COMPUTER MODELING IN CLIMATE SCIENCE: EXPERIMENT, EXPLANATION, PLURALISM

by

Wendy S. Parker

BA, BS, Northern Illinois University, 1997

MA, University of Pittsburgh, 2002

Submitted to the Graduate Faculty of

Arts and Sciences in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2003

UNIVERSITY OF PITTSBURGH

FACULTY OF ARTS AND SCIENCES

This dissertation was presented

by

Wendy S. Parker

It was defended on

August 11, 2003

and approved by

John Earman

Sandra Mitchell

Laura Ruetsche

John Norton
Dissertation Director

COMPUTER MODELING IN CLIMATE SCIENCE: EXPERIMENT, EXPLANATION, PLURALISM

Wendy S. Parker, PhD

University of Pittsburgh, 2003

Computer simulation modeling is an important part of contemporary scientific practice but has not yet received much attention from philosophers. The present project helps to fill this lacuna in the philosophical literature by addressing three questions that arise in the context of computer simulation of Earth's climate. (1) Computer simulation experimentation commonly is viewed as a suspect methodology, in contrast to the trusted mainstay of material experimentation. Are the results of computer simulation experiments somehow deeply problematic in ways that the results of material experiments are not? I argue against categorical skepticism toward the results of computer simulation experiments by revealing important parallels in the epistemologies of material and computer simulation experimentation. (2) It has often been remarked that simple computer simulation models—but not complex ones—contribute substantially to our understanding of the atmosphere and climate system. Is this view of the relative contributions of simple and complex models tenable? I show that both simple and complex climate models can promote scientific understanding and argue that the apparent contribution of simple models depends upon whether a causal or deductive account of scientific understanding is adopted. (3) When two incompatible scientific theories are under consideration, they typically are viewed as competitors, and we seek evidence that refutes at least one of the theories. In the study of climate change, however, logically incompatible computer simulation models are accepted as complementary resources for investigating future climate. How can we make sense of this use of incompatible models? I show that a collection of incompatible climate models persists in part because of difficulties faced in evaluating and comparing climate models. I then discuss the rationale for using these incompatible models together and argue that this climate model pluralism has both competitive and integrative components.

FOREWORD

I owe many thanks the for the financial assistance, intellectual guidance, and emotional support that I have received in recent years. In my first two years in the program, a Graduate Research Fellowship from the National Science Foundation freed me from teaching responsibilities and allowed me to concentrate on my coursework. During my last academic year, a Mellon Fellowship from the University of Pittsburgh afforded me the opportunity to focus my attention fully on the completion of this dissertation. I owe special thanks to my advisor, John Norton, who has been supportive and helpful throughout the entire dissertation process. The comments and advice that I received from other members of my committee—John Earman, Sandra Mitchell, and Laura Ruetsche—facilitated substantial improvements to several chapters and were greatly appreciated. The project also benefited from conversations with Daniela Bailer-Jones, Gualtiero Piccinini, James Risbey, and John Roberts. I am especially grateful to Francis Longworth for his constant help, enthusiasm, and encouragement. Finally, I would like to thank my parents for continuous support of every kind.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1
Introduction

Only fifty years ago, computer simulation modeling was a new activity undertaken in only a handful of scientific disciplines. Since then, it has come to play an increasingly important role in scientific practice. In fact, it is not an exaggeration to say that computer simulation modeling is nearly ubiquitous in science today. Scientists in fields as diverse as physics, evolutionary biology, and economics engage in computer simulation modeling on a regular basis, and in some fields, such as meteorology, the introduction of computer simulation modeling has led to truly revolutionary advances.

Despite the importance of computer simulation modeling in contemporary science, philosophers have had surprisingly little to say about it. A common theme in the work that has been done is that computer simulation modeling is somehow intermediate between theorizing and experimenting (see e.g. Rohrlich 1991; Humphreys 1994; Galison 1997; Dowling 1999; Sismondo 1999). But there is much more to be said about computer simulation modeling. Precisely because it does not fit neatly into either methodological category, computer simulation modeling raises a host of interesting philosophical—and especially epistemological—questions and puzzles. These questions and puzzles are worthy of philosophers' attention both because addressing them will contribute to a more complete understanding of science as it is now practiced, and because (as we will see) some scientific issues that may have profound societal impacts cannot be responsibly attended to without a relatively clear understanding of the nature, strengths, and limitations of computer simulation modeling. Put succinctly, the practice of computer simulation modeling generates its own important philosophical questions and puzzles, and there has not yet been enough systematic and sustained effort to address them.

This dissertation will contribute to the filling of this lacuna in the philosophical literature. The particular questions and puzzles to be considered are epistemological and methodological in nature and are drawn from a scientific arena in which computer simulation modeling plays an especially central role: the study of Earth's weather and climate. Widespread concern over the possibility of global warming has fueled research in this arena in recent years. Because computer simulation models are involved in almost every aspect of global warming research, it is not surprising that disagreement over the epistemic status of such models has been the centerpiece of both the scientific and policy-related debates over global warming. No doubt some of this disagreement is politically driven, but genuine confusion and uncertainty about the epistemic value of climate models has left plenty of room for divergent views to flourish. This project thus aims not only to contribute to the philosophical literature on computer simulation modeling, but to do so in a way that is useful to scientists and policymakers who are struggling to understand the nature and value of climate models. The core of the dissertation addresses three sets of questions:

(1) Computer simulation experiments play a central role in climate change research and in the study of the climate system more generally. One tactic that is available to (and is sometimes used by) global warming skeptics is to declare that it is "only a model" behind the results of these computer simulation experiments. The implicit contrast is with results obtained from "real" material experiments, like those that take place in the chemistry laboratory using real chemicals and test tubes. But are the results of computer simulation experiments somehow deeply problematic or questionable in a way that the results of material experiments are not? In Chapter 4, I will argue against blanket or categorical skepticism toward the results of computer simulation experiments. I will suggest that part of the skepticism that does exist comes from a failure to appreciate important parallels in the epistemologies of material and computer simulation experimentation, and I will highlight some of these parallels.

(2) Beginning early on in discussions of weather and climate modeling and continuing up to the present, we find scientists remarking that simple computer simulation models—but not complex ones—contribute substantially to our understanding of the atmosphere and climate system. This is a puzzling claim. It implies that the descriptive accuracy of a model and its value in promoting understanding are not aligned. How could simple models, which patently "get it wrong" relative to more complete theoretical descriptions nevertheless play an important part in helping us to understand the world? Is this view of the relative contributions of simple and complex models

tenable? In Chapter 5, I will show that this view appears more or less plausible depending on the account of scientific understanding that one adopts. I will argue that ultimately the view is untenable: both simple and complex models contribute to our understanding of weather and climate in ways that I will illustrate.

(3) In the study of climate change, there are not only simple and complex computer simulation models—there are many simple models and many complex models. Furthermore, many of the models at a given level of complexity are logically incompatible with one another. Even more puzzling is the fact that these incompatible models are used together as complementary resources for investigating future climate change. By contrast, when two logically incompatible theories are available, they typically are viewed as competitors, and scientists seek evidence that refutes one theory and supports the other. Why is the situation different in climate modeling? How should we characterize this climate model pluralism? In Chapter 6, I will give a non-instrumentalist explanation of the persistence of logically incompatible climate models, and I will discuss the rationale for using such models together to investigate climate change. I will suggest that this interesting use of incompatible models involves two kinds of pluralism: ontic competitive pluralism and pragmatic integrative pluralism.

The remaining chapters will provide philosophical and scientific background and summarize conclusions. In Chapter 2, I will present and defend the view of scientific models (in general) and computer simulation models (in particular) that will be assumed in the rest of the dissertation. I will endorse an intentional, pragmatic account of models in science according to which a model is a representation designated as such in virtue of perceived relevant similarities between the model and that which it is chosen to represent. Chapter 3 will introduce the reader to the study of climate change and will discuss the variety of models used to investigate climate change as well as the several ways in which they are used in such investigation. Finally, in Chapter 7, I will provide some concluding remarks.

2

Models and Computer Simulation Modeling

2.1 Introduction

This chapter provides a general discussion of models and computer simulation as background for later chapters. In Section 2.2, I describe the general view of models in science that I endorse; this view is intentional and pragmatic and has much in common with that of Giere (1988, 1999a, 1999b) and Teller (2001). In Section 2.3, as a transition to the topic of computer simulation modeling, I characterize some types of models that are commonly associated with but should be distinguished from computer simulation models. In Section 2.4, I turn to the topic of computer simulation modeling. In the process of examining accounts of simulation that have been offered, I present and defend my own characterization of computer simulation and clarify related terminology that will be important to keep in mind in later chapters. Then, in Section 2.5, I consider the most detailed recent work on the practice of computer simulation modeling, which was put forth by Winsberg (1999b, 1999c); he argues that, in the study of complex physical systems, computer simulation modeling is an activity that typically involves the use of a number of particular kinds of models. I agree with Winsberg that a series of different kinds of models is typically used, but I argue that the series of models differs from that described by Winsberg. Section 2.6 includes a summary of the most important conclusions.

2.2 Models in science: a pragmatic view

Although references to what would now be called models can be found in natural philosophical works dating to the 17th century (if not long before), the topic of models in science did not receive much philosophical attention until the mid-20th century. Since then, a substantial but quite disjointed philosophical literature has developed. I will not attempt to chronicle the development of this literature—it would be a substantial undertaking and one not especially relevant to the present project, since the philosophical literature on computer simulation has developed largely independently of that on models in science. However, it will be useful to provide some discussion of the view of models that I endorse, since this will provide a larger framework within which to situate the topic of computer simulation models (i.e. the particular type of model that is of interest in this project).

Before describing the view of models that I endorse, I want to emphasize that it is not at all easy to give a general yet interesting account of models in science, at least not if one wants to respect the actual usage of the term by scientists and philosophers of science. As noted by Goodman already several decades ago (see Goodman 1968), the term "model" is used to refer to an incredibly diverse collection of entities. Consider some examples. There is the solar system model of the atom. There are model organisms—physical entities such as worms and rats—that are used to find out about humans and other organisms. There are mathematical models of ideal and real pendulums. There are computer simulation models used to forecast the weather. There are scale models studied in wind tunnels. The list could go on. The challenge of giving a general account of models is to say something non-trivial about what these various "models" have in common. Even attempting to give a relatively comprehensive and detailed account would be a dissertation project in itself; here, I just want to discuss the main features of the kind of view that I endorse.

Like many other contemporary philosophers of science (e.g. Bailer-Jones (2003), Giere (1999a), Hughes (1997, 1999), Morrison (1998), Morrison and Morgan (1999), Suarez (1999)), I take the fundamental feature of models to be that they represent—models are *representations*. A central question that a representation view of models must answer is: in virtue of what are models said to represent?

To begin to answer this question, I will suggest that an entity does not become a model in virtue of its *actually* standing in some particular relation to that which it is said to model. Rather, it is considered a model because it is *perceived* or *believed* to stand in some relation to that which it is said to model. A scientist or scientific community designates one entity to be a model of some other entity because there is some perceived relation between the two entities. In this sense, my view of models is an intentional view—entities are models because they are designated as such by us (see Teller 2001 for a recent related view). This does not mean that models can be chosen or constructed arbitrarily or that just anything will be chosen or constructed to be a model of any other thing. It also does not mean that models must be artifacts as opposed to natural objects or systems. It simply means that we decide rather than discover that something is a model of something else. We can discover that something is a good model of something else or that one model is better for some purpose than another, but this is discovering something *about* a model rather than *that* something is a model in the first place.

It is obvious that in practice scientists do not just arbitrarily designate one entity to be a model of another—as suggested above, there must be some perceived relation between the two entities. What is the nature of that relation? Here is where the sheer variety of models creates difficulties; if we are not careful, our answer to this question is bound to exclude from the category "models" many of those things that scientists actually (and legitimately, it seems) identify as such. For example, perceived isomorphism between a model and that which it models is obviously too strong. By isomorphism, I mean a one-to-one correspondence between the elements and relations of the model and that which it models. If perceived isomorphism were the appropriate relation, then schematic diagrams of natural phenomena that reflect only the most prominent or salient features of the phenomena (and only to a limited degree) would fail to count as models, and so would many other "models." But schematic diagrams do qualify as a type of model in scientific practice. We need a less restrictive characterization of the relation between the model and that which it models.[1]

I suggest that the relation be characterized as *perceived relevant similarity*. One entity is chosen to represent another entity because the former is perceived to be similar in relevant respects and to relevant degrees to the latter. Which respects and degrees are considered relevant depends upon the purposes for which the model is intended to be used. For example, suppose we want a model organism that we can use to study how a new cancer drug will affect brain tumors in humans. In this case, we want to select a model organism whose brain tumors are similar in various ways and to various degrees to those in humans, but we do not care whether the model organism is dissimilar to humans

---

[1] Giere (1999a, 46) also rejects isomorphism as the sense in which models represent. It is interesting that Suarez (1999, 77-79) argues that isomorphism is not sufficient for representation.

with respect to hairiness, number of toes, mating behavior, etc. The purposes for which a model is to be used help to determine in which respects and to which degrees a model should be similar to its target. Of course, we rarely know for sure exactly in which respects and to which degrees a model should be similar to its target if our intended purposes are to be fulfilled, but consideration of these purposes almost always helps us to conclude some things about those respects and degrees. Particular entities are selected or constructed to be models because they are perceived to be *relevantly similar* to that which they are chosen to represent.

This view has much in common with that of Giere and of Teller, who also take similarity to be constitutive of the representation relation (see Giere 1999b, 123 and Teller 2001, 398-402).[2] What is "similarity"? Like Hesse (1966), Giere and Teller seem to take similarity to be a kind of primitive notion. They explicitly deny that a general account of similarity is needed. Giere argues that what it means to say that a model is "similar" to some other thing is context dependent but not vacuous—in a given context with particular modeling goals, one can specify (1) which aspects of the model should be similar to which aspects of the thing to be modeled, (2) the way in which these aspects should be similar, and (3) the degree to which they should be similar (Giere 1999a, 45-46). Teller claims that no general account of similarity can be given but argues that such a demand is misguided to begin with (2001, 401-402), because the details of the particular modeling situation at hand specify what counts as (relevant) similarity, i.e. specify the respects in which and degrees to which a proposed model should be similar to that which it models. I am sympathetic to this analysis, although I would express it slightly differently as follows: to say that a model is (relevantly) similar to that which it models is to say that, in particular respects and to particular degrees, the model is like that which it is chosen to model. Similarity is context dependent in that context (including modeling purpose) helps us to determine the respects and degrees of likeness considered important.[3]

---

[2] However, while Giere is often concerned with what he calls "theoretical models"—abstract objects such as the simple harmonic oscillator which are defined by the equations of a theory (see e.g. 1988, 78-79)—I intend for my discussion of models to apply much more broadly, as indicated above, to concrete objects, computer simulation models, etc.

[3] This way of thinking about similarity does leave us with "likeness" as a primitive notion, but given the aims of this chapter, it does not seem worthwhile to push the analysis much further.

R.I.G. Hughes (1997, 1999) issues a challenge to the idea that similarity is an important part of an account of models in science. He denies that models in science typically are similar to that which they represent (1997, S329 and 1999, 126). While he admits that some types of models (e.g. scale models) do resemble their subjects, he claims that more typical abstract (especially "theoretical") models are not similar to that which they model. As a supposed illustration of this, he argues that the abstract "ideal pendulum" is not similar to any actual material pendulum:

> We may even be tempted to say that in both cases [ideal and material] the relation between the pendulum's length and its periodic time is approximately the same, and that they are in that respect similar to each other. But the ideal pendulum has no length, and there is no time in which it completes an oscillation. It is an abstract object, similar to material pendulums in no obvious sense. (1997, S330)

This view seems mistaken, however. Surely it is possible to present the situation a little more precisely and clearly so that the problem disappears, i.e. so that there *is* a similarity of the kind that Hughes denies. In particular, one can say that the mathematical relation between the quantity that is called "length" and the quantity that is called "period" for the ideal pendulum is similar to the relation between quantities that are found when the length and period of the material pendulum are measured. On this description, the ideal and material pendulums do have some properties that are similar. I would further assert that it is in large part *because* of this perceived similarity that the ideal pendulum is taken to be a model of the material one.[4]

Thus, on the intentional, pragmatic account of models that I endorse, a model is a *representation* designated as such in virtue of *perceived relevant similarities* between the model and that which it is chosen to represent. The *purposes* for which the model will be used provide direction in determining in which respects and to which degrees this similarity should hold. This account makes clear the dual or hybrid nature of models: models represent in virtue of their perceived relationship to a modeling target, but the nature of the required relationship is shaped by the purposes for which the model is to be

[4] French (2002) also seems to reject Hughes' argument, though French suggests that partial isomorphism (rather than similarity) might be a promising way to characterize the representation relation. He does not provide any extended argument that similarity is an inappropriate characterization. It may be that at bottom the partial isomorphism approach and the similarity approach amount to largely the same thing, but I will not pursue this here, since it would make little difference to the rest of my project.

used. Put succinctly, models are both representations and tools, and these aspects are interrelated.[5] This account of models is general but informative and yet succeeds in recognizing as models those entities that are actually designated as such in scientific practice. An especially attractive feature of the account is that it puts no a priori constraints on the "ingredients" that can relied upon in constructing models, and in particular does not require that models stand in any close relation to theories.

2.3 Distinguishing some important types of models

Before examining in some detail the type of model that is of special interest in this project, it is worthwhile to distinguish some closely-related types of models that are often confused with computer simulation models. Perhaps because they all typically involve equations, mathematical models, numerical models (a particular species of mathematical model), and computer simulation models are often spoken of interchangeably. However, there are differences among these types of models, and overlooking these will only lead to confusion (and sloppy philosophy) when it comes time to make and evaluate claims about computer simulation models. To help avoid this, in this section, I want to characterize mathematical models in general and numerical models in particular in order to distinguish them from computer simulation models, which will be examined in detail in the next section.

Mathematical models in science are sets of mathematical relations that are taken to describe relations among particular properties of a system of interest. For example, the ideal gas law describes the relations among the pressure, temperature, and volume of an ideal gas. Not all of the variables or terms that appear in a mathematical model must be assumed to stand for properties of the physical system of interest. For example, we might find that tacking on a particular "fudge factor" results in better predictions from our mathematical model, and we might have no idea about the relation (if any) that the fudge factor has to any physical features of the modeled system. But at least some of the mathematical relations constituting the model are thought to be similar to the mathematical relations that would hold between quantities in the system of interest if

[5] Morrison and Morgan (1999) present a similarly dualistic picture of models as "mediating instruments." I want to emphasize more strongly than they do the fact that the instrumental dimension of models influences their representational dimension.

those quantities were to be measured (as in the pendulum example described in the last section).

Mathematical models may rely heavily on accepted background theory, but they need not. For example, the following mathematical model might be determined primarily from empirical measurements of the two variables (through regression):

$$Z = 240R^{1.5}$$

where Z stands for the reflectivity measured by a radar for a particular volume of atmosphere and R stands for the rain rate at the Earth's surface below that volume of atmosphere. This relation can be used to estimate the rain rate, given a reflectivity measurement. Of course, it is not the case that no theoretical knowledge went into obtaining the Z-R model. Theoretical knowledge concerning the behavior of the emitted radiation, the reflectivity of water droplets, etc. all will be relied upon at least implicitly in obtaining such a model. Such knowledge underwrites our belief that our radar can tell us about the reflectivity of raindrops in the first place. However, the Z-R relation itself is not obtained via derivation from some theory concerning the relation of Z and R, but rather is obtained simply by fitting a curve to particular measurements of Z and R that are made.[6]

The mathematical relations that constitute a mathematical model may or may not be solvable, either in theory or in practice, whether by us or by some artificial device. In fact, it has turned out that many of the mathematical models of interest in science are constituted by equations that are extremely difficult, if not impossible, to solve analytically. Such is the case, for example, for the equations that describe the large-scale motion of the Earth's atmosphere. Morton (1993) calls this the problem of unsolvability/uncomputability. In order to make progress in such situations, scientists

[6] An example of a type of mathematical model that is closely tied to theory is Giere's (1999a) *theoretical model*. Giere describes theoretical models as abstract objects that are defined using theoretical principles such as Newton's Laws, the Principle of Relativity, and the laws of Mendelian Genetics (1999a, 51). Although he describes theoretical models not as consisting of mathematical equations themselves but rather as abstract objects described by equations, I still classify them as mathematical models, since the mathematical relations (whether expressed using a *particular* set of equations or not) are what defines the models and since when we communicate their content we typically do use equations. Examples of theoretical models include the ideal pendulum, the ideal harmonic oscillator, and a generic two-body system (in the context of Newtonian mechanics).

often resort to the use of a particular species of mathematical model, namely, the numerical model.

Numerical models also consist of mathematical relations, but these models are specifically designed in order to be solvable by numerical computation. This means that the construction of numerical models requires knowledge of available numerical methods. For instance, one may need to know methods for discretizing continuous equations in such a way that the solutions of the discrete equations approximate those of the continuous ones. Simplification of equations may also be required in order to get them into a form that can be solved numerically. For instance, terms whose impact on the solution is small (as shown e.g. by scale analysis of the equations) but whose form is very complicated may be replaced by simpler terms or even discarded entirely. When numerical models are used, the aim is to find solutions to the numerical model equations that approximate to some desired degree of accuracy the solutions to a mathematical model whose equations (probably continuous) we have been unable to solve.

Neither mathematical models in general nor numerical models in particular incorporate explicitly any information about the means by which their equations are to be solved; these models are constituted by equations (and other mathematical relations, such as inequalities) but not by instructions for solving those equations. Even though numerical models are designed to be amenable to solution by numerical methods, the models themselves do not specify precisely how their equations are to be solved (e.g. do not specify time-step magnitudes, boundary conditions). In this respect, mathematical models and numerical models differ from computer simulation models, which, as we will see, do incorporate instructions for solving equations that they include. Although these types of models should be distinguished, it is true that mathematical models and especially numerical models often play an important role in the practice of computer simulation modeling, which will be discussed in Section 2.5.

2.4 Characterizing simulation and computer simulation modeling

If the topic of models is a relatively recent one to arrive on the philosophy of science scene, the topic of computer simulation modeling is even more recent. This is not surprising. Because digital computers were not really in use (except in a few contexts) until the 1960s, it is to be expected that the philosophical literature on simulation modeling will have developed only in the past few decades. However, given that through the course of those few decades computer simulation has become an incredibly widespread practice in both natural and social science, it is also to be expected that computer simulation would have become a topic of real interest to philosophers of science. After all, this is a new aspect of scientific practice that is different, at least ostensibly, from what has come before.

Yet philosophers of science have paid surprisingly little attention to the topic of computer simulation. Although Bunge made some early contributions to the philosophical literature on simulation (e.g. Bunge 1969), and some methodologically-minded scientists offered what might be deemed philosophical discussions on the topic of simulation (e.g. Guetzkow et al. 1972, Shannon 1975), philosophers of science in general only began to consider the topic of computer simulation in the natural sciences in the 1990s.[7] Even then, the volume of philosophical literature was not very large, given the established and growing importance of simulation in the natural sciences. I will review some of this literature below and use it as a guide in presenting my own views on the nature of computer simulation modeling.

Humphreys 1991 is an important general contribution to the philosophical literature on computer simulation in the natural (and especially physical) sciences. Humphreys offers the following working definition of computer simulation:

> A computer simulation is any computer-implemented method for exploring the properties of mathematical models where analytic methods are unavailable. (1991, 501)

He is careful to argue that numerical methods should be distinguished from computer simulations, which (unlike numerical methods) pertain to specific scientific problems and are implemented on concrete machines that can carry out the needed computations (Humphreys 1991, 502). Put differently, numerical methods may be used to solve a variety of types of equations that may not be at all related to any particular scientific problem, while computer simulations are carried out on particular machines making use

---

[7] I will not address the history of philosophical interest in the use of computers to model human cognition or the human brain. This is a large literature that is almost entirely separate from what has been written on computer simulation in the natural sciences.

of numerical methods in order to investigate specific problems (e.g. the effect of carbon dioxide on the radiative properties of the Earth's atmosphere).

Humphreys' working definition of computer simulation is criticized by Hartmann (1996), who raised two apt objections. First, he points out that the definition does not emphasize the "dynamic character" of computer simulation models and the mathematical models that typically underlie them (Hartmann 1996, 84). That is, the definition does not make any explicit reference to the fact that simulations have to do with the *temporal* evolution of systems. In this context, Hartmann makes a useful distinction between static and dynamic models. A static model only includes assumptions about a system that is not changing through time, while a dynamic model also includes assumptions about the time-evolution of the system (Hartmann 1996, 82). Second, Hartmann finds the definition too restrictive in its focus on analytically-intractable mathematical models (Hartmann 1996, 84). This restriction appears in Humphrey's definition in the last clause, according to which computer simulation is used only where analytic methods are unavailable. Hartmann suggests that computer simulation is also used (and useful) even when analytic methods are available, because "visualizing the result of a simulation on a computer screen...may increase our understanding of the system more than complicated formulas written down on a paper would ever do" (Hartmann 1996, 84).[8]

Hartmann also presents his own alternative characterization of computer simulation. His definition of a computer simulation relies upon his prior definition of simulation more generally. His definition of simulation can be presented as follows: a *simulation* imitates one object or system whose state changes in time by another object or system whose state changes in time (1996, 83).[9] This kind of imitation results, according to Hartmann, "when the equations of the underlying dynamic model are solved" (Hartmann 1996, 83). He then characterizes a *computer simulation* as a simulation run on

[8] Hughes (1999) offers an additional criticism of Humphreys' working definition. According to Hughes, the definition "blurs the distinction...between computer simulation and the use of the computer to solve intractable equations" (1999, 131). While I can imagine how a distinction between these two uses of the computer might be made, I do not find a clear account of the distinction in Hughes 1999, despite the examples that are offered to illustrate (see Hughes 1999, 129-130).

[9] Hartmann presents this definition in two steps, one defining simulation in terms of process and a second defining process in terms of a state-changing system. This account of simulation is adopted by Hughes (1999, 130) and Guala (2002, 62), except that Guala makes the slight change of defining a process in terms of a time-ordered sequence of states rather than in terms of the system displaying those states.

a computer (Hartmann 1996, 83). This is a more inclusive account of computer simulation than that offered by Humphreys (1991), since it includes cases in which analytic methods are available. Hartmann suggests that this broader account fits better with the variety of ways in which computer simulation is used in science.

I think that a slightly clearer general account can be given. I define a *simulation model* as a dynamic model whose temporal evolution is supposed to be similar (in some respects and to some degrees) on some level of description to the temporal evolution of the system that it models.[10] When the simulation model actually evolves temporally (naturally or through human intervention), this is a *simulation* of the system that it models. Some simulation models can be transformed into computer code whose execution results in the solution of the simulation model equations.[11] This computer code is a *computer simulation model.* The execution of this code (the "running" of the computer simulation model) is a *computer simulation.*

An easily overlooked difference between my definition of a simulation model and Hartmann's definition of simulation is my explicit indication that the requisite similarity must be supposed to exist on at least one level of description. I include this requirement not to restrict the definition; rather, this requirement is intended to ensure that the definition is *not* interpreted in a more demanding way, e.g. as requiring that a simulation model must be similar on *all* levels of description to the system that it models. I do not mean to imply that Hartmann thinks that simulation involves this all-level similarity; I think it unlikely that he holds such a view. But his definition of simulation does not make this clear enough. I want to make explicit just how weak I intend the similarity requirement to be. This requirement must be weak if the associated definition of computer simulation is to fit with actual scientific usage of the term. This is because the actual execution of the computer code often involves step-by-step calculations whose products, on some levels of description, *cannot* be supposed to be similar (in certain important respects) to the actual objects or systems of interest.

