical claims rather than on the *extension* of theoretical knowledge into new areas—constructing new knowledge about a system assumed to be under the domain of the theory.

Probably the first philosopher of science to emphasize the importance of *theory articulation* as a creative process was Thomas Kuhn. In "The Function of Measurement in Modern Physical Science," Kuhn illustrates the point that much of "normal science" consists of theory articulation. He notes that most novel theories, including Newton's theories of motion and gravitation, are capable, upon their first presentation, of few novel predictions.

The new order provided by a revolutionary new theory in the natural sciences is always overwhelmingly a *potential order*.

Much work and skill, together with occasional genius, are required to make it *actual*. (Kuhn, 1977, p. 188)

"Complex mathematical manipulation," together with "essential approximations," is often necessary in order to get theories to yield experimental predictions. (Kuhn, 1977, p. 190) Kuhn also argues that none of this ought to be construed as attempts at confirmation, because he points out, failure does not count as disconfirmation. Accordingly, a failure in constructing a model of the storm does not count against whatever theory or theories we use to model such phenomena.

Nevertheless, Kuhn resists going so far as to say that theory articulation is productive of new knowledge. " His [the scientist's] success [at opening up new areas of application for a theory] lies only in the explicit demonstration of a *previously implicit* agreement between theory and the world. No novelty in nature has been revealed." (Kuhn, 1977, p.192)

Kuhn's insights are helpful, yet somewhat insufficient. The concept of theory articulation, the idea that bringing theory into contact with the world is often a nontrivial and nondeductive process, is very helpful in understanding what goes on in complex situations involving simulation. His assumption, however, that "no novelty in nature is revealed" misses the mark. Simulation modeling is clearly a case of theory articulation in the spirit of Kuhn. It is a nontrivial process of bringing a theoretical structure into resonance with some phenomena that is 'implicitly' in its domain. But simulation modeling, when successful, does reveal novel aspects of nature. Often simulation will enable us to produce a representation of an aspect of nature that is extremely difficult to observe. Even if the system in question can be observed in detail, often the simulation will bring a level of mathematical order where before there was only seemingly random detail. In Kuhn's own words, there is *actual order* where before there was only *potential order*.

This blind spot in Kuhn, the reluctance to see theory articulation as a form of knowledge production, is not a minor one; if a process provides no new knowledge, then it requires no epistemology. Because of the blind spot, it is natural for Kuhn not even to express any concerns

about justification. He talks about mathematical manipulation, approximation, and idealization, but there is not one worry about when or why or how these steps might be *justified*.

This is a persistent problem in the philosophy of science literature on theory articulation. Most philosophers of science assume that approximations, idealizations, and other such transformations serve only to allow us to compare theoretical predictions with observed results, usually in order to test the theory. As such, they have not been concerned with the autonomous assessment of the reliability of these inductive steps; autonomous in the sense of being based on the means by which the steps were carried out, rather than merely on how successfully the outcome matches the world.

Consider, for example, the work of the philosopher of science Ronald Laymon. Laymon has written extensively on the use of approximations and idealizations in deriving conclusions from theories, including one paper that specifically addresses their role in computer simulation. Laymon's work in this area, however, has focused primarily on epistemological issues relating to *theory confirmation* under the influence of approximations and idealizations. If a theoretical structure requires the application of approximations in order to make predictions about the world, then Laymon worries about what impact the (then only approximate) success of the prediction has on the relative confirmation of the underlying theory. (Laymon, 1985, 1990)

For our purposes, the important thing to note about Laymon's work is that for him, an epistemology of approximation is ultimately about the truth of the theory, not the justification of the approximation itself. Laymon's epistemology is an epistemology of theories and not really an epistemology of approximations and idealizations. It is simply worth noting here another example of a philosopher of science who has viewed the epistemological importance of approximations as relating exclusively to the justification of theoretical knowledge, and not to the results of the approximations themselves.

Jeffrey Ramsey has made some telling criticisms of prevailing philosophical accounts of

approximation. In "An expanded epistemology for approximations,"(Ramsey, 1992) Ramsey criticizes Laymon, as well as Cartwright (1983,1989), Giere (1988), and other philosophers for this persistent problem in their treatments of approximation. (Ramsey, 1992, p.156) First, he argues that all of these philosophers have a static conception of approximation, in which an approximation is a static relation between two structures: the theoretical structure on the one hand and the real-world empirical structure on the other. Ramsey wants to emphasize the dynamic nature of approximation as an act, instead of looking at approximations as fixed things. In other words, he wants to focus his epistemological lens on the context-dependent act of approximating, and not on the relation between the end product of the act and the world.

