METAPHYSICAL AND EPISTEMOLOGICAL ISSUES IN COMPLEX SYSTEMS

Robert C. Bishop

1 INTRODUCTION

There are multiple senses of complexity. For instance, something can be considered complex because it is complicated or intricate (think of an engine or a watch with many moving parts and interrelated systems). However, systems that have a complicated set of interacting parts may actually exhibit relatively simple behavior (this is the case for engines). In contrast, the notion of complexity of interest in "complexity studies" centers on systems whose behavior is nonlinear and typically exhibits self-organization and collective effects. Such behavior appears to be anything but simple (see sec. 2.5), yet it is the case that systems with relatively few interacting components can exhibit this kind of intricate behavior. For instance, phase changes in Ising models and other systems with large numbers of components are examples of systems exhibiting complex self-organizing behavior. But even three-particle billiard-ball-like systems can exhibit the requisite complex behavior. On the other hand, many n-body systems, where n is large, do not exhibit complexity (e.g., Brownian motion of molecules) because the interactions among the constituents are not of the right type. An important feature of the systems of interest in complexity studies is that the interactions of the system components be nonlinear. Still, characterizing the resulting behavior, and the complex systems which exhibit it, is one of the major challenges facing scientists studying complexity and calls for the development of new concepts and techniques (e.g., [Badii and Politi, 1997]).

A number of metaphysical and epistemological issues are raised by the investigation and behavior of complex systems. Before treating some of these issues, a characterization of nonlinear dynamics and complexity is given. Along with this background, some folklore about chaos and complexity will be discussed. Although some claim that chaos is ubiquitous and many take the signal feature of chaos to be exponential growth in uncertainty (parameterized by Lyapunov exponents, see sec. 2.4), these examples of folklore turn out to be misleading. They give rise to rather surprising further folklore that chaos and complexity spell the end of predictability and determinism. But when we see that Lyapunov exponents, at least in their global form, and measures for exponential divergence of trajectories only apply to infinitesimal quantities in the infinite time limit, this further folklore also

Handbook of the Philosophy of Science. Volume 10: Philosophy of Complex Systems.
Volume editor: Cliff Hooker. General editors: Dov M. Gabbay, Paul Thagard and John Woods.

turns out to be misleading. Instead, the loss of linear superposition in nonlinear systems is one of the crucial features of complex systems. This latter feature is related to the fact that complex behavior is not limited to large multi-component systems, but can arise in fairly simple systems as well.

The impact of nonlinearity on predictability and determinism will be discussed including, briefly, the potential impact of quantum mechanics. Some have argued that chaos and complexity lead to radical revisions in our conception of determinism, namely that determinism is a layered concept (e.g., [Kellert, 1993]), but such arguments turn on misunderstandings of determinism and predictability and their subtle relations in the context of nonlinear dynamics. When the previously mentioned folklore is cleared away, the relationship among determinism, predictability and nonlinearity can be seen more clearly, but still contains some subtle features. Moreover, the lack of linear superposition in complex systems also has implications for confirmation, causation, reduction and emergence, and natural laws in nonlinear dynamics all of which raise important questions for the application of complex nonlinear models to actual-world problems.

2 NONLINEAR DYNAMICS: FOLKLORE AND SUBTLETIES

I will begin with a distinction that is immediately relevant to physical descriptions of states and properties known as the ontic/epistemic distinction tracing back at least to Erhard Scheibe [1964] and subsequently elaborated by others [Primas, 1990; 1994; Atmanspacher, 1994; 2002; d'Espagnat, 1994; Bishop, 2002). Roughly, ontic states and properties are features of physical systems as they are "when nobody is looking," whereas epistemic states and properties refer to features of physical systems as accessed empirically. An important special case of ontic states and properties are those that are deterministic and describable in terms of points in an appropriate state space (see secs. 2.3 and 3.2 below); whereas an important special case of epistemic states and properties are those that are describable in terms of probability distributions (or density operators) on some appropriate state space. The ontic/epistemic distinction helps eliminate of confusions which arise in the discussions of nonlinear dynamics and complexity as we will see.

2.1 Dynamical systems

Complexity and chaos are primarily understood as mathematical behaviors of *dynamical systems*. Dynamical systems are deterministic mathematical models, where time can be either a continuous or a discrete variable (a simple example would be the equation describing a pendulum swinging in a grandfather clock). Such models may be studied as purely mathematical objects or may be used to describe a target system (some kind of physical, ecological or financial system, say). Both qualitative and quantitative properties of such models are of interest to scientists.

The equations of a dynamical system are often referred to as dynamical or evolution equations describing the change in time of variables taken to adequately describe the target system. A complete specification of the initial state of such equations is referred to as the *initial conditions* for the model, while a characterization of the boundaries for the model domain are known as the *boundary conditions*. A simple example of a dynamical system would be the equations modeling a particular chemical reaction, where a set of equations relates the temperature, pressure, amounts of the various compounds and their reaction rates. The boundary condition might be that the container walls are maintained at a fixed temperature. The initial conditions would be the starting concentrations of the chemical compounds. The dynamical system would then be taken to describe the behavior of the chemical mixture over time.

2.2 Nonlinear dynamics and linear superposition

The dynamical systems of interest in complexity studies are *nonlinear*. A dynamical system is characterized as linear or nonlinear depending on the nature of the dynamical equations describing the target system. Consider a differential equation system $d\boldsymbol{x}/dt = \boldsymbol{F}\boldsymbol{x}$, where the set of variables $\boldsymbol{x} = x_1, x_2, ..., x_n$ might represent positions, momenta, chemical concentration or other key features of the target system. Suppose that $\boldsymbol{x}_1(t)$ and $\boldsymbol{x}_2(t)$ are solutions of our equation system. If the system of equations is linear, it is easy to show that $\boldsymbol{x}_3(t) = a\boldsymbol{x}_1(t) + b\boldsymbol{x}_2(t)$ is also a solution, where a and b are constants. This is known as the *principle of linear superposition*.

If the principle of linear superposition holds, then, roughly, a system behaves such that any multiplicative change in a variable, by a factor α say, implies a multiplicative or proportional change of its output by α. For example, if you start with your television at low volume and turn the volume control up one unit, the volume increases one unit. If you now turn the control up two units, the volume increases two units. These are examples of linear responses. In a nonlinear system, linear superposition fails and a system *need not* change proportionally to the change in a variable. If you turn your volume control up two units and the volume increases tenfold, this would be an example of a nonlinear response.

2.3 State space and the faithful model assumption

Dynamical systems involve a *state space*, an abstract mathematical space of points where each point represents a possible state of the system. An instantaneous state is taken to be characterized by the instantaneous values of the variables considered crucial for a complete description of the state. When the state of the system is fully characterized by position and momentum variables (often symbolized as q and p, respectively), the resulting space is often called a *phase space*. A model can be studied in state space by following its trajectory, which is a history of the model's behavior in terms of its state transitions from the initial state to some

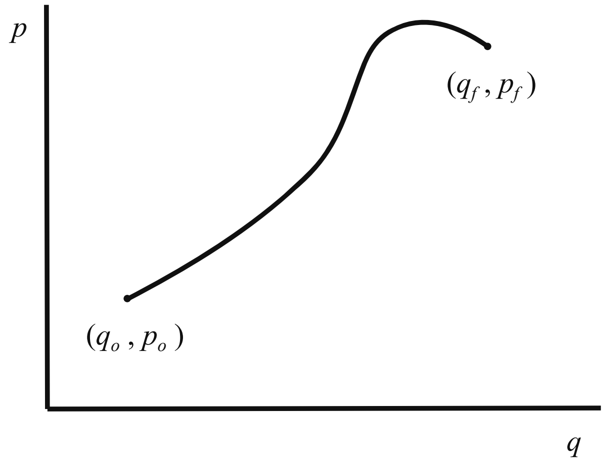

Figure 1.

chosen final state (Figure. 1). The evolution equations govern the path — the history of state transitions — of the system in state space.

There are some little noticed yet crucial assumptions being made in this account of dynamical systems and state spaces. Namely, that the actual state of a target system is accurately characterized by the values of the crucial state space variables and that a physical state corresponds via these values to a point in state space. These assumptions allow us to develop mathematical models for the evolution of these points in state space and to consider such models as representing the target systems of interest (perhaps through an isomorphism or some more complicated relation). In other words, we assume that our mathematical models are faithful representations of target systems and that the state spaces employed faithfully represent the space of actual possibilities of target systems. This package of assumptions is known as the *faithful model assumption* (e.g., [Bishop, 2005a; 2006]). In its idealized limit — *the perfect model scenario* [Judd and Smith, 2001]— it can license the (perhaps sloppy) slide between model talk and system talk (i.e., whatever is true of the model is also true of the target system and vice versa).

2.4 Sensitive dependence and Lyapunov exponents

One striking feature of chaos and complexity is their *sensitive dependence on initial conditions*: the property of a dynamical system to show possibly extremely different behavior with only the slightest of changes in initial conditions. A very popular measure of this sensitive dependence involves the explosive growth of the smallest uncertainties in the initial conditions of a nonlinear system. This explosive growth is often defined as an exponential parameterized by the largest *global*

Lyapunov exponent. These exponents arise naturally out of linear stability analysis of the trajectories of nonlinear evolution equations in a suitable state space. The infinite time limit plays an important role in this analysis, indicating that global Lyapunov exponents represent time-averaged quantities so that transient behavior has decayed. The existence of this limit is guaranteed by Oseledec's [1969] multiplicative ergodic theorem, which holds under mild conditions.

