ECOLOGY AS HISTORICAL SCIENCE

Bryson Brown

1 HISTORY IN SCIENCE AND HISTORICAL SCIENCES

There are two long-term issues about the nature of science that I want to address in this essay. The first concerns the overall *shape* of our scientific understanding of the world—an issue that was once central to philosophy of science, but has been largely eclipsed during the twentieth century. The second concerns the different subject matters of the sciences we distinguish today, and the implications of these for the forms of scientific understanding we seek.

The shape of science as a whole was once a central preoccupation of philosophers of science, who developed systematic views about the subject matters of the different sciences and their relations, about their relative standing in terms of authority and fundamentality and their special methodological characters. Taxonomies of human knowledge about the natural world traditionally distinguished *natural history* from *natural philosophy*, with natural history conceived as taxonomic and descriptive while natural philosophy dealt with causal relations, and aimed to produce not just descriptions but *explanations* of its phenomena.

However, during the eighteenth and nineteenth centuries, causal questions began to arise in fields that had been part of natural history. In biology, taxonomic work by Linnaeus, John Ray and others led to a plethora of newly recognized species, giving rise to puzzles about the distribution of different forms of life around the world, about the nature of species (arising especially from the difficulty of distinguishing well-marked subspecies from closely resembling but distinct species), about relations between different species (Linnaeus suggested some species had arisen by hybridization of other species), about extinction (arising from a growing recognition that numerous fossil species seemed to lack living representatives), and about the origin of species (as growing knowledge of the fossil record suggested the familiar species of today did not exist during earlier periods in the earth's history) (see [Young, 76f; Rudwick, 2007, 349f.]). In geology, the description of formations and their spatial relations (driven in large part by the practical concerns of miners) led scientists to an increasingly historical vision of their origins (see [Rudwick, 2007 181f.]), consciously developed in parallel with antiquarian history, in which fossils and other traces of the earlier earth served in place of ancient monuments and buildings to illuminate a distant, unrecorded past.

Since then, the sciences of natural history have become both explicitly causal and truly historical. As a result, the kind of explanation that we find in these sciences is notably different from the ideal of explanation inherited from the western

Handbook of the Philosophy of Science. Volume 11: Philosophy of Ecology.
Volume editors: Kevin deLaplante, Bryson Brown and Kent A. Peacock. General editors: Dov M. Gabbay, Paul Thagard and John Woods.

philosophical tradition. Ancient ideas about epistemology focused on what is fixed and unchanging as the proper objects of knowledge; the ever-changing world of individuals and their particular stories, inconstant and imperfect as they are, were regarded, at best, as second-class subjects of knowledge. Thus geometry, mathematics, biological and other species (conceived as fixed kinds to which various concrete and imperfect individuals belong), and physics (as a universal science of motion and change) provided true material for knowledge. This kind of knowledge is fully understood because it is grounded in the unchanging, necessary nature of things.

There is a special sort of explanation that seems a credible goal for such sciences, and unachievable for a truly historical science. On the assumption that fundamental principles are self-explanatory,[1] such explanations are *closed*, that is, they appeal to nothing that is not itself explained. But historical explanations are always open, appealing to conditions and circumstances and sequences of events as boundary conditions that themselves remain unexplained. Such explanations have traditionally been regarded as *incomplete* (hence various regresses of explanation, central to some forms of cosmological argument). But if we regard an explanandum as truly contingent, we cannot expect it to be explained except by appeal to other contingencies.

In what follows we will examine the distinction between the *natural sciences*, generally regarded as including physics, chemistry, and some aspects of biology and even ecology when directed towards present life and considered aside from the historical/evolutionary origins of living things and long-term changes in ecosystems, and the *historical sciences*, including cosmology, geology, earth systems science, evolutionary biology and ecology. This division, which is familiar in outline but perhaps not in detail, along with the motives for drawing a line separating these sciences will be examined carefully here. I will argue that there are indeed close parallels connecting biology and especially ecology to the historical sciences, some with important methodological implications, although the most important parallels are not really about history at all.

As the eighteenth century French savant the Comte de Buffon understood it, ecology represented an interesting middle stage in the emergence of historical science. Buffon envisaged a natural course of development for the earth and for life on earth—so his vision is clearly causal. And this natural course of development provides a narrative for a history of life on earth. According to Buffon, life formed as soon as the temperature of the cooling earth was low enough to allow it, arising first at the poles. As the earth continued to cool, the first forms of life migrated towards the equator while new forms, adapted to colder conditions, arose at the poles. In general, climatic and soil conditions determine all the rest: "Ainsi la terre

[1]See Aristotle, *An Post.* A.2: "We suppose ourselves to possess unqualified scientific knowledge of a thing, as opposed to knowing it in the accidental way in which the sophist knows, when we think that we know the cause on which the fact depends, as the cause of that fact and of no other, and, further, that the fact could not be other than it is." The intellect ('nous') is involved in the special grasp we have of principles, which cannot be demonstrated; this grasp assures us that the principles are among the things that 'could not be other' than they are.

fait les plantes, la terre et les plantes font les animaux, la terre, les plantes et les animaux font l'homme" (Thus the land makes the plants, the land and the plants make the animals, the land, the plants and the animals make man) [Buffon, 1756, p. 58]. Later biogeographical research made it clear that climate and soil were not enough to determine flora and fauna. But Buffon's notion of a tight link between physical conditions and life forms persisted in Lyell's extreme uniformitarianism, when he proposed that should the climate revert to Mesozoic conditions, it might bring back the prehistoric beasts of those times: "The huge iguanodon might reappear in the woods, and the ichthyosaur in the sea, while the pterodactyle might flit again through umbrageous groves of tree-ferns" [Lyell, 1830, p. 123].

But Buffon's narrative is not fully historical in the sense set out above: on Buffon's account, there is only one possible course for the history of life, only a single possible ecology for any specific geographical region (soil and climate). This is closely related to an attractive explanatory ambition, viz. to arrive at an account of things that (at least in principle) *rules out* alternatives, not conditionally, but absolutely; such an explanation is only possible if the present state of things is regarded as somehow inevitable, rather than contingent. Admittedly, there *is* one dramatic contingency at the very outset of Buffon's theory of the earth, as the material of the planets is dragged out of the sun by the gravitational force of a passing comet. But given that one cosmic accident (itself, perhaps, bound to occur *somewhere* in the vastness of our universe), the subsequent course of events is firmly fixed—something entirely foreign to the richly contingent narratives of the historical sciences.

The goal of developing an over-view of the entirety of human knowledge, and the special place and contribution that each science has in it, faded in importance for twentieth century philosophers of science. Nevertheless, some questions about relations between different sciences have remained important. First, there are logical and metaphysical questions about the relations between different sciences, revolving closely around the ideas of *reduction* and *supervenience*. Second, there are questions about evidence and methodology, with implications for the special methods and the relative authority of different sciences (an issue that becomes particularly important when conflicts arise, but also contributes to debates about, and misunderstandings of, the scientific world view). Finally, there are questions about the nature of the *explanations* that the different sciences provide for the phenomena they address.

Immense and wide-ranging effort is required to produce, evaluate and extend theories like evolution by natural selection or plate tectonics. Great effort is needed even to establish important empirical regularities, such as William Smith's insight into fossils and their regular appearance in certain geological formations, in turn related by superposition and thus by time. Similarly, great intellectual insight was required to connect the layered structure of rock formations to the temporal ordering of the processes that deposited them (and the principles of stratigraphy that follow from that connection between spatial structure and temporal sequence) [Albritton, p. 34, Rudwick, pp. 203–214]. The historical sciences have been able to transform our understanding of the natural world so dramatically because their practitioners have built a detailed and richly connected body of knowledge about the natural world and the processes by which many of its features have developed over time.

But the contribution the historical sciences have made to our understanding of the world is notably different from that of the more theory-centered sciences. Rather than focus on theory, which provides an account of basic concepts and their inferential relations[2] (an account that is intuitively freighted with a kind of necessity, encouraging metaphysical notions about natural laws and essences), the historical sciences aim first at *applications*, via the construction of narratives, contingent from the start, which coherently encompass and account for many patterns and features of life and the earth. The coherence of this narrative is not a matter of having demonstrated that it follows from first principles, and it is often largely insensitive to the details of those principles; it arises instead from the fact that the processes invoked in the narrative are grounded in and cohere with a rich and ever-growing variety of observations. Though these observations and the narrative we embed them in are contingent in themselves, they often fit together so intricately that it would be difficult (at best), and impossible (practically speaking) to invent an equally rich and coherent fiction.

In this essay I argue that we should think of ecology as an historical science, despite the fact that ecological models do not, in general, appeal to long periods of time as part of the story they tell about the populations, communities or ecosystems they represent. Ecology shares important features with evolutionary biology, geology and other historical sciences—features that illuminate the epistemic contact between ecological models and the phenomena we apply them to, the limitations of ecological models as predictive tools and the kinds of explanation we can expect from ecological models. By these criteria, ecology fits with the historical sciences—more generally, it emerges that these epistemological and explanatory characteristics that it shares with the historical sciences provide a more interesting dividing line within science than the element of historicity itself (in its contemporary sense), which turns out to be less central in our taxonomy. Finally, our conclusions have implications for what we should expect of ecology, and even for how ecological research ought to be done.