Ramsey's second criticism is that, because of their focus on approximations as static relations between two structures, traditional philosophical accounts of approximation and idealization can all be seen as "comparison" accounts. Ramsey's argument is that all prior accounts take as the sole criteria of epistemic adequacy the degree to which the *end result* of the approximation *resembles* some empirical structure.

When we look specifically at the techniques employed in simulation, we find this point to be especially on the mark. It is precisely on this point that an epistemology of simulation modeling must focus: What *are* the factors that contribute to the notion that a computational model constructed on a set of approximations and idealizations is valid? Here, Ramsey's analysis of approximations is important for us. The realization that techniques of approximation are in fact inferences that require an epistemology that looks at the means of the methods, and not just their ends, is exactly the jumping off point that we need.

The transformations involved in simulation cannot be judged solely by comparing their results to the world. Since simulations are used to generate representations of systems for which data are conspicuously sparse, the transformations they make use of need to be justified internally; that is, based on their own internal form, and not solely on the basis of what they

produce.

§5 Simulation and Experiment.

I should make it clear from the outset that I believe it unlikely that there is a logic to the epistemology of simulation modeling. What I propose to offer is something that is inspired by what Allen Franklin offered in "The Epistemology of Experiment," a chapter of his book, *The Neglect of Experiment*. (Franklin, 1986) Franklin asks: "How do we come to believe rationally in an experimental result?" His answer is that there are various strategies used by scientists which philosophers ought to be able to see as providing grounds for rational belief in experimental results. While Franklin provides a list of twelve or so such strategies, he is careful to note that this list is neither exclusive nor exhaustive, and that no subset of the list is a necessary or sufficient condition for rational belief. He also notes that grounds for rational belief do not guarantee certain knowledge. Sometimes we may rationally believe something which we all later regard as wrong. (Franklin, 1986, p.190-91)

The epistemology of simulation is an analogous project. It is the study of the means by which we sanction belief in the outcome of simulation studies, despite their motley methodology. I will argue that in order really to understand the relationship between models of phenomena and scientific theories, we need to understand the processes by which these results get sanctioned. One of the benefits of gaining some insight into the epistemological foundations of the simulation practice is that it facilitates a better understanding of this relationship.

One of the central themes of Franklin's work is that experimenters are constantly preoccupied with scrutinizing experimental setups to uncover possible sources of artifact. Then they can work to eliminate the impact of these disturbances on experimental results. The same is certainly true of simulationists. Naturally, the complexity of this task is proportional to the complexity of the methodology of simulation. This methodology includes, among its components:

- A calculational structure for the theory.
- Techniques of mathematical transformation.
- A choice of parameters, initial conditions, and boundary conditions.
- Reduction of degrees of freedom.
- Ad hoc models.
- A computer and a computer algorithm.
- A graphics system.
- An interpretation of numerical and graphical output coupled with an assessment of their reliability.

A thorough epistemology of simulation requires a detailed analysis of the role of each of these components and of how a skilled simulationist can manage each of their potential contributions as sources of error. Here, I will have to be satisfied to note a few crucial features of the process of sanctioning.

To begin with, there is a diverse range of elements that need to be subjected to epistemological scrutiny. Each of the elements listed above mediates between our theoretical models and our simulation results. If the influences and possible pitfalls of each element are not properly understood and managed by the simulationist, then they represent a potential threat to the credibility of simulation results. Understanding and managing these diverse factors requires reliance upon an equally diverse range of sources of knowledge and skills. Much of this knowledge is not contained in the theoretical structure that formed the original basis for the simulation.

A simple example for comparison will help to this point. In celestial mechanics, there exist no general closed form solutions for the orbits of three massive gravitational bodies. There do, however, exist certain restricted solutions. Lagrange, for example, was able to write down a solution for two large masses orbiting around each other, orbited in turn by a third body of infinitesimal mass. (Moulton, 1970, 277-321) Using techniques of qualitative analysis, such as linear perturbation techniques, it is possible to study certain dynamical models, in this case certain celestial systems, for which there is no closed form solution, by studying perturbations of systems for which there do exist such solutions. These techniques study the orbits “nearby” to

orbits that have closed form solutions. This allows us to write down a mathematical function that bears a mathematically provable degree of asymptotic similarity to the actual solution of our dynamical model, even though that model itself has no closed form solution.