Imagine a small ball of points in state space around the initial conditions $\boldsymbol{x}(0)$. For any number $\delta > 0$ and for every slightly different initial condition $\boldsymbol{y}(0)$ in this small ball, exponential growth means the initial uncertainty, $|\boldsymbol{x}(0) - \boldsymbol{y}(0)| < \delta$, will evolve such that $|\boldsymbol{x}(t) - \boldsymbol{y}(t)| \approx |\boldsymbol{x}(0) - \boldsymbol{y}(0)|e^{\lambda t}$.[1] Here, λ is interpreted as the largest global Lyapunov exponent and is taken to represent the average rate of divergence of neighboring trajectories issuing forth from points very nearby $\boldsymbol{x}(0)$. If $\lambda > 0$, then the growth in uncertainty is exponential (if $\lambda < 0$, there is exponential convergence of neighboring trajectories). In general, such growth cannot go on forever. If the system is bounded in space and in momentum, there will be limits as to how far nearby trajectories can diverge from one another.

In most all philosophy literature and much physics literature, sensitive dependence as parameterized by global Lyapunov exponents is taken to be a distinguishing mark of chaotic dynamics. That is to say, exponential growth in the separation of neighboring trajectories characterized by λ is taken to be a property of a particular kind of dynamics that can only exist in nonlinear systems and models. However, there are problems with this folklore for defining sensitive dependence (and, hence, characterizing chaos and complexity using Lyapunov exponents).

One problem is that the definition of global Lyapunov exponents involves the infinite time limit. Strictly speaking, λ only characterizes growth in uncertainties as t increases without bounds, not for any finite time. At best, this would imply that sensitive dependence characterized by a global Lyapunov exponent can only hold for the large time limit. And this would further imply that chaotic phenomenon can only arise in this limit, contrary to what we take to be our best evidence. Furthermore, neither our models nor physical systems persist for infinite time, but an infinitely long time is required to verify the presumed exponential divergence of trajectories issuing from infinitesimally close points in state space.

The standard physicist's assumption that an infinite-time limit can be used to *effectively* represent some large but finite elapsed time will not do in the context of nonlinear dynamics either. When the finite-time Lyapunov exponents are calculated, they do not usually lead to on-average exponential growth as characterized by the global Lyapunov exponents (e.g., [Smith, *et al.*, 1999]). This is because the propagator — an operator evolving the uncertainty in some ball of points in state space forward in time — varies from point to point in state space for any finite times. The propagator is a function of the position $\boldsymbol{x}$ in state space and only approaches a constant in the infinite time limit. So local finite-time Lyapunov exponents vary from point to point in state space (whereas global Lyapunov ex-

[1] Technically, this kind of measure is taken to be valid in an appropriate state space for "almost all" points in the region around $\boldsymbol{x}(0)$ except a set of measure zero.

ponents do not). Therefore, trajectories diverge and converge from each other at various rates as they evolve in time, so that the uncertainty usually does not vary uniformly in the chaotic region of state space [Smith, *et al.*, 1999; Smith, 2000]. In contrast, global Lyapunov exponents are on-average global measures of uncertainty growth and imply that uncertainties grow uniformly (for $\lambda > 0$). Such uniform growth rarely occurs outside a few simple mathematical models (e.g., the baker's map).

For instance, the Lorenz, Moore-Spiegel, Rössler, Henon and Ikeda attractors all possess regions dominated by decreasing uncertainties in time, where uncertainties associated with different trajectories issuing forth from some small ball in state space shrink for the amount of time trajectories remain within such regions (e.g., [Smith, *et al.*, 1999, 2870-9; Ziehmann, *et al.*, 2000, 273-83]). What this means is that on-average exponential growth in uncertainties is not guaranteed for chaotic dynamics. Linear stability analysis indicates when nonlinearities can be expected to dominate the dynamics. And local finite-time Lyapunov exponents can indicate regions on an attractor where these nonlinearities will cause *all* uncertainties to decrease — cause trajectories to converge rather than diverge — so long as trajectories remain in those regions.

To summarize this first problem, the folklore that trajectories issuing forth from neighboring points in some ball in state space are guaranteed to diverge on-average exponentially in a chaotic region of state space is false in any sense other than for infinitesimal uncertainties in the infinite time limit.

The second problem with our folklore is that there simply is no implication that finite uncertainties will exhibit an on-average growth rate characterized by *any* Lyapunov exponents, local or global. As pointed out above, the linearized dynamics used to derive global Lyapunov exponents presupposes infinitesimal uncertainties. But when uncertainties are finite, linearized dynamics involving infinitesimals does not appropriately characterize the growth of finite uncertainties aside from telling us when nonlinearities should be expected to be important (this latter information is extremely useful however). Infinitesimal uncertainties can never become finite in finite time except through super-exponential growth. And even if infinitesimal uncertainties became finite after a finite time, that would presuppose the dynamics is unconfined; however, the interesting features of nonlinear dynamics usually take place in subregions of state space (e.g., on particular energy surfaces or in regions where attractors exist). Presupposing an unconfined dynamics, then, would be inconsistent with the features we are typically trying to capture.

One can ask whether the on-average exponential growth rate characterized by global Lyapunov exponents can ever be attributed legitimately to diverging trajectories if their separation is no longer infinitesimal. Examining simple models like the baker's map might seem to indicate yes. However, answering this question requires some care for more complicated systems like the Lorenz or Moore-Spiegel attractors. It can turn out to be the case that the rate of divergence in the finite separation between two nearby trajectories in a chaotic region changes character numerous times over the course of their winding around in state space, sometimes faster, sometimes slower than that calculated from global Lyapunov exponents, sometimes contracting, sometimes diverging [Smith, *et al.*, 1999; [Ziehmann, *et al.*, 2000]). In the long run, will some of these trajectories *effectively* diverge *as if* there was on-average exponential growth in uncertainties as characterized by global Lyapunov exponents? It is conjectured that the set of initial points in the state space exhibiting this behavior is a set of measure zero. This means that although there are an infinite number of points exhibiting this behavior, this set represents zero percent of the number of points composing the attractor. The details of the kinds of divergence (convergence) uncertainties undergo depend crucially on the detailed structure of the dynamics (i.e., it is determined point-by-point by local growth and convergence of finite uncertainties and not by any Lyapunov exponents).

In practice, however, all finite uncertainties saturate at the diameter of the attractor. The uncertainty reaches some maximum amount of spreading after a finite time and is not well quantified by measures derived from global Lyapunov exponents (e.g., [Lorenz, 1965]). So the folklore — that on-average exponential divergence of trajectories characterizes chaotic dynamics and complexity — is misleading for nonlinear systems. Therefore, drawing an inference from the presence of positive global Lyapunov exponents to the existence of on-average exponentially diverging trajectories for a dynamical system is shaky at best.

2.5 Complexity measures

Although several formal definitions of complexity have been proposed for characterizing random, chaotic and other forms of complex behavior, there is no consensus on which is the best definition, nor do these different definitions agree in picking out the same categories of behavior [Grassberger, 1989; Wackerbauer, *et al.*, 1994; Badii and Politi, 1997]. There is some evidence to suggest that different measures are useful for characterizing interesting behaviors of different systems for different purposes [Wackerbauer, *et al.*, 1994]. Perhaps this is not too surprising as it can be argued that complexity is just the kind of feature requiring a complex suite of tools and measures [Badii and Politi, 1997]. However, most of these complexity measures provide no intuitive access to the issues of emergence and causation at work in complex systems (some dynamical measures applicable in particular circumstances are exceptions). This is because most measures of complexity are formalized in terms of probabilities with no explicit reference to physical system variables (again, dynamical measures are an exception). Physical variables are implicitly involved in probabilistic measures because such variables are required to define the state space over which probability measures are defined.

Often it is more informative to characterize complex systems phenomenologically. Some of the most important features in these characterizations are:

- *Many-body systems.* Some systems exhibit complex behavior with as few as three constituents, while others require large numbers of constituents.

- *Broken symmetry.* Various kinds of symmetries, such as homogeneous arrangements in space, may exist before some parameter reaches a critical value, but not beyond.

- *Hierarchy.* There are levels or nested structures that may be distinguished, often requiring different descriptions at the different levels (e.g., large-scale motions in fluids vs. small-scale fluctuations).

- *Irreversibility.* Distinguishable hierarchies usually are indicators of or result from irreversible processes (e.g., diffusion, effusion).

- *Relations.* System constituents are coupled to each other via some kinds of relations, so are not mere aggregates like sand grain piles.

- *Situatedness.* The dynamics of the constituents usually depend upon the structures in which they are embedded as well as the environment and history of the system as a whole.

- *Integrity.* Systems display an organic unity of function which is absent if one of the constituents or internal structures is absent or if relations among the structures and constituents is broken.

- *Integration.* Various forms of structural/functional relations, such as feedback loops couple the components contributing crucially to maintaining system integrity.

- *Intricate behavior.* System behavior lies somewhere between simple order and total disorder such that it is difficult to describe and does not merely exhibit randomly produced structures.

- *Stability.* The organization and relational unity of the system is preserved under small perturbations and adaptive under moderate changes in its environment.

- *Observer relativity.* The complexity of systems depends on how we observe and describe them. Measures of and judgements about complexity are not independent of the observer and her choice of measurement apparatus [Grassberger, 1989; Crutchfield, 1994].

Such features of complex systems make the development of context-free measures of complexity unlikely (e.g., aside from observer relativity, the sense of order invoked in defining behavior as "intricate" depends on context). This can be illustrated by focusing on the nature of hierarchies in complex systems.

2.6 Hierarchies and sensitive dependence

The concept of *hierarchy* in the context of complex systems is of particular note. In some systems the hierarchy of physical forces and dynamical time scales (e.g., elementary particles, molecules, crystals) provide ontologically distinguishable levels of structure. In some cases the lower-level constituents may provide both necessary and sufficient conditions for the existence and behavior of the higher-level structures. In complex systems, however, levels of structure are often only epistemically distinguishable in terms of dynamical time scales. Furthermore, these levels are coupled to each other in such a way that at least some of the higher-level structures are not fully determined by, and even influence and constrain, the behavior of constituents in lower-level structures. That is, the lower-level constituents provide necessary but *not* sufficient conditions for the existence and behavior of some of the higher-level structures (cf. [Bishop, 2005b; 2008a; Bishop and Atmanspacher, 2006]). Moreover, the lower-level constituents may not even provide necessary and sufficient conditions for their own behavior if the higher-level structures and dynamics can constrain or otherwise influence the behavior of lower-level constituents (e.g., [Bishop, 2008a]). This latter kind of hierarchy is called a *control hierarchy* [Pattee, 1973, 75-9; Primas, 1983, 314-23]. Control hierarchies are distinguished from merely hierarchical structure like sand grain piles through the kinds of control they exert on lower-level structures and dynamics.