---

[10] This characterization makes use of Hartmann's notion of a *dynamic model*, as defined above. All simulation models are dynamic models, but not all dynamic models need be simulation models.

[11] In this case, the simulation model will need to be a numerical model, as described in the previous section. The process of arriving at a computer simulation model may be a complicated, multi-step process, as described in the next section.

One way to see this is to consider the situation that arises in computer simulation of the Earth's atmosphere for purposes of weather prediction. The computer simulation generates values of temperature, pressure, humidity, etc. for a large number of grid points intended to correspond to volumes of atmosphere at or above the Earth's surface. The computer simulation produces these grids (or matrices) of values for a series of times $t_1 \ldots t_n$ that are supposed to correspond to actual times in the evolution of the atmosphere. For example, the computer simulation might predict what the temperature will be in the volumes represented by the grid points at 10am next Saturday, at 11am next Saturday, at noon next Saturday, and so on. However, even though a grid of values is supposed to describe the state of the atmosphere at a *particular* time t, this grid of values for time t is generated in the computer simulation by calculating a value for each grid point individually *in succession* (through time). In the real atmosphere, the new state of the system at time t emerges *simultaneously* at all of the locations to which the model grid points correspond. Thus, on a very detailed level of description that follows step-by-step through the execution of the computer code, the temporal evolution of the model atmosphere (the one described by the simulation model) differs qualitatively from what happens in the real atmosphere and in an important respect is not similar to what happens in reality. However, on a different level of description, the temporal evolution of the model atmosphere in the computer simulation can be supposed to be similar to the evolution of the real atmosphere. This will be the case, for instance, if the simulation is described in terms of the grids of values produced *for* times $t_1 \ldots t_n$.

I do not want to belabor this point unnecessarily. I just want to make clear that the account of computer simulation that I am proposing above is not a very restrictive one; it does *not* require similarity on all levels of description. Like Hartmann, I want my account to be inclusive enough to recognize as computer simulation models many of those entities that are so designated in scientific practice.

2.5 The practice of simulating complex physical systems

A recent, in-depth examination of computer simulation is offered by Winsberg (1999a, 1999b, 1999c, 2001). He is concerned with a very particular type of computer simulation, namely, computer simulation of complex physical systems for which there are established theories of the processes that make up the system.[12] This makes his analysis of special interest to the present project, since the climate system is a complex physical system comprised of many processes, some (though not all) of which are well-understood theoretically. Because of its relevance, it is worth examining Winsberg's account in some detail.

Winsberg presents simulation of complex systems as a practice involving the construction of a hierarchy of models of different types (1999c, 279 and 1999b, 256). The hierarchy is constructed as follows. First, it is necessary to identify the theory or theories that are applicable to the phenomena of interest (1999c, 279). One then develops a *mechanical model*, which Winsberg describes as "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" (Winsberg 1999c, 279).

After one has a mechanical model, it is necessary to take into account the parameters, boundary values, and initial conditions that will transform the general mechanical model into a family of *dynamical models* that apply specifically to the class of phenomena of interest (Winsberg 1999b, 258). As an example of a dynamical model, Winsberg (1999a, 8) describes a set of partial differential equations which include the variables relevant to the development of a severe storm (as presented in Wilhelmson et al. 1990).

Since Winsberg is specifically concerned with cases in which the dynamical models involve sets of equations that are mathematically intractable, the next step in the simulation process is to construct a *computational model*. This involves translating the equations (usually by discretizing them) into a form amenable to numerical solution on a computer and then actually writing/building the computer code/program that the computer will follow in order to solve the equations. According to Winsberg (1999c, 282), this also typically requires the use of *ad hoc* modeling assumptions, which help to make the dynamical model computationally tractable via the use of approximations and other simplifications that are not necessarily theory-driven.

[12] Apparently because Winsberg is concerned with this particular type of computer simulation, he does not give a general account of simulation.

I would suggest that Winsberg's computational model actually involves two types of models. First, there is what I have called a *numerical model* in the previous section; this consists of equations amenable to solution by numerical computation. Second, there is the actual computer code whose execution will result in the repeated solution of the equations comprising the numerical model. The computer code, of all of the models discussed so far, seems most deserving of the label *computer simulation model.* This model incorporates not just a set of equations amenable to numerical solution but also the techniques and methods that will actually be used to solve the equations—it indicates how the solution of the equations should proceed, in what order they are to be solved, to what degree of accuracy, etc.

The last model in Winsberg's hierarchy of models is what he calls the *model of the phenomena.* This is a "manifold representation that embodies the relevant knowledge, gathered from all relevant sources, about the phenomena" (Winsberg 1999c, 283). The model of the phenomena can include pictures, text, and equations considered relevant to the characterization of the phenomena from the point of view of theory, observation and/or modeling results. Winsberg considers this synthesis of all available (relevant) information about the phenomena to be the final goal of a simulation study (Winsberg 1999c, 283).

Again, I think that a further distinction can be made. After the computer program has been run, there is a data set waiting to be analyzed, interpreted and displayed, as Winsberg recognizes. (Sometimes this data is even displayed as part of the running of the program.) This data set can be displayed in various ways to make it more cognitively accessible to scientists, but the data set itself can also be considered a model. I propose that this data set, which is the immediate product of running a computer simulation model (because that is what the model was designed to produce), be included in the hierarchy proposed by Winsberg. Perhaps it could be called a *time-series model*, since the data set is a time series of states of the system described by the numerical model, intended to model the real system of interest. In other words, the data represent successive "snapshots" of the state of the system of interest in terms of the chosen variables. However, this time-series model can be distinguished from the model of the phenomena as identified by Winsberg. The model of the phenomena incorporates much more than (and may not include much information gleaned from) the data set produced by the simulation model run. It seems more appropriate to view the model of the phenomena as a typical desideratum of scientists but as something that need not be closely tied to any particular simulation modeling study.

Ultimately, I am in agreement with Winsberg that simulation modeling of complex physical systems is an activity that typically involves several different kinds of models. This is borne out by an examination of actual scientific practice. However, I would propose an amended series of models that can accommodate situations in which well-confirmed background theories describing all aspects of the system of interest are not available. My amended series of models proceeds as follows.[13]

(1) The process begins with the construction of a *descriptive model.* This is very much like Winsberg's "mechanical model". The descriptive model describes in a general, typically qualitative way what the important components and processes are that need to be considered in simulating the system. The model is constructed using whatever background knowledge is available to the model builder, whether theoretical or empirical. It may also involve some guesswork. The main goal is to identify the basic components and processes that will need to be represented in the computer simulation model. For example, for the climate system, the descriptive model might include such components as the ocean, the atmosphere, the cryosphere, and the biosphere, and such processes as (among others) evaporation, freezing/melting of ice sheets, precipitation, the carbon cycle, and radiative transfer (see Chapter 3). In addition, the model might include some rough characterization of which components and processes influence which others, and in which ways. The descriptive model often is presented as a diagram, or perhaps in some very simple mathematical form, indicating on which other variables a particular variable is likely to depend.

(2) The descriptive model will serve as an important preliminary guide in the construction of a *dynamical model* of the system of interest. This is a type of mathematical model that consists of a set of mathematical relations (usually equations) that are supposed to describe the time evolution of the system of interest. As with the

---

[13] Note that this series of models is not intended to describe what goes on in every modeling situation; for example, in some cases, modeling activity begins from an already existing dynamical model.

descriptive model, the dynamical model may incorporate theoretical or empirical information as well as guesswork—there may be aspects or processes that simply are not well enough understood for them to be confidently characterized using our background knowledge but that are thought to be important enough that some attempt at representing them must be made. In the dynamical model, however, this information will be presented in the form of equations or other mathematical relations.

(3) Third, the dynamical model is translated into a *numerical model*, a set of mathematical relations amenable to solution by numerical computation. Accomplishing this may require, as Winsberg suggests, the use of *ad hoc* modeling assumptions, e.g. in the form of simplification, approximation, or reduction of degrees of freedom. These *ad hoc* modeling assumptions may be necessary in order to render tractable the equations of the dynamical model. Accomplishing tractablility may also require consideration of the initial and boundary conditions as well as the particular parameter values that might be used.

(4) Then, the numerical model is translated into a *computer simulation model*, which is the computer code that is programmed into and executed on some computer (in conjunction with the initial, boundary and parameter values just mentioned). The computer simulation model indicates how the equations of the numerical model will be solved, in what order, to what degree of accuracy, etc. The execution of this program on a computer in conjunction with a set of initial, boundary, and parameter values (as needed) is called a "computer simulation" or a "run" of the computer simulation model.

(5) Finally, the output from a run of the computer simulation model is a *time-series model* of the system of interest under the specified initial and boundary conditions and in terms of the chosen variables. This is a data set that can be further manipulated for display, analysis and interpretation.

This, I claim, is a better characterization of the typical sequence of models that is employed in the simulation of complex physical systems. Unlike Winsberg's account, it (i) does not require that the system under consideration be very well understood from a theoretical point of view, (ii) recognizes that there is an important difference between numerical models and computer simulation models, (iii) identifies which model in the sequence of models is the computer simulation model (and thus what ought to be the referent of claims concerning the quality, etc. of computer simulation models), and (iv) maintains a distinction between the immediate product of the simulation (a time-series model) and the larger framework into which the simulation results are incorporated (Winsberg's model of the phenomena).

2.6 Conclusions

Models are entities that represent other entities in virtue of there being *perceived relevant similarities* between the two entities. The *purposes* for which a model is to be used inform decisions about the respects in which and the degrees to which the model should be similar to that which it models; the relevance of similarities is purpose-dependent. A *simulation model* is a dynamic model whose temporal evolution is supposed to be similar on some level of description to the temporal evolution of the system that it models, and a *computer simulation model* is a type of simulation model whose temporal evolution is brought about using a computer. In the practice of computer simulation of complex physical systems, scientists often make use of a sequence of models of different types, one of which is a computer simulation model.

3

Climate Change and Climate Modeling

3.1 Introduction

This chapter provides a background discussion of climate science and climate modeling. In Section 2, I introduce the notion of climate and discuss in general terms the types of influences that can affect Earth's climate. The Earth's climate system is typically characterized in terms of several interacting component systems, which are described in Section 3. In Section 4, I give a sketch of the development of climate modeling and then describe the collection of climate models currently in use. I indicate some differences among so-called "simple" and "complex" climate models and provide examples of these types of models. Section 5 makes evident the central role played by climate models in the study of climate change. In Section 6, a summary of important information from the chapter is given.

3.2 Climate and the Earth's radiative balance

Climate can be characterized in at least two ways. Traditionally, "climate" has referred to average weather conditions over some period of time, typically on the order of decades or centuries. The average conditions of interest include both mean values (e.g. annual mean temperature) and the frequency and magnitude of deviations from the mean values. In other words, describing the climate of a region involves describing the statistical characteristics of its weather conditions. A more recent characterization of climate focuses on the climate system as a whole, rather than on weather conditions alone; on this new characterization, "climate" refers to the statistical characteristics of the state of the climate system. The climate system will be discussed in more detail below.

The ultimate source of energy driving the Earth's weather and climate is the sun, which provides energy in the form of radiation. When the Earth's climate is stable, the energy received from the sun is balanced (over some period of time, such as a year) by the radiation energy emitted back to space by the Earth. Anything that can alter either the average amount of solar energy absorbed or the average amount of terrestrial energy emitted can upset this "radiative balance" and lead to a change in the Earth's climate.

This can happen as a result of changes in a variety of factors that are considered "external" to the climate system itself. For instance, variations in the output of the sun and variations in the orbital characteristics of the Earth affect how much radiation is received on average at the top of the atmosphere. Likewise, volcanic eruptions and movement of the continents can alter the amount of radiation absorbed and emitted by the Earth's atmosphere and surface. Human activities are also included in this category of external climate forcing factors. The emission of carbon dioxide and changes in land usage (e.g. clearing of forest for city expansion) are examples. In addition to climate change due to external forcing factors, there can be changes within the climate system itself that alter the radiation budget. To understand how this might occur, it will be helpful first to examine the basic components and processes that make up the Earth's climate system.

3.3 The climate system

The Earth's climate system is typically characterized in terms of several component systems, including the atmosphere, hydrosphere, cryosphere, land surface, and biosphere (see IPCC 2001, 87). These complicated component systems interact with one another in a myriad of ways, resulting in a very complex climate system. The component systems and some of their interactions are described briefly below. [14]

3.3.1 Atmosphere

The atmosphere is the layer of air that surrounds the Earth. It is here that the infamous greenhouse effect occurs. The greenhouse effect is a consequence of the radiative properties of the gases found in the Earth's atmosphere. In particular, some of these gases are relatively transparent to solar (shortwave) radiation but relatively opaque to terrestrial (longwave) radiation. These "greenhouse gases" allow solar radiation to pass through the atmosphere and heat the Earth's surface, but they absorb and then re-emit (in all directions) some of the infrared radiation that is emitted from the Earth's surface. The

[14] The discussion of non-atmosphere components of the climate system draws on that given in IPCC 2001, 87-89.

net effect is a warming of the Earth's surface due to the re-emission of infrared radiation back toward the surface. Because of this "greenhouse effect," the Earth is maintained at an average temperature that is about 35 K warmer than it would be without the presence of the atmosphere.[15] The most important greenhouse gases are water vapor, carbon dioxide, methane, nitrous oxide and tropospheric[16] ozone. These occur only in trace amounts in the atmosphere, but they exert a strong influence on the near-surface climate of the Earth. It is important to realize that this greenhouse effect is a natural consequence of the basic constitution of the Earth's atmosphere. Without it, the Earth would be a much colder place.

Current concern about global climate change focuses on the possible *enhancement* of the natural greenhouse effect as a consequence of human activities. If human activities result in increased concentrations of greenhouse gases in the atmosphere, then it is possible that even more of the infrared radiation emitted from the Earth's surface will be re-emitted back toward the surface, resulting in further warming near the surface.[17] There is observational evidence that atmospheric concentrations of some greenhouse gases, including carbon dioxide, have increased dramatically since pre-industrial times, in large part due to the burning of fossil fuels (see IPCC 2001, Chapter 3). However, it is not a foregone conclusion that increased concentrations of these greenhouse gases will result in a substantial "global warming" of the Earth's near-surface climate, because the climate system is so complicated, involving many processes that interact and exhibit strong feedback behavior.

A feedback occurs when the effects of some process come to have an influence on the continued occurrence of the process itself, i.e. when the effects of a process "feed back" into that process. Feedbacks are described as negative or positive, according to

---

[15] Note that "greenhouse effect" is a misnomer; the heating in a greenhouse occurs primarily due to suppressed convection, while the heating near the Earth's surface due to "greenhouse gases" occurs because of the re-emission of terrestrial radiation back toward the Earth. It is not correct that "the glass panes of a greenhouse function in this manner exactly analogous to the atmosphere in maintaining high greenhouse temperatures" (Huschke 1959, 261), though this misconception seems to be the origin of the name.

[16] The troposphere is the lowest layer of the Earth's atmosphere. It extends from the surface to an average height of about 10 kilometers, although it extends higher in the tropics and not as high in the polar regions. The layer above the troposphere is the stratosphere.

[17] Changes in the land surface due to human activities also might contribute to climate change, but this is not currently the focus of research on climate change.

whether they weaken or strengthen the process (respectively). Before giving examples of feedbacks in the atmosphere, it is useful to consider examples from everyday life. For example, in winter, a furnace thermostat is part of a negative feedback process. Whenever the temperature in the house goes above or below the temperature at which the thermostat is set, the furnace turns off or on (respectively) to help bring the room temperature back to the desired temperature. Positive feedback occurs when a microphone is held too close to a speaker that is transmitting the signal from the microphone. An amplification cycle occurs: the microphone initially picks up some weak signal (e.g. a bit of background noise) which is amplified and emitted from the speaker; but the microphone is close to the speaker and so picks up its own amplified signal, which is then further amplified and emitted from the speaker, and so on. The result is a runaway process that produces an increasingly loud signal from the speaker.

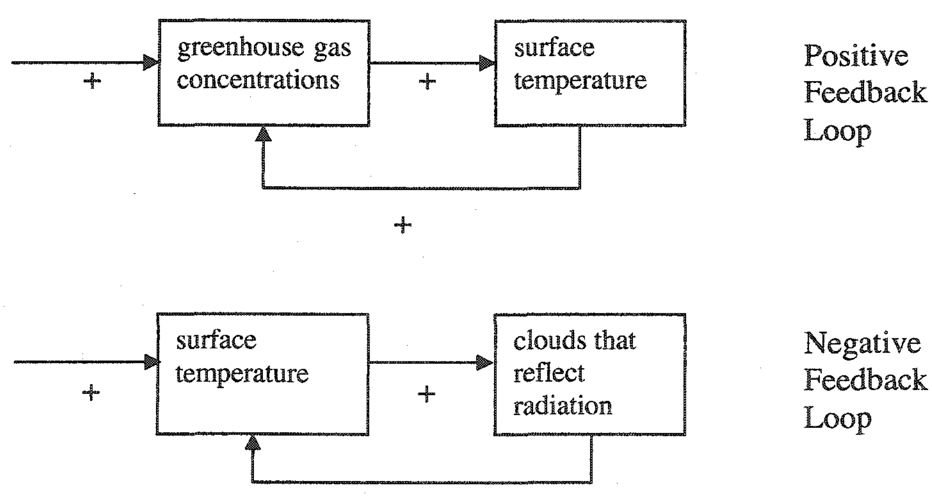

Figure 1. Examples of feedbacks within the climate system

Feedbacks can occur in the atmosphere as well. For instance, if increased greenhouse gas concentrations result in warming that in turn results in an increase in the amount of water vapor present in the atmosphere (e.g. due to increased evaporation), then the presence of more water vapor (a greenhouse gas) can further enhance the warming, leading to even further increases in water vapor concentration, and so on. This kind of process exhibits a positive feedback (see Figure 1). There is some evidence that this kind of water vapor feedback does occur and that the way that it is represented in climate models strongly affects the amount of warming predicted to result from increased greenhouse gas concentrations, in part because it can amplify other feedback processes in the climate system (see IPCC 2001, 423-427). By contrast, if increased warming near the surface leads to the formation of many more clouds that reflect more incoming solar radiation and thereby prevent it from reaching the surface, then the initial warming will be checked. This kind of process exhibits a negative feedback (see Figure 1). There is currently a great deal of uncertainly concerning cloud feedback processes like this one. Part of the complication comes from the fact that, while some clouds have an overall cooling effect, other types of clouds are thought to actually contribute to further warming. At present, it is not clear whether any changes in cloud development due to increased greenhouse gas concentrations will have a net warming or cooling effect (see IPCC 2001, 427-431). In general, the existence of feedback processes like these and the uncertainties surrounding their strengths contribute to the difficulty that scientists have in determining what will be the ultimate effects of increased concentrations of greenhouse gases.

There are many other atmospheric processes that play a role in shaping the Earth's climate on a regional and global scale, but there is not enough room to discuss them in detail now. In general terms, they are processes that transport energy within the climate system or lead to a change in the amount or distribution of radiatively active gases or particulate matter within the atmosphere. Some of these processes are associated with the large-scale circulation of the atmosphere, while others occur on smaller scales, as in the case of the formation of clouds and precipitation. The formation of clouds is an especially important process, since clouds can exert strong cooling and warming effects on climate by reflecting, absorbing and emitting radiation. These smaller-scale processes pose some difficulty for climate modelers, as will be explained below. In addition, there are processes that involve interaction of the atmosphere with other components of the climate system. Some of these will be mentioned as the other components are examined.

3.3.2 Hydrosphere

The hydrosphere consists of liquid water at and below the Earth's surface, including oceans, rivers, lakes and aquifers (IPCC 2001, 88). The oceans play an especially important role in shaping climate. They store and transport a huge amount of energy. Because of the large heat capacity of water, the oceans change temperature much more slowly than the atmosphere above, acting as a moderating force in the face of rapid changes in atmospheric conditions. The oceans also serve as a sink for atmospheric carbon dioxide, which is stored in the oceans in very large amounts. This is just one of the many ways in which the oceans interact with the atmosphere. For instance, water is evaporated from the ocean to the atmosphere and precipitated back to the ocean from the atmosphere, and winds exert a force on the surface of the ocean, affecting its circulation.

3.3.3 Cryosphere

The cryosphere is comprised of all of the frozen water at and below the surface of the Earth. This includes glaciers, ice sheets, sea ice, snow (on land), and permafrost (IPCC 2001, 88). This frozen water influences the radiation budget of the Earth in several ways, most obviously by reflecting a great deal of incoming solar radiation; typically, snow/ice surfaces reflect much more solar radiation than would be reflected by the underlying surface (land or sea) in the absence of the snow/ice. The cryosphere also interacts with the atmosphere and hydrosphere. For instance, it plays a role in driving deep ocean water circulation, and it affects exchanges of mass and momentum between the atmosphere and the ocean (IPCC 2001, 88-9).

3.3.4 Land surface

The land surface is another important part of the climate system. Several of the ways in which the land surface influences climate are discussed in IPCC 2001. For instance, the land surface reflects and absorbs solar radiation and emits infrared radiation. In addition, evaporation from the land surface transports water vapor and heat to the atmosphere. The land surface also interacts with the atmosphere via frictional drag on air movements. Dust and particulate matter blown from the land surface into the atmosphere can affect the transfer of radiation in the atmosphere. A multi-component interaction involving the land surface can be seen in the following process: precipitation (atmosphere) carries soil and materials from the land surface to rivers and lakes (hydrosphere).

3.3.5 Biosphere

The biosphere is comprised of living matter at and below the surface of the Earth. Like the land surface, the biosphere reflects, absorbs and emits radiation. In addition, evapotranspiration from plants transports substantial amounts of water vapor and heat to the atmosphere. Plants also store large amounts of carbon dioxide, some of which is returned to the atmosphere when they decay (e.g. in the clearing of forests). Biospheric processes are only now beginning to be represented in climate models.

3.4 Climate modeling

Although only a few of the interactions among the component systems were listed in the previous section, it is easy to see that the components are strongly interactive. Exchanges of heat, momentum and mass constantly occur among the component systems. As IPCC 2001 points out, there is “a virtually inexhaustible list of complex interactions some of which are poorly known or perhaps even unknown” (89). This means that climate modelers face a substantial challenge in trying to develop computer simulation models of the Earth’s climate system.

3.4.1 A sketch of the development of climate modeling

Numerical simulation of the climate system grew directly out of efforts to numerically forecast the weather, which began in the early 1950s with the development of the first electronic (and digital) computers. In March and April of 1950, a series of numerical weather calculations were made using one of the first electronic computers, the ENIAC. These forecasts were the culmination of several years of work by a group of researchers gathered at Princeton’s Institute for Advanced Study as part of the “Meteorology Project,” a government-funded research project proposed by John von Neumann. He saw in the complicated, nonlinear problems of dynamic meteorology an opportunity to demonstrate the great potential of the electronic computer. The proposal for the project, which von Neumann submitted to the Navy Office of Research and Inventions in May 1946, identified as its objective: “the investigation of the theory of dynamic meteorology in order to make it accessible to high speed, electronic, digital, automatic computing, of a type which is beginning to be available, and which is likely to be increasingly available in the future" (see Thompson 1983/1990, 106).

Over the next several years, a number of numerical weather forecasting models were developed and tested (see Edwards 2000 for more details). With the reasonable success of these relatively simplified models in simulating large-scale atmospheric motions over limited spatial domains, it was not long before interest grew in simulating the global circulation of the atmosphere. A 2-layer hemispheric model developed in 1955 is generally considered to be the first general circulation model (see Phillips 1956). Over the next decade or so, modelers worked to develop improved atmospheric general circulation models—among other things, they replaced some very simplified equations with more realistic ones, increased the number of vertical levels for which the models made predictions, and modeled the atmosphere in terms of three spatial dimensions (rather than just one or two). It was not really until the 1970s that climate modelers began to represent in any detail more than just the atmospheric component of the climate system. They then began to develop models of other components of the climate system and to couple these other models to the already existing atmospheric general circulation models. By the early 1990s, general circulation models (also sometimes called "global climate models") had been coupled to ocean and sea-ice models and land surface models, and thus were beginning to include some representations of most of the major components of the climate system (as identified above). Present day climate models often also include some representation of aerosols and the carbon cycle; dynamic models of vegetation are not yet generally incorporated (see IPCC 2001, 48). Because they are now becoming so comprehensive, they are sometimes referred to as "climate system models." While this sketch focuses mainly on the development of general circulation models/climate system models (or "complex models" as they will be considered below), simplified models of the climate system were also being developed throughout the latter half of the 20th century.

3.4.2 Present-day climate modeling

The study of weather and climate currently involves the use of a collection of computer simulation models of differing complexity. Reference is often made to simple models, models of intermediate complexity, and complex models, as if there were three well-defined and distinct model types. In reality, there are no hard and fast lines dividing such categories of models. In fact, different authors mean different things when they speak of these model categories, especially when they speak of simple models. Some authors mean to refer to the simplest of models, while other authors mean to refer to models that are *relatively* simple when compared to the most complex models; the latter "simple" models may still be rather complex. Still, models that are among the simplest do differ in systematic ways from those that are among the most complex. These differences will be described below, along with examples of simple and complex models. Models of intermediate complexity will not be discussed; as one would suspect, their characteristics fall in between those described below for simple and complex models.

3.4.2.1 Characteristics of simple and complex models

Simple and complex models of weather and climate typically differ from one another in one or more of the following ways.

Spatial dimensions

Simple models often describe the atmosphere or climate system in terms of only one or two spatial dimensions. For instance, simple climate models might represent the climate system in terms of a single, multi-layer column of ocean and atmosphere that can be construed as representing the average state of the climate system across all latitudes and longitudes. By contrast, complex models typically represent the atmosphere or climate system in three spatial dimensions. Unlike simple models, they do not average over (or omit) one or more spatial dimensions.

Spatiotemporal resolution

Simple models have relatively coarse spatial and temporal resolutions. For example, a simple climate model might give calculations for only a few points (or even a single point) intended to represent the average conditions over all land mass in a hemisphere. On the other hand, complex models have relatively fine spatial and temporal resolutions.[18] Typical horizontal spatial resolution (the represented distance between grid

[18] I say "relatively fine" because there are models that have much finer resolution. For instance, models of atmospheric convection might have a spatial resolution on the order of tens of meters. Among models intended to simulate the global climate system, however, the resolution of general circulation models/climate system models is the finest—certainly much finer than that of simple climate models.

points) in complex climate models is on the order of a couple hundred kilometers. In the vertical, the resolution in the ocean part of a complex climate model is on the order of a few hundred meters, while in the atmosphere part of a complex climate model the vertical resolution is on the order of one kilometer (except for levels very near the surface in the ocean and the atmosphere, which have finer resolution). The equations in the climate model are usually solved in time steps on the order of a half hour, meaning that predictions of the state of the climate system are made for every half hour of its evolution (i.e. half-hour "snapshots"). The spatial and temporal resolution is limited by available computing power. Increasing the spatial resolution of the model dramatically increases the number of grid points and hence the number of calculations that must be done as part of a simulation.

Comprehensiveness

Simple models represent only a subset of the processes that are represented in complex models. In other words, they are less comprehensive than complex models. For example, many simple climate models do not represent sea ice, while many complex models do. Likewise, simple early numerical weather prediction models did not have any representation of precipitation, while today's complex models do. Some of the very simplest climate models in use represent only a handful of processes (or even fewer). These models are designed to isolate particular processes thought to be important in the functioning of the larger climate system or atmosphere. By contrast, complex models represent many components and processes thought to be important in shaping the behavior of the atmosphere or climate system. These models explicitly calculate the large-scale dynamical evolution of the momentum, heat and moisture/salinity fields in the atmosphere (and in the ocean in climate models). Complex climate models also include detailed representations of land surface processes.