Notice, though, that in the celestial mechanics example the resources that we need to establish the validity of our results come from within the theoretical structure of celestial mechanics proper. In contrast, as the issues concerning the tractability of a dynamical model grow more complex, and as qualitative analysis techniques begin to fail, more and more resources get drawn into the attempt get a handle on the empirical consequences of the dynamical model. As this happens, we begin to depend on more and more outside resources in order to warrant the validity of our results. With all of its mediating elements, simulation techniques represent, in a sense, the extreme case.

For example, extra-theoretical knowledge of storm dynamics is needed to sanction Wilhemson's confidence in his claim that "a simple rain model" is capable of providing the "key storm-driving forces." We could go through each of the different mediating transformation of the storm simulation, and talk about what sources of knowledge are required in order to rationally justify them, and we would find that most of them lay outside of any of our general theories of fluid mechanics and thermodynamics. The general point is that there are a large number of mediating influences in simulation that need to be subjected to epistemological scrutiny, and that many of the resources needed to carry out this scrutiny cannot be located inside the original theoretical structure. In the case of *ad hoc* models, often these derive their support from empirical findings not incorporated into fundamental theory.

It is also important to note the pivotal role of the 'observer' in the sanctioning of simulation results. On the one hand, simulations are often performed to learn about systems for which data are sparse. As such, comparison with real data can never be the autonomous criterion by which simulation results can be judged. Nevertheless, calibration does play a very important

role. Simulation models can be calibrated in three different ways: by comparing their results to experiments, to analysis, and to other simulations. The first criterion that a simulation must meet is to be able to reproduce known analytical results. Even for complex systems, there often exist, under highly constrained conditions, limited analytical results for the full equations of a mechanical model. Typically, these results apply to highly symmetrical, equilibrium-state instances of the system, or from instances that can be studied as small deviation, or perturbations, of such instances.

Results can also be compared to the output of another simulation in if it uses a different algorithm, or even better, if it is of a small local region within the broader system and makes use of a more complete dynamical model.

Simulation results are also calibrated against experimental results. Unfortunately, this kind of comparison is often not as easy as it might seem, since data from these different sources may come in different forms. For example, simulation data and experimental data are not always obtained at the same spatial position within a system; the grid points of the simulation do not often correspond to the location of probes within the experimental setup. Moreover, it is typically less interesting to compare detailed data— piece by piece—than it is to compare the characteristic features of the simulation results and empirical results, especially when the empirical results come in the form of flow visualization. (Shirayama and Kuwahara, 1990, p.67)

Most of all, the observer is crucial because, absent an observer who can compare images against images, there is no metric of similarity between the different data sets that need to be compared. Visualization is by far the most effective means of identifying characteristic features out of complex dynamical data sets, and so it is the most, if not only, effective means of judging the degree of calibration a simulation enjoys with other data sets and with analytic results. Thus, visualization plays a crucial role in sanctioning as well as in analyzing simulation results. Not only does the epistemology of simulation call upon resources that are empirical, and that come

from outside of the theory, it also calls upon the faculties of the observer.

§6 What are Models?

Once we understand that simulation has its own epistemology, and that that epistemology necessitates the scrutiny of all of the elements I have listed, we are forced to take a fresh look at what we take to be the nature of the relationship between scientific theories and their empirical consequences.

In the philosophy of science, there are basically two competing accounts of the relationship between theories and their real-world consequences. One comes from the syntactic, or "received," view of theories and the other from the semantic view. According to the 'syntactic' view, specifying the syntax, axioms, and correspondence rules of a formal theory specifies a scientific theory.² The empirical consequences of a theory are exactly those sentences that are logically derivable from the formal theory with the help of the appropriate correspondence rules. According to the 'semantic' view, a theory is specified by specifying the semantics; that is, by specifying a family of models, in the set-theoretic sense. A class of phenomena is said to be an empirical consequence of the theory if it is 'isomorphic' to a sub-structure of one member of the family of models. (Suppes, 1962, Suppe, 1974, p.221-30, VanFraassen, 1970, 1980, p.64-70)

In other words, under the syntactic view, the empirical consequences of a theory bear a strictly logical relation to the theory itself. They are simply deduced from the theory. In the semantic conception, this is not the case. The empirical consequences of a theory can only be determined once an isomorphism has been established between a model of the data in question and a model of the theory. The establishment of this isomorphism, as well as the construction of the model of the data, can require many extra-theoretic resources.