In complex systems, control hierarchies affect lower-level constituents primarily through constraints. The most important examples of constraints actively change the rate of reactions or other processes of constituents relative to the unconstrained situation (e.g., switches and catalysts). These constraints control lower-level constituents without removing all the latter's configurational degrees of freedom (in contrast to simple crystals, for instance). These top-down constraints may be external, due to the environment interacting with the system. Or such constraints may arise internally within the system due to the collective effects of its constituents or some other higher-level structural feature. Typically fundamental forces like gravity and electromagnetism are not explicitly identified with these latter internally generated constraints.

The notions of hierarchy and sensitive dependence allow us to formulate a more qualitative distinction between linear and nonlinear systems (though this characterization can also be made empirically precise — see [Busse, 1978], and [Cross and Hohenberg, 1993] for examples). Linear systems can be straightforwardly decomposed into and composed by subsystems (a consequence of the principle of linear superposition). For a concrete example of the principle of linear superposition, consider linear (harmonic) vibrations of a string which can be analyzed as a superposition of normal modes. These normal modes can be treated as uncoupled individual oscillators. The composition of the string's vibration out of these component vibrations is then analogous to aggregating these parts into a whole ("the whole is the sum of its parts"). The linear behavior of such systems in these cases is sometimes called *resultant* (in contrast with emergent).[2]

In nonlinear systems, by contrast, this straightforward idea of composition fails (a consequence of the failure of the principle of linear superposition). When the behaviors of the constituents of a system are highly coherent and correlated, the system cannot be treated even approximately as a collection of uncoupled individual parts ("the whole is *different* than the sum of its parts"). Rather, some particular global or nonlocal description[3] is required taking into account that individual constituents cannot be fully characterized without reference to larger-scale structures of the system. Rayleigh-Bénard convection, for instance, exhibits what is called *generalized rigidity* — the individual constituents are so correlated with all other constituents that no constituent of the system can be changed except by applying some change to the system as a whole. Such holistic behaviors are often referred to as *emergent* (in contrast with resultant).

The tight coupling between constituents in nonlinear systems is related to the nonseparability of the Hamiltonian. The latter is a function which corresponds to the total energy of the system and is related to the system's time evolution. Roughly, a Hamiltonian is separable if there exists a transformation carrying the Hamiltonian describing a system of N coupled constituents into N equations each describing the behavior of an individual system constituent. Otherwise, the Hamiltonian is nonseparable and the interactions within the system cannot be decomposed into interactions among only the individual components of the system.

In summary, linear systems can be decomposed into their constituent parts and the behavior of each component can be changed independently of the other components (which will then respond to the change introduced). Nonlinear systems often exhibit collective behavior where an individual system component cannot be isolated and its behavior changed independently of the rest of the system. Modifications of behaviors in a nonlinear system may have to take place at some higher hierarchical level or even at the level of the system as a whole.

2.7 Identity and individuation and a classical measurement problem

The interplay among hierarchical levels in nonlinear systems exhibiting complexity blur distinctions like part-whole, system-environment, constituent-level and so forth (e.g., cases where hierarchies are only distinguishable by differing time scales rather than by ontologically distinct features). The mathematical modeling of physical systems requires us to make distinctions between variables and parameters as well as between systems and their environments. However, when linear superposition is lost, systems can be exquisitely sensitive to the smallest of influences. A small change in the parameter of a model can result in significantly different behavior in its time evolution, making the difference between whether

[2]See [McLaughlin, 1982] for a discussion of the origin and history of the terms 'resultant' and 'emergent.'

[3]A nonlocal description in nonlinear dynamics denotes a description that necessarily must refer to wider system and environmental features in addition to local interactions of individual constituents with one another.

the system exhibits chaotic behavior or not, for instance. Parameters like the heat on a system's surface due to its environment may vary over time leading to wide variations in the time evolution of the system variables as well as temporal change in parameters. In such cases, the distinction between model variables and parameters tends to break down. Similarly, when a nonlinear system exhibits sensitive dependence, even the slightest change in the environment of a system can have a significant effect on the system's behavior. In such cases the distinction between system and environment breaks down. For instance, even the behavior of an electron at the 'edge' of the galaxy would affect a system of billiard balls undergoing continuous collisions [Crutchfield, 1994, p. 239], so the system/environment distinction becomes more a matter of pragmatically cutting up the 'system' and 'environment' in ways that are useful for our analysis.

All these subtleties raise questions about identity and individuation for complex systems. For instance, can a complex system somehow be identified as a distinct individual from its environment? Can various hierarchies of a complex system be individuated from each other? Asking these questions presupposes both that a distinct entity can be identified as well as individuated from other entities. Consider the so-called butterfly effect. Earth's weather is a complex system, but its potential sensitivity to the slightest changes of conditions leave its boundaries ill-defined if the flapping of a butterfly's wings in Argentina can cause a tornado in Texas three weeks later. Is the butterfly's flapping an internal or external source of wind and pressure disturbance? Turning towards space, is the magnetosphere surrounding the earth, which exhibits 'space weather' and shields the earth from lethal solar radiation, a distinct weather system or a qualitatively different extension of the Earth's weather system?

Traditional questions about identity and individuation revolve around numerical identity and the criteria for individuation and identity through time. It certainly seems plausible to consider butterflies, the Earth's weather and the Earth's magnetosphere (with its space weather) as numerically distinct systems (or as numerically distinct subsystems of a larger system). After all, according to Leibniz's principle of the identity of indiscernibles, these different "systems" do not share all their properties. On the other hand, systems are generally composed of subsystems that differ in properties, so given the lack of absolute boundaries between them, perhaps the best way to conceive of the butterfly-weather-magnetosphere system is as one very large complex system. As suggested by the phenomenological properties of complex systems (sec. 2.5), it is often the case that distinctions between parts and wholes, hierarchies and the like are pragmatic rather than absolute. There are further problems with identifying the boundary between the butterfly-weather-magnetosphere system and its solar system/galactic environment (e.g., electromagnetic fields and gravity extend over enormous distances in space).

Classical views of identity and individuation based on Leibniz's principle might be of some use in the pragmatic project of identifying complex systems and their components. However, these would only yield identification and individuation based on the kinds of questions scientific and other forms of inquiry raise and not a kind of objective ontology of distinct things (hence, many of our judgements about identity and individuation in nonlinear dynamics are epistemic rather than ontic). Whether the kinds of features described in secs. 2.5 and 2.6 imply there are no rigid designators, and hence complex systems represent a case of contingent identity and individuation [Kripke, 1980], is an open question.

Such features also raise questions about our epistemic access to complex systems. Obviously, some kind of cuts between observer and observed, and between system and environment have to be made. Along with this difficulty, there are clear epistemic difficulties confronting the measurement and study of complexity.

One epistemic difficulty is the mismatch between the accuracy or level of fine-grained access to the dynamics of a complex system and its underlying states and properties (i.e., the ontic/epistemic distinction). If a particular measurement apparatus only samples some even relatively fine-grained partition of the dynamical states of a complex system, the result will effectively be a mapping of (perhaps infinitely) many system states into a much smaller finite number of measurement apparatus states (e.g., [Crutchfield, 1994]). Such a mapping produces an apparent complexity — epistemic dynamical states in the measurement apparatus' projected space — that may not faithfully represent the complexity (or simplicity) of the system's actual dynamics — ontic states.

Another epistemic difficulty is that any measurement apparatus used to ascertain system states necessarily will introduce a small disturbance into complex systems that, in turn, will be amplified by sensitive dependence. No matter how fine-grained the measurement instrument, no matter how tiny the disturbance, this perturbation will produce an unpredictable influence on the future behavior of the system under study, resulting in limitations on our knowledge of a complex system's future. Along with the disturbance introduced to the complex system being measured, there is also a small uncertainty in the measurement apparatus itself. So the apparatus must also measure both itself and its disturbance perfectly for a full accounting of the exact state of the complex system being studied. This, in turn, leads to an infinite regress of measurements measuring measurements requiring the storage of the information of the entire universe's state within a subsystem of it, namely the measurement instrument. Because a system exhibiting sensitive dependence is involved and any measurement uncertainty will be amplified, an infinite amount of information stored in the measurement apparatus is required, which is physically impossible.

3 METAPHYSICAL AND EPISTEMOLOGICAL IMPLICATIONS

Complex systems, then, have rich implications for metaphysics and epistemology. Some of these implications for determinism, prediction, confirmation, causation, reductionism and emergence, and laws of nature will be surveyed here. I will begin with some epistemological implications that lead naturally into metaphysical issues.

3.1 Predictability and confirmation

So long as the uncertainty in ascertaining the initial state of a nonlinear system remains infinitesimal, there are no serious limitations on our ability to predict future states of such systems due to rapid growth in uncertainties.[4] In this sense, it is not the presence of a positive global Lyapunov exponent that signals predictability problems for nonlinear systems per se; rather it is the loss of linear superposition that leads to possible rapid growth in *finite* uncertainties in the measurement of initial states.[5] When the disturbance of the initial state due to the act of measurement is included, rapid growth in the total uncertainty in the initial state places impossibly severe constraints on the predictability of individual trajectories of systems or their components over various time scales.