Level of parameterization

Many of the processes that simple models do represent are parameterized, because of the models' reduced spatial dimension and coarse spatial scale. The notion of a parameterization needs some explication. If a process occurs on a spatial and temporal scale smaller than that for which the model performs calculations, then that process cannot be explicitly represented in the model, i.e. the detailed occurrence of the process cannot be represented. If this "sub-grid" process influences what happens on the larger scale for which model calculations are performed, then it will be important to represent the effects of that sub-grid process in some way. Often this is done using a parameterization—a relationship between the time-averaged or area-averaged effects of the sub-grid process and one or more of the larger-scale variables in the model. These parameterizations can be physically based, drawing on theoretical understanding of the relationship between the unresolved (i.e. sub-grid) and resolved (larger-scale) quantities. Parameterizations also can be constructed with the help of observed empirical relationships between the unresolved and resolved quantities. In some cases, they also incorporate guesswork. An example of a parameterization in the simple climate model to be described below is that for the thermohaline circulation, an important circulation in the ocean. Instead of explicitly calculating the dynamics of this circulation, the circulation is represented using a single parameter that specifies the rate of upwelling of cold water as a simple function of the change in temperature of the ocean mixed layer.[19]

It is important to recognize that nearly all models—simple or complex—currently incorporate parameterizations, since there are always some processes that occur on scales smaller than those resolved by the model. For example, in complex climate models parameterizations are often employed in representing (among other things) the radiative effects of greenhouse gases, the microphysical and radiative properties of clouds, the occurrence and effects of moist atmospheric convection, and the transport of momentum and moisture near the Earth's surface. So, as in simple models, parameterization is needed. The difference is that complex models resolve some processes that are not resolved in simple models, so that parameterizations in complex models need to be used only for processes that occur at a "deeper" or smaller-scale level in the system.[20]

3.4.2.2 Examples of simple and complex models

For brevity's sake, I will only present examples of simple and complex climate models (not weather forecasting models); these should sufficiently illustrate the kinds of differences described in the last section.

[19] The ocean mixed layer is the layer nearest to the surface; its thickness is on the order of 50 meters.
[20] In principle, this is the case, but limitations on computing time sometimes create the need for parameterizations of some relatively large-scale processes in complex climate models, too.

A Simple Climate Model

As an example of a simple climate model, consider an upwelling-diffusion energy balance model (UD/EBM) of the type described in IPCC 1996 and 1997 and Raper et al. 2001 (see Figure 2 for a schematic illustration). This model treats the atmosphere as a single well-mixed volume (a box). It does not explicitly calculate any dynamics for the atmosphere—everything that happens there is parameterized. The same is true for the land surface, which is represented by two boxes—one each for the Northern and Southern Hemispheres. Two multi-layer vertical columns represent the oceans, again one for each of the two hemispheres.

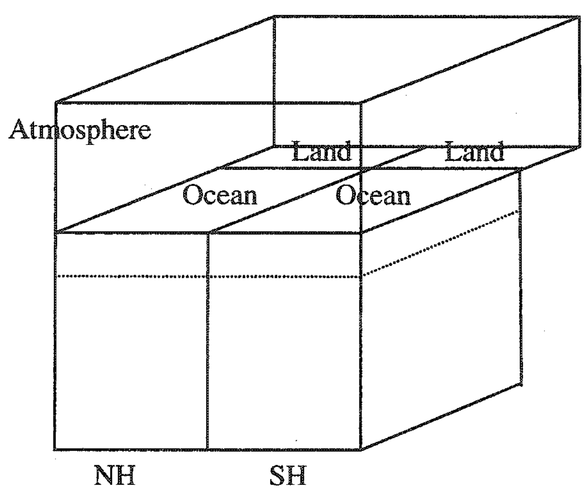

Figure 2. A schematic illustration of a simple UD/EBM climate system.

The uppermost layer of each ocean column is the mixed layer, whose thickness is specified as a model parameter. In addition, there are parameters representing the vertical transfer of heat within the one-dimensional ocean column; downward transfer of heat from the warm ocean surface is represented as a diffusion process, while the rate of upward transfer of cold water from near the ocean bottom is represented as a simple function of the change in the temperature of the ocean mixed layer. (In the real ocean, relatively cold polar near-surface water sinks, leading to an upwelling of cold bottom water in lower latitudes.) Other important parameters in the model include the climate sensitivity parameter, which specifies the equilibrium global-mean temperature change for a doubling of carbon dioxide concentration, and a parameter that specifies how the temperature of the cold water sinking near the poles changes when the global mean temperature changes. Previously, simple models like these were used as stand-alone tools for investigating the climate system and climate change; now, they are also being used as malleable tools whose parameters can be changed so that a particular run of the model produces output that approximately matches that of a particular complex model.

A Complex Climate Model

The complex climate model to be described is the National Center for Atmospheric Research (NCAR) Community Climate System Model (CCSM), version 2.0, which was released in May 2002. This type of model is among the most comprehensive and complex types of models currently used to study the Earth's climate. CCSM will be discussed in rather limited detail here, since the details can become rather tedious; the interested reader is referred to the Appendix for a slightly expanded description. The CCSM incorporates several distinct sub-models that correspond roughly to four of the components of the climate system discussed in above: an atmospheric general circulation model (GCM), an oceanic GCM, a land surface model, and a sea-ice model. These component models communicate and interact via a "flux coupler" that passes information among the components, calculates some of the fluxes of energy and moisture to be exchanged, and coordinates the spatiotemporal evolution of the model components, some of which incorporate longer time steps or finer spatial scales than others (Boville and Gent 1998, 1115). The CCSM thus includes four types of component models plus the flux coupler (see Figure 3).[21]

[21] This modular configuration allows for different component models (e.g. land surface models) to be "plugged in" to the larger model structure without changing the rest of the model (Boville and Gent 1998, 1116). This is a nice feature, since it allows one to investigate the effects of alternative models of some component system without having to make compatibility adjustments in all other component models—any necessary adjustments can be made in the flux coupler.

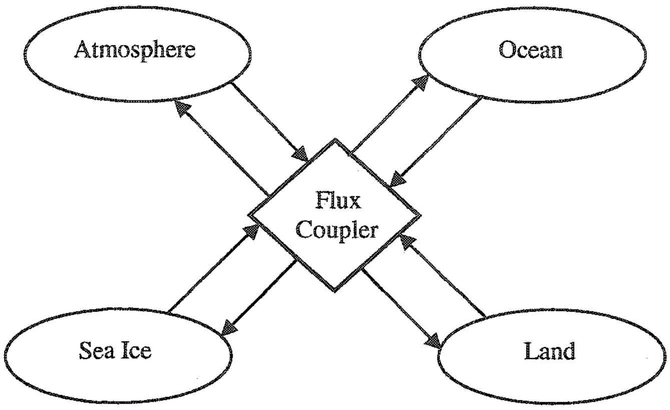

Figure 3. The major component submodels of the NCAR CCSM 2.0.

Both the atmosphere and ocean submodels are primitive equation dynamical models, which means that they rely on fundamental physical equations in calculating the large-scale dynamical evolution of the atmosphere and ocean.[22] The atmosphere submodel also includes representations of the transfer of radiation within the atmosphere (whether cloudy or clear), the formation of clouds and precipitation, and the complicated evolution of the momentum, heat and moisture fields near the Earth's surface (in the boundary layer), where turbulence plays an important role. The land surface submodel specifies different surface types (e.g. wetlands, lake, desert) for each grid cell in contact with the overlying atmosphere. These surface types define vegetation and soil types whose properties, along with information from the atmospheric model, are used to determine fluxes of energy, momentum, moisture and carbon dioxide to the atmosphere. The sea ice submodel represents sea ice in terms of sea ice area, sea ice volume, sea ice internal energy, snow volume, surface temperature of snow/ice, sea ice velocity, and stress components (Briegleb et al. 2002, 6). These state variables evolve during the

[22] These fundamental equations for the atmosphere typically include: horizontal equations of motion, the hydrostatic law (which defines pressure change with height), a mass continuity equation, the first law of thermodynamics, the equation of state for air, and a balance equation for water vapor (see Peixoto and Oort 1992, 454-455). The equations for the ocean are similar in type to those for the atmosphere, although the water vapor balance equation is replaced by a balance equation for salinity (see Peixoto and Oort 1992).

simulation (in part based on information received about the changing oceanic and atmospheric conditions), which in turn results in changes in the fluxes of momentum, heat, and moisture to the ocean and atmosphere. Each of these component models is itself rather complex, and the collection of component models, along with the flux coupler, together constitute the climate system model.

3.5 The role of climate models in the study of climate change

Climate models play a central role in the study of climate change. They are used not only to make projections of future climate changes, but also to determine whether any climate change is currently occurring and, if it is occurring, what might be the cause of this change. In addition, climate models are used to conduct a variety of computer simulation experiments designed to help answer basic questions about the climate system. Without climate models, the study of climate change would be severely limited.

3.5.1 Estimating internal variability and detecting climate change

Even in the absence of human activities, the average state of the climate system varies somewhat from year to year. This *natural variability* can result from factors internal or external to the climate system itself. Fluctuations can result from natural cycles in the climate system, such as El Nino-Southern Oscillation (ENSO), whose periods are longer than one year. There are also small random variations internal to the climate system that can contribute to year-to-year variability. Natural variability that results from internal fluctuations is referred to as *internal variability*. Natural external factors can also contribute to natural variability. As discussed above, these external factors include variations in solar output and the occurrence of volcanic eruptions. Both of these external factors can alter the amount of shortwave radiation received at the Earth's surface.[23]

Climate change can occur when there are changes in natural external factors and/or when there are non-natural (i.e. human-caused) activities affecting the climate system. Detecting such climate change requires establishing that observed recent climate

[23] Variation in the Earth's orbital parameters is an additional external factor that contributes to natural variability, but since this variation occurs on very long time scales, it typically is not considered relevant to recent fluctuations in climate (i.e. those in the last two centuries).

trends are unusual (in a statistical sense) relative to what is expected due to internal variability alone. Thus, in order to determine whether observed recent climate trends are unusual, it is necessary to have some estimate of the internal variability of the Earth's climate system.

Unfortunately, this information is difficult to obtain. It might seem that observations are the obvious source for this information. However, there are several problems with using observations to determine internal variability. First, the direct observations that are available are limited in spatiotemporal coverage and in the physical quantities that have been measured. They are available primarily for land locations and for conditions very near the Earth's surface, and they have been made systematically over wide areas only in the last century or so. Second, these observations provide information not about internal variability alone, but about internal variability in combination with both natural variability due to external sources and variations due to human activities. Reconstructions of paleoclimatic[24] conditions from proxy indicators, such as tree rings and ice cores, might seem a more promising source of information about internal variability, since these do not include the effects of human industrial activities (assuming that they are for periods prior to ~1800). However, variability due to natural external factors is still present in paleoclimate data, so even these data are not ideal. In addition, although some progress has been made in improving the quality of paleoclimate reconstructions, there are still substantial uncertainties involved, and spatial coverage is limited. Thus observations, whether direct or reconstructed, provide only limited assistance in estimating global features of the internal variability of the climate system.

As a consequence, climate models have become the main source of information about internal variability. They are used to carry out simulations of the evolution of the climate system in the absence of human-related factors and in the absence of changes in solar output and volcanic eruptions. These simulations are carried out for hundreds of model-years, and the statistical properties of the simulated global climate conditions (on time-scales from decades to centuries) are then used as an estimate of the internal variability of the real climate system. Such simulations have been carried out using

[24] Paleoclimatic conditions are those that occurred a long time ago; depending on the research context, paleoclimatic conditions might refer to those occurring anytime from billions of years ago to a few hundred years ago.

several different complex climate models. From these, estimates of the variation in global mean temperature due to internal variability have been extracted.

Estimates of internal variability differ somewhat from model to model, but overall they compare reasonably well to the estimates of internal variability that have been obtained from observations (IPCC 2001, 702). Based on these model-derived estimates of internal variability, it is currently judged to be unlikely that recent observed temperature changes are due to internal variability alone, since the observed trend in near-surface temperature (since ~1850) is statistically unusual even when the various model estimates of variability are increased by a factor of two or more (IPCC 2001, 730).

3.5.2 Attribution of climate change

Attribution of climate change to some set of causes requires showing that observed changes are (1) consistent with the expected response of the climate to those causes and (2) not consistent with the response expected as a result of any other physically plausible combination of causes (IPCC 2001, 700). There are numerous forcing factors—factors that can affect the Earth's radiation budget—that might have contributed to the observed recent climate changes, but current studies only consider the factors thought most likely to have had a substantial impact. These forcing factors include increased concentrations of greenhouse gases, increased concentrations of sulfate aerosols, changes in solar output, volcanic eruptions, and (sometimes) changes in concentrations of tropospheric and stratospheric ozone.[25]

The process of attribution relies on complex climate models for estimates of the expected response of the climate system to various combinations of these forcing factors. Computer simulations are carried out in which these forcing factors are allowed to act alone or in combination, and then characteristics of the simulated response are compared to those of the observed climate.[26] Various characteristics can be compared. A standard comparison is that between simulated and observed global mean near-surface temperature

---

[25] Greenhouse gases and sulfate aerosols are anthropogenic forcing factors resulting primarily from the burning of fossil fuels. Ozone is affected by both human and solar influences, although it is not emitted directly. Solar output changes and volcanic eruptions are natural but external forcing factors.

[26] Ideally, the model climate system would respond uniquely to different forcing factors, so that each forcing factor would bring about unique spatiotemporal changes in the temperature fields, precipitation fields, etc. That is, ideally each forcing factor would have its own "signature" or "fingerprint" in the simulations. In fact, this is not the case—some forcing factors share certain response characteristics, and this "degeneracy" complicates the attribution process to some extent.

(annual, decadal, or multi-decadal). More sophisticated comparisons focus on spatial and temporal patterns in near-surface temperature changes. Similar comparisons are performed for changes in temperature in the upper troposphere and stratosphere. The methods used to carry out these comparisons range from simple overlaid time-series plots of the simulated and observed temperatures (where one can easily "eyeball" the similarity) to rather complicated and abstract statistical techniques that focus on the aspects of the simulated temperature response that differ most markedly from fluctuations expected due to internal variability, i.e. which try to maximize signal-to-noise ratios (see IPCC 2001, Chapter 12 for details).

Although complex climate models differ somewhat in their simulated responses to the various forcing factors, a number of conclusions have been drawn by the IPCC regarding attribution of climate change. Relevant to (2) above, the modeling studies have found that simulations involving only natural external forcing factors produce results that are consistent neither with late 20th century observed changes in the vertical structure of atmospheric temperature nor with late 20th century observed changes in near-surface temperature, although there is some evidence that natural factors may have played a role in warming that occurred earlier in the 20th century (IPCC 2001, 730-731). Relevant to (1) above, it has been concluded that simulations involving anthropogenic forcing factors including greenhouse gases and aerosols produce response patterns that are detectable in the 20th century near-surface temperature record (IPCC 2001, 730). Again relevant to (1), the best fit to observed global mean temperature changes seems to result from simulations involving both natural and anthropogenic forcing factors (IPCC 2001, 699). The most recent summary conclusion drawn by the IPCC, based on these and many other modeling results, is that "most of the observed warming over the last 50 years is likely to have been due to the increase in greenhouse gas concentrations" (IPCC 2001, 699).[27]

[27] One modeling result that has been called upon in order to challenge this conclusion concerns simulated versus observed changes in the temperature of the lower troposphere during the last twenty years; simulations suggest that this region should have warmed faster than the near-surface region, but observations show the opposite. This discrepancy is sometimes cited as evidence of the inadequacy of complex climate models in order to call into question other results obtained using these models. The IPCC, while admitting that substantial uncertainties still exist in multiple aspects of the detection and attribution process, does not seem to think that this discrepancy undermines its results and conclusions.

3.5.3 Projections of future climate

Both simple and complex climate models play a role in estimating future climatic conditions. Environmental policymakers (and citizens) want to know what kind of climate changes are likely to occur under various greenhouse gas emissions scenarios, so that informed policy decisions can be made. To address this matter, a large number of possible emissions scenarios have been developed. These are intended to reflect greenhouse gas emissions under various social, economic and political conditions that might obtain in the next century. Climate models are used to estimate what kinds of climate changes would be likely to occur under each scenario. These estimates are referred to as "projections" rather than "predictions" to reflect the fact that they are based on possible future greenhouse gas emissions, which are not necessarily the emissions that will actually occur.

Until recently, only a very few projections of future climate had been carried out using complex climate models, primarily because of the substantial computing resources required. In the last few years, however, with further model development and greater computing resources, there has been an organized worldwide effort to generate and compare complex climate model projections for a few different scenarios. Some of these emissions scenarios are intended to be rather realistic, in the sense that they are considered plausible or possible scenarios for the next century. For these, simulations are carried out in which the hypothesized greenhouse gas emissions are prescribed at each model time step, in order to see what kind of effects on climate emerge by the end of the 21st century. Other computer simulation experiments are conducted for more idealized greenhouse gas emissions scenarios, e.g. for a scenario in which greenhouse gas concentrations increase by 1% (compounded) per year, until the concentration has doubled. Although this is an idealized scenario, it is useful for model comparison purposes and is within the range of possible emissions scenarios, in the sense that the overall emissions increase is not wildly implausible, even if the smooth rate of increase is implausible (see IPCC 2001, 527).

Simple climate models also are used to make projections of future climate, because computing resources are still limited enough that it is not feasible to use complex models to investigate the many emissions scenarios that have been developed (~35 in the most recent set of scenarios). As a consequence, most of the emissions scenarios are investigated using simple models. Although the simple models were used as stand-alone models in the past for this purpose, the current approach is to tune a simple model so that it replicates the projections that have been made by an individual complex model. This is accomplished by adjusting the simple model parameters until the simple model results match, to some degree of accuracy, the available complex model results (typically for global mean parameters, such as global mean surface temperature).

The most recent results from these modeling experiments are available in IPCC 2001 (see especially Chapter 9). For the limited number of scenarios for which complex climate models have made projections, the estimates of global mean temperature change by the year 2100 range from about 0.9 to 4.5 degrees Celsius (IPCC 2001, 527).[28] For the wider range of scenarios for which simple climate models have made projections, the estimates of global mean temperature change in the year 2100 range from about 1.4 to 5.8 degrees Celsius (IPCC 2001, 527).[29]

3.5.4 Other uses: process and sensitivity studies

Climate models are used in other ways that are not directly part of the detection and attribution process. Two examples of this type of usage can be found in the modeling of individual processes and in studies investigating the sensitivity of models to various changes in their formulations.

In the context of climate studies, the least comprehensive models are those that aim to represent a single process in isolation from the rest of the climate system. For example, detailed models are constructed to represent types of atmosphere-biosphere exchanges, the indirect effects of aerosols, and atmospheric convection. Why would such models be constructed? Aside from basic scientific interest, one reason is that understanding how a process works in isolation may help scientists to gain a qualitative understanding of the contribution that the process is likely to make to the overall climate. Of course, this practice is at least partially foiled if the process of interest interacts in complicated ways with other components of the climate system. Another reason for

---

[28] These values are based on 30-year averages for the period ending in the year 2100, relative to the global mean temperature for 1961-1990 (IPCC 2001, 527).

[29] These values are based on the global mean temperature at the year 2100, relative to the global mean temperature for 1990 (IPCC 2001, 527).

constructing detailed models of individual processes is that studying their behavior can help to reveal which kinds of simplified representations of the process might capture the most important features of the process. More comprehensive climate models often must incorporate these simplified representations of component processes, since computational limitations can preclude representing all processes in the most advanced and detailed way permitted by current understanding.

Sensitivity studies aim to determine the effects of changing the formulation of a climate model, whatever its degree of complexity. This can involve "turning off" some process in the model to see how the model results differ in comparison to results obtained when the process is included. Alternatively, a sensitivity study can involve changing the way that some process is represented in the model, again in order to see how this change affects the simulation results. This type of sensitivity study often is carried out to investigate the role of particular parameterizations in producing simulation results. For some processes, there may be several different parameterizations in use in the climate modeling community, and it may be unclear (from the point of view of physical theory or empirical studies) whether one of these parameterizations more accurately describes the process than the other parameterizations. It is important for scientists to determine whether simulation results are especially sensitive to the way in which some process is parameterized. For one thing, this sensitivity might turn out to explain many of the differences in results obtained from different climate models. In addition, knowing that simulation results are sensitive to a particular parameterization will direct scientists to focus their efforts on developing the best parameterization possible, which in turn will (hopefully) increase the utility of climate simulations.

Process and sensitivity studies like these are somewhat peripheral to the study of climate change in that these modeling studies typically are not designed to answer questions about detection and attribution. However, these studies play an important role in answering other questions about the functioning of the climate system and in the development of the climate models that are used in answering questions about detection and attribution, so their importance should not be overlooked.

3.6 Summary

"Climate" refers to the statistical characteristics of weather conditions, including means and deviations from means. The Earth's climate can be affected by changes internal or external to the climate system. The climate system is typically characterized in terms of several interacting component systems, including the atmosphere, hydrosphere, cryosphere, land surface, and biosphere. Modeling of the climate system grew out of numerical weather prediction modeling. The earliest climate models were atmosphere models only; in recent decades, other component systems have come to be represented in climate models. Today, there exists a collection of computer simulation models of differing complexity. Very simple models tend to differ from very complex models with respect to: spatial dimensions, spatiotemporal resolution, comprehensiveness, and level of parameterization. Climate models play a central role in the study of climate change. They are used to estimate natural variability and detect climate change, to attribute climate change, to provide projections of future climate change, and to conduct process and sensitivity studies.

4

Experimental Validity in Computer Simulation and Material Experimentation

4.1 Introduction

In the scattered pockets of philosophical work on computer simulation, one of the views most commonly expressed is that computer simulation is somehow intermediate between theory and traditional material experimentation (e.g. Galison 1997; Rohrlich 1991; Humphreys 1994; Dowling 1999; Sismondo 1999). Very recently, philosophers have begun to consider in some detail just how computer simulation does relate to theorizing and to material experimentation (see Norton and Suppe 2001, Guala 2002, Morgan 2002 and 2003, Winsberg 2003). This chapter continues in this vein, taking a closer look at the issue of experimental validity in both material experimentation (ME) and computer simulation experimentation (CSE).

CSE involves the manipulation of a computer simulation model in order to investigate the effects of that manipulation. In the study of climate and climate change, computer simulation experiments are carried out on a regular basis. For example, as discussed in the last chapter, computer simulation models are used to investigate what might happen if various climate system processes are "turned off" or if greenhouse gas concentrations change in a particular way. One tactic that is available to (and is sometimes used by) climate skeptics is to declare that it is "only a model" behind the results of these experiments. The implicit contrast is with results obtained from "real" material experiments, like those that take place in the chemistry laboratory using real chemicals and test tubes. Such material experiments are accepted by many scientists and philosophers as a standard and reliable way of finding out about the empirical world. By contrast, as the climate skeptics' approach illustrates, computer simulation experimentation frequently is viewed as a suspect methodology whose value as a source of empirical knowledge is questionable.

The ultimate goal of the discussion that follows is to combat this kind of blanket skepticism about the results of CSE. I suggest that part of the skepticism that does exist comes from a failure to appreciate important parallels in the epistemologies of material and computer simulation experiments. My strategy will be to highlight some of these parallels in order to narrow the perceived gulf between these methodologies. The key similarity explored in this chapter is as follows: in *both* types of experimentation, scientists must be concerned with both the internal and external validity of their experiments. In Section 4.2, I introduce the notions of internal and external validity as components of experimental validity. In Section 4.3, I consider the issue of internal validity in CSE and show that the strategies offered by Allan Franklin (1989) for sanctioning belief in the internal validity of the results of material experiments have analogues in CSE. I find that, with respect to internal validity, CSE may even enjoy certain advantages over ME. In Section 4.4, I turn to the topic of external validity. I maintain that in both CSE and ME arguing for the external validity of results is a matter of justifying simulation-to-target inferences. I then show by example that it is not necessarily more difficult to make good arguments for the external validity of results in CSE than for the external validity of results in ME, even though ME does have certain advantages over CSE when it comes to external validity. I conclude in Section 4.5 that CSE and ME enjoy different but limited advantages with respect to experimental validity and that the validity of the results of both types of experimentation should be judged on a case by case basis, rather than categorically.

4.2 Experimentation and experimental validity

For purposes of this discussion, a very detailed account of experimentation is unnecessary. The view of experimentation that will be adopted as a working account is as follows: an experiment is any activity that involves the manipulation of some entity or system in order to investigate its properties and behavior. Many experiments are goal-oriented—they are conducted with one or more particular questions in mind. This does not mean that experiments cannot be conducted without very specific questions in mind (i.e. "just to see what will happen"), but the present chapter is concerned with the large number of experiments that that do aim to provide particular information.[30]

A specific question that an experiment is designed to help answer will be referred to as a *target question*. An example of a target question is: How much weight will be lost on average by obese adults who daily ingest 10 milligrams of sibutramine for one year? Target questions are identified as such because they indicate the *target domains* of experiments.[31] The target domain is that about which the experiment ultimately is supposed to provide information. It includes both objects and conditions. The target domain might consist of a particular system, such as the Earth's atmosphere, subjected to some conditions of interest (e.g. CO_2 concentrations set to 400 parts per million). Or it might consist of a type of object, such as this make of car, under some conditions of interest (e.g. wet pavement). Or it might consist of a group of entities under some conditions of interest, as in the case of the target question above. In that example, the target domain consists of all obese adults who will take sibutramine as described. The target domain of an experiment must be distinguished from the *experimental domain*, which is made up of the particular entities and conditions that are actually manipulated as part of the experiment. In the weight loss example, the experimental domain might be particular obese adults who participate in a study by ingesting sibutramine (e.g. 1000 such adults).[32]

Making use of this vocabulary, a *material experiment* can be described as one whose experimental domain consists of material entities subjected to particular material conditions. Similarly, a *computer simulation experiment* is an experiment whose experimental domain is composed of one or more computer simulation models. Computer simulation models can be manipulated either by making changes to their constitutive parameters and relations or, perhaps more commonly, by altering their initial and boundary conditions. Much of the activity of simulation modeling involves this kind of manipulation of the models in order to investigate their properties and behavior. This

[30] Experiments that are not designed to help answer very particular questions have been called "exploratory" experiments. See Steinle (1997) and Ribe and Steinle (2002).
[31] Guala (2002) adopts similar terminology: he uses "system" where I use "domain."
[32] Such an experiment likely would involve a control group of obese adults as well. The adults in the control group would not be part of the experimental domain (because they would not ingest sibutramine), but they would be part of the experiment.

kind of activity qualifies as a form of experimentation according to the working account described above and is what is meant by "computer simulation experiment" in this discussion.[33]

Of utmost importance in goal-oriented experimentation is that the experimental results be "valid." This validity is not the kind spoken of in the context of deductive argumentation; rather, it concerns the rationale for drawing inferences about a target domain. Following Campbell and Stanley (1963), *experimental validity* can be broken down into two sub-species of validity, which are individually necessary and jointly sufficient for a goal-oriented experimental result to be deemed valid.

Internal validity ultimately concerns whether the experimental outcomes of interest can be attributed to the manipulation that the experiment was designed to investigate, e.g. whether giving a drug to mice in the treatment group was a cause of their recovery from illness. The question of internal validity is the question of whether something went wrong in the execution of the experiment that allowed other confounding factors to influence the results of interest. This can happen due to poor experimental design—i.e. the very plan of the experiment allowed for unintended factors to influence the result—or due to ordinary problems in the execution of a well-designed experiment, e.g. when an instrument malfunctions or an experimenter accidentally bumps a sensitive apparatus. To argue that an experiment is internally valid, one must argue that the experimental outcomes of interest can in fact be attributed to the particular, local manipulation that the experiment was designed to investigate. Some of such argumentation will aim to show that it is unlikely that the experimental outcome is an artifact brought about by improper functioning of the experimental apparatus. This aspect of internal validity is the focus of Franklin's work on the epistemology of experiment (see e.g. Franklin 1989), which will be considered in detail in the next section.