However, in the semantic conception, while the exact empirical consequences of a theory are not logically determined by the theory, the sub-model, which is a candidate for being isomorphic to the data model, is. The sub-model is said to be embeddable in the semantic

structure of the theory. It follows directly from the theory by semantic implication.

In most cases, the equations that form the theoretical basis of systems of interest to simulationists are analytically unsolvable. That is, there is no mathematically expressible function that is the solution to these equations. Therefore, representations of these systems cannot possibly come in the form of linguistic entities derivable from a linguistic theoretical structure. The syntactic view won't cut it.

Moreover, deductive inferences by definition confer certainty on their conclusions (provided the premises are true!) The inferences that take place in simulation modeling confer no such thing. At their best, they confer reasonable warrant for believing the conclusions reached, and this only when painstaking steps are taken to insure success.

Nevertheless, we must keep in mind that what we are after here is not the human *means* by which results are acquired, but the *relation* that they *actually bear* to our best theories of the physical world. Might it not be argued then, that while simulation results are not *actually deducible* from the theory, they do bear a logical relation to the theoretical structure; a relation of *ideal deduction*. Surely, the argument goes, the actual relation between theory and model is independent of the means we use to arrive at the model, and it is equally independent of the trust we put in the model.

Before we consider this possibility, we need to look carefully at what we might be proposing. One thing that we would probably not want to propose is that our models of phenomena stand in an actual logical relation to our theories. Our models are approximate, and they are often rich with detail that we know is not robust.

What one might propose though, is that if all has gone well, the results of a particular simulation closely resemble an ideal structure that stands in a purely logical relation to the theory. The structure itself and the exact nature of the relationship are unfortunately inaccessible to us, but we should be confident that it is there and that we have successfully approximated it. The

fact that we lack the certainty of deduction merely reflects the fact that there is some chance that we are wrong and that our results do not bear any relation to the ideal structure. In other words, there is some structure that is a logical consequent of our theories, and what we have is something which we strongly believe approximates it. Is this a view that we should adopt?

The semantic view of theories provides us with a good framework for articulating this picture. In the semantic view, one way of characterizing a dynamical theory is a phase diagram. In this picture, the theory specifies what are the allowable state transition trajectories for any system under the domain of the theory. But in non-linear theories, the notion of “specifying” is quite weak. Even though there is a restricted class of trajectories which are logically determined by the theory, there may, even in principle, be no way of extracting or displaying what these trajectories are.

In cases, like ours, where there are no known solutions to the equations provided by the theory, we have no direct way of displaying what the appropriate family of trajectories is. But we might express the hope that our models of phenomena can play a surrogate role. On this view, models of phenomena represent the best attempts of scientists to qualitatively and approximately characterize a family of state transition trajectories that are semantically implied by the theory for a given range of parameters.

Even though this way of characterizing the situation may strike some as overly formal, it does provide us with a means for asking a central question here. Should we indeed think of the models of phenomena that simulationists provide to us as approximate qualitative characterizations of set theoretic ‘solutions’ of the theoretical equations? Simply put, do we have grounds to believe that models of phenomena are close approximations to structures that are models of our theories in the strict formal parlance of the semantic conception.

I would argue that we do not. Given the rich complexity of the process of deriving warrant for simulation results, and the extent to which this process focuses on elements *external*

to the theoretical structure itself, it would be unrealistic to interpret this warranting process as being about the results' relationship to some formal model. That is, if we think carefully about the epistemological steps that go into warranting the reliability of simulation results, we find that they have everything to do with ensuring that the results match up well with the world, and little to do with ensuring that they resemble some ideal structure.