The case of forecasting individual trajectories of complex systems receives the most attention in discussing the implications of chaos and complexity for predictability. For example, even under the perfect model scenario (sec 2.3), no matter how many observations of a system we make, we will always be faced with the problem that there will be a set of trajectories in the model state space that are indistinguishable from the actual trajectory of the target system [Judd and Smith, 2001]. Indeed, even for infinite past observations, we cannot eliminate the uncertainty in the epistemic states given some unknown ontic state of the target system. One important reason for this difficulty is traced back to the faithful model assumption (sec. 2.3). Suppose the nonlinear model state space (of a weather forecasting model, say) is a faithful representation of the possibilities lying in the physical space of the target system (Western European weather, say). No matter how fine-grained we make our model state space, it will still be the case that there are many different states of the target system (ontic states) that are mappable into the same state of the model state space (epistemic states). This means that there will always be many more target system states than there are model states.[6]

The constraints nonlinear systems place on the prediction of individual trajectories do not spell doom for predictability of systems exhibiting complex behavior, however. There are other statistical or probabilistic forms of prediction that can

[4]At least this is the case for chaos. In the more general context of nonlinear dynamics, such as the three-dimensional Navier-Stokes equations, it remains a grand prize challenge question as to whether there are solutions where infinitesimal uncertainties blowup on finite time scales. If answered in the affirmative, the loss of linear superposition would pose more potent challenges to prediction than the much ballyhooed chaos.

[5]More precisely, the loss of linear superposition is a necessary condition for rapid growth in uncertainties. Since nonlinear systems do not always exhibit rapid uncertainty growth, the detailed character of the actual parameter values and nonlinear dynamics has to be investigated for the conditions governing uncertainty dynamics.

[6]At least this is the case for any computational models since the equations have to be discretized. In those cases where we can develop a fully analytical model, in principle we could get an exact match between the number of possible model states and the number of target system states. Such analytical models are rare in complexity studies (many of the analytical models are toy models, like the baker's map, which, while illustrative of techniques, are misleading when it comes to metaphysical and ontological conclusions due to their simplicity).

be effectively applied to such systems (though these also have their limits; see [Smith, 2003]). For instance, one can apply various techniques for forming ensembles of initial states surrounding the assumed actual state of the target system and evolve these ensembles forward in time to forecast the behavior of the target system with specified measures for the uncertainties (e.g., [Judd and Smith, 2001] and references therein).

Moreover, faithfulness problems for our nonlinear mathematical models are problematic for standard approaches to confirming such models. We typically rely on the faithfulness of our mathematical models for our confirmation or verification of their efficacy in capturing behavior of target systems, but when the models are nonlinear and the target systems complex, faithfulness turns out to be inadequate for these standard confirmation practices.

Given a target system to be modeled (e.g., the weather over Western Europe), and invoking the faithful model assumption, there are two basic approaches to model confirmation discussed in the philosophical literature on modeling, focusing on individual trajectories and following a strategy known as piecemeal improvement.[7] These piecemeal strategies are also found in the work of scientists modeling actual-world systems and represent competing approaches vying for government funding (for an early discussion, see [Thompson, 1957]).

The first basic approach is to focus on successive refinements to the accuracy of the initial data fed into the model while keeping the model fixed (e.g., [Laymon, 1989, p. 359]). The intuition lying behind this approach is that if a model is faithful in reproducing the behavior of the target system to a high degree, refining the precision of the initial data fed to the model will lead to its behavior monotonically converging to the target system's behavior. This is to say that as the uncertainty in the initial data is reduced, a faithful model's behavior is expected to converge to the target system's behavior. Invoking the faithful model assumption, if one were to plot the trajectory of the target system in an appropriate state space, the model trajectory in the same state space would monotonically become more like the system trajectory as the data is refined. The second basic approach is to focus on successive refinements of the model while keeping the initial data fed into the model fixed (e.g., [Wimsatt, 1987]). The intuition here is that if a model is faithful in reproducing the behavior of the target system to a high degree, refining the model will produce an even better fit with the target system's behavior given good initial data. This is to say that if a model is faithful, successive model improvements will lead to its behavior monotonically converging to the target system's behavior. Again, invoking the faithful model assumption, if one were to plot the trajectory of the target system in an appropriate state space, the model trajectory in the same state space would monotonically become more like the system trajectory as the model is made more realistic.

What both of these basic approaches have in common is that piecemeal monotonic convergence of model behavior to target system behavior is a means of con-

[7] I will ignore bootstrapping approaches as they suffer similar problems, but only complicate the discussion (e.g., [Koperski, 1998]).

firming the model. By either improving the quality of the initial data or improving the quality of the model, the model in question reproduces the target system's behavior monotonically better and yields predictions of the future states of the target system that show monotonically less deviation with respect to the behavior of the target system. In this sense, monotonic convergence to the behavior of the target system is a key mark for whether the model is a good one. If monotonic convergence to the target system behavior is not found by pursuing either of these basic approaches, then the model is considered to be disconfirmed.

For linear models it is easy to see the intuitive appeal of such piecemeal strategies. After all, for linear systems of equations a small change in the magnitude of a variable is guaranteed to yield a proportional change in the output of the model. So by making piecemeal refinements to the initial data or to the linear model only proportional changes in model output are expected. If the linear model is faithful, then making small improvements "in the right direction" in either the quality of the initial data or the model itself can be tracked by improved model performance. The qualifier "in the right direction," drawing upon the faithful model assumption, means that the data quality really is increased or that the model really is more realistic (i.e., captures more features of the target system in an increasingly accurate way), and is signified by the model's monotonically improved performance with respect to the target system.[8]

However, both of these basic approaches to confirming models encounter serious difficulties when applied to nonlinear models exhibiting sensitive dependence [Koperski, 1998; Bishop, 2008b]. In the first instance, successive small refinements in the initial data fed into nonlinear models is not guaranteed to lead to any convergence between model behavior and target system behavior. Due to the loss of linear superposition, any small refinements in initial data can lead to non-proportional changes in model behavior rendering this piecemeal convergence strategy ineffective as a means for confirming the model. Even a refinement of the quality of the data "in the right direction" is not guaranteed to lead to the nonlinear model monotonically improving in capturing the target system's behavior. The small refinement in data quality may very well lead to the model behavior diverging away from the system's behavior.

In the second instance, keeping the data fixed but making successive refinements in nonlinear models is also not guaranteed to lead to any convergence between model behavior and target system behavior. Due to the loss of linear superposition, any small changes in the model, say by adding additional higher-order terms into the equations, can lead to non-proportional changes in model behavior for the same initial data, again rendering the convergence strategy ineffective as a means for confirming the model. Even if a small refinement to the model is made "in the right direction," there is no guarantee that the nonlinear model will monotonically improve in capturing the target system's behavior. The small refinement in the model may very well lead to the model behavior diverging away from the system's

[8] If one waits for long enough times piecemeal confirmation strategies will also fail for linear systems if there are imperfections in the data or models.

behavior.

So whereas for linear models piecemeal strategies might be expected to lead to better confirmed models (presuming the target system exhibits only stable linear behavior), no such expectation is justified for nonlinear models exhibiting sensitive dependence deployed in the characterization of nonlinear target systems. Even for a faithful nonlinear model, the smallest changes in either the initial data or the model itself may result in non-proportional changes in model output, an output that is not guaranteed to "move in the right direction" even if the small changes are made "in the right direction."[9]

Sticking with the individual trajectory approach, one might consider alternatives to these piecemeal confirmation strategies. One possibility is to turn to a Bayesian framework for confirmation, but similar problems arise here for nonlinear models exhibiting sensitive dependence. Given that there are no perfect models in the model class to which we would apply a Bayesian scheme and given the fact that imperfect models will fail to reproduce or predict target system behavior over time scales that may be long or short compared to our interests, there again is no guarantee that some kind of systematic improvement can be achieved for our nonlinear models.[10] Another approach is to seek trajectories issuing forth from the set of initial conditions in the model space — presumably the actual state of the target system has been mapped into this set — that areis consistent with all observations of the target system over the time period of interest. Given a faithful model, choose an initial condition consistent with the observational uncertainty that then yields a model trajectory passing within the observational uncertainty of the desired future observations. Models can then be judged as better or worse depending on the length of their shadowing times. Finding such trajectories that consistently shadow target system observations for longer and longer times under changes in either the initial data or the models themselves may be quite difficult, however. Furthermore, it is possible to construct models that can shadow any set of observations without those models having any physical correspondence to the target system (e.g., [Smith, 1997, p. 224-225]).

It seems that probabilistic models utilizing ensembles of trajectories or ensembles of probability distributions would allow for a clearer sense of confirmation. Yet, similar problems can crop up here as well. Ensemble forecasting models can give unique, but incorrect indications of the target system's future behavior or such models can give no unique indications of expected future behavior. And even when an ensemble model gives a relatively unique indication that tracks with the outcomes of a target system over a shorter time scale, its indications may diverge significantly from that time point forward.[11] Again, we face difficulties with formulating a systematic confirmation scheme.

[9]Of course, this lack of guarantee of monotonic improvement also raises questions about what "in the right direction" means, but I will ignore these difficulties here.

[10]Here I leave aside the problem that having no perfect model in our model class renders most Bayesian schemes ill-defined.

[11]See [Smith, 1997, pp. 236-237]) for an illuminating example of this in the context of weather forecasting.

While the above difficulties with determining when a nonlinear model is good causes problems for philosophers' desires to produce a systematic theory of confirmation for all models, this situation does not impede physicists and others from finding ways to improve their models and make determinations about how good their imperfect models are. However, it is also the case that these model builders do not follow some universal scheme for improving or confirming their models, but use a variety of techniques (e.g., [Smith, 1997]).