External validity concerns the generalization of experimental results. Arguing for the external validity of an experiment involves justifying the application of internally

[33] Morgan (2003) and Guala (2002) discuss what they call "hybrid" forms of experiment. The experimental domain of a *hybrid experiment* includes both material entities and computer or mathematical models. Morgan also labels as "hybrid" those computer simulation experiments that are designed to produce data that mimic data that would be collected by making observations of the empirical world (see 2003, 224). I do not see why these experiments should be considered hybrids, but I will not examine the issue here. My focus in what follows will be on material and computer simulation experiments.

valid experimental results to situations outside of the experimental setting (i.e. to the target domain in goal-oriented experiments). That is, it involves supplying reasons for thinking that the results obtained via the particular manipulations, objects, and conditions in the experiment tell something about what will happen to other objects under other (perhaps related) conditions and manipulations. External validity is thus a matter of the representativeness of the experimental domain for purposes of answering the target question. In the context of material experiment, the issue of external validity often arises (though perhaps not by name) in discussions of the "artificiality" of the laboratory setting—the controlled conditions of the laboratory setting may be unlike those of the real world situations to which the experimental results are hoped to be applied. For example, when an experiment is performed on animals in captivity, reasons typically must be given to support the application of the results to animals in the wild.

This analysis of experimental validity seems to have been used first over forty years ago, as a way of thinking about potential problems in experimental tests of educational techniques (see Campbell 1957, Stanley and Campbell 1963). I mention this as a reminder that the terminology of internal and external validity was introduced in the context of ME. As will become clearer below, there seems to be a tendency today to associate the question of internal validity with ME and the question of external validity with CSE. The perception seems to be that in ME the main challenge is to ensure controlled environmental conditions and distinguish real effects from artifacts of the experimental apparatus, while in CSE the main challenge is to ensure that the experimental domain is representative of the target domain for purposes of answering the target question. I will emphasize in what follows that *both* ME and CSE typically involve questions of *both* internal and external validity.

4.3 Internal validity

Part of the evidence that external validity is not fully recognized as an important epistemic obstacle for ME comes from discussions of the epistemology of experiment. In particular, I have in mind the work done by Franklin (e.g. 1989) concerning strategies for sanctioning belief in the "validity" of an experimental result. Franklin indicates that he is discussing the epistemology of ME, but the validity under discussion seems to be internal validity only, and primarily only one aspect of internal validity, namely, the separation of real effects from artifacts created by the experimental apparatus. We may be tempted to infer from this that the epistemology of material experiment *just is* the study of strategies for determining whether an experiment is internally valid, but this would be a mistake: as I will show in Section 4.4, the epistemology of ME needs to pay close attention to issues of external validity as well.[34] The present section focuses on internal validity, using Franklin's strategies as a resource. Franklin appears to offer them as a kind of representative sample of the rational strategies that scientists use in defending the internal validity of their results in ME, and so it is of some interest to determine to what extent analogous strategies exist for offering evidence for the internal validity of results in CSE.[35] Before investigating the possibility of such analogous strategies, I first want to clarify what is at issue when it comes to the internal validity of computer simulation experiments, since this is not a commonly discussed topic.

4.3.1 Understanding internal validity in computer simulation experimentation

In ME, arguing for internal validity is a matter of separating real effects from artifacts of the experimental set-up and thus is concerned with such things as proper functioning of experimental apparatus and successful controlling for other confounding factors. In CSE, arguing for internal validity is also a matter of separation of real effects from artifacts, but the "real" effects are whatever our mathematical model of the target domain entails should occur as a result of the manipulation of interest.

To see this, it helps to review some of the kinds of models involved in the practice of computer simulation (see also Chapter 2). We formulate a mathematical model that we

---

[34] I want to empasize that Franklin does not say that he is giving a complete epistemology of experiment. For my purposes, it is helpful to draw attention to the fact that Franklin chooses to focus on internal validity, but I do not mean to suggest that Franklin actually believes that this is all there is to the epistemology of experiment.

[35] Eric Winsberg remarks of Franklin's strategies that "it is a straightforward exercise to go through this list and see that many, if not all, of these techniques apply directly or by analogy to the sanctioning of simulation results" (2003, 121). He recommends Weissart 1997 and his own 1999a for more details. Weissart has indeed discussed some of Franklin's strategies in the context of an historical case study on a particular problem in computer simulation, but he is concerned with only a subset of Franklin's strategies and leaves room for more discussion even with respect to this subset (as we will see). Winsberg compares the spirit of his own discussion of error management in computer simulation to that of Franklin's, but I do not find that he actually compares Franklin's strategies to those that he identifies as important in the context of computer simulation, although I think that such a comparison could be made. Upon close inspection, it does not seem that the application of Franklin's strategies to the context of CSE is quite so straightforward, and I think it worthwhile to work through the details.

take to describe the behavior of the target domain. (Whether this model is an accurate enough description of target domain behavior is a matter of external rather than internal validity.) We then transform the mathematical model equations into equations that are more amenable to numerical solution by computer, i.e. we transform the mathematical model into a numerical model. We must devise an algorithm to solve the equations of the numerical model under the conditions of interest.[36] We hope that the algorithm will yield solutions for the numerical model equations that approximate to some desired degree of accuracy the solutions of the computationally intractable mathematical model equations.[37] If we succeed in this approximation endeavor, then the computer simulation experiment can be considered internally valid. A computer simulation experiment is internally valid when its relevant results approximate (to some desired degree of accuracy) those entailed by the chosen mathematical model of the target domain; internal validity fails to hold when the relevant results instead reflect computational artifacts brought about either by the way in which we formulated and/or implemented our numerical model and our solution algorithm or by some truly external interference, such as a power surge that causes our computer to malfunction. Though I will not provide any extended argument for it here, I think it pretty clear that today the former path for undermining internal validity is of much greater concern than the latter—power surges and computing device malfunctions are not common suspects when the results of a computer simulation experiment are other than expected, but computational methods are a regular topic of concern. In what follows, I set aside such concerns about power surges and the like and consider internal validity to be a matter of the adequacy of the computational methodology employed.

How could a poorly-implemented computational methodology lead to results that do not approximate the outcomes that (in principle) are entailed by our mathematical model of the target domain? There are several ways. Computational instability produces "blow ups" in the simulation when the space and time scales of a finite difference scheme do not satisfy certain conditions. Truncation error arises because, in order to approximate

[36] The implementation of this algorithm in the form of computer code *is* the computer simulation model.

[37] The numerical solutions need not approximate the real solutions with respect to information that is not relevant to the answering of our target question; the important thing is for the answer to our target question to be close enough (according to our standards) to the answer that we would have obtained if the mathematical model equations could have been solved directly.

the continuous equations of our mathematical model by finite difference equations, we truncate a Taylor series approximation to the derivatives (see Holton 1992, 446). If the truncation error is too large, then the results of our computer simulation experiment may differ too much from the results entailed by the mathematical model directly. Round-off error results because, for each number stored, the computer can store only a finite number of significant digits. These round-off errors can sometimes accumulate if the computer simulation involves many iterative-loop calculations. Finally, there is simple programming error, which occurs when we make mistakes (e.g. "typos") in implementing an algorithm for solving our numerical equations.[38]

4.3.2 Strategies for obtaining prima facie evidence of internal validity

I now turn to the nine strategies offered by Franklin for obtaining evidence that an observed result is a real one, rather than an artifact of the experimental apparatus. Some of these strategies are concerned with evidence that the experimental apparatus is functioning properly (and thereby giving us results that are not artifacts), while others are more directly concerned with the plausibility of the result of interest. It will help to keep in mind throughout the discussion that in CSE the analogue of the experimental apparatus is the computational methodology—the way the numerical equations are produced and then actually solved. It is also important to keep in mind that Franklin's strategies are acknowledged by him to be neither individually sufficient for separating real effects from artifacts nor exhaustive of strategies that one might use to argue that a result is not an artifact. I will not discuss which of Franklin's strategies are likely to provide more or less compelling evidence for internal validity; a Bayesian analysis of these strategies in the context of ME has been given by Franklin and Howson (1988). I have the simpler goal of showing that, for each of Franklin's strategies, an analogous strategy is used in the context of CSE. I will discuss Franklin's strategies in three groups: strategies involving prediction, strategies involving theory, and miscellaneous strategies. The table below previews the analogue strategies that I will identify.

[38] Winsberg (1999a, 27-28) has also given a short discussion of these "computational/mathematical sources of error" in his typology of error in computer simulation.

Table 1. Summary of analogous strategies for ensuring internal validity.

Franklin's Strategies for Material Experimentation	Analogous Strategies for Computer Simulation Experimentation
Apparatus gives results that match known results	Relevant output of simulation matches with either (a) analytic solution or (b) natural phenomenon thought to occur in real world
Expected artifacts of technique are seen in experimental results	Expected artifacts of computational technique are seen in simulation output
Expected effects of intervention on experimental domain are seen to occur	Expected effects of intervention on either (a) model parameter values or (b) the computational algorithm itself are observed
Observed result of experiment can be explained by existing theory of the phenomenon	Outcome of computer simulation experiment can be explained by existing theory of the phenomenon
Proper functioning of experimental apparatus depends on well-corroborated theory	Adequacy of computational technique is underwritten by sound mathematical theorizing
Results of experiment replicated in other experiments using different kinds of apparatus	Results of computer simulation experiment confirmed by other simulation experiments using different kind of numerical technique
Results are too coherent and natural looking to be artifacts of apparatus	Results are too coherent and natural looking to be artifacts of computational technique
Results are not consistent with any plausible alternative explanation	Computational errors that might contribute to result can be ruled out
Results are unlikely to be the results of random measurement errors	Results are unlikely to be due to truncation or round-off error

4.3.2.1 Strategies involving prediction

(a) *Reproducing known phenomena.* The idea behind this strategy is that "the ability of the apparatus to reproduce already known phenomena argues both for its proper operation and in favor of the results obtained" (Franklin 1989, 447-448). Thus, when a thermometer accurately registers the temperatures of volumes of liquid whose temperatures we know in advance, we have evidence both that the thermometer is working properly and that it has registered the correct temperature for the volume of liquid whose temperature we sought to determine in our experiment. The same type of strategy is used in CSE in at least two different ways. First, we can show that our algorithm gives results that match those obtained analytically from the mathematical model (for the small set of such analytic results that might be available).[39] Second, we can point to the fact that the computer simulation output displays features that we associate with known real-world phenomena. For example (and speaking more plainly), we can point to the fact that our weather simulation model accurately predicts the occurrence of weather phenomena like extratropical cyclones. The assumption is that it is very unlikely that the computer simulation model would predict these known phenomena if there were major problems with its computational methodology. Thus, it is not only facts about the relation between the computer simulation results and the solution of the mathematical model equations that can serve as prima facie evidence of internal validity; facts about the relation between the computer simulation results and our understanding of the real world are invoked as well.

(b) *Producing expected artifacts.* A closely related strategy involves detecting artifacts already expected to be present (Franklin 1989, 449). In ME, we sometimes know that our techniques will lead to particular artifacts in our results. Franklin describes a case in which a substance of interest could only be studied in mixture with other, known substances. When the substance of interest was studied via infrared spectroscopy, and the measured absorption lines were seen to reflect the expected effects of the presence of the other substances in the mixture (i.e. peaks and valleys in anticipated locations), this constituted evidence that the instrument for making the measurements was working correctly (Franklin 1989, 449). An analogous strategy is used in CSE. Here, the

[39] This analogous strategy is noted by Weissart; he calls it "calibration" (1997, 123).

anticipated artifacts are ones that result from the limitations of the methods that we have used to solve the numerical equations at hand. For instance, when using weather forecasting models, we may be able to predict in advance the rate of propagation of errors into the forecast region due to the use of simplified boundary conditions; we thus can predict when the forecast accuracy will deteriorate for various parts of the forecast region. The fact that we can predict this development of artifacts counts as prima facie evidence that our computational method is working as designed. (It is the fact that we know of certain shortcomings of the design that we are able to have such evidence.)[40]

(c) *Producing expected outcomes of interventions.* A third strategy involving prediction concerns intervention: we predict "what will be observed after the intervention if the apparatus is working properly or as expected" and "when the predicted observation is made we increase our belief in both the proper operation of the apparatus and in its results" (Franklin 1989, 440). For example, if we heat a volume of liquid and observe that the temperature registered by our thermometer increases by an expected amount or at an expected rate, we then have prima facie evidence that the thermometer is working properly and therefore that it registers the correct temperature of our heated volume of fluid. An analogous strategy in CSE involves something like manipulation of parameter values in the computer simulation model. For instance, we might increase the heat capacity of the ocean in our computer simulation model of the climate system, and if the rate of increase of global annual mean temperature then slows for a period (as expected), we have some evidence that our computational scheme is not seriously malfunctioning. A not strictly analogous but closely related strategy would be to predict the effects of intervening on the computational method itself, e.g. by decreasing the time step of our calculations. If the results of the simulation experiment do not change substantially after

[40] Weissart (1997) seems to misinterpret the strategy of anticipating artifacts. His analogue in CSE involves verifying that energy is conserved at each step in a computer simulation of a system in which energy is supposed to be conserved (1997, 122-123). He apparently takes such conservation of energy to be an artifact, but this does not seem to match the common conception of artifacts as erroneous results. On the other hand, Weissart's suggestion that we might anticipate errors in our simulation results due to round-off errors that accumulate during iterative-loop calculations (1997, 122) seems more on the mark, though he does not pursue this in much detail.

decreasing the time step, we have prima facie evidence that the computational scheme is working properly.[41]

4.3.2.2 Strategies involving theory

(a) *Results explained by theory of phenomena.* The idea here is that confidence in the legitimacy of observed phenomena (as real effects as opposed to artifacts) increases when the observations can be explained using an existing, accepted theory that applies to the observed phenomena (see Franklin 1989, 442-446). Thus, suppose we find that the addition of certain chemicals to a lake is followed by marked decrease in the observed number of catfish in that lake. The observed decrease might be real or it might be an artifact of our technique for estimating the number of catfish present in the lake. If we can argue on a theoretical basis (e.g. making appeal to what we know about the biochemistry of catfish) that catfish are likely to die in the presence of the chemicals that were added to the lake, then we increase our confidence that the observed decline in the catfish population is real and not an artifact of our estimation technique. In the case of computer simulation experiments, there is a close analogy to this strategy. By showing that an interesting result of a computer simulation experiment can be explained using accepted theory, we have prima facie evidence that the interesting result is really entailed by our mathematical model and is not just an artifact of our computational scheme.[42]

(b) *Well-corroborated theory of the apparatus.* A related strategy involves appeal to a well-corroborated theory of the apparatus used in an experiment: "if…the proper operation of the apparatus depends on such a theory, then it can be argued that the evidence supporting the theory also gives reasons to believe the observations" made using that apparatus (Franklin 1989, 440). Franklin gives the example of the electron microscope: since its proper operation depends upon well-corroborated theory, we have reason to believe that the electron microscope generally shows us real effects rather than artifacts. An analogous strategy in CSE involves arguing that the computational

[41] Weissart (1997, 123-124) characterizes this kind of test involving the decrease of the computational time step as a form of what Franklin identifies as the use of a "different experimental apparatus" to provide confirmation of a result. I will suggest below that this involves using a different experimental apparatus only in a very weak sense and that a stronger sense is better aligned with another example that Weissart gives.

[42] It is worth emphasizing the "prima facie" nature of the evidence here. The theory-simulation match is only evidence of the proper functioning of the computational scheme if the match is not the result of "tuning" the model for the express purpose of producing such a match. See Chapter 6 for more on tuning.

techniques used in the computer simulation experiment are based on sound mathematical theorizing. For example, one could argue that finite difference techniques have their basis in the calculus and that theorems have been proven concerning the stability of these techniques.

4.3.2.3 Miscellaneous strategies

(a) *Independent confirmation.* This strategy involves showing that the results of an experiment can be replicated in different experiments using different experimental apparatus (Franklin 1989, 438). The example discussed by Franklin follows Ian Hacking's argument that when something can be observed using two or more different types of microscopes, this counts as strong evidence that the observations are revealing something real, rather than an artifact. Two experiments can be considered "different" experiments when they involve apparatus based on different theories, and perhaps even when they simply involve differences of size, geometry or personnel (Franklin 1989, 438). In CSE, a direct analogue of confirming a result using different apparatus involves changing the *kind* of numerical technique used to transform and solve the equations of the mathematical model, e.g. one might switch from a finite differencing scheme to a spectral approach in solving the equations of a weather simulation model.[43] If we also accept that experiments are different even when they involve only variations in "size," then we can have a different experiment just by re-running a simulation with a different time step in its finite differencing scheme. This is perhaps analogous to changing the degree of magnification used when studying a slide under a microscope. The change might reveal some artifacts (e.g. those that resulted because the previous lens was dirty) but any artifacts resulting from something other than the lenses (e.g. a distortion in the eyepiece) can remain undetected. This kind of "different" experiment would seem to confer relatively weak evidence of internal validity. Perhaps the "most different" experiment that we might conduct in order to confirm that a CSE result is not an artifact would be a *material* experiment that we take to capture the situation represented in our computer simulation experiment. If this material experiment gave the same result as the computer

[43] A spectral model represents the mass of the atmosphere in terms of a series of waves of differing frequencies (see Holton 1992 for a brief discussion).

simulation experiment, we would have independent confirmation that the result was not just an artifact of our computational methodology.

(b) *Internal coherence of results.* This strategy involves arguing that it is unlikely that the instrument would report such coherent-looking observations if it were malfunctioning. Franklin's example concerns the use of Galileo's telescope to observe the moons of Jupiter (see 1989, 440-441). The suggestion is that it hardly seems possible that behavior as complicated and coherent as that exhibited by the specks of light seen with the telescope (i.e. behavior that looked like the regular motions of a small planetary system complete with eclipses) could be the result of an artifact of the measuring instrument (see Franklin 1989, 441). The assumption is that the appearance of such complicated and coherent behavior only is plausible if some real phenomena are being observed with the telescope. An analogous strategy in CSE involves arguing that some result simply looks "too natural" to be a computational artifact. It is true that some computational artifacts (e.g. blow ups due to computational instability) have distinctive features that we do not typically observe in the natural systems being simulated.[44]

(c) *Elimination of plausible alternative explanations.* The strategy of eliminating plausible alternative explanations involves arguing that we can rule out specific, alternative explanations of results. Franklin's example concerns observations of electrical discharges in the rings of Saturn by a passing spacecraft (Franklin 1989, 446-447). He shows how several plausible alternative explanations of the observed electrical discharges in the rings of Saturn were eliminated: among other factors, such things as a poor telemetry link between the spacecraft and the Earth and the presence of electric discharges near the spacecraft (rather than in the rings of Saturn) due to environmental phenomena and/or dust particle-spacecraft interaction were each ruled out. The fact that all of these plausible alternative explanations could be ruled out counted as evidence that the observation of electrical discharges from Saturn's rings reflected a real effect rather than an artifact. In CSE, the analogue is to argue that some result should be seen as a "real" result because we can rule out all of the plausible computational errors that we can imagine might contribute to such a result. Thus we might offer specific evidence that the

[44] A related argument concerning features of the results themselves focuses on consistency of results (see Weissart 1997, 124 for an application to CSE).

result is not the product of computational instability, round-off error, truncation error, propagation of boundary effects, or programming error.

(d) *Statistical arguments*. This strategy is the most difficult to encapsulate. The idea seems to be, at least in part, that a surprising result obtained from a material experiment is not considered a real result as long as it is plausible that it arose due to small, random measurement errors (see Franklin 1989, 455-458). The discovery of the top quark might be a good example of the use of statistical arguments in this way: the existence of the top quark was accepted because it was thought that the data obtained from many, many measurements were very unlikely to have been as they were if top quarks did not exist (i.e. the data were unlikely to be the result of small, random measurement errors). It is not entirely clear whether a direct analogy to this strategy exists in CSE. Perhaps the analogue concerns results that differ a little bit from what is expected: because we know that a computer simulation will involve some small unavoidable errors (e.g. due to truncation and round-off), we hesitate before considering "real" any results that might have occurred due to those errors. For example, it might be that energy is not exactly conserved from one time-step to the next in our computer simulation, but rather than accepting this as a "real" consequence of our mathematical model equations we assume that it is an artifact due to truncation and round-off error. For CSE, this strategy thus appears to have something in common with the previous strategy.

4.3.3 Conclusions concerning internal validity: the computational advantage?

The foregoing analysis shows that the strategies that Franklin identifies for ME have analogues in the context of CSE. The analogues discussed concern the adequacy of the computational methodology employed, rather than the proper functioning of some physical experimental apparatus. It is worth showing in detail the parallels between these strategies in order to emphasize the fact that internal validity is not only an issue in both ME and CSE—analogous strategies are often employed in arguing for such validity in both types of experimentation.

I would argue that CSE actually enjoys certain advantages over ME when it comes to internal validity. In ME, a major concern is to ensure that conditions in the laboratory are precisely as desired—controlling lab conditions is big part of the job of the experimentalist, and this is probably one reason why internal validity is a matter of such concern in ME. But in CSE, controlling these (virtual) conditions is a straightforward matter; such control is achieved simply by setting parameter values or changing the form of equations.[45] If one wants there to be no friction, one simply sets the friction variable to zero in the computer simulation model; if one wants wind velocity to increase exponentially during some time interval, one simply includes the appropriate function in the computer simulation model. Things can go wrong only if one makes a mistake in setting these controls (e.g. a "typo"). In principle, we have complete control over the (virtual) conditions that interest us. In this respect, at least, computer simulation experiments have an advantage over material experiments when it comes to arguing for the internal validity of results. Obviously, this advantage is not decisive—the discussion above suggests numerous ways in which a computer simulation experiment might nevertheless go wrong (e.g. due to computational instability, etc.)—but it is an advantage nonetheless.

4.4 External validity

As discussed in Section 4.2, arguing for the external validity of an experiment involves justifying the application of experimental results to situations outside of the experimental setting. The issue of external validity (even if not mentioned by name) is the root of much criticism of CSE. This section will take a closer look at external validity in both CSE and ME. I will first discuss what is at issue when it comes to external validity, although in this case it is ME that requires extended discussion and clarification. Then, I will show by example that it is not necessarily more difficult to argue convincingly for the external validity of results in CSE than for the external validity of results in ME.

4.4.1 Understanding external validity in material experimentation

In CSE, arguing for the external validity of some experimental result typically requires arguing that the right relation holds between the computer simulation model and the system being represented. If the results of the experiment are to apply to the target system, then the behavior of the model system (i.e. the simulation) must be similar in relevant respects and to relevant degrees to the behavior of the target system. As I

[45] This advantage is also noted by Christie (2001, 163).

suggested in Chapter 2, which respects and degrees are considered relevant depends upon the purpose for which the model is to be used, i.e. upon the question that the computer simulation experiment is supposed to help answer. A computer simulation experiment is externally valid (relative to a target domain) if the relevant similarities obtain.

External validity in ME is also a matter of the relation between the experimental domain and the target domain. It is important to recognize that, as in CSE, the experimental domain of a material experiment often serves as a model of the target domain. A common misconception is that computer simulation experiments involve the use of models in this way, while material experiments do not. For example, consider this statement from a recent textbook on simulation in the social sciences:

> The major difference is that while in an experiment, one is controlling the actual object of interest (for example, in a chemistry experiment, the chemicals under investigation), in a simulation one is experimenting with a model rather than the phenomenon itself. (Gilbert and Troitzsch 1999, 13)

For this view to be correct, most (or perhaps all) material experiments would need to be ones in which the experimental and target domains were identical. That is, they would need to be experiments designed to answer questions about what happens to some particular thing at some particular time under specific conditions. For example, such an experiment might involve investigating the effects of closing a particular traffic route at a particular time by actually closing that route at that time.

However, even brief reflection on the practice of science reveals that a great number of material experiments do not have identical experimental and target domains. In many cases, scientists want to answer general questions about types of entities and groups, and in these cases it usually is not feasible to have identical experimental and target domains. Some of the sample experiments described in Table 2 illustrate this. For instance, returning to the weight loss example again, it simply is not feasible to perform a study in which all of the relevant obese adults take part. The experimental domain consists instead of a randomly selected subset of the target domain. Similarly, consider the experiment in Table 2 concerning samples of elemental substance *S*. It is not feasible

Table 2. Sample target questions and their associated domains.

Target question	Target domain	Experimental domain	Experiment Type
How will traffic patterns change if route R is closed at time T?	Traffic patterns near closed route R after time T	Traffic patterns near closed route R after time T	Material
At what temperature will elemental substance S melt, under pressure P?	All masses of S under pressure P	A small number of samples of S under P	Material
How much weight will be lost on average by obese adults who daily ingest Z?	All obese adults who will ingest Z	A particular group of obese adults who ingest Z	Material
How much stress will be put on the tail of jets of type J when at cruising altitude with a K knot crosswind?	All jets of type J when at cruising altitude with a K knot crosswind	A miniature version of J in a wind tunnel subjected to crosswind W	Material
What will be the effects on humans of daily ingestion of X amount of saccharin?	All humans who will daily ingest X amount of saccharin	A particular group of rats who daily ingest M amount of saccharin	Material
What is the range of a projectile with properties E if it is fired with initial force F and initial angle A under environmental conditions C?	Projectiles with properties E under conditions under conditions F, A and C	A computer simulation model of the motion of such a projectile, with E, F, A, C represented	Computer Simulation
How will increasing the concentration of substance Y in the Earth's atmosphere affect the mean annual global temperature?	The Earth's atmosphere with increased concentration of Y	A computer simulation model of the Earth's atmosphere, in which the increased concentration of Y is represented	Computer Simulation

to test all samples of *S*, since they presumably are scattered far and wide in unknown locations; instead only a few samples of *S* are examined as part of the experiment. Material experiments may have non-identical experimental and target domains for other reasons as well. Suppose that one wants to know how much stress will be put on the tail of a particular type of jet when the jet is at cruising altitude and experiencing a strong crosswind (see Table 2). If there is some risk that the stress may be so great that it will cause the tail to break off or will cause other serious damage to the jet, then ethical and financial reasons may prohibit engineers from using real jets to answer the target question. Instead, a miniature material model of the jet—adjusted in particular ways in accordance with the guidelines of scale modeling—is studied in a wind tunnel.[46]

It is uncontroversial in the wind tunnel case that the experimental domain is serving as a model of the target domain. But I suggest that in the other cases as well (when the experimental and target domains differ), the experimental domains are serving as models of the target domains. Returning to Table 2, the sample of obese adults who take substance *Z* serves as a model of the larger population of such adults, the samples of substance *S* serve as models of all other masses of that substance, and the rats ingesting saccharin serve as models of humans ingesting saccharin. Each of these experimental domains is taken to represent the respective target domain because it is thought to be similar to the target domain in respects relevant for the purposes of experiment. The experiments are conducted on models.

If experimental domains are models in both CSE and ME, then carrying out an experiment of either type amounts to carrying out a simulation. A simulation involves the modeling of some system as it evolves through time (see Chapter 2), and this is just what happens in an experiment. The sequence of events that constitutes the experiment—the change/evolution of the experimental domain as a consequence of some manipulation—is intended to simulate the target domain as it evolves under the conditions of interest. This seems obvious for the case of CSE, but it is often true in the case of ME as well. It is most easily seen for the wind tunnel and rats/saccharin examples in Table 2.