2 STATUS AND AUTHORITY AMONG THE SCIENCES

In terms of status, the historical sciences, including geology, paleontology, physical anthropology, taxonomy and ecology, have often had to take a back seat to the natural sciences and especially physics. To choose a particularly egregious example of a physics-centered view of science, Ernest Rutherford once famously remarked, "[i]n science there is only physics; all the rest is stamp collecting." This is obviously

[2]Scientific theories structured in this way provide a logico-mathematical framework within which we can state observations and make inferences from them.

a case of (perhaps deliberately exaggerated) physics chauvinism. It is true that the historical sciences lack the formal unity and elegance of mathematical physics—but they make up for their looser, more eclectic conceptual structure in the breadth of their scope, the rich variety of concepts, principles and processes they invoke, and the beautiful and subtle inferences they are able to make. Further, the historical vision of the earth and of life on earth that has emerged from geology and biology since the eighteenth century constitutes as important a change in our world view (and especially our understanding of our own place in the world) as the Copernican revolution.

For a long time many philosophers of science sided with Rutherford—to the point that Karl Popper once claimed that evolutionary biology was a pseudoscience.[3] There are obvious reasons, both internal and external, for this preference. First, where science is conducted through the use of a clear set of rich and unified theoretical principles, logically-minded philosophers find a happy hunting ground. Second, there is a sense that many philosophers share, that physics gives us insight into fundamental ontological questions about the make-up of the natural world. This makes the claims and evidence of physics particularly interesting from a metaphysical point of view. Third, the great success of physics in illuminating, revealing and producing a wide range of striking phenomena makes physics extremely interesting from an epistemic point of view as well. Finally, from an external point of view, the great prestige enjoyed by physics (the 'queen of the sciences') since Newton—especially in the English-speaking world—made it a natural focus for philosophical studies of science.

More recently this imbalance in favour of physics has been set at least partly right. Since the mid-twentieth century philosophy of science has become a recognized specialization in philosophy, often pursued by scholars with both scientific and philosophical training. At the same time, the attention of many philosophers of science has turned towards a wider view of the sciences, including detailed and careful studies of the history of science, and work on a wide range of specific sciences including chemistry, biology and geology. The historical sciences are now recognized as clearly worthy of philosophical study in their own right. As a result, we are now in a position to appreciate more fully what distinguishes historical sciences from the more theoretical sciences.

Some have even tried to reverse Rutherford's invidious ranking, arguing that while physics may be able to claim authority over the most basic principles governing nature as a whole, biology (and ecology) are more *comprehensive* because they deal with a much richer variety of processes that require the *fullest* collection of natural principles to be understood. This debate, obviously enough, points towards the long and complex literature on reduction, supervenience, emergence

[3]It's particularly interesting that it is Popper, with his rigorous insistence on falsification as the touchstone of science, who took this position. I argue below that there are some important methodological differences between the historical sciences and the natural sciences, that some of these differences might be mistaken for flaws in the methodology of the historical sciences, but that the testability of claims in the historical sciences is still robust.

and related concepts; here we will treat these topics, if only in passing, from an *inferentialist* point of view, that is, by focusing on the inferences that scientists make as they work with theories, models and observations.

Concern about the comparative authority of different sciences is especially acute when tensions arise between them. One striking example is the late 19th century debate over the age of the earth. Thermodynamic calculations by William Thompson (later Lord Kelvin), assuming a solid earth with an initial temperature at the melting point of its main constituents and a gravitational theory of the sun's energy, suggested that the earth was between about 40 and 400 million years old, and that the sun could not radiate energy at its present rate for more than roughly 100 million years. Despite arguments by Perry showing that a molten interior with convection currents could allow a much greater age for the earth, Kelvin insisted on his own model (arguing that transverse earthquake waves demonstrated the earth's solidity), and later tightened the limits on the earth's age to between 10 and 40 million years, based on new data for the heat capacity and melting temperature of various kinds of rock.[4]

The tension between these arguments of Kelvin's (and Kelvin's stature as a leader in physics) and the views of geologists (even some like Croll, who had been content to live within the limits of Kelvin's earlier results) became quite sharp in the last years of the nineteenth century. These geologists were convinced by their own evidence, most dramatically in the sheer thickness of past sedimentary formations, accumulated slowly as erosion wore down previous rocks and accumulated new beds of sediment, of a much longer history for the earth. Still, there was something *elastic* about the rough measures of time the geological evidence provided. In response to Kelvin's sophisticated calculations, geologists could offer only the crudest of hour-glass equations:

$$\text{minimum time elapsed} = \frac{\text{minimum accumulation (of sedimentary rock, erosion or other)}}{\text{maximum average rate of accumulation.}}$$

Certainly judgments about the minimum total accumulation and the maximum average rate of accumulation were variable—still, that this relation between accumulations and time holds is incontestable. This connection between the geological evidence and a minimum age for the earth is extremely robust (even young-earth creationists' 'flood geology' only raises—to absurd levels—the maximum average rate of accumulation). By contrast the relation between Kelvin's calculations and the age of the earth (and the sun) depended critically on the details of Kelvin's models. Perry's model vastly extended the age of the earth by invoking convection currents to speed up the transfer of deep heat from the earth's core to the surface, maintaining a higher temperature gradient at the surface. More radically, T. C. Chamberlin suggested that atoms might be 'seats of enormous energies' [Chamberlin, p. 18], able to replenish the energy radiated by the sun—this suggestion, subsequently borne out, breaks the connection between Kelvin's evidence and the age of the sun altogether.

[4]See [Burchfield, 1975, especially chapter II].

The evidence that geologists relied on was (and is) robustly connected to the age of the earth, while the evidence Kelvin appealed to depended on his particular (solid) model of the earth and on the assumption that there was no source of energy that could replace (a significant part of) the heat radiated into space by the earth and the sun. However, the inferences connecting Kelvin's assumptions and his models to the conclusions he drew from them provided much tighter constraints on the ages of the earth and the sun than the geological evidence could give, and the theoretical sophistication of his calculations combined with the general prestige of physics added still more weight to Kelvin's argument. A simple illustration of the authority Kelvin wielded is that in his 1903 essay, 'Was the earth made for man,' Mark Twain took it for granted that Kelvin's (early) figure of approximately 100 million years was the best science could offer.

This contrast suggests that our evaluations of the status and authority of different sciences depend on a multi-dimensional comparison that is by no means easy to reduce to a one-dimensional measure, even with respect to a single question. The natural sciences, exemplified by Kelvin's calculations, provide well-tested, mathematically powerful models for a wide range of phenomena. But their application in particular cases depends crucially on whether the models applied really fit the case, and whether basic theoretical assumptions that have been successful to date can be reliably applied in contexts where basic parameter values are extreme and/or where so-far undetectable levels of violations of the assumptions would be sufficient to invalidate the model. The application of natural science models in such cases might be described as *brittle*, because it can be shattered by new evidence demanding distinct models with very different implications and by new phenomena that occur only at low frequencies or under extreme conditions.

The historical sciences are generally less vulnerable to shifts in the detailed models of various natural processes. This is partly because there are so many coherence checks that can be applied to test and confirm their conclusions, and partly because often the fine details of processes, such as the mechanics and chemistry of surfaces, weathering, frost cycles and so on, or of burial, decay, permineralization, etc., don't threaten to transform the broader observable effects of erosion or fossilization. That a river valley was excavated after a volcanic eruption is sufficiently demonstrated by noting that the valley cuts through a flow of lava from the eruption, regardless of the details of how the water flowing through the river managed to cut through the rock or whether the valley was cut by steady, gradual river flow or by one or a few massive floods; that a small, three-toed animal with some characteristics now found only among horses lived during a certain period is demonstrated by the fossil remains and the formation they were found in, regardless of the details of how the remains were preserved until the present or the precise chemical processes involved in cementing the rock it was found in. The coherence of these inferences with familiar and straightforward observations about how rivers flow and alter the landscape,[5] about how objects with the shape (and

[5]See [Twain, 1883, especially chapters VI-XIII], in which Twain describes learning to 'read the river' as a pilot.

other characteristics) of bone or shell or wood come to be and what can happen to them after an organism dies, convinced scientists that these basic inferences of geomorphology and paleontology were correct long before refined accounts of the details of these processes became available. Further, the bare possibility that these inferences might be mistaken is extremely remote: no credible alternative model will revise these conclusions, even if it substantially alters our understanding of the detailed physical and chemical processes involved. Only a thoroughly radical large-scale transformation of our understanding of the world could lead to the surrender of these basic principles of geology and paleontology.[6]

This reliance on coherence checks to ground our inferences directs those inferences towards the past, tracing backwards towards Reichenbachian 'forks': in general past events leave multiple traces of different kinds, which we can compare against each other now to test historical hypotheses. Predictive power in this context is typically limited to a kind of *retrospective* prediction—that is, traces of a process (say, the iridium-rich layer at the KT boundary construed as a trace of the impact of an asteroid or comet) allow us to predict other traces as well (such as the possibility of finding an impact crater, evidence of a massive tsunami along fossil coastlines if the point of impact was in a sea or ocean, shocked quartz crystals in the boundary layer, and evidence of widespread fires in the boundary layer). As more of these other traces were found, the impact hypothesis became practically certain. Moreover, the resulting establishment of various phenomena as reliable indicators of certain past events makes further inferences stronger: systematic study of such traces both refines our understanding of how they are produced and the special, detailed features they display, providing still more secure ways of making the case for (or against) similar events having occurred in other cases.