Last, but not least, there are ramifications here for the use of nonlinear models in the development and assessment of public policy. Policy formation and assessment often utilizes model forecasts and if the models and systems lying at the core of our policy deliberations are nonlinear, then policy assessment will be affected by the same lack of guarantee as model confirmation due to the loss of linear superposition. Suppose government officials are using a nonlinear model in the formulation of economic policies designed to keep GDP ever increasing while minimizing unemployment (among achieving other socio-economic goals). While it is true that there will be some uncertainty generated by running the model several times over slightly different data sets and parameter settings, assume policies taking these uncertainties into account to some degree can be fashioned. Once in place, the policies need assessment regarding their effectiveness and potential adverse effects, but such assessment will not be merely a function of looking at monthly or quarterly reports on GDP and employment data to see if targets are being met. The nonlinear economic model driving the policy decisions will need to be rerun to check if trends are indeed moving in the right direction and are on the right course with respect to the earlier forecasts. But, of course, the data fed into the model have now changed and there is no guarantee that the model will produce a forecast with this new data that fits well with the old forecasts used to craft the original policies. How, then, are policy makers to make reliable assessments of policies? The same problem that small changes in data or model in nonlinear contexts are not guaranteed to yield proportionate model outputs or monotonically improved model performance also plagues policy assessment using nonlinear models.

3.2 Determinism

Intuitively, one might think that if a system is deterministic, then it surely must be predictable, but the relationship between determinism and predictability is much too subtle to support this intuition [Bishop, 2003]. Predictability of systems has much to do with epistemic states while determinism has to do with ontic states. And while the characteristics of ontic states should have some implications for the character and behavior of epistemic states, it is difficult at best to draw any conclusions about the ontic states of a system based on our access to epistemic states. This is a fundamental reason why the often discussed unpredictability of chaotic and complex systems by itself does not undermine the determinism of the underlying ontic states of nonlinear systems in classical mechanics. So arguments like Karl Popper's [1950] to the effect that a breakdown in predictability implies a breakdown in determinism trade on confusing epistemic conclusions for ontic ones.

The trend since the Scientific Revolution has been to support the belief in metaphysical determinism by appealing to the determinism of theories and models from physics, though this strategy is not without its subtleties and surprises [Bishop, 2006]. A standard way of characterizing a mathematical model as deterministic is through the property *unique evolution*:

> Unique Evolution: A given state of a model is always followed (preceded) by the same history of state transitions.

The basic idea is that every time one returns the mathematical model to the same initial state (or any state in the history of state transitions), it will undergo the same history of transitions from state to state and likewise for the target system if the faithful model assumption holds. In other words the evolution will be unique given a specification of initial and boundary conditions.[12] For example the equations of motion for a frictionless pendulum will produce the same solution for the motion as long as the same initial velocity and initial position are chosen.

It is not uncommon to find arguments in the literature that purport to show that chaos and complexity tell against determinism. For example, in several books John Polkinghorne has pushed the claim that the kind of sensitive dependence exhibited by complex dynamical systems should lead us to view even the deterministically rigid world of classical mechanics as ontologically indeterministic. Here is an instance of this line of reasoning:

> The apparently deterministic proves to be intrinsically unpredictable. It is suggested that the natural interpretation of this exquisite sensitivity is to treat it, not merely as an epistemological barrier, but as an indication of the ontological openness of the world of complex dynamical systems [Polkinghorne, 1989, p. 43].

He attempts to make this line of thought plausible through demanding a close link between epistemology and ontology under a critical realist reading of the two.

If we remain at the level of dynamical systems — i.e., mathematics — then clearly there is a serious problem with this line of reasoning. Namely, the mathematical equations giving rise to the exquisite sensitivity and attendant predictability problems are deterministic in exactly the sense of unique evolution described above. So our ontic description in terms of these equations push in precisely the opposite direction that Polkinghorne pursues. Although it is true that apparent indeterminism can be generated if the state space one uses to analyze chaotic behavior is coarse-grained, this produces only an epistemic form of indeterminism. The underlying equations are still fully deterministic.

[12]Note that as formulated, unique evolution expresses state transitions in both directions (future and past). It can easily be recast to allow for unidirectional state transitions (future only or past only) if desired.

Instead, to raise questions about the determinism of real-world systems, one has to pursue the nature of these complex models and their implications as well as examine their presumed connection with target systems via the faithful model assumption. As pointed out in sec. 2 above, the mathematical modeling of actual-world systems requires us to make distinctions between variables and parameters as well as between systems and their boundaries. As we saw, these distinctions become problematic in the context of complex systems, where linear superposition is lost and such systems can be exquisitely sensitive to the smallest of influences. Such features raise questions about our epistemic access to systems and models in the investigation of complex systems, but they also raise questions about making sense of the supposed determinism of target systems. As an example, consider applying a deterministic mathematical model to forecasting the weather over Western Europe, where the identity and individuation of that system is questionable (sec. 2.7). What all do we have to include in this model to be able to make some reasonable pronouncement about whether Western Europe's weather is deterministic or not? Do we need only include a particular fluid mass over this particular continent, or over the earth's surface or that plus the stratosphere and magnetosphere, or .... And do we have to include every butterfly flapping its wings to get an identifiable target system?

There is a further problem in our application of deterministic models to actual-world complex systems and our belief that those systems are deterministic. Although the faithful model assumption appears fairly unproblematic in some simple contexts, if the system in question is nonlinear the faithful model assumption raises serious difficulties for inferring the determinism of the target system from the deterministic character of the model. For example, there is the problem that there will always be many more target system states than there are model states as described above (sec. 3.1).

More fundamentally, there is the problem of the mapping between the model and the target system itself. Even for a faithful model, we do not have a guarantee that the mapping between the model and the target system is one-to-one as we customarily assume. The mapping may actually be a many-to-one relation (e.g., several different nonlinear faithful models of the same target system as is the case with competing weather forecasting and climate prediction models) or a many-to-many relationship. For many classical mechanics problems — namely, where linear models or force functions are used in Newton's second law — the mapping between model and target system appears to be straightforwardly one-to-one with plausible empirical support. By contrast, in nonlinear contexts where one might be constructing a model from a data set generated by observing a system, there are potentially many nonlinear models that can be constructed, and each model may be as empirically adequate to the system behavior as any other. For the inference from the deterministic character of our mathematical model to the deterministic character of the target system to hold appears to require either a one-to-one relationship between a deterministic model and target system or that the entire model class in a many-to-one relation be deterministic.[13]

A different approach attempting to call the ontological determinism of the macroscopic world into question via complexity is the research on far-from-equilibrium systems by Ilya Prigogine and his Brussels-Austin Group [Antoniou and Prigogine, 1993; Petrosky and Prigogine, 1996; 1997; Prigogine, 1997].

Conventional physics describes physical systems using particle trajectories as a fundamental explanatory element of its models. If a system of particles is distributed uniformly in position in a region of space, the system is said to be in thermodynamic equilibrium (e.g. cream uniformly distributed throughout a cup of coffee). In contrast, a system is far-from-equilibrium (nonequilibrium) if the particles are arranged so that highly ordered structures appear (e.g. a cube of ice floating in tea). This means that the behavior of a model is derivable from the trajectories of the particles composing the model. The equations governing the motion of these particles are reversible with respect to time (they can be run backwards and forwards like a film). When there are too many particles involved to make these types of calculations feasible (as in gases or liquids), coarse-grained averaging procedures are used to develop a statistical picture of how the system behaves rather than focusing on the behavior of individual particles.

In contrast the Brussels-Austin approach views these systems in terms of models whose fundamental explanatory elements are distributions; that is to say, the arrangements of the particles are the fundamental explanatory elements and not the individual particles and trajectories.[14] The equations governing the behavior of these distributions are generally *irreversible* with respect to time. Moreover, focusing exclusively on distribution functions opens the possibility that macroscopic nonequilibrium models are irreducibly indeterministic, an indeterminism that has nothing to do with epistemic access to the system. If true, this would mean that probabilities are as much an ontologically fundamental element of the macroscopic world as they are of the microscopic.

One important insight of the Brussels-Austin Group shift away from trajectories to distributions as fundamental elements is that explanation also shifts from a local context (set of particle trajectories) to a global context (distribution of the entire set of particles). A system acting as a whole may produce collective effects that are not reducible to a summation of the trajectories and subelements composing the system [Petrosky and Prigogine, 1997; Bishop, 2004]. However, there are serious open questions about this approach, for instance what could be the physical source of such indeterminism[15] and what is the appropriate interpretation of the probabilistic distributions? Thus, the Brussels-Austin approach remains quite

[13]That this last requirement is nontrivial is exemplified in that different modeling teams will often submit proposals for the same project, where some propose deterministic models and others propose nondeterministic models.

[14]This does not imply, as some have erroneously thought (e.g., [Bricmont, 1995, 165-6]) that the Brussels-Austin Group argued there was no such thing as individual particle trajectories in such complex systems.

[15]One possibility is that this kind of indeterminism is ontologically emergent from the underlying dynamics (see sec 3.4 below).

speculative.

Another possible implication of complexity for determinism lies in the sensitivity of such systems to the smallest of disturbances. Some have argued that the sensitive dependence found in macroscopic chaotic systems opens such systems to the influence of quantum effects (e.g., [Hobbs, 1991; Kellert, 1993]). The line of thinking in these sensitive dependence arguments is that nonlinear chaotic systems whose initial states can be localized to a small patch of state space, because of quantum fluctuations, will have future states that can only be localized within a much larger patch of state space. For example, two isomorphic nonlinear systems of classical mechanics exhibiting sensitive dependence, whose initial states differ only in the quantum fluctuations within their initial conditions, will have future states that will differ significantly at later times. Since quantum mechanics sets a lower bound on the size of the patch of initial conditions, unique evolution must fail for such nonlinear chaotic systems.

This provocative line of argument is beset with a number of subtleties and difficulties, however. For example, there are difficult issues regarding the appropriate version of quantum mechanics (e.g., von Neumann, Bohmian or decoherence theories), the nature of quantum measurement theory (collapse vs. non-collapse theories), and the selection of the initial state characterizing the system that must be resolved before one can say clearly whether or not unique evolution is violated [Bishop, 2008c]. Just because quantum effects might influence macroscopic systems exhibiting sensitive dependence does not guarantee that determinism fails for such systems. Whether quantum interactions with such systems contribute *indeterministically* to the outcomes of these systems depends on the currently undecidable question of indeterminism in quantum mechanics, a resolution of the measurement problem, and a decision as to where to place the boundary between system and measurement.