[46] Norton and Suppe (2001, 70) identify this type of experiment (scale models of airplanes in wind tunnels) as a paradigmatic example of traditional experimentation. Sterrett (2003) discusses the methodology of scale modeling.

The fact that the experimental domains of both computer simulation experiments and material experiments typically serve as models of their respective target domains has not been emphasized enough in the literature. Norton and Suppe (2001) argue that both ME and CSE involve the use of models, but they emphasize the presence of models in the context of instrumentation and data processing, not as the very entities on which experiments are performed.[47] Although Morgan (2002) does describe mathematical model experiments as experiments *on* models, she does not characterize material experiments (as a group) in terms of models at all. She does identify different types of "representing relations" that might obtain between experimental and target domains, which suggests that her view of experimentation may be similar to the one presented here (see Morgan 2003, 227-232).[48] The view of experiment and simulation that seems to conflict most directly with my own is that offered by Guala (2002). He draws an ontological distinction between "genuine" experiments and "mere" simulations, suggesting that experiments involve "deep" "material" correspondence between experimental and target domains, while simulations involve only "abstract" and "formal" similarity (see Guala 2002, 67). It seems possible that our views could be reconciled by saying that material experiments typically involve material similarities between experimental and target domains, while computer simulation experiments typically involve formal similarity instead. (I will not pursue that reconciliation in detail here.)

To return to the purported difference between ME and CSE indicated above, it is not that CSE involves experimentation on models while ME does not; the difference is that material experiments are conducted on material models, while computer simulation experiments are conducted on computer simulation models. If this difference is put in terms of simulation, then the difference is that material experiments are material simulations (i.e. simulations using material entities), while computer simulation experiments are symbolic simulations. When it comes to external validity, in both types of experimentation the task is to justify inferences about targets on the basis of results

[47] Instruments embody models that translate the physical properties of that which they probe into dial and meter readings; this is a kind of physical dependence on models (see Norton and Suppe 2001, 72-73). In addition, the data that are the end product of experiment typically have a computational dependence on models, because they are the result of further processing guided by assumptions (Norton and Suppe 2001, 72-73). For instance, corrections may have been applied to account for known deficiencies or biases in the instruments used in the experiment.

[48] Morgan does not actually use the terminology "experimental domain" and "target domain."

obtained via simulations. Thus, in *both* ME and CSE, arguing for external validity typically involves justifying simulation-to-target inferences.

4.4.2 Arguing for external validity

Yet we can still distinguish *computer*-simulation-to-target inferences from *material*-simulation-to-target inferences. This raises the question of whether these different types of inferences are equally easy to justify. If it were more difficult to justify computer-simulation-to-target inferences than to justify material-simulation-to-target inferences, then this would provide some basis for the kind of blanket skepticism toward CSE that I discussed above.

Guala (2002, 70) points out that the kind of information needed to justify the application of experimental results to systems outside of the laboratory (or model world) will be *different* in ME and CSE: in the case of ME, one will need good reason to think that the experimental and target domains have relevant material similarities, while in the case of CSE one will need good reason to think that the target domain has been adequately described by the equations and values in the model and that these equations have been solved using a trustworthy method.[49] The issue at hand is whether it is more difficult to have the required good reasons in CSE than in ME. Guala hints that it is more difficult when he says that material experiments "do not require as much" in the way of background knowledge as computer simulation experiments do (2002, 69). Morgan (2003, 231) addresses the issue more directly, remarking that mathematical model experiments "present much greater difficulties" than material experiments when it comes to justifying inferences about target domains. But the examples below illustrate that this is not always the case. For some computer simulation experiments, it is relatively easy to justify inferences about the target domain, and for some it is quite difficult. The same is true of material experiments. For both types of experimentation, the difficulty involved in justifying simulation-to-target inferences varies from case to case across a broad spectrum.

---

[49] Morgan expresses the difference like this: one needs "accurate replication" of the target domain in the case of material experiment but "accurate representation" of the target domain in the case of computer simulation experiment (see 2002, 57).

4.4.2.1 Material experiments

The second experiment listed in Table 2 is an example of a material experiment in which it is not very difficult to justify the relevant inference about the target domain, i.e. relevant for answering the target question about melting temperature of elemental substance *S*. Assuming internal validity (i.e. assuming that the pressure in the experimental apparatus is very close to *P*, that the instruments used to measure temperature are functioning properly and are sensitive enough to register the result, and that one has a reliable method for determining when *S* begins to melt), one needs to know only that small deviations from *P* will not make much difference to the experimental results and that the samples of *S* are pure or else impure in ways that allow for corrections to be made to the experimental results. Realizing these conditions is not excessively difficult. Suppose that the experiment then involves a total of ten measurements of melting temperature on ten different samples. The average of these temperatures will provide a good estimate of the temperature at which other samples of *S* will melt when under pressure *P*. Experiments like these have been carried out time and time again, and their results are part of the accepted body of science.

By contrast, the study of model organisms provides examples of material experiments in which it may be unclear to what extent, if any, the results of an experiment apply to the target domain. An interesting real-life example of this kind of situation can be found in attempts to determine whether saccharin is a human carcinogen. Studies found an increased incidence of cancerous tumors in some types of male rats who were given saccharin. On the basis of this finding, precautionary measures were taken. Saccharin was classified as a chemical "reasonably anticipated to be a human carcinogen" by the U.S. National Toxicology Program (NTP), and products containing saccharin began to carry warning labels indicating that saccharin had been shown to cause cancer in laboratory rats. But saccharin never was listed as "known" to be a human carcinogen, and saccharin-containing products did not carry labels saying that saccharin causes cancer in humans. The problem was uncertainty concerning the relation between biological processes in rats and humans, i.e. concerning whether the experimental and target domains were similar in respects relevant for answering the target question about cancer in humans. Although rats and humans have certain biological similarities—this is one reason for using rats rather than some other type of organism in these experiments—there are also known differences, and there is much about the relation between rats and humans that is unknown. At the time of the early rat studies, it was unclear whether the unknown mechanisms producing the cancer in the rats could also act in humans.[50] So, although the experiments were material, they did *not* allow for the answering of detailed questions about the effects of saccharin on humans.[51]

In the case of the rats, there are good material models available, i.e. human beings, but there are ethical reasons prohibiting their use in the experiments. There are also cases in which it is just very difficult to build or find an adequate material model of the system of interest. A good example of this is found in the study of the Earth's climate. A recent textbook explains that:

> ...there exist no physical models which can simulate the complex behavior of the climate system in a laboratory environment in an adequate way. For example, the nonlinear interactions between the various subsystems are impossible to reproduce, even partially, in any laboratory experiment. The laboratory analogs (dishpan or annulus experiments) obtained so far can only yield some general dynamical characteristics of the rotating ocean-atmosphere system in a very preliminary and coarse way. (Peixoto and Oort 1992, 450)

The detailed evolution of the climate system is the result of a variety of interacting processes that are not easy to recreate in a laboratory setting. Laboratory analogues are used (as the rats are in the experiments described above) in the hope that the experiments will tell us something about the real climate system. But since these analog experiments simulate only some very general features of the climate system, it is difficult to know to what extent they can be used to answer questions about the real atmosphere. In these types of cases, too, it can be very difficult to justify material-simulation-to-target inferences.

---

[50] LaFollette and Shanks (1995) provide an interesting discussion of the limitations of experiments on model organisms for drawing inferences about causal mechanisms in humans. Schaffner (2001) also discusses the issue of extrapolating from animal models to other animals and humans.

[51] After 20 years, the NTP removed saccharin from the list of chemicals reasonably anticipated to be carcinogenic to humans. According to the NTP report, the removal from the list was "based on the perception that the observed bladder tumors in rats arise by mechanisms not relevant to humans, and the lack of data in humans suggesting a carcinogenic hazard" (see NTP 2001, Appendix B, 7).

4.4.2.2 Computer simulation experiments

Consider a computer simulation experiment in which the release angle of a shot put is altered by a moderate amount in order to study the effects of this change on the shot trajectory. The target question might concern the peak height and range of the shot put, given that it is thrown with a particular initial force and a particular initial release angle, under particular environmental conditions (see Table 2). In principle, a number of factors could influence the trajectory of the shot put, including gravity, aerodynamic drag, air density, wind effects, rotational effects and other factors (see e.g. De Mestre 1990). But for many cases, only the effects of gravitational and aerodynamic drag forces will need to be included in order to answer the target question at hand.[52] In these cases, it will be a rather simple exercise to construct and employ a computer simulation model that predicts the shot put trajectory. Physical theory, empirical data, and mathematical analysis all will be employed in constructing the model. For example: Newtonian gravitational theory will be appealed to in representing the gravitational force; empirical data concerning the drag coefficients for spherical objects will be employed in representing the drag force; the equations of motion that incorporate these representations of gravitational and drag forces (in accordance with Newton's 2nd Law of Motion) will be discretized and solved numerically, e.g. according to the Euler-Cromer method (see e.g. Gould and Tobochnik 1987, 46-48). This is a case in which one has well-established background knowledge about the causal factors relevant to the trajectory of the shot put, about how to represent these factors mathematically, and about the adequacy of the mathematical techniques used to solve the mathematical equations describing the action of these factors.[53] Because one has good reasons for trusting this background knowledge (e.g. it has been confirmed time and time again), one has good reasons for thinking that

[52] Sometimes it is known that the other factors are not present at all. In addition, even if it is known that they are present, it is sometimes possible to estimate, e.g. through scale analysis, that they affect the trajectory so minimally that their exclusion will not prevent us from answering our target question to the desired degree of accuracy.

[53] One may even be able to predict how closely the simulated trajectory will match a given real trajectory by estimating the order of magnitude of the effects of the omitted factors (e.g. rotational effects) and the errors introduced by the computational method.

the results of the computer simulation experiment can be applied rather directly to the target domain in order to answer the target question at hand.[54,55]

For a good example of a case in which it can be difficult to justify computer-simulation-to-target inferences, I again return to the study of the Earth's climate system. Many computer simulation models of the climate system have been developed, and many of these are used to conduct computer simulation experiments. Often, the goal is to find out how the Earth's climate will change if atmospheric greenhouse gas concentrations change in some way. However, there is uncertainty about the capacity of these models to answer questions about what will happen in the future. The uncertainty arises in part because the climate system is not entirely well understood. It is not clear that all of the important processes controlling climate have been identified, much less adequately specified in the computer simulation models. In addition, some phenomena that are thought to play an important role in shaping climate are believed to be relatively poorly represented in the computer simulation models, for a variety of reasons. The formation of clouds is the prime example of this type of phenomenon. Clouds develop on a spatial scale that is smaller than the scale resolved by computer simulation models of climate, but clouds play an important role in the climate system. As a consequence, the effects of clouds must somehow be accounted for in climate models, even though individual clouds are not simulated. These effects are complicated, however, and it is difficult to represent them adequately in terms of other, resolved model parameters. Currently, modeling

---

[54] I do not mean to suggest that all simulation-to-target inferences must be justified in this way, i.e. by appealing to well-established background knowledge of causal factors. In some cases, there may be other strategies for justifying such inferences, e.g. by making an inductive argument based on past predictive successes had by the model. So, I am not offering a general account of what constitutes adequate justification for simulation-to-target inferences. I am only presenting an example of what could constitute such justification in a particular case. This is all that is required for my argument. While it is interesting and important to consider what a more general account would look like, it goes beyond the scope of the present discussion.

[55] A similar type of example can be found in computer simulation experiments designed to answer questions about a real pendulum. The factors that might be taken into account in the computer simulation model are discussed in some detail in Morrison (1999, 48-53). As Oreskes (2000, 78) has recognized, relatively simple cases like these may be the exception rather than the rule in scientific practice. But I am not arguing that justifying computer-simulation-to-target inferences is always or even typically easy to do. I simply want to remind those who are skeptical that it is *not always very difficult* to do. The fact that computer simulation often is used as a resource for studying complicated, relatively poorly-understood systems indicates that we should be cautious in applying the results of some particular computer simulation studies, but it says nothing about the general soundness of the methodology of computer simulation.

studies do not agree even on whether clouds will reinforce or hinder any warming that might otherwise occur due to changes in greenhouse gas concentrations (see IPCC 2001, 427-431). Uncertainties like these, concerning the extent to which important processes have been adequately represented in climate models, translate into uncertainty concerning the extent to which the results of computer simulation experiments involving climate models provide information about the real climate system.[56] This means that there may very well be room for climate skeptics to question some climate modeling results. But, as the examples just discussed show, such questioning ought to be a consequence of the details of the climate modeling case rather than the mere fact that climate models *are* computer simulation models.

4.4.3 Conclusions regarding external validity: the material advantage?

If we conceive of experiments as simulations, we see that the problem of external validity is really the problem of justifying simulation-to-target inferences and that it is an issue in both CSE and many material experiments. The examples above illustrate that that it is not necessarily more difficult to justify computer-simulation-to-target inferences than material-simulation-to-target inferences. There can be computer simulation experiments for which one can provide good reasons for thinking that the results do allow one to answer the target question at hand, and there can be material experiments for which one has real difficulty in determining whether the results apply to the target domain at all. The difficulty involved in justifying simulation-to-target inferences varies substantially from case to case and is not determined by whether the experiment is material. The fact that an experiment is material (or not) is not decisive when it comes to justifying experimental results as externally valid.

However, it is true that, when it comes to external validity, there are epistemic advantages that can be had in ME that cannot be had in CSE. In material experiments in which parts of the experimental and target domains are made of the same (kind of) material, there is the epistemic advantage that these domains are thereby known to have some "built-in" similarities. For example, if humans are used to represent human beings in an experiment, one can be confident based on that fact alone that the experimental and

[56] It is interesting to note that, despite the uncertainty involved, computer simulation models of climate (rather than the material analog models described above) are currently widely used in research on climate change.

target domains will have many similarities; presumably it is in virtue of some of these similarities that the entities are classified as "humans" in the first place. But in computer simulation experiments these built-in similarities are not present, and so it is up to the experimenter to ensure that all of the required relations between the experimental and target domains obtain. This kind of advantage seems to be at least part of what leads Guala and Morgan to suggest that computer-simulation-to-target inferences are more difficult to justify than material-simulation-to-target inferences (see Guala 2002, 65 and Morgan 2002, 54).

Nevertheless, it is important to recognize the limitations of this advantage. For one thing, there is no guarantee that the built-in similarities will be the ones that are important for the experiment at hand; there are differences among humans, and it is easy to imagine experiments in which these differences really matter. For instance, clinical trial results obtained by studying male humans may not apply to female humans. In addition, it is not clear that one will actually know in advance what many of the built-in similarities are. One may know what some of the similarities are likely to be—for example, humans are likely to have two kidneys—but in the end one may have to check to see that the expected similarities really are present (since there are humans who have only one kidney). Furthermore, having parts of the experimental and target domains made of the same material type is not *decisive* when it comes to justifying simulation-to-target inferences. Simply having human test subjects take part in psychology experiments will not guarantee that the experimental results apply to a particular real-world situation—obviously other details of the experimental set-up are important. Finally, it is important to remember that the advantage is enjoyed only by a subset of material experiments, since not all material experiments have experimental and target domains made of the same kinds of material things.

But a second epistemic advantage *is* enjoyed by all material experiments. All material experiments have the advantage that they are constrained by the empirical world in a way that computer simulation experiments are not. When experiments involve the manipulation of material objects, the characteristics of these objects place some limits on the possible outcomes of the experiments. For instance, when a material fluid flow experiment is carried out, mass and energy will be conserved. By contrast, in a computer simulation experiment, it is possible for these conservation laws to be violated. In computer simulation experiments, there is the possibility of ascribing to systems characteristics that they could not display in actuality. In material experiments one at least does not have to worry about the experimental domain behaving in a way that violates what can actually happen in the material world. However, knowing that the experimental domain does not violate what can happen in the material world does not take one very far in justifying material-simulation-to-target inferences. For instance, there is no guarantee that the relations that hold in the experimental domain are the ones that are relevant for understanding what will happen in the target domain. This is illustrated by the saccharin example above—it eventually was determined that what was happening in the rats does not happen in humans.

4.5 Conclusions

Material experiments and computer simulation experiments both involve questions of internal and external validity. Failure to recognize this can lead to misconceptions about the relative epistemic difficulties faced in CSE and ME. CSE seems to enjoy certain advantages when it comes to overcoming the problem of internal validity (e.g. control of "environmental" conditions is a straightforward matter), while ME has some advantages with respect to external validity (as just discussed). However, the respective advantages enjoyed by CSE and ME are generally not sufficient to justify either the internal or external validity of experimental results. Thus, the experimental validity of results in both CSE and ME should be judged on a case by case basis, rather than categorically.

5

Explaining Weather and Climate: What Role for Simple Models?

5.1 Introduction

The study of weather and climate (hereafter "SWC") currently involves the use of a collection of computer simulation models whose complexity spans a broad spectrum. While there is no hard and fast line dividing so-called "simple" models from "complex" models, models near the simple end of the spectrum do tend to differ in systematic ways from models near the complex end of the spectrum. As we saw in Chapter 3, simple climate models generally (i) represent only one or two spatial dimensions of the climate system (ii) have relatively coarse spatial and temporal resolutions and (iii) represent a markedly reduced set of physical processes, many of which are parameterized.[57] By contrast, complex models generally (i) represent all three spatial dimensions of the climate system (ii) have relatively fine spatial and temporal resolutions and (iii) represent a rather large set of interrelated physical processes, with many of these representations grounded in accepted physical theory. Insofar as the climate system is a three-dimensional system whose behavior is the result of the action of numerous physical processes interacting in complicated ways on both large and small spatiotemporal scales, simple models are relatively poor characterizations of the climate system when compared to complex models. So why do we need the simple models at all?

An obvious reason has to do with the limits of our computing power. Complex models require vast computing resources if simulations of long time periods are to be carried out, as is the case in the study of climate change. As discussed in Chapter 3, one important part of current research on climate change is concerned with generating projections of future climatic changes under a variety of possible greenhouse gas emissions scenarios. It simply has not been computationally feasible to use complex

[57] See 3.4.2.1 for more details on simple and complex models.

models to make these projections for more than a few scenarios. Because of their computational efficiency, simple models have been used to generate the bulk of global mean annual temperature projections thus far (see IPCC 1996, 1997, 2001). However, the predictive value of simple models is limited, most obviously because they do not make any predictions (and therefore certainly not good ones) for one or more spatial dimensions of the three-dimensional climate system. For example, nobody thinks that simple one-dimensional models are likely to be very useful for investigating regional climate change, because such change is thought to depend upon processes that are not represented in detail (and sometimes not represented at all) in such simple models. In general, when it comes to prediction, simple models seem to be resorted to when pragmatic reasons prohibit the use of more complex models. Ultimately (and not surprisingly), complex models are viewed as the most promising predictive tools.

Are there other reasons for developing simple models when more comprehensive models are available? Surely simple models would not be considered the best tools for arriving at explanations of phenomena in SWC, since their predictive limitations would seem to limit their explanatory value as well. If we want to explain some phenomenon, we might turn to computer simulation models to provide us with information about that phenomenon. Presumably, if such a model cannot even simulate the occurrence of the phenomenon of interest, then it cannot provide us with much direct information about that phenomenon. For instance, if a simple model cannot simulate the occurrence of anything like an El Nino event, then it cannot provide us with much direct information about the El Nino phenomenon. At the very least, such a model would not seem to have the potential to tell us as much about the El Nino phenomenon as a model that could simulate the occurrence and detailed structure of the phenomenon. Since complex models have the potential to supply us with direct and more detailed information about many more phenomena than simple models, ultimately the complex models would seem to be the more valuable resource when it comes to explaining (and therefore understanding) phenomena in SWC.

But precisely the opposite often is claimed! A commonly expressed view in SWC can be summarized as follows:

(SU) Simple models—but not complex models—contribute substantially to our understanding of the atmosphere and climate system.

This is a puzzling claim. It says that models that are relatively inaccurate characterizations of their target systems nevertheless contribute substantially to our understanding, while models that are relatively accurate characterizations nevertheless do not contribute much. In other words, it says that the descriptive accuracy of a model and its value in promoting understanding are not aligned.[58] How could models that patently "get it wrong" relative to more complete theoretical descriptions nevertheless play a central role in helping us to understand the world?[59] And why would complex models, which constitute our best representations of the world, not be good tools for helping us to understand the world?

The discussion that follows aims to make sense of and evaluate SU. In order to stave off worries from the outset, I want to make clear that in general claims like SU are not about "understanding" in a psychological sense, i.e. some kind of pleasant "feeling of satisfaction" that one develops. Scientists in SWC seem concerned instead with something like "scientific understanding" as characterized by Salmon (1998, Ch.5). According to Salmon, scientific understanding comes in at least two sub-varieties that involve, respectively (1) developing a world-picture into which we can fit natural phenomena that we encounter and (2) obtaining knowledge of how things work in the world, i.e. of the mechanisms (typically causal) that produce the phenomena that we encounter (Salmon 1998, Ch.5). Deductive-nomological (and unificationist) explanations on the one hand and causal-mechanical explanations on the other hand contribute to

[58] The trade-off that is at issue here sounds much like that identified by Cartwright (1983) between truth and explanatory power. However, she is concerned with the relative explanatory power of fundamental and phenomenological *laws*, whereas the present chapter is concerned with the contributions of simple and complex *models*. Since both simple and complex climate models typically incorporate Cartwright's phenomenological laws, there does not seem to be a straightforward mapping of these models to Cartwright's laws. In future work, I would like to further explore the relation between Cartwright's analyses and those of the present (and next) chapter.

[59] Wimsatt (1987) considers a related question. However, he focuses on the use of "false" models in the search for better models and theories. The present discussion focuses instead on the use of simplified models when "better" models (using Wimsatt's terminology) and accepted background theory are already available. Still, some of the ways in which so-called false models are used in the search for better models (according to Wimsatt) seem closely related to the ways in which simple models are used in order to promote understanding in SWC. I will point out these similarities below.

understanding in the senses of sub-varieties (1) and (2), respectively. As we will see, both sub-varieties of understanding are invoked in SWC.[60]

In Section 2, I give evidence that one view of understanding common in SWC is similar to the causal sub-variety identified by Salmon. In Section 3, I show that there are ways in which simple models can outperform complex models in contributing to understanding in this causal sense; this constitutes preliminary evidence in favor of claims like SU. I argue in Section 4, however, that SU is ultimately untenable under the causal view of understanding, because complex models also contribute substantially to understanding in the causal sense, in ways that I indicate. In Section 5, I suggest that SU fares even worse under an alternative/secondary view of understanding in SWC according to which understanding increases when phenomena are explained in something like Hempel's (1965) Deductive-Nomological (DN) sense. I conclude in Section 6 that SU is an oversimplification, and I offer a more nuanced view of the value of simple and complex models in promoting understanding in SWC.

5.2 Understanding and explanation in SWC: causal knowledge

In this section, I present evidence that one conception of understanding in SWC is closely related to learning about "how things work" in the world. In particular, the evidence suggests that in SWC understanding is increased when we obtain information about or "insight into" causal dependencies and causal mechanisms. In line with more general remarks on understanding and explanation by Salmon (1998, Ch.5), I will then suggest that this causal information is thought to increase understanding in SWC at least in part because it aids in the construction of causal explanations of phenomena exhibited by the atmosphere and climate system.

5.2.1 A causal view of understanding in SWC

Some of the most direct statements of SU come from Edward Lorenz, an atmospheric scientist who is probably most famous for his work in chaos theory. Lorenz discusses the role of simple and complex models in several places. For instance, in a

[60] I will not attempt to make much more precise Salmon's sub-varieties of understanding. The present paper aims to make sense of what goes on in SWC, rather than to critique possible accounts of understanding. For purposes of this paper, working with Salmon's somewhat loosely characterized notions of understanding is adequate.

1960 paper in which he attempts to find maximally simplified mathematical models of large-scale atmospheric behavior, he explicitly identifies prediction and explanation of atmospheric phenomena as the goals of the dynamic meteorologist (1960, 243), but he suggests that these goals are achieved using different types of models. According to Lorenz, even producing perfect forecasts via perfect equations and exact initial conditions "would not by itself increase our understanding of the atmosphere, no matter how important it might be from other considerations" (Lorenz 1960, 243). For example, if we predicted the occurrence of a particular hurricane using such a perfect complex model:

> We might still be justified in asking why the hurricane formed. The answer that the physical laws required a hurricane to form from the given antecedent conditions might not satisfy us, since we were aware of that fact even before integrating the equations. (Lorenz 1960, 243-244)[61]

According to Lorenz, while such complex models can make valuable predictions, "it is only when we use systematically imperfect equations or initial conditions that we can begin to gain further understanding of the phenomena which we observe" (Lorenz 1960, 244). This is because by comparing the results obtained using systematically imperfect models—ones that fail to represent features of the atmosphere that might be causally relevant to the production of the phenomena of interest—with what occurs in the real world, we "gain some insight concerning the relative importance of the retained and omitted features" (Lorenz 1960, 244). That is, we learn something about the causal contributions of the retained and omitted features in the production of the phenomena of interest.[62]

Lorenz considers the scenario of something like a perfect model again in several later discussions. He asks what would be gained by obtaining a "super-model" of the climate system (1970, 329) or a perfect model of the general circulation of the atmosphere (1967, 134). He suggests that solving the equations of the climate super-model so as to simulate the past climate "may give us little insight as to why the [climatic] changes took place," because we would not be able to tell which feature or set of features might have produced the changes (1970, 329). Likewise, reproducing the general circulation of the atmosphere in all of its relevant details:

[61] Regarding the last sentence of the quotation, Lorenz seems to be assuming that we also *know* that the perfect model is a perfect model.
[62] As indicated below, these "features" might include conditions, processes, and/or forces.

> …would not, however, necessarily increase our physical insight. The total behavior of the circulation is so complex that the relative importance of various physical features, such as the Earth's topography and the presence of water, is no more evident from an examination of numerical solutions than from direct observations of the real atmosphere (1967, 134).

In both cases, Lorenz points out that it would be a shortcoming of the very complex model that it would not give us much physical insight into the causal contributions made by various factors in producing observed phenomena. It is only by using simplified models that we begin to discern these causal contributions.

Schneider and Dickinson (1974) also discuss modeling methodology and the roles of models of different complexity. They distinguish "mechanistic" models and "simulation" models, which are varieties of "simple" and "complex" models, respectively.[63] According to Schneider and Dickinson, mechanistic models typically "investigate a single mechanism or a small number of simply coupled mechanisms, taking as being given…all the other components of climate" with the goal of "understanding the dependence of the particular mechanism on the other parameters of the problem" (1974, 456). Simulation models, by contrast, "include as internal variables as many interacting physical processes as possible," but these models are so complicated that it is "difficult to trace cause and effect relationships" (Schneider and Dickinson 1974, 456). These complex models "usually require a great deal of analysis and computer time in order to provide much understanding of the individual mechanisms and their dependence on each other" (Schneider and Dickinson 1974, 456). Schneider and Dickinson go on to claim that, while constructing comprehensive, realistic models is one "ultimate objective" of climate modeling, this construction probably cannot be done successfully "without the understanding derived from simpler models of individual processes" (1974, 456)—that is, without knowledge of causal dependencies among the various features and mechanisms that might be represented in the complex models.

A more recent anthology (Newton and Holopainen 1990) confirms—without explicitly discussing the nature of understanding—that understanding in SWC often is closely related to learning about causal dependencies. This collection of review papers,

[63] As is also the case with the distinction between simple and complex models, Schneider and Dickinson emphasize that the distinction between mechanistic and simulation models is not sharply defined and does not "do justice to the variety of modeling approaches that have been and will be used" (1974, 456).

which focuses on the topic of extratropical cyclones, contains two papers whose titles indicate that they are about "advances in understanding" of cyclones and cyclone development, respectively (Reed 1990 and Anthes 1990). In each paper, substantial space is devoted to discussing what has been learned about the ways in which various conditions and processes, either singly or in combination, influence the structure and development of cyclones. The implication is that the advance in understanding is constituted by this knowledge of causal dependencies and the causal roles of conditions and processes.