The prediction of a future impact is much harder; of course what has happened once may well happen again, but the evidence we would use to predict a particular impact (as opposed to merely evaluating the likelihood of such an impact occurring within some interval of time based on the historical record of impacts) has to draw on the theories and inferences of celestial mechanics to detect an asteroid or comet on an orbit that will intersect the earth's. The observation of such a body really can provide a reliable prediction, but only because we are able to exclude as highly unlikely any dramatic alteration of its expected orbit within the time frame of the prediction: the principle gravitational influences of the sun and planets are well-understood, and the probability of some other body coming close enough to substantially change its orbit is extremely low given the prevailing conditions in our solar system. Further, there is enough uncertainty about the details of these orbits that, for any timescale greater than some tens of years, the prediction of an event as precisely constrained as a collision becomes effectively impossible—so,

[6] As an illustration of just how far such an hypothesis has to go, consider *Darwinia*, by Robert Charles Wilson. (Spoiler alert!) In this imaginative story, Darwinian evolution is undermined by the sudden replacement of Europe with a new continent inhabited by forms of life utterly unlike anything else on earth. In the end, this is explained by the fact that the world is really a kind of cosmic computer program.

while we can retroactively establish that a collision occurred millions of years in the past, we have no means by which we could hope to predict a collision as little as a thousand years hence.

For the purposes of confirmation, one advantage of predictive inferences is that there is little chance that the prediction is actually *ad hoc*.[7] But a retrodictive success could just be a disguised bit of ad hocery. Nevertheless, it's not difficult to find cases where this suggestion is extremely fanciful: for example, consider fossils showing traits intermediate, in various respects, between modern humans and the great apes. The probable existence of some such fossils follows from our evolution from a common ancestor with the apes, a hypothesis which in turn has been massively confirmed by the wealth of hominin fossils discovered since Darwin first claimed that we share a recent common ancestor with the great apes. It would be silly (at best) to suggest that Darwin actually had access to such fossils and deliberately shaped his theory of evolution to ensure that it predicted such fossils are likely to be found—silly both because the fossils were unknown at the time and because there is no room or need for such adjustment of Darwin's theory.

Aside from the often dismissible risk of such *ad hoc* manoeuvres, the epistemology of the historical sciences is on a very firm footing; in fact, they are arguably better supported by their evidence than the theories of the natural sciences, since, as we've already seen, the narratives that we arrive at in the historical sciences tend to survive changes in detail that drastically alter the theoretical principles of our physics and chemistry.[8] Still, it's worth pausing here to respond to an objection that is often heard, though rarely in academic circles. The objection concerns the special role of laboratory results in the natural sciences; in particular, some creationists have argued that the absence of replicated laboratory tests undermines the evidence for the historical narratives of both evolution and geology.[9]

My response is two-pronged. The first prong points out that this is just wrong. Many laboratory tests of both evolutionary and geological ideas have been conducted. Long-term experiments with bacteria have demonstrated evolution by natural selection over thousands of generations, including the development of new metabolic capabilities.[10] Laboratory experiments have explored the properties of many kinds of rock, including details of mineral composition, structure, melting temperatures etc., as well as the processes involved in sedimentation, earthquake dynamics and many other central issues in geology. Laboratory work on multiple forms of radiological dating has confirmed the ordering of formations and forms of life that emerged from stratigraphic work beginning in the 18th century.

[7]Someone could 'gin up' a prediction of some sort of dramatic event deliberately, on the outside chance that it might come true. The familiar 'psychic's' strategy of making multiple predictions and then publicizing only the successful ones comes to mind here.

[8]This independence is symmetrical, of course: changes in accepted historical narratives can occur without requiring changes in basic physics or chemistry. But it is more striking in the other direction, because the historical narratives are ultimately grounded on processes that are described in terms of these sciences.

[9]See the list of creationist claims at `talkorigins.org`, for references.

[10]For a recent and dramatic example see [Blount *et al.*, 2008].

But this response is unconvincing for most of those who raise the objection. It's tempting, and partly right, to diagnose this response as a purely defensive refusal to understand the evidence for evolution and geology. But that isn't all there is to it. There is a real difference here between the natural sciences and the historical sciences, and, though the difference doesn't undermine the evidence for the historical sciences at all, one can see why some would be tempted to think that it does. The temptation arises from the fact that many of the central principles of the natural sciences are directly tested in laboratory: we can, after all, precisely measure many kinds of basic physical interactions and their results there. Moreover, we have also successfully applied the results to predict the behaviour of many systems, both in the lab and in nature.

There certainly are important lab results in the historical sciences—lab work on genetics and biochemistry has been central to the development and refinement of evolutionary biology since the nineteenth century, and lab work has been similarly central to many geological questions as well. However, what people tend to think of as the *main principles* of these historical sciences are broad, long-term historical claims that aren't open to direct testing, in the lab or outside of it. What underlies these challenges to historical science (though it is generally not made explicit) is the notion that the distant past is a proper subject for skeptical worries, while what happens in laboratories is not. Consequently, while the laboratory tests of various processes and principles are taken to establish those processes and principles as reliable aspects of how the natural world operates, their application to unravel the distant past is regarded as dubious at best.

This concern combines with the relatively weak predictive powers of the historical sciences: while we can predict that living things will go on changing over time, that various geological processes, including the movements of tectonic plates, slippage and occasional earthquakes on active faults and various forms of erosion and sedimentation will continue, detailed predictions of specific events (the emergence of new adaptations, or the timing and exact locations of earthquakes, eruptions, etc.) are extremely difficult to make, and appeal to longer term processes and predictions about them (such as dramatic shifts in geography over millions of years) are treated with the same skepticism as claims about the deep past, despite their elegant fit with so much current evidence. The upshot is that laboratory science, celestial mechanics, and immediately testable claims are seen as far better supported by the evidence than any science whose principle claims concern the course of events in deep time could be.

The second prong of my response addresses this challenge directly, asking what justifies this special skepticism about distant periods of time and the long integration that accumulates small changes of the kinds observed over short periods into the dramatic changes that make up the history of our planet. As Charles Lyell argued in his *Principles of Geology*, 'drafts on the bank of time' are far less troublesome, when it comes to understanding their significance and implications, than the invocation of processes (whether natural or not) that can't be observed in detail now because they no longer occur.

Any narrative about the past (or anticipation of the future) is justified, if at all, by its coherence with the various traces we can find now together with our understanding of the processes linking those traces to the various events they can tell us about. As we've already noted here, the historical sciences have produced many remarkably rich, detailed and coherent narratives describing the development of the universe, the solar system, the earth and life on earth. The fact that other narratives can be imagined (including fanciful 'false-past' narratives) is no more evidence for skepticism about the past than the fact that other courses of events (in and outside of laboratories) can be imagined is evidence for skepticism about the present workings of the natural world. The many observational tests these narratives have passed, as the present traces of the processes they invoke were tracked down and documented, make them convincing parts of the natural history of our world.

Ecology shares many of these characteristics of historical sciences. It has been subject to criticism from partisans of the 'hard' sciences. Its subject matter involves rich and complexly connected processes, and predictions in ecology are well-known to be difficult at best. Finally, the retroactive construction of explanatory narratives plays a central role in ecological investigation. So, in both its epistemology and its methodology, ecology groups naturally with the historical sciences.

3 AIMS OF EXPLANATION

Another important contrast between the historical and natural sciences is the focus of the historical sciences on particular applications as opposed to general principles. Of course both principles and applications are part of every science. But in the natural sciences particular applications are typically concerned with how to account for some phenomena using specific theoretical principles—it is generally presumed that the features to be modeled can be expressed within the language of that theory, and that the principles of the theory should provide all the necessary constraints to make the model 'work' (if it emerges that they don't, the theory is in trouble). Further, we expect that the phenomena will be *the same* in any similar case: a successful account of the phenomena will be, in that sense, entirely general.

By contrast, applications in the historical sciences are not theoretically *pure*, often involving rich interactions between processes that are described in terms of different collections of basic principles. Further, they are not treated as *closed*; we expect the historical processes of geological and biological change to be interrupted, altered, even transformed by outside influences such as the eruption of a volcano many kilometres away, a planet-wide climatological shift that gives rise to an ice age, or the sudden impact of an extra-terrestrial body. The task is to unravel a particular sequence of events, not to identify a type of process that will be regularly repeated in every similar case. Differences over preferred models or conceptual outlooks often have more to do with preferences concerning starting points and which factors are treated as intrinsic features of the models and which as exogenous influences altering the course of events. See, for example, C. Eliot's discussion of Clement and Gleason's views of ecological succession, in this volume.