Moreover, the possible constraints of nonlinear classical mechanics systems on the amplification of quantum effects must be considered on a case-by-case basis. For instance, damping due to friction can place constraints on how quickly amplification of quantum effects can take place before they are completely washed out [Bishop, 2008c]. And one has to investigate the local finite-time dynamics for each system because these may not yield any on-average growth in uncertainties (sec 2.4).

3.3 Causation

The analysis of determinism in complex systems is complicated by the fact that there are additional forms of causation arising in such systems that must be taken into account. Indeed, understanding whether a process is deterministic or not often depends upon understanding the underlying causal mechanism(s). There is no consensus account of what causes are; rather, there is a set of accounts — e.g. counterfactual, logical, probabilistic, process, regularity, structural (e.g., see [Sosa and Tooley, 1993]) — that each have strengths and weaknesses (and, perhaps, like definitions of complexity, have different applicability for different purposes). These accounts of causation require rethinking in the face of the richness of nonlinear dynamics. As indicated in section 2 above, chaos, complexity and self-organization are behaviors where complex wholes play important roles in constraining their parts. Such inter-level causation more generally has received little philosophical attention relative to bottom-up efficient modes of causation.

Immanuel Kant was one of the first to recognize the peculiarities of what we now call self-organization in living systems. He classifies such phenomena as *intrinsic physical ends* [Kant, 1980, p. 18] because they are in some sense both cause and effect of themselves. For instance, according to Kant, a tree "in the genus, now as effect, now as cause, continually generated from itself and likewise generating itself, . . . preserves itself generically" [Kant, 1980, p. 18]. An entity is an intrinsic physical end if "its parts, both as to their existence and form, are only possible by their relation to the whole" [Kant, 1980, p. 20]. Self-organizing systems, particularly organisms, are produced by, and in turn, produce the whole. Each part — such as they are distinguishable — exists in virtue of the activity of the other parts and the whole. Furthermore, each part exists for the sake of the other parts as well as the whole. "Only under those conditions and upon those terms can such a product be organized and self-organized *being*, and as such be called a physical end" [Kant, 1980, p. 22].

In Kant's view self-organization "has nothing analogous to any causality known to us" [Kant, 1980, p. 23] because the dominant concrete conception of causation available to him was that of external forces acting on systems generally through contact as exemplified in the model systems of Newtonian mechanics. Given his recognition that self-organizing systems required some kind of time-irreversible processes and that Newtonian dynamics was fully time-reversible, he relegated our lack of understanding how self-organization comes about to a limitation of reason [Kant, 1980, pp. 22-4]. Counterfactual, logical and regularity analyses of causation fare no better at penetrating this lack of understanding. While process and structural accounts each appear to have some pieces of the puzzle for understanding self-organization, process theories lack an adequate account of the structural constraints of wholes on parts, while structural theories lack an adequate account of processes.

Causation in complex systems has been given very little sustained analysis in the philosophy literature relative to causation in general ([Juarrero, 1999] is a notable exception). Probably this lack of attention is largely due to a widely shared assumption that causal analysis in complex systems is no different in kind than in typical metaphysics literature (save it is obviously more complex than the usual examples on which philosophers tend to focus). However, complexity raises difficult questions for thinking about causation when nonlinear inter-level relationships, rapid amplification of the smallest perturbations and so forth are present. For example, how are we to identify the causes at work in systems exhibiting sensitive dependence? What to do if quantum effects possibly can play causal roles in such systems (sec 3.2) or electrons dancing about in a distant galaxy possibly can play a causal role in such systems here on Earth (sec 2.7)? How far down do we have to go to identify all the causes at work in a complex macroscopic system (e.g., to butterfly wing flaps, to the atomic level or beyond)? Or how far do we have to extend a complex system to capture all of its causes (e.g., weather near the earth's surface or must we include the stratosphere, troposphere, magnetosphere, solar system, etc.)? There is a real problem here of specifying just what the causes in complex systems are aside from the trivial answer: everything! Finding principled ways to draw the boundary around the "crucial" or "dominant" causes at work in complex systems is difficult, to say the least because one of the lessons that nonlinear dynamics teaches us is that "small" causes are not insignificant (e.g., [Bishop, 2008a; 2008c]).

Hierarchies also raise questions about causation in complex systems (secs. 2.6 and 2.7). Typical metaphysical analyses consider all causation as "bottom up," where the system components are the only causal actors and systems as a whole have causal power in virtue of the causal powers of their constituents (e.g., [Kim, 2007]). But when control hierarchies act to limit or govern the causal influence of system components, is this not "top down?" Instances where lower levels provide necessary but insufficient conditions for the total behavior of wholes are rather routine in complexity systems. Moreover, when higher levels in a hierarchy and wholes act to constrain or direct the causal powers of constituents, even lower-level constituents in and of themselves turn out to not have necessary and sufficient conditions for governing all of their behavior (sec. 3.4). To repeat some key ideas of secs. 2.6 and 2.7 in slightly different language, in complex systems the formation of control hierarchies often comes about when a new form of dynamics arises that exhibits downward constraint on system constituents and is self-sustaining (e.g., [Hooker, 2004, pp. 449-477; Bishop, 2008c]). This kind of dynamical top-down constraint has a character more resembling Aristotle's notion of formal cause than efficient cause and has been largely unanalyzed by analytic philosophers (who tend to focus on logical and formal relationships among efficient causes in bottom-up constructions than on dynamics and dynamical relations).

3.4 Reduction and emergence

The issues of identity and individuation as well as causation in complex systems lead naturally to a discussion of reduction and emergence in complex systems. In rough outline, reductionist lore maintains that properties and behavior of systems as a whole are completely determined by the states and properties of their parts (ontic claim) or are explainable in terms of the states and properties of their parts (epistemic claim). Defenders of emergence deny one or both of these claims. The property of linear superposition plays an interesting role in the concepts of resultant and emergent forces in such systems. However, the loss of superposition and the possibilities for holism and constraining causation lead to the need to consider an alternative to the received views.

For instance, the lack of necessary and sufficient conditions for the behavior of lower-level constituents in complex systems directly challenges reductive atomism (e.g., control hierarchies). One of the core principles of atomistic physicalism as identified by Robert van Gulick is that "The only law-like regularities needed for the determination of macro features by micro features are those that govern the interactions of those micro features in all contexts, systemic or otherwise" [van Gulick, 2001, p. 18]. However, in complex systems control hierarchies and other inter-level causal relations are crucial to determining the behavior of system constituents. The behavior of constituents in such systems is *conditioned* by contexts in which the constituents are situated.

Recall the characterization of nonlinear systems and the linear superposition principle (sec. 2.2) and the relationship of the failure of that principle to hierarchical and large-scale structure behavior (sec. 2.6). When linear superposition holds, a system can be decomposed into its constituent parts and the behavior of each component is independent of the other components. This is the typical way in which philosophical proponents of reductionists tend to conceive of all systems. In contrast, when linear superposition breaks down, as it does for complex systems, such systems often exhibit behaviors reflecting the fact that individual system components are not independent of each other. Moreover, the behavior of individual system components are not even independent of the wholes (and various hierarchies in between). Hierarchies and wholes act to enable or constrain various possibilities for component behavior relative to what would be possible for the components if the hierarchies and wholes were absent.[16]

The interplay between parts and wholes in complex systems leads to the self-organization observed in such systems. Their sensitive dependence on the smallest of changes at the component level is partly guided by the inter-level causal relations in such systems (e.g., determining the characteristic features of convecting cells in Rayleigh-Bénard convection due to initial perturbations and instabilities in the system). This kind of behavior may be fruitfully captured by the concept of *contextual emergence* (see [Bishop, 2005b; Bishop and Atmanspacher, 2006]):

> The properties and behaviors of a system at a particular level (including its laws) offer necessary but not sufficient conditions for the properties and behaviors at a higher level.

The reference to necessary conditions at the lower level means that properties and behaviors of components at the higher level of a system may imply the properties and behaviors of components at the lower level. However, the converse is not true as the lower-level properties and behaviors do not offer sufficient conditions for the properties and behaviors of higher-level components. Contingent conditions specifying the context for the transition from the lower to the higher level of properties and behaviors are required to provide such sufficient conditions. In complex

[16] [Juarrero, 1999; Silberstein and McGeever, 1999; Bishop, 2004; Bishop and Atmanspacher, 2006; Ellis, 2006; Bishop, 2008a]. In the language of Lagrangian mechanics the failure of the laws and conditions at the lower level to serve as both necessary and sufficient conditions for higher-level behavior is due to the failure of the constraints to be holonomic (see [Symon, 1971], sec 9.4, for a discussion of holonomicity).

systems, such contingent contexts are not given by the lower-level properties and behaviors alone. Moreover, the conditions for specifying a candidate reduction are not well defined until an appropriate contingent context is properly specified [Hooker, 2004, pp. 467-468].

In this sense, complex systems seem to validate intuitions many emergentists have that some form of holism plays an important causal role in complex systems that is missed by reductionist analyses. However, care is needed with typical emergentist slogans such as "The whole cannot be predicted from its parts," or "The whole cannot be explained from its parts" when it comes to contextual emergence like that exemplified in complex systems. Relevant information about the lower-level properties and behaviors of constituents *plus* the specification of an appropriate contingent context allows for the (in principle) prediction or explanation of higher-level properties and behaviors in many cases (e.g. [Primas, 1998; Bishop, 2005b; Bishop and Atmanspacher, 2006]). So complexity holds surprises for both proponents of reductionism and emergence.

3.5 Laws

While most metaphysicians focus on the "upward" flow of efficient causation from system components to system behavior as a whole, the possibilities for inter-level causal relationships in the dynamics of complex systems like convecting fluids present plausible examples of a "downward" flow of causation constraining the behavior of system components. Such behaviors clearly raise questions about the nature of laws in complex systems.