In present-day climate modeling literature, SU continues to surface in the context of discussions about the value of models of different complexity. These models have been organized and presented in terms of hierarchies or pyramids, which purport to show some of the relations among the models as well as their individual features. Shackley et al. (1998) take issue with such hierarchies and pyramids, which typically have simple models at their bottoms, claiming that the hierarchies and pyramids are interpreted normatively and thus contribute to an undervaluing of simple models in the context of SWC. Shackley et al. provide quotations from interviews with various unidentified present-day climate modelers in order to illustrate what they (i.e. Shackley et al.) consider to be standard views about the utility of simple and complex models. One such modeler is quoted as saying:

> Simpler models are useful for isolating and illustrating points in an easily understandable manner. ...They are essential for understanding but useless for prediction. (1998, 165)

When Shackley et al. summarize prevailing sentiment about simple models, they conclude that "simpler models are portrayed as useful *heuristics* to generate insights about the processes involved in climate change..." (1998, 165, italics in original).[64] As in previous discussions, simple models are said to help us obtain information about the causal contributions of features of the atmosphere or climate system; in this case the features of interest are those involved in the production of climatic change.

[64] Shackley et al. are describing what they take to be the generally perceived role of simple models in SWC, not what they think the role of simple models ought to be or must be.

5.2.2 A causal view of explanation in SWC

In philosophical analyses, it is rather common for understanding and explanation to be closely related as follows: explanation promotes or increases understanding. This relation is assumed, for instance, in the discussion given by Salmon (1998, Ch.5) mentioned above—causal explanations promote the sub-variety of understanding that involves knowing how things work in the world. In accord with this idea, I suggest that the causal information described in the previous section is thought to advance understanding in SWC in part because it helps us to construct causal explanations. The evidence that causal explanation is a preferred mode of explanation is primarily indirect, since discussion of such philosophical matters is rather rare in published works in SWC. I offer the following as preliminary evidence.[65]

Turning first to the early history of numerical weather prediction, a causal view of explanation appears to be assumed in several places by Jule Charney, who was centrally involved in the first attempts to simulate atmospheric motions using electronic computers. For instance, in a paper on numerical methods in dynamical meteorology, he notes that if, after surveying atmospheric phenomena, one

> ...were to conclude that methodologically the first task of the dynamical meteorologist should have been the explanation of the large-scale transient flows [e.g. extratropical cyclones], it may come as a considerable surprise that until the 1930's meteorologists did not understand the cause of the motion of the migratory vortices, or even why they move most often in an easterly direction, not to speak of the cause of their origin.... (1955, 799)

Here, the implication is that explanation involves knowledge of causes and the telling of "why" some phenomenon occurs. A similar implication is made by Charney in a paper published the following year:

[65] Winsberg (1999a) argues that explanatory practice associated with the simulation of complex physical systems does not fit with causal-mechanical explanation as it has been characterized by people like Salmon. While I question some of the reasons that Winsberg gives, I agree with him that in many cases in the study of complex physical systems, explanations are not given strictly in terms of the fundamental variables that appear in the equations of the associated theory; the explanations are not always "micro-level" causal explanations (as I claim below). But, despite this, I do think that the explanations given in many of these cases should be considered causal explanations. As Machamer et al. (2000) have suggested, causal-mechanical explanations can be given on a number of different levels of description. In any case, I am not arguing that causal explanations in SWC fit precisely with any extant philosophical account of causal explanation.

> Two problems that have occupied the major attention of dynamic meteorologists in recent years have been the explanation of (a) the generation and motion of the great migratory cyclonic and anticyclonic vortices of middle latitudes, and (b) the mechanism by which the mean zonally averaged circulation of the atmosphere is maintained against friction. (1956, 323)

For other historical evidence of a preference for a causal view of explanation, we can again turn to the published work of Lorenz. He remarks in more than one place that "we" or "some people" may not be satisfied with explanations of phenomena that simply tell us that the physical laws required the phenomena to form from the antecedent conditions (see e.g. 1960, 243 and 1966, 418). What is lacking in this type of explanation "...is a real physical insight into the mechanism through which" the phenomenon is produced (Lorenz 1966, 418). Speaking of such a missing mechanism for one particular phenomenon, he remarks that "if there is a simple process which could readily be described in a qualitative manner, it has so far been obscured by the complexity of the total problem" (Lorenz 1966, 418).

More generally, textbooks and research monographs in SWC provide evidence for something like a causal-mechanical view of explanation in SWC; these are full of discussions of causal dependencies and causal stories (clarified below), and in practice it is often these discussions that are cited as the "explanations" of phenomena in SWC, not the computer simulations which successfully produce as part of their output statements describing the phenomena.

In light of this evidence and drawing on my own experience in SWC, I would characterize the causal view of explanation in SWC as one according to which (a) an explanation is an answer to a "why" or "how" question and (b) adequate answers to these questions typically involve identifying causal dependencies and, preferably, telling causal stories.[66] Let us consider (a) and (b) in more detail. (a) is meant to indicate that causal explanations are linguistic things (perhaps including equations) and are produced by

[66] I want to emphasize that I am not trying to offer a definitive philosophical account of scientific explanation or even a definitive account of causal explanation. It is well known that these projects are onerous and fraught with difficulties. In fact, I am not even trying to give a definitive account of causal explanation *in SWC*. Scientists in SWC do not pretend to be working with a well-defined, detailed, and consistent account; they simply make remarks here and there that have to do with explanations involving causes, as illustrated above. My modest aim is to extract from these remarks and from my observation of practice in SWC some general features of those explanations (as identified by practitioners in SWC) that involve causes. I do not need more than this in order to carry out my evaluation of claims like SU.

humans.[67] (b) is meant to indicate what kind of content should appear in these linguistic constructions if they are to be deemed acceptable by practitioners in SWC. In particular, they at least should identify which causal factors play a role in the production of the phenomenon of interest (causal dependencies) and hopefully also should tell *how* it is that those factors bring about the phenomenon to be explained (causal stories). To explicate further: by a causal story, I mean a discussion, whether purely qualitative or involving quantitative information, of how the members of a set of causal factors work together, whether all at once or in succession, to bring about the phenomenon to be explained.[68] By causal factors, I mean physical conditions (e.g. a type of temperature field), forces (e.g. pressure gradient and gravitational forces), and micro-scale or macro-scale processes (e.g. evaporation, tilting of an updraft) that are part of the accepted ontology of SWC (i.e. not astrological powers, human desires, etc.). These causal factors need not be micro-level factors; they can include such things as hurricanes, cold fronts, updrafts, etc.[69] A causal story likely will involve mechanism descriptions as articulated by Machamer et al. (2000), but a causal story can also be something as simple as a description of a force analysis that shows why a system exhibits some characteristic of interest, as the following example illustrates.

Suppose that an explanation is requested as follows: Why is there a narrow jet stream near the tropopause in the middle latitudes of the northern hemisphere? An answer to this question that provides an explanation in the form of a causal story might proceed as follows:

> The general circulation of the atmosphere can be divided into three zones in each hemisphere—polar, midlatitude/subtropical, and tropical—with somewhat steeper horizontal temperature gradients at the boundaries between the zones than within the zones. At the boundary between the polar and subtropical regions, especially strong horizontal temperature gradients can exist. Because the pressure is decreasing more slowly per unit height on the (less dense) warm air side of the gradient than on the

[67] It seems possible in principle that verbal discussions constituting explanations could be generated by computers, but this has not yet been the case in SWC.

[68] This does not seem too different from Cartwright's notion of causal story (see e.g. 1983, Chapter 4).

[69] That is, in SWC, causal stories can include elements that do NOT appear explicitly in the equations included in an associated computer simulation model; causal stories might include aggregate/larger-scale processes. This is an interesting feature of explanatory practice in SWC that deserves further attention but is beyond the scope of the present discussion (see Christie 2001 for a preliminary discussion and also comments in Winsberg 1999a).

> (more dense) cold air side of the gradient, there is a pressure difference horizontally across the temperature gradient at a given height above the ground, and this pressure difference increases with height until the temperature gradient begins to diminish near the tropopause. Thus, a pressure gradient force whose magnitude increases with height is exerted in the cold direction on air in the region of the temperature gradient. When the temperature decreases to the north (as it does on average from equator to pole in the northern hemisphere), this force is a northward force. In addition, the Coriolis force (due to the Earth's rotation) is acting and deflects air flow toward the right in the northern hemisphere. The result is a narrow region of accelerated eastward flow that reaches a maximum near the tropopause above the polar/subtropical boundary—this is the subtropical jet stream.

This causal story explains the observed movement of air (i.e. the jet stream) as the result of the simultaneous action of forces whose presence depends upon physical conditions in the atmosphere and the properties of fluids (e.g. temperature gradients, density relations). It is the kind of causal story that might appear in an introductory meteorology textbook; in more advanced contexts, such an explanation of the jet stream might incorporate equations (e.g. the thermal wind equation) and/or omit information assumed to be already known. As indicated in (b) above, a causal story explaining some other phenomenon might involve a succession of causally connected events whose termination is the phenomenon to be explained. The factors mentioned in a causal story may themselves require explanation in other contexts, but they are taken as basic or foundational in that story. Thus, a causal story explaining something other than the presence of the jet stream may make reference to (and treat as basic) the jet stream as a factor that can contribute to the production of the phenomenon of interest.

If, as I and many others assume, explanation contributes to understanding, then the "explanatory value" of a computer simulation model is an indicator (perhaps imperfect) of how much that model can aid our quest for understanding. On my view, the explanatory value of a model in some domain is defined by its usefulness as a resource for the construction of explanations in that domain. Put somewhat differently, a model has explanatory value to the extent that it can help us to obtain explanatory information. In the spirit of Railton (1981, 240), by "explanatory information" I mean statements that reduce our uncertainty about the form and content of a sought-after explanation. When causal explanations that answer why/how questions are sought, explanatory value is a function of both facts about the model and facts about the people constructing the explanations, including their background knowledge and cognitive capacities. That is, because the causal explanations are ultimately of our own construction, whether a model is a useful resource for the construction of these explanations depends on both what the model can provide (e.g. true and relevant data about causal factors) and what we are able to do with what the model provides. For instance, if model output comes in the form of long lists of numbers, but we are not able to organize or manipulate those lists in ways that make salient any relevant causal information (e.g. before advanced visualization tools were developed), then the model may not have much explanatory value for us at that time; it will not be very helpful as we seek answers to our why/how questions. Thus, under the sketched causal view of explanation that I claim is common in SWC, the explanatory value of a computer simulation model is relative to a community, a time, and a set of questions to be answered. Keeping this in mind will help us to see why simple models can outperform complex models in certain ways when it comes to advancing understanding in the causal sense. This is the topic of the next section.

5.3 Evidence in favor of SU

Let us first review the sense in which simple models of weather and climate are relatively poor characterizations of their target systems (i.e. the atmosphere and climate system). These models generally fail to represent one or more spatial dimensions of their target systems, omit factors thought to be relevant to the functioning of the target systems, and often include very oversimplified characterizations of factors that they do include. Compared to complex models, which represent all spatial dimensions and a larger subset of the factors that are thought to be important, with many of these representations based soundly on physical theory, the simple models are relatively poor characterizations of their target systems.

How then might simple models like these nevertheless help us to understand phenomena, while complex models do not? That is, how could these simple models serve as valuable resources for identifying causal dependencies and constructing causal stories, in ways that complex models could not? I want to discuss two ways in which simple models can outperform complex models as resources for these purposes. This discussion will allow us to see why scientists in SWC might be led to claim something like SU.

5.3.1 The basics: identifying causal dependencies and essential causal processes

The climate system is made up of numerous nonlinear and interactive processes, which makes it difficult to infer causal relations simply by observing it in action. Even if one has some idea of which causal factors might be acting, it is another matter to figure out which combinations of factors are responsible for which of the phenomena observed to occur. The same is true of complex model behavior. As many modelers have remarked, determining why things happen as they do in complex model simulations is often very nearly as difficult as figuring out why things happen as they do in the real atmosphere (e.g. North et al. 1981, von Storch and Navarra 1999). Complex model simulations give us evidence that the set of factors represented in the model is sufficient for producing the phenomena that appear in the simulated atmosphere or climate system, but we often do not know, just by looking, which represented members of the set actually contribute to the production of which features of the simulation. It may be that only a subset of the represented factors actually has any role in producing a phenomenon of interest, but which subset it is may not be at all obvious.

Simplified models can help us to identify these causal dependencies in complex models and the real systems they model. Simple models can be used to explore which subset of factors is necessary and sufficient for producing something like the phenomenon of interest. A good example of this use of simple models can be found in the early history of computer simulation of the atmosphere and climate system. One of the main goals of the first attempts to simulate the atmosphere and climate system was to gain insight into the mechanisms by which various gross, salient features of these systems were produced (see e.g. 1946 report in Appendix of Thompson 1983/1990). By gross, salient features, I mean such things as the transient large-scale vortices that populate the atmosphere in the middle latitudes (i.e. extratropical cyclones and anticyclones), the general increase of potential temperature with height in the atmosphere, the magnitude of the poleward decrease of temperature in the troposphere, the distribution of mean zonal wind, etc. (see Thompson 1983/1990 and Phillips 1956, 124).[70]

The explicit approach taken was to formulate extremely simplified models and then build up incrementally to more and more complex models (see e.g. Charney 1949, Phillips 1956, and Lorenz 1960).[71] In this way, one could observe how the behavior of the model system changed with each addition of a new physical factor (e.g. frictional drag or a primitive hydrologic cycle). One could observe what happened in the absence of particular factors and then see how the addition of those factors—one at a time—changed the way the model system behaved. In this way, one could begin to determine which causal factors might be essential to the production of which features of the real atmosphere or climate system (see e.g. Charney 1951, 253). One could also gain some confidence that some small set of causal factors that one hypothesized to be essential for the production of various phenomena was at least sufficient for producing something like those phenomena (see e.g. Charney 1972, 122).

Using simple models in this way can yield surprising results. Sometimes it turns out that the simulations produced with simple models include features that previously were thought to result from the action of factors not included in those simple models. Such a situation arose in the early history of computer simulation of the atmosphere. Contrary to expectations, a very simplified model produced simulations that displayed a kind of primitive "development" of the type associated with extratropical cyclones. As one modeler put it:

> If these forecasts had been made instead from the complete equations of motion, it probably would not have been at all clear that these developments were essentially barotropic and quasi-geostrophic in nature. (Phillips 1951, 393)

[70] Explaining characteristic features of a system seems to be a primary concern of Batterman in his recent work on understanding and explanation (2000, 2001). He focuses on "asymptotic analysis" (an analytical methodology) as a method for achieving or promoting understanding of gross, salient features of a system. Batterman takes himself to be arguing against the adequacy of the causal-mechanical account of explanation (at least for some purposes), since he thinks that the causal-mechanical view calls for revealing mechanisms in their full, overwhelming detail. He wants to highlight what he takes to be a different kind of explanation and understanding—one in which the "basic features" of the system or phenomenon of interest are "transparently exhibited by the mathematics" as a consequence of asymptotic analysis (2000, 252).

[71] A similar function of "false" models is identified by Wimsatt: "An oversimplified model may act as a starting point in a series of models of increasing complexity and realism" (1987, 30).

The precise meanings of "barotropic" and "quasi-geostrophic" are not important; the point is that using complex models would have obscured the fact that these atmospheric developments might be produced under conditions other than those previously thought essential to their production. Of course, there is no guarantee that the behavior of the real atmosphere or complex model does result from the factors that produced similar behavior in the simple model atmosphere (as e.g. Phillips (1956, 154) was quick to point out), but at least one has a starting point.

It should be noted that very simple models can take us only so far in this way, since they can produce simulated atmospheres and climate systems that resemble the real systems only in rather limited ways. There are many phenomena whose essential causal ingredients go far beyond those included in very simple models and/or that require representations of those ingredients in a more complete way than is done in very simple models (see Andrews 2000, 211 for a similar point and examples). Even primitive forms of these phenomena will not show up in very simple model simulations. However, as one builds up to more and more complex models, one can keep in mind what was learned about causal dependencies from simpler models and thereby will be in a better position to hypothesize about the causal factors responsible for features exhibited in the more complex model simulations.

Thus, there was good reason for employing simple models in the early computer simulation studies of the atmosphere and climate system: the modelers wanted to identify which factors played a role in the production of salient atmospheric and climatic phenomena, and employing the most comprehensive and complex model that could have been developed (and much more complete hydrodynamic equations *were* available at the time) would not have helped much in reaching this goal, even if such a model would have been a useful predictive tool. If such a complex model had been used right away, it would not have been clear to the modelers which causal factors contributed to the production of which predicted phenomena, because these dependencies would have been obscured by the detail of the simulation. At that time, given their limited knowledge of the workings of the atmosphere and climate system and given the set of questions that the modelers wanted to answer, simple models were more useful resources for learning about the relevant causal dependencies than complex models.

This use of simple models did not end with the early days of computer simulation of the atmosphere and climate system. For example, making direct reference to the hierarchical approach to modeling pursued in those earlier decades, a modeler remarked some thirty years later:

> But as models become more complex, it is difficult, with highly nonlinear and interactive processes, to say why we obtain a given result. ...One must also be prepared to go backward, hierarchically speaking, in order to isolate essential processes responsible for results observed from more comprehensive models. (Smagorinsky 1983, 37)

Going backward in the hierarchy does not necessarily mean going all the way back to the simplest models. It might involve selecting a few causal factors suspected to be important in the production of the result and running simulations in which only these factors (perhaps in various combinations) are represented (or represented in detail). Because they are not too demanding computationally, simple models can be used to run many simulations, exploring many combinations of causal factors. Simple models are still used in this way to gain insight into causal dependencies that might obtain in more complex models and in the real systems being modeled.[72] This use of simple models may be part of what motivates claims like SU.

5.3.2 The next step: cognitive accessibility and the construction of causal stories

A second way in which simple computer simulation models can outperform complex models in helping us to learn about causes relates closely to cognitive accessibility. When the goal is to explain some feature of a phenomenon, it generally is not enough to simply identify causal dependencies in the form of input/output pairs, e.g to show that if factors *A*,*B*,*C* are present at time t_0, then feature *F* emerges at time t_n. It is better if one can tell, in a step-by-step fashion, how these factors interact and work together to produce the feature of interest. That is, it is better if a causal story detailing events between t_0 and t_n can be told. It is worth briefly discussing why having a causal story is preferable. The main reason, it would seem, is that one is better prepared to predict the effects of interventions on the system at times between t_0 and t_n. For example, suppose we know that warmer-than-usual ocean temperatures off the coast of New

[72] This use of simple models seems similar to the following function of "false" models as identified by Wimsatt: "An oversimplified model may provide a simpler model for answering questions about the properties of more complex models that also appear in the simpler case, and answers derived here can sometimes be extended to cover the more complex models" (1987, 31).

Zealand are typically followed one month later by colder-than-usual ocean temperatures near San Diego. Suppose that we then observe the warmer-than-usual temperatures near New Zealand. Unless we know something about how the warmer temperatures eventually lead to the colder temperatures (assuming that there is some causal connection), we may have no idea whether changing ship routes near Japan would prevent the occurrence of cold temperatures next month near San Diego. Although a causal story may not always provide the information needed to make good predictions about counterfactual interventions, in general a causal story will give us more of the relevant information than mere input/output pairs. Put succinctly, having a causal story generally expands the range of interventions about which one can reason counterfactually. A causal story typically is a richer inferential tool than a set of input/output causal dependence relations.

How could simple models be especially helpful when it comes to constructing causal stories? When factors represented in the model are few in number and take a rather simplified form (e.g. have been linearized or include only a few terms), it is easier on average to construct a causal story about why particular features of the simulation were present than when the factors represented are many, interrelated, and nonlinear. This is because, for a very simple model, we may be able to imagine without too much difficulty the likely results of interactions among the small number of represented factors, which will make it easier for us to "follow" the behavior of the model. As we observe the behavior of the model, we may be able to tell ourselves a plausible story about how the behavior of the model through time results from the represented factors. It may be exceedingly difficult to do this (with any confidence) when observing the behavior of a complex model, at least for many features of complex model behavior. Thus, simple models can be of more value when it comes to constructing causal stories than complex models, because their behavior is (on average) more cognitively accessible than that of complex models.

In fact, simple models are often used in order to construct causal stories not only about the behavior of real-world systems of interest, like the climate system, but also about the behavior of complex models. By studying the behavior of the simple models, and constructing causal stories about how physical processes represented in the simple models might work together to bring about various features of the simulations, scientists can form some idea of how these mechanisms might be acting in more complex models (see e.g. Brovkin et al. 1998 and comments in Andrews 2000, 204). The causal stories developed on the basis of simple model behavior can serve as a starting point for interpreting complex model behavior.

We now have seen two ways in which simple models can outperform complex models in advancing understanding in the causal sense, namely, by helping us to identify which causal factors are key in the production of which basic features of the atmosphere and climate system and by helping us to construct causal stories about how the members of a set of causal factors might work together to produce phenomena of interest. The fact that simple models can help in these ways may be part of what leads to claims like SU. [73]

5.4 Complex models and causal understanding

A more thorough look reveals, however, that complex models also contribute to understanding in the causal sense. I will describe two ways in which complex models outperform simple models in increasing understanding in the causal sense. This suggests that SU is ultimately untenable under the causal view of understanding in SWC.

5.4.1 Surrogate observational resources

One way in which complex models can have explanatory value (and thus help to advance understanding in the causal sense) is by serving as sources of "observational" data. In many situations in SWC, we would like to know something about the magnitude of various physical quantities at particular times and locations (in the real world). For instance, we would like to know what the pressure field looks like both near and high above the ground in the vicinity of a front and how that field changes over the course of a few days. Or we would like to know what the vorticity field looks like in a supercell thunderstorm before a tornado develops. In practice, it may be very difficult to obtain this information observationally—it is simply not easy, for a variety of reasons, to make measurements that reveal the structures of these fields. Yet obtaining this information

[73] The foregoing discussion should remind us that simple models have had a role to play in promoting causal understanding not merely in classrooms (i.e. as pedagogical tools) but at the frontiers of research in SWC. This fact is sometimes overlooked.

might be especially valuable when it comes to constructing causal explanations. For instance, our ability to construct an explanation of the development of the tornado (as a phenomenon) might depend crucially on knowing what the vorticity field typically looks like just before tornadoes form.

Complex models have the potential to provide us with such information. For any variable included in the complex model, we can view its value at every spatial location and every time represented in a simulation. In this way, we can "look inside" a simulated phenomenon in a way that we are not able to do in the real world. Visual display tools may be especially useful in this regard. The data generated by the simulation thus becomes a surrogate for observational data; the simulated vorticity field serves as a surrogate for the actual vorticity field. Of course, there is no guarantee that the information that the simulation provides will be accurate; complex models are not perfect, and there will be some distortions in complex model simulations. We need to devise tests of hypotheses that we form on the basis of the simulated data, but this is true of any hypothesis.

In general, complex models outperform simple models in providing such surrogate data. This is both because complex models generally incorporate fewer and less severe distortions than simple models (and thus can be seen as more reliable sources of information on average) and because complex models explicitly calculate values for many more variables—and with higher spatiotemporal resolution—than simple models do. For example, a simple climate model might not calculate any dynamics for the atmosphere, might represent the ocean as a one-dimensional column, and might prescribe (rather than calculate) such things as the climate sensitivity. By contrast, a complex climate model might explicitly calculate values for a very large number of variables and does so with relatively fine spatiotemporal resolution. The output of a complex model can be a rich source of information that aids us in the construction of causal explanations and thereby increases understanding in the causal sense.

5.4.2 More on causal dependencies: sensitivity studies and hypothesis testing

As noted in Section 5.3.1, simple models can only take us so far when it comes to identifying causal dependencies. This is because simple models can produce simulated climate systems that resemble the real system only in rather limited ways, so many phenomena will not appear even in a primitive form in simple model simulations. Simple models may outperform complex models when it comes to identifying which causal factors contribute to the production of some basic features of the climate system, but complex models can be more useful for investigating phenomena whose production relies upon the complicated interaction of several causal factors. This is not to say that investigating such phenomena will be a straightforward and easy task using complex models. But there are ways in which complex models can be especially useful for learning about causal dependences in situations where many causal factors are acting.

One way is via so-called "sensitivity studies" in which various causal factors are turned on or off in turn in order to see the effects on the simulated phenomena (see also Section 3.5.4). This practice is not unrelated to that described in Section 5.3.1 above, but the models used in these sensitivity studies are generally much more complex than those referred to in Section 5.3.1, and the idea is not to isolate a few processes thought to be important to the production of a phenomenon but rather to take a very complex model that does a relatively good job of simulating the phenomenon of interest and successively turn off or manipulate individual causal factors thought to influence the production of the phenomenon.[74] This can be costly from a computational point of view, so there are practical limitations on the extent to which the properties of a model can be explored in this way, but sensitivity studies do have the potential to provide valuable information. Referring to the study of past climatic changes, Lorenz notes that, by carrying out sensitivity studies on the (hypothetical) super-model of climate mentioned above, we eventually might be able to "say what features or combinations of features *could have* produced the changes" (1970, 329, italics in original). Likewise, Reed (1990, 38) notes the importance of sensitivity studies in understanding the physical processes that contribute to the formation of extratropical cyclones.

In a similar way, complex models can be used to test hypotheses about phenomena. If we hypothesize that the structure of the vorticity field in a thunderstorm is a critical causal factor in determining whether a tornado will form, we can test our

[74] This practice seems to correspond roughly to the following function of "false" models identified by Wimsatt: "An incorrect simpler model can be used as a reference standard to evaluate causal claims about the effects of variables left out of it but included in more complete models, or in different competing models to determine how these models fare if these variables are left out" (1987, 31).

hypothesis by changing the structure of the vorticity field and checking to see if the tornado still forms in the simulation. In general, we cannot perform these types of experiments in the real atmosphere or climate system. And, in many cases, simple models will not be appropriate tools for testing such hypotheses, because they may be known to omit potentially important causal factors and/or because they may be unable to simulate the phenomenon of interest at all.

Thus, it seems that SU is untenable under the causal view of understanding that is common in SWC. Complex models can outperform simple models in advancing understanding in the causal sense by serving as sources of surrogate data and by allowing us to investigate and test hypothesized causal dependencies in complicated situations that cannot be simulated by simple models. In these ways, complex models can provide us with explanatory information that aids us in the construction of causal explanations and thereby advances understanding in the causal sense.

5.5 Understanding via DN-like explanation: more trouble for SU

The causal view of explanation discussed in Section 5.2 is not the only view that is at work in SWC, even if it a rather common one. In fact, the very authors who sometimes seem to assume a causal view of explanation will, at other times, acknowledge at least implicitly a view of explanation that resembles Hempel's DN account. For example, Lorenz remarks that "mathematical solutions do constitute acceptable explanations for many physical phenomena" (1966, 418). The phenomena are explained because statements describing their occurrence are shown to follow from the system of equations assumed to govern atmospheric motions (and these are incorporated approximately in the computer simulation model). Likewise, claims made by Charney hint at something like a DN account of explanation. For instance, he says that it is possible to "show by mathematical means" that the processes conjectured to drive the general circulation of the atmosphere do "in fact explain the gross character of the general circulation" (1959, 1650). This showing by mathematical means involves the generation of a computer simulation in which the large-scale motions in the simulated atmosphere look something like those of the real atmosphere.