Of course the processes modeled by the natural sciences are also, as concrete individual processes, subject to such external influences. But the aims of the natural sciences don't include a systematic account of such external influences and their roles in the development of particular systems—when our interests do turn to the particular, we will certainly seek out the particular circumstances that explain what happens in an unusual case, but in the natural sciences our interests are not typically focused on the particular. We notice and try to explain unusual cases precisely because we see them as exceptions to the rule, and consequently important tests of it: the usefulness of the basic principles is first illustrated in relatively pure cases, but ultimately every case has to be reconciled with the principles.

Consider the erosion of a particular geological formation. In general, many kinds of processes will be involved, from small-scale mechanical and chemical goings-on to large-scale meteorological and climatological phenomena. The results will have effects on water conditions in the watershed, the soils in the region, and on plants and animals; in turn, plants and soils will alter erosional processes. More significantly, the processes will be *local* and *contingent*: the results will depend on the detailed history of that particular formation and the broader context (climatological, geological and biological) in which its erosion took place. Small features, such as the location and orientation of cracks in the bedrock, can have substantial influence on the direction of water flows and the subsequent development of a drainage system. Not only do the details matter here, but also the course of events outside the region, which often intrude on and alter the processes under study. Any explanation of the erosion of this formation will draw on many contingencies, both in the detailed interaction and feedback processes influencing events, and in the impacts of external events. As a result, detailed and reliable predictions are difficult, if not impossible.

This is partly a matter of complexity. Complex feedback interactions can occur in any science, but in the historical sciences as well as in biology and especially ecology, they are inevitable: even very simple population models can generate chaotic behaviour.[11] The upshot is that, even if we begin with very similar circumstances, the results we obtain from our models, and very probably the outcomes in the natural world, can vary widely. But it is also a matter of the focus on the particular: rather than begin with a universal course of events that will be characteristic of such situations in general though it may, in particular cases, be interrupted or altered, we aim to identify the particular sequence of events that has produced the erosional effects observable in this particular case. Finally, it is also due to

[11] For example, consider the simple logistic equation, $P_{n+1} = r(P_{\max} - P_n)$, where $P_{\max}$ represents the maximum population that can be sustained, r the rate of growth and P_n an initial population. If we normalize the population measure by setting $P_{\max}$ to 1, the result is chaotic for r greater than about 3.57.

the general openness of such systems, and the infrequency of useful explanatory narratives driven by a single isolated process that we can characterize by appeal to theoretically pure principles.

4 THE ROLE OF PREDICTION

Many events and processes studied by the natural sciences allow for powerful models that produce highly constraining, reliable predictions; in fact, the predictive success of such models is often taken as a paradigmatic example of (and a principal type of evidence for) good work in those sciences. This is closely related to the role of general principles used to produce the predictions: in the natural sciences, these principles are believed, at least in many cases, to specify all the relevant quantities and how they influence the system: they aim to be closed, in this sense. Even if the systems being described are not strictly closed and may, in some cases, be disrupted by external interventions, such external disruptions don't undermine the model, although predictive failure in the absence of external disruption does.

This sort of closure leaves aside the challenge of determining the right values to assign to the relevant quantities, as well as the often very difficult problem of calculating or inferring the consequences of such a set of conditions for a system. It also sets aside the fact that the systems we are describing are vulnerable to external influences that we may not be able to anticipate even when we take those external influences to be subject to and describable in terms of the same fundamental principles and quantities.

Nevertheless, the natural sciences do manage to assign values and perform reliable, detailed calculations predicting the behaviour of some important real systems.[12] This success in isolating[13] the course of certain kinds of processes underwrites some of the more metaphysical elements in scientific thought—what we see here is the 'natural' development of such systems in the absence of external interference (though even on a billiard table, a standard illustration of basic mechanical processes, the subtle effects of gravity together with rapid amplification of deviations in the motions of the colliding balls ensure that after more than a few collisions the state of the table is dramatically different from what it would be without the minuscule gravitational influence of the moon). Still, while a closed, predictive account of events on the table is not entirely possible, the external influences that affect it can be expressed, in principle, in terms provided by our theory

[12]See [Cartwright, 1999], *The Dappled World*, for some limits and interesting comments on this issue—the result is often that what we model is the behaviour of very special systems which are developed by experimentalists precisely to isolate/demonstrate certain basic processes and features. See also [Sellars, 1963] "Scientific Realism or Irenic Instrumentalism" for remarks on the role of metaphors as involving second-order similarities in science.

[13]Nature must be 'put to the question', Bacon infamously suggested, in order to achieve this kind of isolation—I see this remark more charitably than some, as drawing a contrast between Bacon and those who, like Descartes (cf. *Principles*) held that science should begin with familiar phenomena in natural contexts, which, Descartes claimed (erroneously, on the evidence) would display the simplest combinations of 'natures'.

of mechanics, and they are small enough not to matter for simple cases involving just a few collisions. This encourages the hope that in principle, a full calculation of all such influences would allow a perfect modeling of the system, a hope beautifully expressed by Galileo when he declared that imperfect results could be obtained from a correct mathematical model only because of the imprecision of our measurements and our own failure to account for all influences on the system:

> Just as the accountant who wants his calculations to deal with sugar, silk and wool must subtract the boxes, bales and other packings, so the mathematical physicist, when he wants to recognize in the concrete the effects which he has proved in the abstract, must deduct the material hindrances; and if he is able to do that, I assure you that matters are in no less agreement than for arithmetical computations.... The sources of error, then, lie not in abstractness or concreteness, not in geometry or physics, but in a calculator who does not know how to make a true accounting. (Galileo, *Dialogue of Two World-Systems*, cited in [Drake, 1970, pp. 68–69].)

Certainly our efforts to apply ever-higher levels of precision in measurement and accounting for small influences on mechanical systems have been well-rewarded, from lens grinding to using pendulums to measure the gravitational attraction of the earth to Hamilton's chronometer and on to today's efforts to detect gravitational waves or the Higgs boson.

5 A CASE IN POINT

Consider the famous debate over the relation between the extinct Dodo and the vanishing Tambalacoque tree, *Sideroxylon grandiflorum* (formerly *Calvaria major*). In a very influential paper, Stanley A. Temple [1977] proposed that the unusually heavy seeds of the Tambalacoque could not germinate unless they had been abraded by passing through a Dodo's gizzard. Temple's argument drew on the apparent absence of young Tambalacoque trees in Mauritius' forests as well as an experiment in which Temple fed fresh Tambalacoque fruit to turkeys (a somewhat smaller bird than the Dodo): three of ten seeds that were either regurgitated or passed whole through the turkeys' digestive tracts did germinate (though seven were crushed). But in a vigorous critique of Temple's work, Mark Witmer and Anthony Cheke [1991] drew on a richer range of evidence to argue that Temple's hypothesis of an obligate mutualism could not be right. Though it's clear that germination is rare in the field, Tambalacoque seeds have been reported to germinate without such treatment and unpublished trials showed no difference in germination rates between abraded and unabraded seeds; though Tambalacoque trees younger than the extinction of the Dodo are rare, they are not unknown; the hard endocarps of Tambalacoque seeds have a natural line of weakness along which the endocarp can split, allowing the seed to germinate (a characteristic Tambalacoque seeds share with *Canarium paniculatum*). Witmer and Cheke's investigations suggest that Tambalacoque seeds are very susceptible to fungal infections, and may (like those of many other tropical trees) need to be cleaned of pulp before the fruit begins to rot, in order to germinate. But other lost or reduced fauna of Mauritius, including an extinct parrot and two species of tortoise, may have eaten and dispersed the seeds. Further losses to introduced species at the vulnerable seedling and sapling stages, along with habitat degradation and competition from newly introduced tree species may also have contributed to the Tambalacoque's decline.

The evidence for this richer account of one ecological change on Maritius closely parallels the kinds of evidence that ground narratives in the historical sciences. The reasoning is clearly retrospective: the decline of Tambalacoque trees is well known; the process(es) that have led to that decline are what is in question. Certain kinds of processes—seeds' failure to germinate through disease or the absence of some helpful factor previously present, seedlings' and/or saplings' failure to survive, are known to be potential factors in such a decline; various tests and signs indicating the importance (or lack of importance) of these factors are explored. An account is supported by our evidence when it coherently connects the results of such studies into an account of the trees' decline; the more fully an account fits details in our evidence, integrating what we can discover about changes in population structure over time and how various processes can affect germination and survival of young trees over the last 300 years, the more satisfactory our explanation of this ecological change. One natural way to construe the reasoning involved is in terms of *eliminative induction*:[14] we accept a particular explanatory narrative when (and only when) the initial constraints on credible types of explanation and the accumulated evidence rule out other narratives; in general, such acceptance leaves open only the possibility that the problem was mis-posed from the outset.

Clearly enough, the result in this case will be a retrospective explanation for what has happened; there is little reason to expect that a prediction of the decline (or of other, similar declines) would be practically possible: much of the evidence used in testing and confirming our explanation wouldn't be available in advance. Moreover, ecologists would have little reason to pursue the evidence that might be available prior to the trees' decline. Many different kinds of ecological changes occur when an island is invaded by so many foreign species and subjected to new forms of agricultural exploitation. Attempting to anticipate them in advance would require extremely high levels of initial information, including detailed and complete models of complex ecological interactions and a rich variety of precise data to apply those models predictively. Finally, the bearing of that evidence would be hard to sort out in advance given the complexity of the web of interactions affecting the survival of various kinds of trees in such circumstances.