A very popular conception of science is that its goal is to discover natural laws often given form as universal statements in scientific theories. Philosophically there have been two main traditions for analyzing the nature of these laws. One tradition is the necessitarian tradition, where laws represent genuine necessities governing the regularities and patterns we find in reality. The other tradition is the regularity tradition, where laws are descriptive of regularities and patterns but no genuine underlying necessities exist. Although these two traditions involve differing relationships to the various theories of causation on offer, a very important class of laws in scientific theories have been causal laws, which govern or specify the history of state transitions a system will make given some initial starting conditions. Such causal laws often play important roles in philosophical analyses of science as well as metaphysical accounts of reality, however features like inter-level causation challenge an exclusive focus on this role.

For example, exclusive focus on causal laws might lead one to worry that fundamental laws of physics — viewed as causal laws — are being violated if the lowest-level constituents of systems (e.g., elementary particles, molecules) plus the fundamental laws are not jointly necessary and sufficient to fully determine the behaviors of these constituents (and, thereby, determine the behaviors of all higher-level constituents). Much is going on in this philosophical worry. Part of what is presupposed in this worry is an understanding of natural laws as being universal in the sense of being context free (as in the characterization of atomistic physicalism given above in sec. 3.4). Another presupposition in this worry is that the laws governing actual-world systems are only or primarily of the causal variety. Although perhaps under analyzed, science also includes structuring laws that govern or structure the range of possibilities, but do not necessarily specify which of those possibilities are actualized.

Rather than giving in to the worry, perhaps the lesson of nonlinear dynamics and complexity is that we should reconceive the primary role laws play in scientific theories. After all, scientific theories are notorious for eschewing language of the efficient causal type with which philosophers are enamored. What fundamental laws primarily do is carry out structuring functions, where the relevant space of possibilities is determined for the behaviors of the constituents guided by them, but such laws do not fully determine which of these possibilities in the allowable space are actualized. That can only be fully determined by concrete contexts into which the laws are coming to expression (which may involve higher-level causal and/or structuring laws). For instance, Newton's law of gravity determines the space of possible motions for an apple falling from a tree, but the concrete context where I reach out my hand and catch the apple actualizes a particular possibility among all those allowed even though the particular context is not included in the law of gravity.

The point, here, is very similar to that for contextual emergence above (sec. 3.4): while fundamental laws establish necessary conditions for the possible behaviors of objects, contingent contexts must be added in to establish jointly necessary and sufficient conditions for actual behaviors. If the fundamental laws play primarily structuring roles in nature, then concrete contexts are as important as the laws.[17] Empirically, this is consonant with the intricate and delicate behavior of complex systems. We should, then, resist the tendency to either appeal to only fundamental causal laws in our explanations or to pit causal laws against structuring laws in competing explanations. Sound explanations of complex systems are likely to involve both appeals to causal mechanisms and ordering/constraining structure via the interrelationships among the lower-level and higher-level dynamics[18]

If fundamental laws primarily structure the spaces of possibility they establish but do not fully determine the outcomes within this space, worries about violations of fundamental laws fade away. Hierarchies and wholes in complex systems act to constrain or direct the possibilities made available by lower-level laws as opposed to somehow violating those laws. Since such laws are part of the necessary conditions for the behavior of higher-level constituents, such laws cannot be violated by the behavior of higher-level constituents. Indeed, it is possible that the structuring function played by control hierarchies and wholes in complex systems are the result of some as yet unknown nonlinear dynamical laws, which may be causal

[17]Similar lessons about laws can be gleaned from continuum and statistical mechanics even in cases where systems are not exhibiting complexity.

[18]See [Chemero and Silberstein, 2008, sec. 4; Bishop, 2008b, sec. 5.2] for some discussion as to how these two levels might be brought together fruitfully in the service of explanation.

or structural as well as emergent with respect to the contexts where nonlinear interactions are dominant. Even so, fundamental laws still would be necessary for structuring the possibility space for emergent nonlinear laws, though particular features of contexts might be required for particular nonlinear laws to arise.

But complexity raises more questions for laws than the relative roles of structural vs. causal laws. It also raises questions about how we are to conceive laws that require us to fix appropriate boundary conditions for their corresponding equations to be well posed.[19] For example, the fluid equations governing Rayleigh-Bernard convection require the imposition of a constant temperature along the bottom plate of a container holding fluid so that a temperature difference can be established. In an idealized mathematical model of this situation, there is a principled choice to make for the boundary. However, when these same fluid equations are applied to model atmospheric weather, all of the boundary problems mentioned above in sec. 2 arise, where there is no longer any obvious choice for where to place the cut between weather system and boundary. Operationally it is fine that we can make pragmatic choices that give us well-posed equations to solve on a computer, but the foundational questions of the status of the laws in question and the determination of their boundaries remains unanswered. Perhaps such laws as philosophers have typically conceived them have no instantiations for complex systems because these systems lie outside the laws' domain of applicability, given the lack of an ontologically distinguishable boundary. Yet, we still have stable patterns where the dynamics governs the outcomes even if our philosophical analyses of laws come up short in characterizing the dynamics.[20]

Moreover, we typically connect laws with physical systems via models, which means that the faithful model assumption (sec. 2.3) is being invoked. Faithful yet imperfect models leave open questions about the applicability of the laws to phenomena exhibiting complexity. Even if we take the faithful model assumption to its extreme limit — the perfect model scenario — we run into problems since there are too many states indistinguishable from the actual state of the system yielding empirically indistinguishable trajectories in the model state space [Judd and Smith, 2001]. Add into the mix that the mapping between our nonlinear dynamical models could be many-to-one or many-to-many, and our philosophical judgements become difficult about the roles scientific laws play in the corresponding actual-world systems and which of these laws are fundamental and which are emergent (if any).

Empirically, it is hard to be sanguine about the necessitarian tradition on laws which, in turn, puts pressure on realist views of natural laws. On the other hand, the regularity tradition does not rest too comfortably either. The difficulties nonlinear dynamical systems raise for the notion of faithful models lead to analogous

[19]For an insightful survey of the delicate nature of the relationship between laws and boundary conditions, see [Wilson, 1990].

[20]Similar questions can be raised about the status of scaling laws and laws involving universality and order parameters, which are ubiquitous in literature on complex systems. The subtleties of such systems warrant care in being too quick to conclude that these laws are "merely epistemic" because we "lack access" to the underlying causal mechanisms at work in such systems.

worries about the role regularities play in actual-world systems, and about whether the regularities in these systems are fundamental and/or emergent. The regularities of complex systems are there, to be sure, and engage intense scientific study, but our current philosophical understanding of laws appears to be inadequate.

4 DISCUSSION

This brief survey of complexity and its implications suggests that there are challenges to our philosophical and scientific lore about the world. Nonlinear dynamics and the loss of linear superposition shed different light than is typically found in philosophical literature on identity and individuation, predictability, confirmation, determinism, causation, reduction and emergence, and natural laws, some of the bread and butter topics of metaphysics, epistemology and philosophy of science. There is much remaining to be explored about how nonlinear dynamics and complexity can challenge and enrich our understanding of metaphysics and epistemology as well as science and its practices.

On the one hand, nonlinear modeling and insights from complexity studies are in widespread use in the sciences (and increasingly in public policy). There is a genuine sense that nonlinear models have led to tremendous advances (e.g., in chemical studies, drug development, weather prediction, plasma physics). On the other hand, extracting firm results from such modeling is nontrivial, given that these results are context-dependent, qualified by some very strong idealizations and difficult to confirm. This is not the kind of picture of science presented in our textbooks and a lot of the philosophical literature. Moreover, the upshot of complex systems modeling for our pictures of the world, whether deterministic, causal or reductionistic, is challenging and nuanced.

Many philosophers and scientists have reflected on science and its methods with a set of assumptions that underestimate the complexity of science itself. Our assumption that faithful models are relatively unproblematic and relatively straightforward methodologically seemed to serve fine in scientific inquiry before the advent and explosion of nonlinear modeling and complexity studies. However, the latter have pushed our philosophical and scientific understanding to their limits as described above.

One of the important challenges to our thinking about science and its methods that complex system make abundantly clear is clarifying our understanding of the strengths and weaknesses of nonlinear modeling. Roughly, on the strength side, there is greatly increased power to model actual-world phenomena that defies our analytical capabilities. In domains like fluid flow, weather forecasting and fusion reactor modeling, tuning our models to the best available empirical data sets has proved quite useful for refining our models. On the limitation side, there is the challenge of understanding how to extract useful, reliable information from imperfect and difficult-to-confirm models that do have some faithfulness to the target systems of interest. For instance, even when we have tuned our nonlinear models to the best data sets, there is still a great deal of inadequacy in those data sets (no data set is ever perfect) that is translated into our models; this is over and above the inadequacy inherent in our faithful models. At the end of the process there still remains a great deal of uncertainty in the results of our models. To what extent do modelers and end users who receive model outputs as deliverables understand the limitations of these models? To what extent do they understand the uncertainties inherent in the modeling results? What strategies can modelers use to extract trustworthy information in the midst of such model inadequacy and uncertainty? What changes do end users, such as public policy officials, need to make so that their reasoning reflects the limitations and uncertainty of the modeling science on which they are relying? These latter questions are some of the most difficult and pressing philosophical and methodological issues raised by complex systems and their models. Efforts by philosophers, scientists, public policy experts and others to answer such questions would be effort well spent.