In other words, the confidence we have in our simulation results depends on several factors being in place, none of which is guaranteed by our theoretical knowledge. It depends on facts we know about our computers, and about our graphical techniques. It depends on the confidence we have in the various *ad hoc* models we use—confidence we derive from laboratory and observational experience. It depends on our ability to calibrate models against empirical results. And finally, it depends on the confidence we have in our tacit abilities as observers to make judgements about the degree of resemblance between different classes of images; often abilities acquired in the role of skilled experimenters and observers, as well as in the role of skilled simulators. The epistemology of simulation, as we have seen at length, is very much an empirical epistemology, and not merely a mathematico-logical one.

§7 Conclusion.

Ultimately, even though all simulation modeling of the kind described in this paper fundamentally begins with a theoretical model, and even though we think of simulation as an attempt to 'solve' the mathematical equations of this theoretical structure, our theoretical knowledge is just one of several ingredients that simulationists use to produce their end product—a model of the phenomena. So the relation that theory bears to these models is simply this: our confidence in our physical theories is one of the things that, in combination with other elements, *warrants* our rational belief that the models we construct are reliable and cognitively useful.

There may indeed be something somewhat unsatisfying about this conclusion—that what is at root an application of theoretical knowledge bears no direct substantive relationship to the theory itself. Unfortunately, this conclusion is simply the result of a fundamental limitation of our cognitive power. When it comes to complex systems, we simply cannot bend our theories to our cognitive will—they will not yield results with any mechanical turn of a crank. The models that we need to construct in order to do our science need to be constructed delicately, and from as many sources as are available. Consequently, these models are no mere ‘solutions’ to our theoretical structures. Though they are the results of a form of calculation, they are rich physical constructs that mediate between our theories and the world.

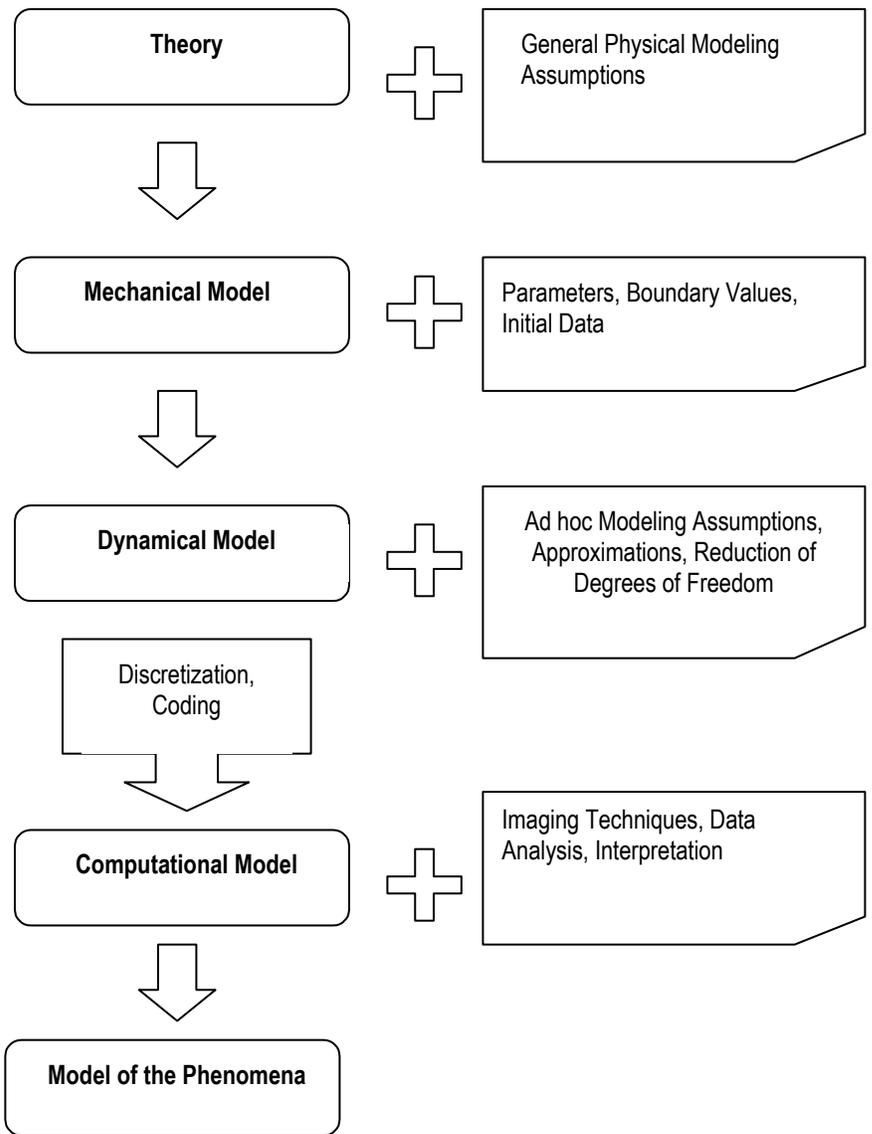


Figure 1. The Hierarchy of Models.