[14]See [Norton, 1993; 2003].

6 REDUCTION AND SUPERVENIENCE

The metaphysical relations between high-level sciences dealing with complex objects and processes and more 'fundamental' sciences have long been an important topic for philosophers of science. The importance of unraveling the relations between the different sciences in order to clarify the sense in which they may be said to collectively represent our best effort at describing the world and at identifying the best methods for producing such descriptions, make this topic worth addressing briefly here. But rather than focus attention on the metaphysical questions, my chief concern will be with the practical constraints that make an account of the world in terms of physics 'all the way up' impossible, and what prospects there are (in absence of this ideally completed unity) for giving substantive expression to the idea of the unity of science.

Ontologically, it often seems intuitively appealing (given the course that our scientific inquiries have taken) to regard the objects of theories applying on larger scales as composed of (and so at least *ontologically* reducible to) the objects of our microtheories. But giving in to the temptations of this metaphysical intuition is light work compared to the hard slogging involved in translating assertions expressed in terms of macrotheory vocabulary into the vocabulary of an ontologically-preferred microtheory. A full theoretical reduction (*cf.* [Bonevac, 1982]), in which the *inferences* made within the reduced theory are captured as special instances of inferences within the reducing theory demands more still. So part of the challenge here is to sort out the different ways in which we might seek to 'reduce' one theory to another.

Wilfrid Sellars distinguishes these types of reduction in "Philosophy and the Scientific Image of Man" ([Sellars 1956], reprinted in [Sellars 1963]), when he comments on the unity of the scientific image: "There is relatively little difficulty in telescoping *some* of the 'partial' images into one image...we can unify the biochemical and the physical images; for to do so requires only an appreciation of the sense in which the objects of biochemical discourse can be equated with complex patterns of the objects of theoretical physics. To make this equation, of course, is not to equate the sciences, for as sciences they have different procedures and connect their theoretical entities via different instruments to intersubjectively accessible features of the manifest world." In this passage two different kinds of reduction are contemplated—equation of the ontology of two sciences, and equation "of the sciences". Sellars elaborates, "[f]or to make this identification is simply to say that the *two* theoretical structures, each with its own connection to the perceptible world, could be replaced by *one* theoretical framework connected *at two levels of complexity* via different instruments and procedures to the world as perceived" [1963, p. 21]. The equation of the two sciences occurs at the level of *vocabulary*, as a function of the 'telescoping' relation: given this replacement, the *reports* we make and conclusions we draw as biochemists come to employ a vocabulary that is based on the vocabulary of theoretical physics.

Sellars further distinguishes this unification of the entities and vocabularies from unification of the *theoretical principles* of the two sciences:

> ...while to say that biochemical substances are complexes of physical particles is in an important sense to imply that the laws obeyed by biochemical substances are 'special cases' of the laws obeyed by physical particles, there is a real danger that the sense in which this is so may be misunderstood. Obviously a specific pattern of physical particles cannot obey different laws in biochemistry than it does in physics. It may, however, be the case that the behaviour of very complex patterns of physical particles is related in no simple way to the behaviour of less complex patterns...There is, consequently, an ambiguity in the statement: The laws of biochemistry are 'special cases' of the laws of physics. It may mean: (a) biochemistry needs no variables which cannot be defined in terms of the variables of atomic physics; (b) the laws relative to certain complex patterns of sub-atomic particles, the counterparts of biochemical compounds, are related in a simple way to laws pertaining to less complex patterns. [p. 21]

Inferential unification is very different from the telescoping unification that arises just from applying the same *language* (*i.e.*, the vocabulary) to report observations and inferences. An inferential reduction would require the basic *inferences* of particle physics to generate the inferences of biochemistry as well; not only the entities and vocabulary, but the science (and language) of biochemistry would then be fully unified with (*i.e.*, reduced to) particle physics.[15]

Distinguishing these different aspects of reduction is particularly helpful because it focuses attention on two separate elements in our use of scientific theories: first, the application of a theory to the world, both in observation, when we respond to situations in the world with assertions in the language of the theory, and in practice, when we use assertions expressed in the language of the theory to guide practical activity, and second, the theoretical inferences that make some assertions in the language follow from others, which provide opportunities to test the theory's ability to coherently represent some situations in the world.

Together, these aspects of reduction engage with the three main elements of Sellars' inferential view of language: norms governing language-entry, language-

[15]The distinction between vocabulary and language drawn here draws on Sellars' ideas about *material inference*; when vocabulary but not language has been 'telescoped' to unite two sciences, the same vocabulary is governed by two systems of material inference rules, one capturing the inferences of each science. The system of the reducing science will be tightly tied to the basic vocabulary, while the other system independently adds inferences applying to the complex systems of basic objects which are described by the reduced science. But with full unification of the language, we will have only one system of material inferences; the inferences of the reduced system will then be understood in terms of certain (generally complex) inferences in the reducing science. It's worth noting as a caveat here that this story does not yet deal with the subtle interactions that arise when not only reduction but also *correction* comes in: the inferences of the reduced science may well be highly reliable even though they are not precisely correct, from the point of view of an acceptable reduction—and they may depend on circumstances that had not yet been identified in the reduced science as required for its success. Consider as an example here the relation between classical and statistical thermodynamics.

language and language-exit 'moves'. Being able to express distinct theories (considered as systems of inference) in a common vocabulary is certainly an advance towards a unified scientific view of things, but the further step of inferential reduction is by far the hardest. Once it is achieved, the reduced science's inferences now appear as a consequence of the reducing science's principles, rather than merely being expressible in the language of the reducing science. Only this much stronger sort of reduction could satisfy those who follow in Rutherford's rhetorical footsteps, holding that only physics makes real contributions to the principles in whose terms we understand our world.

Sellars' account points towards the important practical limits of reductive efforts. We strongly suspect that there will often be no 'simple relations' between laws governing basic entities and laws governing various complex systems of basic entities. This suspicion is particularly well-founded when it comes to organisms and ecologies. Even at the large-scale level, when we try to model interactions between predators, food supplies, population density and disease to capture how a population changes over time, the resulting models are extremely sensitive to the details of these factors. Absent some radical breakthrough in our inferential capacities, it's obvious that there is no serious prospect of an inferential reduction of any of the other sciences, including ecology, to physics.

This observation underscores the independence of explanations in the historical sciences from changes in the fundamental principles of the natural sciences. Many of the inferential links that unify the narratives of the historical sciences have been formed independently of these fundamental principles, and they connect the observations of the historical sciences in ways we cannot replicate using only these fundamental principles: our understanding of the different processes involved in producing the phenomena of the historical sciences, and of their signs and symptoms, developed in a very empirical way from studies of these phenomena. However, this is far from saying that physics does not constrain the processes that the historical sciences describe and explain—it is a very modest sort of emergence that we are discussing. Further, at the micro-level, physics and chemistry do illuminate processes like erosion, fossilization and glaciation, while at the macro-level principles like conservation of mass and energy and the laws of thermodynamics often provide important constraints on our models. For illustrations from ecology, consider the importance of isotope-based measurements in efforts to track ecosystem productivity in the past, and models that trace flows of energy and materials. These refinements, drawing on the principles and applications of the natural sciences, have greatly extended our ability to measure these processes and learn from the traces and patterns they produce.

Here we find a practical kind of unity in the sciences. While the narratives of the historical sciences are well-grounded in their own evidence and the understanding of a wide range of various kinds of complex processes that has emerged from that evidence, they have also been substantially refined and extended, and (not coincidentally) more stringently tested by the application of principles and observational methods drawn from the natural sciences. The confirmation of the established geological column in the light of a multiplicity of radiological dating techniques illustrates both the advances that can be made with the help of observational techniques grounded in the natural sciences and the reliability of well-grounded results in the historical sciences; by way of contrast, the history of continental drift and plate tectonics illustrates a much more complex interaction, in which geophysics at first motivated wide resistance to Wegener's ideas, but later physical measurements (including magnetic surveys around mid-ocean ridges) confirmed the reality of plate motions, establishing a mechanism for 'drift' that escaped the traditional objections. Of course plate motions have since been richly integrated into geological narratives, including much of the evidence Wegener first identified as well as immense amounts of data from subsequent studies of the histories of continents and ocean basins.

7 MODELS AND EVIDENCE

The need for each science to articulate and apply its own inferential structure brings us to a discussion of models. The importance of models as intermediaries between theories and the phenomena we apply them to has been a hot topic in recent philosophy of science. A number of authors, including Nancy Cartwright [1983; 1999], Margaret Morrison [1999] and Naomi Oreskes [2003] have argued that science requires some such intermediary—the logical notion of a theory, *i.e.*, a set of sentences closed under the consequence relation, can't carry the load in practice, both because it does nothing to indicate how the theory is to be applied to the world,[16] and because, even assuming that we know what actual phenomena we want to apply the theory to and how to connect observations of those phenomena to assertions in the language of the theory, providing a full description of the phenomena and determining the implications of that description according to the theory are, in general, beyond us. Simplified descriptions, approximations and selective inferences are inevitable elements in the actual account of the world that our theories inform. *Models* are supposed to embody these descriptions, approximations and inferences.