Any insights gleaned in exploring these questions will no doubt further our thinking in metaphysics and epistemology about determinism, causation, reduction/emergence, natural laws and confirmation in the context of complex systems. One example that is relevant to the latter issues is raised by the possibilities for there to be a many-to-one relationship between different nonlinear models and a single target system: Are such mathematical models simulating the target system or merely mimicking its behavior? To be simulating a system suggests that there is some actual correspondence between the model and the target system it is designed to capture. In contrast, if a model is merely mimicking the behavior of a target system, there is no guarantee that the model has any genuine correspondence to the actual properties of the target system. The model merely imitates behavior. This question becomes particularly important for modern techniques of building nonlinear dynamical models from large time series data sets (e.g., [Smith, 1992]), for example the sunspot record or the daily closing value of a particular stock for some specific period of time. In such cases, after performing some tests on the data set, modelers set about their work constructing mathematical models that reproduce the time series as their output. When multiple such models, each conceptually and mathematically different, reproduce the target system behavior with roughly equal accuracy (and inadequacy!), are such models simulating target systems, or only mimicking them? Here, all the questions about understanding model limitations and uncertainties, strategies for extracting useful information and how to reason about public policy or other implications based on model output are raised. In addition, for the philosophers further questions about what these models are telling us about our world and our access to that world are also raised with their attendant implications for metaphysics and epistemology.

The metaphysical and epistemological implications of complex systems is very rich, indeed.

ACKNOWLEDGEMENTS

Insightful discussions with Cliff Hooker, Jeffery Koperski, Michael Silberstein and Leonard Smith are all gratefully acknowledged. Remaining weaknesses are my sole responsibility.

BIBLIOGRAPHY

[Antoniou and Prigogine, 1993] I. Antoniou and I. Prigogine. Intrinsic Irreversibility and Integrability of Dynamics. *Physica A*, 192, 443-464, 1993.

[Atmanspacher, 2002] H. Atmanspacher. Determinism Is Ontic, Determinability Is Epistemic. In H. Atmanspacher and R. C. Bishop, eds., *Between Chance and Choice: Interdisciplinary Perspectives on Determinism*. Imprint Academic, Thorverton, UK, 2002.

[Badii and Politi, 1997] R. Badii and A. Politi. *Complexity: Hierarchical Structures and Scaling in Physics*. Cambridge University Press, Cambridge, 1997.

[Bishop, 2002] R. C. Bishop. Deterministic and Indeterministic Descriptions. In H. Atmanspacher, R. Bishop (eds.), *Between Chance and Choice: Interdisciplinary Perspectives on Determinism*. Imprint Academic, Thorverton, UK, 2002.

[Bishop, 2003] R. C. Bishop. On Separating Prediction from Determinism," *Erkenntnis* 58, 169-188, 2003.

[Bishop, 2004] R. C. Bishop. Nonequilibrium Statistical Mechanics Brussels-Austin Style," *Studies in History and Philosophy of Modern Physics* 35,1-30, 2004.

[Bishop, 2005a] R. C. Bishop. Anvil or Onion? Determinism as a Layered Concept," *Erkenntnis* 63, 55-71, 2005.

[Bishop, 2005b] R. C. Bishop. Patching Physics and Chemistry Together," *Philosophy of Science* 72, 710-722, 2005.

[Bishop, 2006] R. C. Bishop. Determinism and Indeterminism. In D. M. Borchert (ed. in chief), *Encyclopedia of Philosophy, Second Edition*. Macmillian Reference, Farmington Hills, MI, 3, 29-35, 2006.

[Bishop, 2008a] R. C. Bishop. Downward Causation in Fluid Convection. *Synthese* 160, 229-248, 2008.

[Bishop, 2008b] R. C. Bishop. Chaos. *Stanford Encyclopedia of Philosophy*, `http://plato.stanford.edu/entries/chaos/`, 2008.

[Bishop, 2008c] R. C. Bishop. What Could Be Worse than the Butterfly Effect?", *Canadian Journal of Philosophy* 38, 519-548, 2008.

[Bishop and Atmanspacher, 2006] R. C. Bishop and H. Atmanspacher. Contextual Emergence in the Description of Properties. *Foundations of Physics* 36, 1753-1777, 2006.

[Bricmont, 1995] J. Bricmont. Science of Chaos or Chaos in Science? *Physicalia Magazine* 17, 159-208, 1995.

[Busse, 1978] F. Busse. Non-Linear Properties of Thermal Convection. *Reports on Progress in Physics* 41, 1929-1967, 1978.

[Chemero and Silberstein, 2008] A. Chemero and M. Silberstein. After the Philosophy of Mind: Replacing Scholasticism with Science. *Philosophy of Science* 75, 1-27, 2008.

[Cross and Hohenberg, 1993] M. Cross and P. Hohenberg. Pattern Formation Outside of Equilibrium. *Reviews of Modern Physics* 65, 851-1112, 1993.

[Crutchfield, 1994] J. Crutchfield. Observing Complexity and the Complexity of Observation. In H. Atmanspacher and G. Dalenoort, eds., *Inside Versus Outside*. Springer, Berlin, 235-272, 1994.

[d'Espagnat, 1994] B. d'Espagnat. *Veiled Reality: An Analysis of Present-Day Quantum Mechanical Concepts*. Addison-Wesley, Reading, MA, 1994.

[Ellis, 2006] G. F. R. Ellis. Physics and the Real World. *Foundations of Physics* 36, 227-262, 2006.

[Grassberger, 1989] P. Grassberger. Problems in Quantifying Self-Generated Complexity. *Helvetica Physica Acta* 62, 489-511, 1989.

[Hobbs, 1991] J. Hobbs. Chaos and Indeterminism. *Canadian Journal of Philosophy* 21,141-164, 1991.

[Hooker, 2004] C. A. Hooker. Asymptotics, Reduction and Emergence. *British Journal for the Philosophy of Science* 55, 435-479, 2004.

[Judd and Smith, 2001] K. Judd and L. Smith. Indistinguishable States I: Perfect Model Scenario." *Physica D* 151, 125-141, 2001.

[Juarrero, 1999] A. Juarrero. *Dynamics in Action: Intentional Behavior As a Complex System.* MIT Press, Cambridge, MA, 1999.

[Kant, 1980] I. Kant. *The Critique of Judgement*, trans. J. C. Meredith. Oxford University Press, Oxford, 1980.

[Kellert, 1993] S. Kellert. *In the Wake of Chaos.* Chicago University Press, Chicago, 1993.

[Kim, 2007] J. Kim. *Physicalism, or Something Near Enough.* Princeton University Press, Princeton, 2007.

[Kripke, 1980] S. Kripke. *Naming and Necessity.* Basil Blackwell, Oxford, 1980.

[Koperski, 1998] J. Koperski. Models, Confirmation and Chaos. *Philosophy of Science* 65, 624-648, 1998.

[Lorenz, 1965] E. N. Lorenz. A Study of the Predictability of a 28-Variable Atmospheric Model. *Tellus* 17, 321-33, 1965.

[McLaughlin, 1982] B. McLaughlin. British Emergentism. In A. Beckermann, H. Flohr, and J. Kim, eds., *Emergence or Reduction? Essays on the Prospects of Nonreductive Physicalism.* Walter de Gruyter, Berlin, 49-93, 1982.

[Oseledec, 1969] V. I. Oseledec. A Multiplicative Ergodic Theorem: Lyapunov Characteristic Number for Dynamical Systems. *Transactions of the Moscow Mathematical Society* 19, 197-232, 1969.

[Pattee, 1973] H. Pattee. The Physical Basis and Origin of Hierarchical Control. In H. Pattee (ed.), *Hierarchy Theory: The Challenge of Complex Systems.* George Braziller, New York, 69-108, 1973.

[Petrosky and Prigogine, 1996] T. Petrosky and I. Prigogine. Poincaré Resonances and the Extension of Classical Dynamics. *Chaos, Solitons & Fractals* 7, 441-497, 1996.

[Petrosky and Prigogine, 1997] T. Petrosky and I. Prigogine. The Extension of Classical Dynamics for Unstable Hamiltonian Systems. *Computers & Mathematics with Applications* 34, 1-44, 1997.

[Polkinghorne, 1989] J. C. Polkinghorne. *Science and Creation: The Search for Understanding.* Shambhala, Boston, 1989.

[Popper, 1950] K. Popper. Indeterminism in Quantum Physics and in Classical Physics. *The British Journal for the Philosophy of Science* 1, 117-133, 1950.

[Prigogine, 1997] I. Prigogine. *The End of Certainty: Time, Chaos, and the New Laws of Nature.* The Free Press, New York, 1997.

[Primas, 1983] H. Primas. *Chemistry, Quantum Mechanics and Reductionism: Perspectives in Theoretical Chemistry.* Springer, Berlin, 1983.

[Primas, 1990] H. Primas. Mathematical and Philosophical Questions in the Theory of Open and Macroscopic Quantum Systems. In A. Miller, ed., *Sixty-Two Years of Uncertainty.* Plenum, New York, 233-57, 1990.

[Primas, 1994] H. Primas. Endo- and Exotheories of Matter. In H. Atmanspacher and G. Dalenroot, eds., *Inside Versus Outside.* Springer-Verlag, Berlin, 163-93, 1994.

[Primas, 1998] H. Primas. Emergence in Exact Natural Sciences. *Acta Polytechnica Scandinavica* 91, 8398, 1998.

[Silberstein and McGeever, 1999] M. Silberstein and J. McGeever. The Search for Ontological Emergence. *The Philosophical Quarterly* 49, 182-200, 1999.

[Scheibe, 1964] E. Scheibe. *The Logical Analysis of Quantum Mechanics.* Pergamon Press, Oxford, 1964.

[Smith, 1992] L. A. Smith. Identification and Prediction of Low Dimensional Dynamics. *Physica D* 58, 50-76, 1992.

[Smith, 1997] L. A. Smith. Past and Present Variability of the Solar-Terrestrial System: Measurement, Data Analysis and Theoretical Models. In G. C. Castagnoli, A. Provenzale (eds.), *Proceedings of the Enrico Fermi International School of Physics, Course CXXXIII.* IOS, Amsterdam, 177-246, 1997.

[Smith, 2000] L. A. Smith. Disentangling Uncertainty and Error: On the Predictability of Nonlinear Systems. In A. Mees, ed., *Nonlinear Dynamics and Statistics.* Birkhauser, Boston, 31-64, 2000.