References

- Cartwright, N. 1983. *How the Laws of Physics Lie*. Essays 7 and 8. Oxford: Oxford University Press.
- Cartwright, N. 1989. *Nature's Capacities and Their Measurement*. Oxford: Oxford University Press.
- Franklin, A. 1986. *The Neglect of Experiment*. New York: Cambridge University Press.
- Gierre, R. 1988. *Constructing Science*. Chicago: University of Chicago Press.
- Goodman, N. 1968. *Language of Arts*. New York: The Bobbs-Merrill Company.
- Hempel C.G. 1966. *Philosophy of Natural Science*. New Jersey: Prentice-Hall.
- Kaufmann, W.J., and L.L. Smarr 1993. *Supercomputing and the Transformation of Science*. New York: Scientific American Library.
- Kuhn, T. 1977. "The function of measurement in modern physical science." In *The Essential Tension: Selected Studies in Scientific Tradition and Change*. Chicago: The University of Chicago Press.
- Laymon, R. 1985. "Idealizations and the Testing of Theories by Experimentation." In *Observation, Experiment and Hypothesis in Modern Physical Science*, edited by in P. Achinstein and O. Hannaway. Cambridge, MA: MIT Press.
- Laymon, R. 1990. "Computer Simulations, Idealizations and Approximations." in *PSA 1990 vol. 2*: 519-34. Edited by East Lansing: Philosophy of Science Association.
- Moin, P. and J. Kim 1997. "Tackling turbulence with supercomputers." *Scientific American* 276, 1: 62-68.
- Moulton, F.R. 1970. *An Introduction to Celestial Mechanics*. New York: Dover.
- Ramsey, J.L. 1990. "Beyond Numerical and Causal Accuracy: Expanding the Set of Justificational Criteria", in *PSA 1990 vol.1*: 485-499. Edited by A. Fine, M. Forbes, and L. Wessels. East Lansing: Philosophy of Science Association.
- Ramsey, J.L. 1992, "An Expanded Epistemology of Approximations." in *PSA 1992*, vol. 1: 154-66. Edited by D. Hull, M. Forbes, and K. Okruhlik. East Lansing: Philosophy of Science Association.
- Shirayama, S., and K. Kuwahara 1990. "Flow visualization in computational fluid dynamics." *The International Journal of Supercomputer Applications*, 4, 2: 66-80.
- Smarr, L. 1985. "An approach to complexity" *Science* 228: 403-408.
- Suppe, F. (ed) 1974. *The Structure of Scientific Theories*. Urbana: University of Illinois Press

Suppes, P. 1962. "Models of Data." In *Logic, Methodology and Philosophy of Science*, edited by E. Nagel, P. Suppes and A. Tarski. Stanford: Stanford University Press.

Van Fraassen, B.C. 1970. "On the extension of Beth's semantics of physical theories." *Philosophy of Science* September, 1970: 325-338.

Van Fraassen, B.C. 1980. *The Scientific Image*. Oxford: Oxford University Press.

Wilhelmson, R.B. et al 1990. "A study of the evolution of a numerically modeled severe storm." *The International Journal of Supercomputer Applications*, 4, 2: 20-36.

[†] I am indebted for discussions and advice to Marcel Boumans, Micheal Dickson, Deb Dowling, Robert Downey, Michael Friedman, Sergio Sismondo, Fred Suppe, and two anonymous reviewers.

¹ *c.f.* (Wilhelmson, 1989) on the simulation of a severe storm, (Smarr, 1985), and (Kaufmann and Smarr 1993) on the simulation of intergalactic gas jets, and (Moin and Kim, 1997) on computer simulation in the study of turbulence.

² The canonical discussion of the "received" view of theories can be found in (Suppe, 1974, introduction)