The notion of a *model* here is intuitively straightforward, though there is room for many subtleties. First, models in this sense are not models in the sense of semantic theory, because they generally involve approximations and simplifications that, strictly speaking, are incompatible with the truth of the theory. Instead, they are attempts to capture or express, usually approximately, some of the *implications* of a theory (or theories) for a particular system or type of system. Models in this sense include familiar models of molecules that use sticks and springs joining different coloured wooden balls to represent some aspects of molecular structure, sophisticated computer programs that attempt to capture long-term change in the earth's climate, attempts to calculate the running speed of a *T. rex* based on a model describing the mass, bone structure and muscle strength of the fearsome

[16] See [Brown, 2004].

predator, and Kepler's elliptical model of the orbit of Mars.

Many questions come up here. Just how seriously should we take these models? Are they just a pragmatic element in our scientific practice, serving to link abstract theory to concrete applications, or do they play a substantive role in the content of our scientific accounts of the phenomena we apply them to? How much variation is there in the purposes we use them for? Do different models play different roles in our representation of the world? What sort of evidence do we need to justify adopting a model for the different uses that we put models to? But our concern here is with the special features and challenges of models in ecology, and their relations to features and challenges of models in the historical sciences.

Ecological models are generally divided into three basic types: population models, community models and ecosystem models. Population models focus on capturing changes in the numbers of individual species over time; they vary in the number of factors that they include, whether the model is deterministic (suited only to large populations where chance fluctuations are small enough to be ignored) or stochastic, and whether the model includes any representation of the different properties of individual members of the population (including, for example, a range of values for fecundity, ages of individuals, and links between these and other factors including probability of death within some time period). Community models treat populations of more than one species, including interactions between them (for example, predator-prey relations). Finally, ecosystem models extend (very ambitiously) to the flows of energy and material that link communities to the surrounding, non-living environment; recently, computational models using geographical information systems that provide a representation of the spatial distribution of conditions in the environment have emerged as an important new class of ecological models [Sarkar, 2005].

Even at the level of population models, substantial difficulties arise. Very small differences in input (or *boundary*) conditions can have large effects on the model's results, as can small differences between models.

Four basic challenges that these models face are the challenge of data, the challenge of model complexity, the challenge of natural complexity and the challenge of openness. For the first, we can't expect to have precise and accurate data on an actual population's numbers, the resources the population depends on, the threats its members face or the distribution of each of these in a particular region. Consequently, assigning values to the parameters of our models involves substantial uncertainty. As to the second, the mathematical analysis of ecological models is extremely difficult. They are often exquisitely sensitive both to details of boundary conditions and to the precise structure of the models. Consequently, the uncertainties arising from the first challenge are, at least in general, important to our ability to rely on the models' predictions. Third, our models don't generally include parameters and interactions rich and detailed enough to provide a true picture of all the elements involved in the development of a population, community or, still less, an ecosystem over time. Fourth and finally, the systems we apply these models to are subject to perturbation by external causes, *i.e.*, causes not included in the model.

These challenges are mutually reinforcing: small differences of structure or input can have substantial impacts on the conclusions we would draw from a model, and increasing the detailed structure of the model to provide an intuitively more realistic account of the actual phenomena only adds to the mathematical complexity of the model and to the difficulty of assigning values to an increased number of parameters based on good empirical evidence while uncertainties about interference due to causes not included in a model weaken our ability to evaluate models, and the complexities entailed by attempts to include them limit how far we can go in modeling even known causal factors while threatening (on the other hand) to allow so much unconstrained flexibility as to render 'agreement with the data' an all-too-easy hoop to jump. The result, when these challenges are summed up, is a high level of uncertainty regarding the relation between the development of parameters over time in ecological models and the actual course of real populations, communities and ecosystems.

When we consider these challenges, pessimism about ecological models seems unavoidable. But this pessimism is only justified to the extent that these challenges make *success* for ecological models unlikely—and we can't settle that issue until we've sorted out just what we want these models to do. Any account of *success* for a model will have to begin with what we actually expect it to do. At the empirical and methodological levels, sometimes we expect models to predict future observations, but sometimes we aim at subtler ends. These might include questions about the models themselves, such as identifying constraints on the states the modeled system can achieve or demonstrating the need for further elements in a modified model if the model is to produce reasonable results, and questions about the theory the model draws on for its inferential structure. These kinds of information can serve as premises in important scientific arguments even without substantial predictive success.

One standard view of the logical status of models [Kyburg, 1983; Oreskes, 2003] treats them as contributing certain *conditional premises* to our account of some phenomena; an argument drawing on such a model is broadly a *modus ponendo ponens* inference, with the model telling us that if certain boundary conditions hold, then certain results will follow. This fits nicely with an inferentialist view of models: the key point is that a model allows us to infer from certain premises to certain conclusions. On this account, models are *inferential machinery.* But the uses we put these inferences to vary widely, and we normally draw a pragmatic distinction between inferences that are considered reliable and inferences that, while endorsed by the model, are not regarded as reliable.

8 THE STANDING OF ECOLOGICAL MODELS

Ecological models have come in for some pretty vigorous criticism. For reasons of space we'll focus here on one critique, due to Naomi Oreskes [2003]. Oreskes' main concerns about our attempts to model complex natural systems are straightforward: she argues that uncertainties in such models are inevitable, as they cannot capture the full complexity of the systems they model. Further to this, she argues that efforts to enrich our models and make them more realistic make testing them harder, as the resulting complexity allows increasing room to adjust the model and input parameters to 'fit' our observations, casting doubt on the value of successful predictions as confirmation of the model. Oreskes concludes that it is a mistake to expect models to provide deterministic predictions of outcomes in the natural systems they are meant to represent, although they can illuminate by serving as 'what if' scenarios, to illustrate possible best and worst-cases, and suggest possible outcomes of different sorts of interventions.

This conclusion is well-taken, though I think it's important to add that it reflects a tension between reasonable scientific aims and the practical, public-policy aims which predominate in the examples Oreskes treats. As we've noted, models are not always used to produce predictions, and not all predictions that models apparently give rise to are regarded as *significant* (*i.e.*, sufficiently reliable to guide practical reasoning).

One alternative use of models begins with a deliberately crude model to launch a process of refinement and correction leading to a modified model that we do regard as predictive in at least some respects (see the discussions of pendulum models in [Morgan and Morrison, 1999]). The inferential process in such cases begins with the crude initial model, but then reflects on the model's limitations, contrasted with a more detailed theoretical understanding of the actual processes involved, and introduces step-by-step modifications meant to improve on the initial model. The result may be a model that really is taken to be predictively reliable or even an approximately accurate representation of the real system (or just guidance for building a better clock). However, the inferential process that leads to that final model depends on the crude initial model as well as the subsequent critique and refinement. Further, even the final model may not be used predictively or regarded as a realistic representation—it may still be aimed at identifying constraints on the system modeled, or at serving as a test bed for still more refinement. It may also be used to explore aspects of the theory the model is based on, as in the paradoxes of mechanics that result from Laraudagoitia's Zeno-style puzzles (see [Earman and Norton, 1998]).

A simple but ecologically interesting example is the exponential growth model of a population with unlimited resources. Such models are predictively useless for established wild populations, though they can be predictively successful over a limited number of generations in specific circumstances, such as the introduction of yeast into a fresh barrel of wort. But for Darwin (and later Wallace) such models led to an obvious Malthusian conclusion: in the long run, most organisms cannot survive and reproduce successfully. Here the predictive failure of the model provides a key premise in a convincing argument for an important conclusion: populations of organisms often undergo substantial selective pressure. This important conclusion can be reached without a predictively successful model of any natural population; it requires only the failure of simple models of populations that include no selective pressure, in the light of the exponential reproductive potential shared by all organisms.

Our focus here, however, is on using models to capture certain implications of a theory for some actual phenomena. The idea is that the model should capture, at least approximately, some of the inferences that a good theory would license. This is separate from whether we believe the theory, and from whether we want the model to 'correspond' in some sense to the natural system responsible for the phenomena. It does suppose that there is a theory in the background here, which is not always the case except in a trivial sense of 'theory'. More generally, we may have only the model along with some rough ideas about important features of the system to be modeled—but even then we can then use the model to produce inferences about a particular system or collection of systems, and reflectively evaluate the model in the light of the resulting inferences and how they work out.

As an example of this, consider a purely phenomenological 'model', meant only to capture dynamic relations between certain observable parameters. In principle such models can be very successful. For example, a model of the stock market would be a brilliant (and extremely profitable) success if it were merely predictively successful, regardless of whether it captured any of the 'real' economic dynamics underlying the changes in stock prices it had predicted. (See Vonnegut, *The Sirens of Titan* for an amusing but silly example; more serious examples include technical models of the stock market based on observed cyclical patterns of market changes.)