Note that on a DN-like account of explanation, explanatory value need not depend much on our cognitive capacity or background knowledge. If, instead of seeking causal answers to why/how questions, we seek explanations in the form of derivations from assumptions plus initial and boundary conditions, then explaining a phenomenon might require showing simply that some statement or set of statements (describing the phenomenon) emerges as a conclusion in a derivation of a certain sort. Explanation need not require that we be able to reconstruct mentally the steps in the derivation or even that we be able to follow the derivation when studying it. We might simply accept a particular set of assumptions, conditions, and rules for derivation and acknowledge the ability of some device (living or machine) to apply the rules correctly to the assumptions and conditions. Then, anything derived by the device would be granted the status of "explained."

If one adopts a view of understanding according to which understanding is increased when phenomena are explained in something like a DN sense, complex models would seem to be of much greater explanatory value than simple models and thus would seem to far outperform simple models in increasing understanding. The easiest way to see this is to recognize that simple models can predict only a small subset of the phenomena that complex models are able to predict.

5.6 Conclusions

SU is an oversimplification when it comes to scientific practice in SWC. The analysis of the preceding sections has shown that a more nuanced characterization of the roles of simple and complex models in promoting understanding in SWC is needed. I offer the following characterization:

> (SU') The extent to which simple models can promote understanding depends upon the view of understanding that is adopted: if a deductive view is adopted, then simple models are far less helpful in promoting understanding than complex models, but if a causal view is adopted, then both simple and complex models can promote understanding of phenomena, although they typically do so in different ways.

It is interesting that under neither view of understanding is the original SU tenable. Why, then, would something like the view expressed in SU be relatively common? As suggested above, it likely has something to do with the fact that simple models can outperform complex models in certain ways when it comes to promoting causal understanding. But perhaps something more can be said. If we look at the ways in which simple models are especially useful in advancing causal understanding, we see that, at least in some cases (and certainly in the case of understanding the large-scale dynamical features of the atmosphere), simple models can play an important role in getting the project of causal understanding off the ground. They can be useful in *starting* the building of a stockpile of information about which causal factors might produce which observed basic features of the system of interest (and how they might produce them). Once we have this basic information, we go on to use more and more complex models to investigate further details of the system of interest. The value of simple models for promoting understanding may sometimes be exaggerated (as in SU) because simple models can play a central role in the foundational steps toward developing causal understanding of a system.

6

Using Incompatible Models Together: A Pragmatic Integrative Pluralism

6.1 Introduction

We have seen in previous chapters that the study of Earth's climate currently involves the use of a variety of computer simulation models. It is noteworthy that a substantial number of these are designed to be models of the climate system (as a whole), rather than complementary models of different components of that system. This raises the question: why so many models? After all, there is but one terrestrial climate system.

As we saw in the last chapter, one reason is that models of different complexities are useful for different modeling tasks. For example, simple climate models are used when computational expense is a constraining factor and when global average parameters are to be predicted. Complex models, on the other hand, must be relied upon in studies of regional climate change, since they represent all three spatial dimensions of the climate system in some detail. While the ability of complex models to simulate such regional change is still being investigated, it is agreed that very simple models will not be of much use for regional climate modeling tasks. Thus one dimension of model pluralism in climate science crosses levels of model complexity and exists primarily because, at a given time, different modeling tasks may be best undertaken using different types of climate models.

A second dimension of climate model pluralism is more puzzling. There are not just simple and complex models—there are many simple models and many complex models. For example, a recent report on model evaluation listed more than thirty coupled climate models of approximately equal complexity (see IPCC 2001, Chapter 8). Why does this pluralism within levels of model complexity exist? Let us focus on complex model pluralism in particular. In part, complex model pluralism reflects the fact that there are different mathematical techniques available to climate modelers. For instance, the atmosphere can be represented as a grid of points corresponding to volumes of atmosphere (as I have been assuming) or in terms of a series of waves of differing frequencies, known as a spectral representation (see Holton 1992, 450). The details of these techniques are not important for this discussion; the point is that there are at least two legitimate ways of setting up the mathematics of a climate model. So, among complex models, we have some with grid-point representations of the atmosphere and some with spectral representations of the atmosphere. These models can be based upon the same physical assumptions about large-scale atmospheric dynamics (and often are), even if they handle the mathematical treatment of physical equations in different ways. But this explains only part of the diversity, since complex models also typically differ in some of the assumptions that they make about climate system processes and hence in some of the predictions and retrodictions that they make. That is, the complex models generally are logically incompatible with one another.

At this point we must ask: what explains the persistence of a plurality of logically incompatible complex models? The explanation given for pluralism across levels of model complexity does not apply here—it is not that the different complex models are used for different purposes. Rather, scientists have been *unable* to select from among these incompatible models those that are most promising for the purpose of investigating future climate change. In Section 2, I will discuss several reasons why scientists have been unable to narrow the field of incompatible complex models: (1) There are difficulties in testing model predictions and retrodictions; (2) It is difficult to define an overall "figure of merit" for the simulations; (3) No model is clearly superior to the rest with respect to measures of simulation quality currently in use; (4) There is genuine scientific uncertainty about how to best represent the climate system.[75] In discussing (1) – (4), we will get some sense of why neither scientists nor philosophers find model evaluation to be a straightforward matter.

After this discussion, we will better understand why a plurality of logically incompatible complex climate models continues to exist, and we will be ready to confront an even more surprising feature of complex model pluralism: these logically incompatible models are used together as complementary resources for investigating

[75] I will be concerned with epistemic reasons, rather than social ones, although no doubt there are some social reasons for the persistence of so many models as well.

future climate change. This is peculiar indeed. When two logically incompatible theories are available, they typically are viewed as competitors, and scientists seek evidence that refutes one theory and supports the other. But logically incompatible climate models are currently regarded not as competitors but as a team of models to be used together. As I will argue in Section 3, this is not because climate scientists consider the models to be purely instrumental tools (in which case their incompatibility might not matter, so long as they helped scientists accomplish their goals). Rather, the attitude taken toward complex climate models is one that involves both instrumentalist and realist components. In Section 4, I will show how this mixed status of climate models, in combination with the difficulties faced in their evaluation (as expressed in (1) - (4)), leads scientists to use incompatible climate models as complementary resources for investigating future climate. I will discuss why the "multi-model ensemble" approach used is advantageous and indicate why we must nevertheless be careful in interpreting results obtained via this approach. Finally, in Section 5, I will suggest that this interesting use of incompatible models involves two kinds of model pluralism: ontic competitive pluralism and pragmatic integrative pluralism.

6.2 Explaining the persistence of a plurality of incompatible models

In this section, I will discuss (1) - (4) in more detail, explaining how they together promote to the continued existence of a plurality of incompatible complex climate models.

6.2.1 Difficulties in testing model predictions and retrodictions

If we are interested in identifying models that are most promising as predictive tools, we will want to consider their histories of predictive successes and failures. Unfortunately, for today's climate models, there are virtually no such track records. Today's models make predictions about what might happen ten or fifty or two hundred years from now under conditions that may or may not actually obtain during the intermediate years. This is why such predictions are typically referred to as "projections" instead; they are projections of what would happen if some set of circumstances were to obtain during the next few decades. Weather forecasting models, by contrast, make predictions about what will actually happen over time periods of hours, days or weeks.

For these models, we can and do compile much information about their predictive strengths and weaknesses. But for climate models, there is almost no such information, since the observational data that we need in order to assess the quality of their predictions will not be available, even in principle, for quite some time. The climate models simply cannot be compared with respect to their predictive track records.

In light of this situation, simulations of past and present climate conditions (i.e. "retrodictions") have become a focus of climate model evaluation. The task is to compare the model output with available observational datasets. One serious problem, however, is that we have data for only a few quantities (e.g. temperature, pressure, precipitation), for only relatively recent time periods, and for primarily land locations and near-surface locations, and even these records are incomplete and of variable reliability. We lack a solid observational foundation against which to compare even the retrodictions of climate models. But this is not the only problem encountered in evaluating the retrodictive capacities of climate models, as the next section illustrates.

6.2.2 Difficulty of defining an overall figure of merit

Another difficulty in evaluating climate model retrodictions stems from the vast amount of model output produced. Climate models generate output for thousands and thousands of grid points for years and years of simulated time and for numerous variables. How are we to judge the quality of this output overall? We do have measures for quantitatively assessing model-data fit for individual fields, but we have multiple such measures. Even if we were to privilege a small number of complementary measures for each individual field, we would still need to decide how to combine the scores received for individual fields into an overall "figure of merit" for each model. It is not at all obvious how this should be best done. According to the most recent IPCC chapter on model evaluation, "it has proved elusive to derive a fully comprehensive, multidimensional 'figure of merit' for climate models" (2001, 475). Thus far, scientists have been unable to narrow the field of incompatible complex models using some measure of the overall quality of their retrodictive performance.

One might wonder why a comprehensive figure of merit is needed at all. After all, if we are interested in predicting future temperature changes, why not just evaluate climate models according to how well they simulate temperature changes up until now?

One reason is that models may have been "tuned" to some degree in order to do a reasonably good job of reproducing the available observational temperature record. Tuning involves the manipulation of adjustable parameters in a model in order to bring its output closer to the available observational data. Tuning need not be informed by what is known about the physics of the system being simulated. In fact, the parameters being adjusted may not have any known correlate in the represented system but rather may be included in an ad hoc fashion expressly for the purpose of improving model-data fit. If a climate model is tuned in an ad hoc manner in order to approximately reproduce the available temperature record, there is no guarantee that it will do a similarly good job in predicting future temperatures, since its past successes may have had little to do with how well it described the physical processes acting in the climate system.

The more general reason that scientists would like to have a comprehensive figure of merit has to do with the nature of the climate system. Climate is thought to result from the interaction of numerous processes acting on a broad range of spatiotemporal scales (see Chapter 3). This means that errors in simulating one process may degrade the quality of other simulated variables. Thus it is desirable, even for the sake of prediction, to have climate models that perform well in simulating a range of climatic variables.[76]

6.2.3 No clearly superior model given present evaluations

Most recently, climate model evaluation has involved large "intercomparison" projects. The Coupled Model Intercomparison Projects (CMIP1 and CMIP2) are perhaps the best known of these. The projects require different modeling groups to carry out comparable simulations (e.g. using the same initialization fields and for the same simulated periods of time) and produce time series data for particular climatic variables of interest. Even if no comprehensive figure of merit has been developed, the data produced for these particular variables can be quantitatively compared with one another and with available observational data as long as some measure of model-data fit is selected. To date, no single measure has emerged as the "gold standard" for comparison, even for individual variables, and a variety of measures are in use. Still, it is possible that

[76] I am assuming here that we do not know much about the details of the interactions among climate system processes. We do have some knowledge of these interactions, and this can guide our decisions about what it is most important to "get right" if we want to make accurate predictions of future temperature changes, but detailed knowledge of these interactions is currently quite limited.

one complex model would outperform all others for most variables and for a wide variety of measures of model-data fit. In practice, however, it turns out that when model output is compared with available climatic datasets, no single model consistently scores best even for the limited set of variables and measures of fit that are selected. Instead, some models perform better for some fields and measures of fit, while other models perform better for other fields and measures of fit (see e.g. Lambert and Boer 2001 and IPCC 2001, 482).

The situation is further complicated by the fact that there are several different "observational" datasets with which model output might be compared. Because our observations of climate are incomplete and of variable reliability, it is not a simple matter to produce global climate datasets for use in model evaluation. As noted by Edwards (2001, 61), recent climate model evaluations have made use of data produced via "reanalysis" projects. These projects synthesize observational data and output from weather forecasting models to produce global datasets for hundreds of variables of interest, for the entire globe on a regular grid, and for regular time intervals. Models are used to fill in gaps in datasets, to interpolate observational data to particular grid points, and to derive non-observed fields (e.g. temperature advection and momentum exchange) from observed ones. Thus, some of the reanalysis data are determined almost entirely by observation, while other data are "completely determined by the model" (Kalnay et al. 1996). Given a particular type of field for comparison (e.g. global annual mean precipitation) and a particular measure of model-data fit (e.g. root mean square error), some climate models score better for one "observationally-based" dataset while other models scores better for another such dataset (see e.g. Figure 8.4 in IPCC 2001).[77] In the most recent IPCC report, the authors of the chapter on model evaluation go so far as to say that they "...do not believe it is objectively possible to state which model is 'best overall' for climate projection, since models differ amongst themselves (and with available observations) in many different ways" (2001, 475).[78]

---

[77] The use of reanalysis data in climate model evaluation should be of interest to philosophers of science. Weather forecasting models have many core assumptions in common with climate models, so a dataset whose content depends in part on such weather forecasting models may not be an appropriate resource for evaluating the quality of climate model simulations. The apparent model-data fit may be artificially inflated as a result of the shared assumptions of the weather and climate models. This issue of circularity in the testing of climate models has other complications and is worth pursuing, though I will not do so here.

[78] I think the intended claim is that they do not believe it possible to *identify objectively* which model is best overall for climate projection. Subjective model assessment is common in climate science. For examples of

6.2.4 Scientific uncertainty with respect to climate system representation

According to the last quote, no "best" model for climate projection can be identified in part because models differ amongst themselves in so many ways. At least some of these differences reflect scientists' uncertainty about the nature of processes acting in the climate system and about how various physical processes should be represented in climate models. Some physical processes are still poorly understood, and some occur on spatiotemporal scales smaller than those resolved by the models. In either case, it may be thought that the physical processes are important enough that we should somehow include them in our models, but there may be substantial uncertainty concerning how they can be best represented, leading to the development of logically incompatible representations of those processes in different climate models.

The case of clouds provides a good illustration of this. As mentioned in previous chapters, clouds play an important role in shaping climate. Initial warming due to increased greenhouse gas emissions might lead to changes in the amounts and types of clouds that form, thereby enhancing or offsetting the initial warming. Scientists would like to somehow include the effects of clouds in their climate models. However, individual clouds occur on scales that the models do not explicitly resolve, and there is genuine uncertainty about how clouds interact with the larger climate system and hence about how the effects of clouds can be best parameterized in terms of large-scale, resolved quantities.[79] As a result, several different parameterizations of clouds have been developed, reflecting different approaches to representing clouds within the bounds of our uncertainties.[80] These parameterizations incorporate conflicting assumptions and generally give somewhat different predictions about the effects of clouds on the larger climate system. In fact, clouds are an extremely problematic case: models incorporating different cloud parameterizations do not even agree on whether changes in cloud

---

subjective assessments see Dai et al. 2001, 515 and IPCC 2001, 479. The prevalence of subjective assessment of numerical models (not just climate models) is noted by Oreskes and Berlitz (2001).

[79] See Section 3.4.2.1 for a discussion of parameterization.

[80] While some differences among cloud parameterizations do reflect scientific uncertainty, some also reflect what might be termed "engineering uncertainty." By this I mean uncertainty concerning how to best incorporate (using large-scale, resolved variables) what we do know about clouds. Different parameterizations often are successful in different ways. For instance, one parameterization might output rather accurate values of average cloudiness for one geographical region, while another parameterization does better for another region.

formation due to increased greenhouse gas concentrations will have a net warming effect or a net cooling effect (IPCC 2001, 427-431).

Uncertainty associated with the representation of clouds and other poorly understood and/or sub-grid scale processes has led to the development of numerous complex climate models that differ from one another in many ways, as indicated in the quote from the last section. Scientific uncertainty is both a root of logical incompatibility among climate models and, in combination with the difficulties identified in (1) – (3), a reason for its continued existence.

6.3 The mixed status of complex climate models

Having explained the persistence of a plurality of incompatible complex climate models, we are left with the task of explaining why these incompatible models are viewed not as competitors but as a team of models to be used together in investigating future climate change. One possible explanation would be that climate scientists simply understand climate models to be purely instrumental tools, so that their logical incompatibility need not be troubling as long as they can be used individually in some effective manner. But this explanation seems most promising if the models are to be used for different purposes, and the situation that we are trying to explain is one in which incompatible models are used together in tackling the same modeling tasks.

Furthermore, it is clear that climate models are not considered by scientists to be purely instrumental tools. This can be seen, among other places, in scientists' own statements concerning the basis for confidence in climate models:

> Confidence in climate models depends partly upon their ability to simulate the current climate and recent climate changes, and partly upon the realistic representation of the physical processes that are important to the climate system. (IPCC 1996, 274)

In evaluating climate models, scientists are concerned with both the simulations that the models produce and the assumptions that the models incorporate. Contrary to what would be expected if the models were viewed purely instrumentally, it is not enough for a model to simulate with some accuracy the past and present climate; the model should give the right results (i.e. accurate simulations) for the right reasons (i.e. because the physics of the situation has been accurately described). So, a model may be praised or faulted either on the basis of how well its assumptions mesh with existing background knowledge about the climate system or on the basis of the perceived quality of its simulations.

An illustration of some scientists' concern over getting the right results for the wrong reasons can be found in the case of so-called "flux adjustments." When scientists began to join complex atmosphere models with complex ocean models (to produce complex coupled climate models), they observed that the climate simulated by the coupled models tended to slowly drift away from the equilibrium that was expected to be maintained. One primary factor contributing to this drift was a mismatch between the fluxes of energy at the atmosphere-ocean interface in the coupled models. To remedy the situation, ad hoc adjustments to the flux values were (and sometimes still are) made in order to keep them in line with one another and prevent the drift in the simulations. But the need for flux adjustments is thought by many scientists to indicate that the assumptions built into the coupled climate models are fundamentally deficient. Even scientists who take a somewhat more pragmatic view toward modeling seem to consider flux adjustments to be something of a "necessary evil" and agree that it is preferable for models to perform well without the need for flux adjustments (see e.g. the analysis in Shackley et al. 1999). If a purely instrumental view were taken, then the need for and implementation of flux adjustments would not be considered problematic.

Of course, as suggested above, not all climate scientists have exactly the same attitude toward climate models. In Shackley et al. 1999, it is argued that among climate scientists there are at least two different epistemological approaches to modeling, which are characterized as "purist" and "pragmatist" approaches. In effect, purists and pragmatists differ with respect to how closely they adhere to the "right results for the right reasons" requirement discussed above. Pragmatists tend to be less disturbed than purists by the introduction of ad hoc adjustments whose sole purpose is to improve the apparent fit between model output and available data. This is in part because purists often view simulation of the climate system as a scientific exercise that might advance theoretical knowledge, while the pragmatists often are concerned with simulating climate for purposes of aiding practical decision-making. This does not mean, however, that pragmatists have a purely instrumental view of models; rather, they simply have a somewhat greater (not infinite) tolerance for ad hoc maneuvers than purists do.

For both groups of modelers, the realistic representation of physical processes is considered important and desirable, and it seems likely that both groups would characterize the ideal modeling situation as one in which climate models were constructed entirely via straightforward application of well-established physical principles. There is a strong realist component to the perceived status of climate models: an ultimate aim is to develop models whose assumptions are approximately true of the real climate system.[81] At the same time, because global warming is perceived as an environmental problem that may require immediate action, and because climate models are the most promising resources available for answering questions about climate change, there is pressure to work around the models' present shortcomings and find a way to use them as tools to help answer these questions. In other words, there is also an instrumentalist component to the perceived status of climate models. Thus, climate models come to have a mixed status: they should incorporate realistic assumptions insofar as this is possible, but they also should be useful tools for addressing particular problems and questions. In the next section, we will see that this mixed status of climate models aligns well with the current usage of logically incompatible complex climate models.

6.4 The multi-model ensemble approach in climate modeling

Let us take stock of the situation in climate modeling, given the discussion in the previous sections. A collection of logically incompatible complex models has been developed, each of which is designed to be a realistic characterization of the climate system, insofar as this is possible. But even state-of-the-art complex climate models currently are constituted by a "balance of approximations" (Lambert and Boer 2001, 105) reflecting genuine scientific uncertainty, modeling preferences, and the desire to produce reasonably realistic-looking simulations of past and present climate. None of these models has emerged as clearly superior for purposes of investigating future climate change.

If no model stands out from among the others as a more promising resource for predicting future climate, how are we to proceed? It would not be very sensible to pick a

[81] I rely on the reader's intuitive understanding of "approximate truth" here.

model randomly and base our conclusions and actions on the results given by that model alone, since it might turn out that one or more of the other models will (unbeknownst to us) give more accurate predictions of future climate. We will have riskily "put all of our eggs in one basket." Instead, scientists are pursuing a "multi-model ensemble" approach, a variant type of Monte Carlo method. The multi-model ensemble approach assumes that members of a set—or "ensemble"—of complex models count as approximately equally plausible representations of the climate system.[82] Put slightly differently, although the climate models differ from one another in various respects, each model is assumed to be a reasonable balance of approximations, given present uncertainties. The entire ensemble of models is then used in investigating future climate change. For a given greenhouse gas forcing scenario, each model will be run individually to generate a projection of its own, but the product of the investigation will be the entire collection of projections. In this way, we can investigate the implications of our uncertainty in representing the climate system. Insofar as the ensemble of models spans that uncertainty, it will be reflected (in the loose sense) in range of projections produced.

To illustrate: suppose that we identify a "most likely" greenhouse gas emissions scenario and then use our ensemble of climate models to make projections of global mean annual temperature for the year 2050 under that scenario. It might happen that the members of our ensemble produce 2050 temperatures that seem to vary almost randomly over a wide range of values—our ensemble indicates that the temperature in 2050 might be somewhat cooler or somewhat warmer or just about anything in between. In this case, we learn that our uncertainty in representing the climate system translates into substantial uncertainty with respect to the result of interest. We must conclude (without any further information) that our present understanding of the climate system does not allow us to say with confidence what the temperature will be like in 2050. On the other hand, it might happen instead that nearly all of the projections of 2050 temperatures cluster rather tightly around one particular value. For example, perhaps nearly all of the models agree that there will be moderate warming by 2050. In this case, our uncertainty in representing the climate system turns out not to matter much; despite the differences in their

[82] There are several variations on the multi-model approach; I describe just one of them here.

assumptions, the models basically tell a univocal story about what will happen in the future.

From the point of view of decision-making and planning for the future, we may prefer the latter situation, in which all of the model-derived evidence points to the same conclusion. However, we must proceed with caution. The fact that the models substantially agree in their temperature projections is no guarantee that the agreed upon projection is an accurate one (even if the emissions scenario is a realistic one). It is possible that the models in the ensemble all systematically underestimate or overestimate future temperature. For present-day climate models, the possibility of this kind of systematic error may not be as unlikely as one might think, because the models have not been developed independently of one another; many of today's models are descendents of a small number of climate models constructed early on (see Edwards 2000) and so are likely to have assumptions (and even computer code) in common. Although the models do differ from one another in important respects, their output may exhibit some systematic errors as a result of what they have in common. In fact, recent intercomparison projects have documented some typical systematic errors found in simulations of past and present climate (see Lambert and Boer 2001). It would seem an important next step for scientists to investigate to what extent these known systematic errors (which may be only a subset of all of the systematic errors in the simulations) are likely to impact projections of regional and global climate changes.

Despite the fact that we must be careful when interpreting the results produced by multi-model ensembles, when it comes to addressing the global warming issue, the ensemble approach seems clearly better than the two most obvious alternatives, i.e. selecting a single model randomly for use in planning for the future and/or making no attempt to use climate models to address questions about future climate until a single "best" model can be identified. An ensemble of models incorporating different parameterizations of key climate processes is currently being used to investigate the range of possible future climatic changes under a variety of possible greenhouse gas scenarios and is considered the best way to proceed in the near future in investigating such changes (see IPCC 2001, 511). Instead of pretending that uncertainties do not exist, the ensemble approach acknowledges them and seeks to determine their implications. In the context of the ensemble approach, the fact that the models are logically incompatible need not be problematic, since we do not conclude that each model result is the one "true" outcome but rather that a set of results indicates a range of plausible outcomes. Drawing such a conclusion requires only that each model (and associated initial and boundary conditions) be considered a plausible representation of the climate system. For the reasons outlined in the preceding sections, that is the way that many scientists currently view a large number of complex climate models.

6.5 Pragmatic integrative pluralism

We now have seen both why a plurality of logically incompatible climate models persists and how these incompatible models nevertheless are being used together as complementary resources for investigating future climate. How ought we to characterize this climate model pluralism?

In the philosophical literature, two primary forms of pluralism have been identified—competitive and compatible pluralism (see e.g. Mitchell 2002). Although it is not always emphasized, these forms of pluralism utlimately are about ontology, i.e. they are concerned with accounts of what the world (or some part of it) is like. I will expand their labels here to "ontic competitive pluralism" and "ontic compatible pluralism." In the context of scientific modeling, ontic *competitive* pluralism exists when two models make conflicting claims about the part of the world that they are intended to model. In other words, as representations of the same target, the models are mutually exclusive. For example, we might have one model of the solar system according to which the planetary orbits all lie in the same plane and another model according to which not all orbits lie in the same plane. Typically, when we have two representations that incorporate or entail conflicting claims about the world (and each representation is a candidate for belief or acceptance), we view the representations as competitors—it does not make sense to accept both of them as true of the world, so they compete for our belief/acceptance. By contrast, ontic *compatible* pluralism exists when we have two or more representations that can be true of the world at the same time. These representations do not incorporate conflicting claims about what the world is like. For example, we might have one model of radiation transfer in the atmosphere and another model of plant respiration, both of which could be used at the same time in constructing a larger representation of the climate system. These can be viewed as compatible sub-models of a larger, more comprehensive model (see Bailer-Jones 2000 for another example and discussion). Alternatively, we might have one model that describes only the aggregate features of some system (e.g. mean global temperature and precipitation) and a second model that describes the system in greater detail (e.g. temperature and precipitation on a fine spatial grid), but if the models closely agree in the nature of their assumptions about the system and in their predictions of the aggregate features, then we may consider the situation to be one of ontic compatible pluralism. We need not believe that either of two ontically compatible models is actually true of the world, but it is at least a logical possibility that they are both true of it.

These forms of pluralism can only get us so far as we try to make sense of the situation in climate modeling. The situation does seem to be one of ontic competitive pluralism, since the climate models make mutually conflicting claims about what the climate system is like. Ideally, scientists would like to choose from among the complex climate models that which does actually incorporate the most realistic assumptions about the physical processes that will shape future climatic conditions (whatever those processes are). As we saw above, for a variety of reasons, scientists simply have been unable to identify such a model. The interesting feature of the climate modeling situation, however, is that scientists are not focusing their efforts on paring down the collection of complex models that they now have. In fact, as we saw in the last section, they are actually using the models *together* to investigate future climate. Are the models somehow compatible after all?

Sandra Mitchell's recent work on "integrative" pluralism in biology (see Mitchell 2002) may be of some help to us in making sense of the situation in climate modeling. Her analysis seems relevant because it is concerned with situations in which apparent competitors end up being used together—or integrated. More specifically, she shows how several idealized causal models that seem to provide competing explanations of some type of phenomenon (e.g. division of labor in social insects) in fact often turn out to be compatible when it comes to explaining a particular, concrete instance of that type of phenomenon (e.g. division of labor in leaf-cutting ants). Mitchell argues that the idealized causal models are not actually in competition, because each applies in a different idealized (non-actual) situation (2002, 64). Competition *can* occur among explanations of a particular concrete phenomenon, but each of those competing explanations typically will invoke several contributing causes. In other words, a plurality of idealized causal models will often be integrated in explaining an actual, complex biological phenomenon. There can be pluralism at the level of theoretical modeling, even though there will be only one "true" integrated explanation of any particular, concrete phenomenon (Mitchell 2002, 67). Mitchell's integrative pluralism thus seems to be a particular kind of ontic compatible pluralism—different possible accounts are brought together in producing a single actual account. We can call this "ontic integrative pluralism."