Nevertheless, in many cases we aim to produce models that do represent, if imperfectly, the systems we apply them to, and sometimes we actually think we've succeeded—further, we can have reasons to suppose that we have succeeded at this aim even while successful prediction remains elusive. After all, as Yogi Berra famously noted, "[i]t's tough to make predictions, especially about the future." This is, in effect, the flip side of the concerns about complexity and sensitive dependence raised above, since they imply that (at least with respect to some features of the phenomena) a model of the phenomena can be extremely realistic and still fail to make successful predictions.

Kyburg [1983] argues that predictive success need not be essential to a model's success (*i.e.*, to fulfilling our intentions for the model): we may aim to explore the implications of certain constraints on a system even though we recognize those constraints may not in fact obtain, and that there are other constraints ensuring that the system will not, in fact, develop as 'predicted' by the model. From Kyburg's perspective the Club of Rome model, widely disparaged as an example of a failed model, is of considerably more interest than that characterization suggests: given that there are natural limits of the kind that the model proposed on certain resources, whether the actual limits assumed for the purposes of the model are accurate, the *kinds* of constraints that the model predicted on the sustainability of economic growth (though not their timing) remain significant. Further, despite the model's failure to consider political and social factors that would certainly become significant as resource shortages begin to affect the economy, the constraints it does include are worth exploring on their own.

In another striking example, climatological modeling does not focus on a single *simulation*, *i.e.*, a global circulation model (GCM) together with a set of initial conditions, but an *initial condition* ensemble, which involves a single GCM together with a wide range of initial conditions. Beyond this, climatologists also employ *multi-model* ensembles combining a number of initial condition ensembles based on different GCMs, which turn out to give a better match to climatological observations than initial condition ensembles do. A *climate* for such a model or model ensemble is defined as the features of a simulation that are constant across the ensemble—typically, averages of certain quantities and other statistical measures, along with relations between certain variables. Model *weather*, on the other hand, is those features of a simulation that differ across members of the ensemble. Given the chaotic behaviour of real weather and the obvious limits on data for setting initial conditions and the models as representations of a much more complex real system, model weather has next to no predictive value (weather prediction efforts are guided by more detailed local models). But model climates do have substantial predictive skill, as retrodictive testing shows.

Moreover, when physical details can be successfully added, producing a more constrained model that has improved skill on such measures, we may have reason to hold that the resulting sequence of models demonstrates the characteristics explored by Morgan and Morrison in the case of the pendulum: we are refining the model in the light of a physical understanding of the processes that are actually occurring, and thereby improving its reliability and applicability. This progression may never produce an entirely realistic model, but it can improve and extend the inferences we can make with the model's help.

It's also important to point out that the inferences we regard as supported by a model do not depend just on the 'if-thens' we can extract by using the model. When a real system is known to be predictively intractable in some respects, we may regard a particular model—for example, a fluid mechanics-based model of a bill's trajectory across a public square—as a realistic depiction of the *kinds* of causes at work in a phenomenon despite the failure of such a model to predict certain observations, such as the path of the bill.[17] There can be different kinds of predictions at stake here—fluid mechanics predicts the very unpredictability of

[17]See [Cartwright, 1999, 27f. Of course this makes me and those who agree with me here 'fundamentalists' in Cartwright's sense [Cartwright, 1999, pp. 24–28]. But is such fundamentalism as unreasonable as Cartwright maintains? My view is that, when a theory provides detailed predictive successes in contexts where, by its own lights, such successes are to be expected, and the natural parameters of some other cases where it *predicts* predictive failure still fall within the range where, when prediction is reasonably expected according to the theory, the theory has been shown to be successful, then it's reasonable, absent further contrary evidence, to hold that the theory offers an acceptable account of what's going on in the predictively intractable case. This point is closely related to another fascinating issue, *viz.* the distinction between the inductive skepticism characteristic of science and Humean skepticism. It's clear that even Cartwright rejects Humean skepticism, accepting as she does that many models provide very reliable predictions of the behaviour of certain kinds of systems. By Humean standards, such predictions are just as questionable as the 'fundamentalist' notion that classical fluid mechanics is a sound basis for understanding in outline, though not predicting in detail, the motion of a bill blowing across a square.

such trajectories, a prediction that, so far, is borne out. If classical mechanics made similar predictions of unpredictability about planetary orbits in our solar system, the mere success of positional planetary astronomy would be a convincing counterexample to classical mechanics. Similarly, any model of turbulence that managed to reliably predict such a bill's trajectory, even a successful phenomenological model, would count as powerful evidence against classical fluid mechanics.

Of course when a theory like fluid mechanics takes such a risk and the risk (so far) pans out, the result is (at least) weakly confirming for the theory. On the other hand, current models of turbulent flow also have shared features that can be tested against real turbulent flows.[18] Once again, claims of predictive failure need to distinguish predictions that have failed from those (perhaps subtler or more general) that have succeeded; they must also distinguish those 'predictions' (*i.e.*, claims inferred from a particular model) that are *robust*, *i.e.*, likely to hold if the model is indeed meeting our expectations, from those that are *frail*, *i.e.*, unlikely to hold in the system being modeled even if the model is as accurate and realistic as we can reasonably expect it to be.

A related point arises in a rarely considered argument for modest realism about our cognitive commitments to science. The argument begins with the combination of *confidence* that scientists often express regarding applying a hitherto successful theory or model under the kinds of conditions it has succeeded in and their *reluctance* to put faith in its success under other kinds of conditions. The empirical parameters involved in distinguishing these types of conditions, such as spatial and temporal scales, velocities, temperatures, intensity of gravitational fields, etc., are believed to be (and have indeed turned out to be) predictive of when our theories will and won't fail at various tasks. This both involves and, as I see it, justifies a modest realism about the *significance* of the quantities measured by these parameters (and the types of circumstances distinguished) to how various processes proceed. While scientists are often skeptical about the *truth* of currently successful scientific theories, which would entail their reliability on any scale at all, they are often confident about their applications within the range of parameter values where they have been successful, as well as about the criteria by which we distinguish those established applications from relevantly new applications. Obviously enough, a Humean skeptic would not share this confidence.

An example from ecology is worth examining here, to see how models can be used in non-predictive ways. In this case, the subject is the ecology of cane toads (*Bufo marinus*). In [Lampo and Leo, 1998], a model of the cane toad population is used to help determine what explains the extremely high population densities of cane toads in Australia, where they are an invasive species, compared to their population densities in native habitats. The model is a simple time-based matrix model, distinguishing juvenile from adult (reproductive) stages and including parameters for fecundity, for successful transition from the juvenile to adult stages

[18] Consider the evidence for Kolmogorov's energy spectrum function (see [Frisch, 1995]) and the evidence against the scale-independence of turbulence in the inertial regime that undermines Kolmogorov's account [Mathieu and Scott, 2000].

and for year-to-year survival of adults. This model is extremely schematic—it does not provide a continuous account of the population's size from day to day, distinguish any specific causes of death, include models of resources that cane toads depend on, or their interactions with each other or with other species. But it embodies some obvious constraints all the same: clearly, the reproductive capacity of a species that breeds seasonally turns, in part, on the adult (reproductively active) population at the beginning of the breeding season, and in part on the fecundity of those adults; equally clearly, the size of that population depends on recruitment of new adults and the survival of animals that are already adult. The authors remark, in their summary, "[t]he model presented here is by no means a predictive, but rather an analytical tool" [Lampo and Leo, 1998, p. 395]. Field data provide evidence constraining fecundity and both recruitment and survival. Analysis of the model shows that, at high population densities, equilibrium densities were much more sensitive to adult survival rates than to variations in the other parameters; field data also show that adult survival is indeed much higher in Australia than in South America. The authors conclude, "... there is a general consistency between predicted and observed patterns. Thus we believe our study brings important insights on factors driving the enormous reproductive success of Australian toads and on strategies to control their rate of increase" [Lampo and Leo, 1998, p. 395]. Here we see a clear example of a non-predictive but still cognitive use for ecological models, as well as a significant role for a background, practical aim. Further proposals regarding the usefulness of false models in biology have been made by Wimsatt [1987, p. 28], who suggests that false models "can (1) lead to the detection and estimation of other relevant variables, (2) help to answer questions about more realistic models, (3) lead us to consider other models as ways of asking new questions about models we already have, and (4) (in evolutionary or other historical contexts) determine the efficacy of forces that may not be present in the system under investigation but that may have had a role in producing the form that it has."

9 EPISTEMIC REMARKS

A simple hypothetico-deductive picture of the epistemic situation doesn't fit either the historical sciences or ecology. In these sciences we rarely have a well-characterized (let alone formally specified) model whose applicability to some phenomena is in doubt until a healthy range of predictions have been successful. Much more common is a situation in which we know that a number of processes play a role in the phenomena we wish to understand. We then try to build a useful model by representing those processes in more or less detailed ways. Such models are typically tested by comparing patterns of behaviour in the resulting model with various patterns in our observations.

Although detailed predictions of outcomes are rarely expected, a broader grasp of patterns that make sense in the light of our modeling efforts can still be attained. The upshot, when these efforts are successful, is a *retrospective* understanding of the patterns and some features of individual events and cases. The growth of these sciences over time provides an increasingly rich range of significant observable patterns along with increasingly refined understanding of the various processes that are responsible for them.