The pluralism in climate modeling is also integrative, but in a different way. Results obtained from several incompatible climate models are used together—or integrated—not in order to produce one "true" description of the climate system but rather as a way of taking into account our uncertainty in representing that system. For purposes of investigating the implications of our uncertainty, logically incompatible (but individually plausible) climate models are complementary resources. In other words, from the point of view of methodology—or, more generally, at the level of practice—the models are compatible. Thus, we can characterize the situation in climate modeling as one of "pragmatic" integrative pluralism. Does pragmatic integrative pluralism require that we view the models involved as purely instrumental tools? I would argue that it does not. When it comes to projecting future climate, it is precisely because we have some faith in climate models' plausibility as representations that we use their results together as we do. As suggested above, this is an arena in which it is considered important to get the right results for the right reasons. Even if complex climate models are not thought to be "perfect" or "true" descriptions of the climate system, scientists by and large do believe that they capture many of the most important processes that shape climate. The pragmatic integration of their results does not render them purely instrumental tools.

In the end, we can conclude that two different types of pluralism coexist in the case of climate modeling: an ontic competitive pluralism and a pragmatic integrative pluralism.

6.6 Conclusions

It is not that climate scientists unknowingly accept and employ logically incompatible models; rather, they are simply unable to select from among these models those that are most promising for purposes of investigating future climate. The persistence of a collection of incompatible complex climate models is a consequence of both scientific uncertainty concerning how to best represent the climate system and the difficulty involved in attempting to evaluate the relative merits of complex models. Given this situation, climate scientists are attempting to move forward with the investigation of future climate by adopting a multi-model ensemble approach. In this approach, incompatible complex models are assumed to be individually plausible representations of the climate system and are used together in order to produce a range of projections that reflects our uncertainty concerning how to best represent the climate system. The situation in climate modeling is one of both ontic competitive pluralism and pragmatic integrative pluralism.

7
Concluding Remarks

When we look carefully at the practice of computer simulation modeling, especially in the fields of meteorology and climate science, we see a side of science that is messy, complicated and often very confusing. Computer simulation modeling is rarely the straightforward application of theory that it is sometimes naively assumed to be. Model development instead involves the skilled synthesis of a variety of ingredients, including background theory, empirical data, educated guesses, and ad hoc assumptions. A great deal of "tinkering" and engineering skills are often required in order to bring together these ingredients in a stable and useful way. And then, when it comes time to evaluate computer simulation models as knowledge-producing resources, it is not easy to know how to handle their motley construction. We may see that a model makes accurate predictions, while at the same time we know that it incorporates some very questionable or even flatly false assumptions. What goes on in computer simulation modeling just does not lend itself to simple and straightforward analysis. Yet it is precisely because things are so messy and confusing in this domain that philosophical expertise is sorely needed. In closing, I would like to review the progress that was made in this project in promoting a better understanding of the practice of computer simulation modeling.

(1) The results of computer simulation experiments do not deserve the blanket skepticism that is often directed toward them. We saw that the epistemology of computer simulation experimentation has much more in common with that of material experimentation than is typically recognized. Both types of experimentation involve questions of internal and external validity, and analogous strategies are sometimes used in attempting to argue for such validity. Appreciating such similarities helps to combat the view that computer simulation experimentation is somehow a radically inferior methodology for investigating the empirical world. This is not to say that these two types of experimentation are exactly the same; it is clear that they are not. But in the end the trustworthiness of the results of both computer simulation experiments and material

experiments must be evaluated on a case by case basis by looking at the details of the experiment. We should not impugn the results of computer simulation experiments *just because* they are the results of such experiments.

(2) Simple computer simulation models that are known to be seriously deficient in a variety of ways can nevertheless play an important role in promoting scientific understanding. In the study of weather and climate, practitioners have gone a step further and claimed that simple models *and not complex models* contribute substantially to our understanding, but further analysis reveals this claim to be untenable. Instead, both simple and complex models can promote scientific understanding, although the extent to which each does so depends upon the view of understanding that is adopted. We saw that at least two notions of understanding have been invoked in the study of weather and climate, including causal and deductive varieties. This richness of modeling (and explanatory) practice is not always fully acknowledged, and we fail to do justice to it if we attempt to partition weather and climate models into complex predictive models and simple explanatory models. A more accurate depiction of the situation recognizes that one model can serve a variety of functions, often both predictive and explanatory.

(3) There are circumstances in which logically incompatible computer simulation models can be legitimately and fruitfully used together as complementary epistemic resources. In the study of climate change, scientists face several (currently insurmountable) difficulties in attempting to make judgments of the relative quality of logically incompatible climate models. As a result, they are unable to pick from among the available complex models those that are most promising for investigating future climate. Given this situation, climate scientists are attempting to move forward by adopting a multi-model ensemble approach in which logically incompatible climate models are used together in order to produce a range of future climate projections that reflects current scientific uncertainty. This use of climate models prompts us to distinguish ontic and pragmatic varieties of pluralism: the situation is one of ontic competitive pluralism and pragmatic integrative pluralism. The acceptance of this use of incompatible climate models is underwritten by the fact that models have a dual status as both representations and tools. Recognition of this dual status of models is likely a key to understanding other apparently puzzling aspects of modeling practice as well.

There is still plenty of important work remaining to be done. The practice of computer simulation modeling continues to generate many other interesting questions and problems, and it is time for philosophers to devote some energy to addressing them. No doubt the work will be challenging, but the potential payoff is substantial. One benefit of such work would be a more complete understanding of contemporary scientific practice. This would be an important contribution to philosophy of science in its own right. But the benefits could reach beyond the bounds of science and philosophy. As policymakers struggle to determine how to address the threat of global warming, it is essential that they have some understanding of the computer simulation modeling that is at the center of climate change research, and philosophers could play an important role in shaping that understanding. There is a real opportunity for philosophers to make a difference here.

APPENDIX

Appendix

The NCAR Community Climate System Model

The following provides some additional detail about the NCAR Community Climate System Model 2.0 (CCSM 2.0).

The atmospheric general circulation model (GCM) is the NCAR Community Atmosphere Model (CAM). This is the latest in a series of atmospheric GCMs that have been developed at NCAR over the past 20 years. CAM has a horizontal resolution on the order of 300 kilometers, and it has 26 vertical layers of various thicknesses. CAM is a primitive equation dynamical model, i.e. it relies on fundamental physical equations in calculating the large-scale dynamical evolution of the atmosphere.[83] CAM also includes representations of the transfer of radiation within the atmosphere (whether cloudy or clear), the formation of clouds and precipitation, and the complicated evolution of the momentum, heat and moisture fields near the Earth's surface (in the boundary layer), where turbulence plays an important role. Each of these aspects of CAM is complicated. To begin to get a sense for this, consider that when it comes to clouds CAM includes parameterizations of: the occurrence of different types of cloud, the horizontal extent of the coverage of these cloud types, the vertical distribution of cloud condensate, the sizes of cloud droplets, the radiative properties of the clouds (shortwave and longwave), and many other factors (see e.g. Kiehl et al 1998, 1132-1133). As indicated above, the atmosphere also makes various exchanges of heat, momentum and moisture with the land, ocean and ice surfaces. These fluxes are handled in the flux coupler; for each grid cell (volume of atmosphere) in contact with the Earth's surface, CAM exchanges data with the flux coupler. Some examples of the state variables that CAM sends to the flux

[83] These fundamental equations typically include: horizontal equations of motion, the hydrostatic law (which defines pressure change with height), a mass continuity equation, the first law of thermodynamics, the equation of state for air, and a balance equation for water vapor (see Peixoto and Oort 1992, 454-455).

coupler include: east-west wind velocity, north-south wind velocity, temperature, pressure, and specific humidity (Kauffman and Large 2002). CAM also sends flux variables, including downward longwave radiation, four different categories of downward shortwave radiation, two categories of liquid precipitation, and two categories of frozen precipitation (Kauffman and Large 2002). CAM receives data from the flux coupler, including (but not limited to) information about the reflectivity of the land and ocean surfaces, surface temperature, snow height, and latent and sensible heat fluxes (Kauffman and Large 2002).

The ocean model is the Los Alamos National Laboratory (LANL) Parallel Ocean Program (POP). Like CAM, this model is a primitive equation model: it uses fundamental physical equations to calculate the large-scale dynamical evolution of the ocean.[84] The horizontal resolution of POP is about 100 kilometers in both the east-west direction and north-south directions. Since its spatial resolution differs from that of the atmosphere, the flux coupler must interpolate and average data received from the two models for purposes of flux calculation and exchange. POP sends to the flux coupler information about salinity, velocity, surface temperature, and other quantities (Kauffman and Large 2002). POP receives from the coupler information about sea level pressure, surface stresses, fluxes of shortwave and longwave radiation, latent and sensible heat fluxes, salt flux, precipitation (rain and snow), evaporation, and other quantities (Kauffman and Large 2002).

The land surface model is the NCAR Community Land Model (CLM). Its spatial resolution (grid cell size) matches that of the overlying atmospheric model (CAM). The land surface model specifies different surface types for each grid cell in contact with the overlying atmosphere. These surface types (e.g. wetlands, lake, desert) are held constant through the model run, but their properties are used along with information from the atmospheric model to determine fluxes of energy, momentum, moisture and carbon dioxide to the atmosphere. In computing these fluxes, the CLM accounts for differences among vegetation and soil types. The CLM sends to the flux coupler information about surface temperature, reflectivity for different wavelengths, snow depth, fluxes of latent

[84] These equations are similar in type to those for the atmosphere, although the water vapor balance equation is replaced by a balance equation for salinity (see Peixoto and Oort 1992, 454).

and sensible heat, evaporation, and other quantities (Kauffman and Large 2002). The CLM receives from the coupler information about atmospheric wind velocities, atmospheric pressure, atmospheric temperature, atmospheric humidity, four types of precipitation, and shortwave and longwave radiation fluxes (Kauffman and Large 2002).

The CCSM2.0 also includes a sea ice model, the Community Sea Ice Model, Version 4 (CSIM4). This model has the same horizontal resolution (and uses the same grid) as the ocean model (POP). CSIM4 represents sea ice in terms of sea ice area, sea ice volume, sea ice internal energy, snow volume, surface temperature of snow/ice, sea ice velocity, and stress components (Briegleb et al. 2002, 6). These state variables evolve during the simulation (in part based on information received about the changing oceanic and atmospheric conditions), which in turn results in changes in the fluxes of momentum, heat, and moisture to the ocean and atmosphere. CSIM4 sends to the flux coupler information about ice area, surface temperature, reflectivity, latent and sensible heat fluxes, evaporated water, atmosphere-ice and ocean-ice stresses, heat flux to the ocean, fresh water flux to the ocean, and other quantities (Briegleb et al. 2002, 10). CSIM4 receives from the flux coupler information about atmospheric wind speed, atmospheric temperature, atmospheric humidity, downward radiation, freshwater fluxes due to precipitation, sea surface temperature, sea surface salinity, ocean current velocity, and other quantities (Briegleb et al. 2002, 9).

The flux coupler obviously plays a central role in the CCSM2.0 climate simulation. In addition to simply passing information from one component model to another, the flux coupler actually performs some calculations of fluxes and flux-related quantities. It calculates some fluxes of momentum, heat, moisture, and radiation at the atmosphere/underlying-surface interface. It also calculates an albedo (reflectivity for shortwave radiation) for the ocean surface and the net absorbed solar radiation for each of the land, ocean and ice surfaces (Kauffman and Large 2002).

BIBLIOGRAPHY

Bibliography

Andrews, David G. (2000), *An Introduction to Atmospheric Physics*. Cambridge: Cambridge University Press.

Anthes, Richard A. (1990), "Advances in the Understanding and Prediction of Cyclone Development with Limited-Area Fine-Mesh Models", in Chester Newton. and Eero O. Holopainen (eds.), *Extratropical Cyclones: The Erik Palmen Memorial Volume*. Boston: American Meteorological Society, 221-253.

Bailer-Jones, Daniela (2000), "Modelling Extended Extragalactic Radio Sources", *Studies in the History and Philosophy of Modern Physics* 31: 49-74.

--- (2003), "When Scientific Models Represent", *International Studies in the Philosophy of Science* 17: 59-74.

Batterman, Robert (2000), "A 'Modern' (=Victorian?) Attitude toward Scientific Understanding", *Monist* 83(2): 228-257.

--- (2001), *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction and Emergence*. Oxford: Oxford University Press.

Boville, Byron, and Peter Gent (1998), "The NCAR Climate System Model, Version One", *Journal of Climate* 11: 1115-1130.

Briegleb, Bruce P., et al. (2002), "Description of the Community Climate System Model Version 2 Sea Ice Model", May 2002, <http://www.cgd.ucar.edu/csm/models/ccsm2.0/csim/SciGuide/ice_scidoc.pdf > (September 24, 2002).

Brovkin, Viktor, et al. (1998), "On the stability of the atmosphere-vegetation system in the Sahara/Sahel region", *Journal of Geophysical Research* 103: 31,613-31,624.

Bunge, Mario (1969), "Analogy, Simulation, Representation", *Revue Internationale de Philosophie* 87: 16-33.

Campbell, Donald T. (1957), "Factors Relevant to the Validity of Experiments in Social Settings", *Psychological Bulletin* 54: 297-312.

Cartwright, Nancy (1983), *How the Laws of Physics Lie*. Oxford: Oxford University Press.

Charney, Jule G. (1949), "On a Physical Basis for Numerical Prediction of Large-Scale Motions in the Atmosphere", *Journal of Meteorology* 6, 371-385.

---, (1951), "Dynamical forecasting by numerical process" in *Compendium of Meteorology*. Boston: American Meteorological Society, 470–82.

---, (1955), "Numerical Methods in Dynamical Meteorology", *Proceedings of the National Academy of Sciences* 41, 798-802.

---, (1956), "Some Basic Problems in Dynamic Meteorology", *Journal of Geophysical Research* 61, 323-328.

---, (1959), "On the Theory of the General Circulation of the Atmosphere", in Bert Bolin (ed.), *Rossby Memorial Volume*. New York: Rockefeller Institute Press, 12-17.

---, (1972), "Impact of Computers on Meteorology", *Computer Physics Communications* 3(Supp.): 117-126.

Christie, Maureen (2001), *The Ozone Layer: A Philosophy of Science Perspective*. New York: Cambridge University Press.

Dai, Aiguo, et al. (2001), "Climates of the Twentieth and Twenty-First Centuries Simulated by the NCAR Climate System Model", *Journal of Climate* 14: 485-519.

De Mestre, Neville (1990), *The Mathematics of Projectiles in Sport*. New York: Cambridge University Press.

Dowling, Deborah (1999), "Experimenting on Theories", *Science in Context* 12(2): 261-273.

Edwards, Paul N. (2000), "A Brief History of Atmospheric General Circulation Modeling", in David A. Randall (ed.), *General Circulation Development, Past Present and Future: The Proceedings of a Symposium in Honor of Akio Arakawa*. New York: Academic Press, 67-90.

Edwards, Paul N. (2001), "Representing the Global Atmosphere: Computer Models, Data, and Knowledge about Climate Change", in Clark Miller and Paul N. Edwards (eds.), *Changing the Atmosphere: Expert Knowledge and Environmental Governance*. Cambridge: MIT Press, 31-65.

Franklin, Allan (1989), "The Epistemology of Experiment", in David Gooding, Trevor Pinch, and Simon Schaffer (eds.), *The Uses of Experiment*. Cambridge: Cambridge University Press, 437-460.

---, and Colin Howson (1988), "It Probably is a Valid Experiment: A Bayesian Approach to the Epistemology of Experiment", *Studies in the History and Philosophy of Science* 19: 419-427.

French, Steven (2002), "A Model-Theoretic Account of Representation", 23 March 2003, <http://philsci-archive.pitt.edu/archive/00001068/> (July 3, 2003).

Galison, Peter (1997), *Image and Logic*. Chicago: University of Chicago Press.

Giere, Ronald (1988), *Explaining Science: A Cognitive Approach*. Chicago: University of Chicago Press.

---, (1999a), "Using Models to Represent Reality", in Lorenzo Magnani, Nancy Nersessian, and Paul Thagard (eds.), *Model-Based Reasoning in Scientific Discovery*. New York: Kluwer, 41-57.

---, (1999b), *Science Without Laws*. Chicago: University of Chicago Press.

Gilbert, Nigel, and Klaus Troitzsch (1999), *Simulation for the Social Scientist*. Philadelphia: Open University Press.

Goodman, Nelson (1968), *Language of Arts*. New York: Bobbs-Merrill.

Guala, Francesco (2002), "Models, Simulations, and Experiments", in Lorenzo Magnani and Nancy Nersessian (eds.), *Model-Based Reasoning: Science, Technology, Values*. New York: Kluwer, 59-74.

Gould, Harvey, and Jan Tobochnik (1987), *An Introduction to Computer Simulation Methods: Applications to Physical Systems*. Reading, MA: Addison-Wesley Publishing.

Guetzkow, Harold, Philip Kotler and Randall Schultz (eds.) (1972), *Simulations in Social and Administrative Science: Overview and Case-Examples*. Englewood Cliffs: Prentice Hall.

Hartmann, Stefan (1996), "The World as a Process", in Rainer Hegselmann, Klaus Troitzsch, and Ulrich Mueller (eds.) *Modelling and Simulation in the Social Sciences for the Philosophy of Science Point of View*. Dordrecht: Kluwer, 77-100.

Hempel, Carl (1965), *Aspects of Scientific Explanation*. New York: Free Press.

Henderson-Sellers, Ann, and Kendall McGuffie (1999), "Concepts of Good Science in Climate Change Modelling", *Climatic Change* 42: 597-610.

Hesse, Mary (1966), *Models and Analogies in Science*. Notre Dame: University of Notre Dame Press.

Holton, James R. (1992), *Introduction to Dynamic Meteorology*. San Diego: Academic Press.

Hughes, R.I.G. (1997), "Models and Representation", *Philosophy of Science* 64(4): S325-S336.

---, (1999), "The Ising Model, Computer Simulation, and Universal Physics", in Mary Morgan and Margaret Morrison (eds.), *Models as Mediators*. Cambridge: Cambridge University Press, 97-145.

Humphreys, Paul (1991), "Computer Simulations", in Arthur Fine, Mickey Forbes, and Linda Wessels (eds.), *PSA 1990, Vol. 2*. East Lansing: Philosophy of Science Association, 497-506.

---, (1994), "Numerical Experimentation", in Paul Humphreys (ed.), *Patrick Suppes: Scientific Philosopher, Vol.2*. Boston: Kluwer, 103-121.

Huschke, Ralph E. (ed.) (1959), *Glossary of Meteorology*. Boston: American Meteorological Society Press.

Intergovernmental Panel on Climate Change (IPCC) (1996), *Climate Change 1995: The Scientific Basis*. Cambridge: Cambridge University Press.

Intergovernmental Panel on Climate Change (IPCC) (1997), *An introduction to simple climate models used in the IPCC Second Assessment Report - IPCC Technical Paper II*. Geneva: IPCC.

Intergovernmental Panel on Climate Change (IPCC) (2001), *Climate Change 2001: The Scientific Basis*. Cambridge: Cambridge University Press.

Kalnay, Eugenia et al. (1996), "The NCEP/NCAR 40-Year Reanalysis Project", *Bulletin of the Amercian Meteorological Society* 77: 437-471.

Kauffman, Brian, and William Large (2002), "The CCSM Coupler Version 5.0", May 2002, <http://www.cgd.ucar.edu/csm/models/ccsm2.0/cpl5/users_guide/users_guide.html> (September 24, 2002).

Kiehl, Jeffrey, et al. (1998), "The National Center for Atmospheric Research Community Climate Model: CCM3", *Journal of Climate* 11: 1131-1149.

Lafollette, Hugh, and Niall Shanks (1995), "Two Models of Models in Biomedical Research", *Philosophical Quarterly* 45: 141-160.

Lambert, Steve, and George Boer (2001), "CMIP1 evaluation and intercomparison of coupled climate models", *Climate Dynamics* 17: 83-106.

Lorenz, Edward N. (1960), “Maximum simplification of the dynamic equations”, *Tellus* 12: 243-254.

---, (1966), “The Circulation of the Atmosphere”, *American Scientist* 54: 402-420.

---, (1967), *The Nature and Theory of the General Circulation of the Atmosphere.* Geneva: World Meteorological Organization.

---, (1970), “Climatic Change—A Mathematical Problem”, *Journal of Applied Meteorology* 9: 325-329.

Machamer, Peter, Lindley Darden, and Carl Craver (2000), “Thinking About Mechanisms”, *Philosophy of Science* 67: 1-25.

Mitchell, Sandra D. (2002), “Integrative Pluralism”, *Biology and Philosophy* 17:55-70.

Morgan, Mary (2002), “Model Experiments and Models in Experiments”, in Lorenzo Magnani and Nancy Nersessian (eds.), *Model-Based Reasoning: Science, Technology, Values*. New York: Kluwer, 41-58.

---, (2003), “Experiments Without Material Intervention: Model Experiments, Virtual Experiments and Virtually Experiments”, in Hans Radder (ed.), *The Philosophy of Scientific Experimentation.* Pittsburgh: University of Pittsburgh Press.

Morrison, Margaret (1998), “Modelling Nature: Between Physics and the Physical World”, *Philosophia Naturalis* 35(1), 65-85.

---, (1999), “Models as Autonomous Agents”, in Mary Morgan and Margaret Morrison (eds.), *Models as Mediators*. Cambridge: Cambridge University Press, 38-65.

---, and Mary Morgan (1999), “Models as Mediating Instruments”, in Mary Morgan and Margaret Morrison (eds.), *Models as Mediators*. Cambridge: Cambridge University Press, 10-37.

Morton, Adam (1993), "Mathematical Models: Questions of Trustworthiness", *British Jounral for the Philosophy of Science*, 44(4), 659-674.

National Toxicology Program (NTP), “9th Report on Carcinogens”, January 2001, <http://157.98.13.104/roc/toc9.html> (22 August 22, 2002).

Newton, Chester, and Eero O. Holopainen (eds.) (1990), *Extratropical Cyclones: The Erik Palmen Memorial Volume*. Boston: American Meteorological Society Press.

North, Gerald R., Robert F. Calahan, and James A. Coakley (1981), “Energy Balance Climate Models”, *Reviews of Geophysics and Space Physics* 19(1): 91-121.

Norton, Stephen, and Frederick Suppe (2001), "Why Atmospheric Modeling is Good Science", in Clark Miller and Paul N. Edwards (eds.), *Changing the Atmosphere: Expert Knowledge and Environmental Governance*. Cambridge: MIT Press, 67-105.

Oreskes, Naomi (2000), "Why believe a computer? Models, Measures, and Meaning in the Natural World", in J. Schneiderman (ed.), *The Living Earth*. San Francisco: W.H. Freeman and Co., 70-82.

---, and Kenneth Belitz (2001), "Philosophical Issues in Model Assessment", in Malcolm G. Anderson and Paul D. Bates (eds.), *Model Validation: Perspectives in Hydrological Science*. Chichester: Wiley & Sons.

Peixoto, Jose P., and Abraham H. Oort (1992), *Physics of Climate*. New York: American Institute of Physics Press.

Phillips, Norman A. (1951), "A simple three-dimensional model for the study of large-scale extratropical flow patterns", *Journal of Meteorology* 8: 381-394.

---, (1956), "The general circulation of the atmosphere: a numerical experiment", *Quarterly Journal of the Royal Meteorological Society* 82: 123-164.

Railton, Peter (1981), "Probability, Explanation, and Information", *Synthese* 48: 233-256.

Raper, Sarah, Jonathan M. Gregory, and Tim Osborn (2001), "Use of an upwelling-diffusion energy balance climate model to simulate and diagnose AOGCM results", *Climate Dynamics* 17, 601-613.

Reed, Richard (1990), "Advances in Knowledge and Understanding of Extratropical Cyclones During the Past Quarter Century: An Overview", in Chester Newton. and Eero O. Holopainen (eds.), *Extratropical Cyclones: The Erik Palmen Memorial Volume*. Boston: American Meteorological Society Press, 27-45.

Ribe, Neil, and Friedrich Steinle (2002), "Exploratory Experimentation: Goethe, Land, and Color Theory", *Physics Today* 55(7): 43-49.

Rohrlich, Fritz (1991), "Computer Simulation in the Physical Sciences", in Arthur Fine, Mickey Forbes, and Linda Wessels (eds.), *PSA 1990, Vol. 2*. East Lansing: Philosophy of Science Association, 507-518.

Salmon, Wesley (1998), *Causality and Explanation*. Oxford: Oxford University Press.

Schaffner, Kenneth (2001), "Extrapolation from Animal Models", in Peter Machamer, Peter McLaughlin, and Rick Grush (eds.), *Theory and Method in the Neurosciences*. Pittsburgh: University of Pittsburgh Press, 200-229.

Schneider, Steven, and Robert E. Dickinson (1974), "Climate Modeling", *Reviews of Geophysics and Space Physics* 12: 447-493.

Shackley, Simon, et al. (1998), "Uncertainty, complexity and concepts of good science in climate change modelling: are GCMs the best tools?", *Climate Change* 38: 159-205.

---, et al. (1999), "Adjusting to Policy Expectations in Climate Change Modeling: An Interdisciplinary Study of the Use of Flux Adjustments in Coupled A/O GCMs", *Climatic Change* 43: 413-454.

Shannon, Robert E. (1975), *Systems Simulation: The Art and Science*. Englewood Cliffs: Prentice-Hall.

Sismondo, Sergio (1999), "Models, Simulations, and Their Objects", *Science in Context* 12(2): 247-260.

Smagorinsky, Joseph (1983), "The beginnings of numerical weather prediction and general circulation modeling: Early recollections", *Advances in Geophysics* 25, 3-37.

Stanley, Julian C. and Donald T. Campbell (1963), *Experimental and Quasi-experimental Designs for Research*. Chicago: Rand McNally.

Steinle, Friedrich (1997), "Entering New Fields: Exploratory Uses of Experimentation", *Philosophy of Science* 64: S65-S74.

Sterrett, Susan (2003), "Physical Models and Fundamental Laws: Using One Piece of the World to Tell about Another", *Mind and Society 2002-2.*

Suarez, Mauricio (1999), "Theories, Models, and Representations", in Lorenzo Magnani, Nancy Nersessian, and Paul Thagard (eds.), *Model-Based Reasoning in Scientific Discovery*. New York: Kluwer, 75-84.

Teller, Paul (2001), "Twilight of the Perfect Model Model", *Erkenntnis* 55: 393-415.

Thompson, Phillip D. (1983/1990), "A History of Numerical Weather Prediction in the United States", *Bulletin of the American Meteorological Society* 64, 755-769, reprinted with revisions in Richard S. Lindzen, Edward N. Lorenz, and George W. Platzman (eds.), *The Atmosphere—A Challenge: The Science of Jule Gregory Charney*. Boston: American Meteorological Society Press, 93-119.

von Storch, Hans, and Antonio Navarra (1999), *Analysis of Climate Variability*. Heidelberg: Springer-Verlag.

Weissart, Thomas (1997), *The Genesis of Simulation in Dynamics*. New York: Springer-Verlag.

Wilhelmson, Robert et al. (1990), "A Study of the Evolution of a Numerically Modeled Severe Storm", *The International Journal of Supercomputer Applications* 4/2, 20-36.

Wimsatt, William (1987), "False Models as means to Truer Theories", in Matthew Nitecki and Antoni Hoffman (eds.), *Neutral Models in Biology*. London: Oxford University Press, 23-55.

Winsberg, Eric (1999a), *Simulation and the Philosophy of Science: Computationally Intensive Studies of Complex Physical Systems,* PhD Dissertation. Bloomington: Indiana University.

---, (1999b), "The Hierarchy of Models in Simulation", in Lorenzo Magnani, Nancy Nersessian, and Paul Thagard (eds.), *Model-Based Reasoning in Scientific Discovery*. New York: Kluwer, 255-269.

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

---, (2001), "Simulations, Models, and Theories: Complex Physical Systems and Their Representations", *Philosophy of Science* 68(3), S442-S454.

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