Of course as Hume argued, there is no general logical license for inferences from the truth of a conclusion to the truth of some set of premises it follows from; to choose a trivial example, when the conclusion is a theorem of the language, it follows from every premise set, but those premises certainly aren't confirmed by the truth of the theorem.[19]

Probabilistic accounts of how observations support hypotheses that we've inferred them from have considerable intuitive appeal, but Bayesian methods provide no help with the initial probabilities such accounts depend on; other theories of probability have made heavy going in their attempts to provide a basis for initial probability assignments. In the absence of a general account of how initial probabilities are arrived at, it is all too easy to explain the intuitive appeal of conditionalization as a simple reflection of our intuitions about what evidence (successful predictions in particular) tells for or against a particular hypothesis, rather than as providing a *justification* for those intuitions. Further, the status of consistency constraints only becomes more difficult on a probabilistic approach: the challenge of maintaining consistency in the set of sentences we endorse becomes far more demanding when we're faced with a demand for *coherence*, the probabilistic generalization of consistency. The calculational burden of maintaining coherence in a large set of commitments is very heavy indeed; this renders downgrading the status of coherence to an ideal rarely met except in very specific contexts even more inevitable than the parallel downgrading of consistency.

Equally well established is the point that it's rare for a hypothesis all on its own to entail something that we can test independently. We need to draw on other premises in addition to the hypothesis to arrive at conclusions we can compare against observation (or some other independent and credible source of information). This gives rise to the familiar Duhem-Quine problem. Making a convincing case against skepticism in this context is extremely challenging, though (as noted above) this problem is not any harder for the historical sciences than it is for the natural sciences. In fact, because the importance of certain basic processes invoked in the historical sciences (if not the particular forms they take in a given model) is taken to be settled, skepticism about hypotheses in the historical sciences can be extremely unattractive even when detailed predictions remain impossible.

Because the historical sciences typically appeal to rich narratives involving multiple, interacting processes to provide an explanation for some phenomena and each process involved in such a narrative is often well understood on independent grounds, many *features* of the resulting models are not hypotheses up for test, but components that, in some form or other, *must* belong to any credible model of the phenomena. What tends to be in doubt are questions including how well

[19]See [Norton, 2003] for an account of *local* induction, rejecting the idea that there is any general formal structure that distinguishes all and only good inductive arguments.

our model captures the relevant features, how to represent their interaction, and whether they constitute a *complete* collection of the basic processes a realistic (or useful) model needs to include.

However, even lacking detailed predictive success with regard to the outcomes of particular cases, retrospective refinement and the exploration of a range of different models can lead us to convincing explanations of some observed patterns. Specific predictions are often not the aim, no more than specific predictions of the course of evolution in populations of organisms are the aim of evolutionary biologists.

Success at constructing such explanations emerges from retrospective testing and refinement of models in the light of ongoing observations. Further input from observation (*e.g.*, evidence of new mutations in a population of bacteria and the processes that lead them to be selected for or against)[20] is often needed to link the models provided by our theory to the outcomes in actual populations. But *patterns* of change in many different populations, and the historical patterns that emerge from those changes over time (such as the patterns of resemblance and difference found in the structurally parallel hierarchical trees of organisms that emerge from taxonomy, embryology, biogeography and paleontology) can be explained with the help of such observations. Scale factors and the *Reynolds number* in fluid mechanics are another illustration—though we cannot, as Cartwright emphasizes, predict the trajectory of a bill being blown across a square by turbulent winds, we can predict when flows will become turbulent (at Reynolds numbers greater than 4,000) and how the scales of vortices and eddies relate to the distribution of kinetic energy amongst them (cf. Kolmogorov's statistical theory of turbulence). Why is success at predicting such general and statistical features of turbulent flow to be discounted?

Of course, things would be different if alternative models offered detailed predictions of the bill's path. But does anyone think this is a likely accomplishment? Even regarding it as credible requires strong skepticism about classical fluid mechanics despite its successes in so many applications, including its account of the general features of turbulent flow and the circumstances in which it occurs. Further, there are concrete practical applications associated with contemporary work on turbulence applying computational fluid dynamics to refine Kolmogorov's account [Moin and Kim, 1997].

Obviously, whether a model is successful depends on what we expect of it. Less obviously, we should not assume that models always aim at the same sorts of ends, even when they belong to the same science and paradigmatically successful models in that science achieve certain ends.

Even when actual observations don't fit such a model (as the famous Club of Rome model of resource depletion failed to fit actual economic developments) the model can be very revealing: that the processes involved in this model correctly captured part of what was going on is not in dispute. Its predictive failure shows that other processes (including the conservation and discovery of new sources of essential resources along with development of alternative resources) are also

[20] See [Blount, *et al.*, 2008].

taking place. Given the constraints we believe are in place (natural limits on resource availability) we may use the obvious failure in a *modus tollendo tollens* instead—and follow that step by modifying the models. As the late Henry Kyburg noted, the 'if-then' inferences embodied in the model are important constraints on the system, despite the falsehood of the antecedent: "The weakness of the data base, the obvious inadequacy of the world model as a mechanism for categorical predictions, the fact that the model takes account neither of political feasibility nor moral acceptability—none of these things keep the model from being informative, none of them suggest we need not take the model seriously, none of them provide an excuse for not getting on with the next step, which is that of trying to discover what alternatives are open to us, what courses of action can in fact be implemented, and by whom, and what ethical, political, and social constraints it is possible to impose on those alternatives without eliminating them all" [Kyburg, 1983, p. 11]. Kyburg concluded, "I suggest that in either case, evidence is evidence and we should attend to it. I suggest that the results of programming world models in computers are relevant to our decisions, even though they are not—and were never intended to be—categorical forecasts of the future".

10 CONCLUSION: ECOLOGY AS A HISTORICAL SCIENCE

To this point we've identified a number of differences between the historical and natural sciences and considered some examples drawn from ecology. On each of these points of difference I think ecology falls more naturally on the side of the historical sciences.

First, ecological explanations generally share in the contingency of historical explanations, turning on a wealth of details that are clearly contingent. Consider as an example here cases of invasive species, where facts including a suitable climate, food sources, lack of predators, diseases and other constraints normally faced by species in their home territory serve to explain the dramatic spread of some newly introduced species. *Bufo maritimus*, the cane toad, reaches densities as much as 100 times those typical of its native habitats in some areas of Queensland Australia; the lack of predators able to cope with its poisonous glands plays a substantial role in the high adult survival that (chiefly) explains these high population densities. Such explanations are, sadly, all too often retrospective—a lack of caution and scientific input is not the only reason why invasive species have been deliberately introduced in so many cases, with such unhappy results: it is far from easy to anticipate such disasters, and perhaps the best lesson we can learn from them is simple caution.

Second, ecological explanations share the narrative structure of explanations typical in the historical sciences. Different processes and contingent features of the circumstances come together to provide an ecological account of (for example) the successful migratory habits of water birds, or the demise of the Dodo. These accounts depend on detailed work, which continues to turn up interesting connections that illuminate the natural interactions—the natural history—that extend and refine our explanations.

Third, our understanding of basic ecological processes is not grounded in theoretically pure systems of principles. Instead, our initial grasp of these processes is largely the result of independent, straightforward observations: that organisms tend to reproduce at rates that lead them to outstrip the available resources; that they have certain needs (metabolic, social, climatological); that they often have predators and suffer from various diseases; that each predator and disease poses different levels of risks under different circumstances; that both resources and threats are distributed in space and time in complex ways, and so on. The role of models is to find useful ways to capture such facts and to gather them together in a single inferential tool. Further, the results of such efforts are often useful in ways other than generating predictions based on the model and some observationally-grounded 'input' parameters.

Fourth, the evaluation of ecological models and hypotheses turns less on fundamental principles and more on the set of processes, modeled in different ways, which are combined to explain features of the phenomena. Data limitations are answered, in large part, by retrospective evaluation, which can provide many separate back-tracking inferential paths that coherently support a narrative explanation.

Fifth, ecological models have (at best) very limited predictive power—they are too complex, the phenomena they model are still more complex, our ability to gather and adequately represent the relevant data are too limited, and the systems they involve are open to many different kinds of processes that appear (in our models) as purely exogenous forcings. Nevertheless, we can be justifiably confident that they do capture important features of ecological processes, and that certain particular explanations arrived at in ecological studies are well-founded, due to the combination of confidence about many of the basic processes involved and the rich cross-checking that retrospective investigations can provide.

Sixth, Sellars' distinction between ontological and inferential reductions makes clear how the explanations offered by historical sciences including ecology can be substantively independent of the detailed principles that underlie various ecological and historical processes. The inferences we make are, and must be, shaped by the demands of particular domains[21] even if we maintain, as ontological reductionists, that the items we are describing are ultimately based on the ontology of physics.

Seventh and last, the use of ecological models displays all the features typical of models in the historical sciences as well, including openness, problems of both model and natural complexity, and substantial limitations on the data we can gather, when compared with the richness and detail of the natural phenomena we are describing. Finding grounds for taking such models seriously turns not on predicting detailed outcomes, but on finding cases in which the models are able to capture various features of the phenomena in ways that are independently credible, because support for the different processes and their interactions are drawn from separate lines of evidence. The result is, as Kyburg emphasized, a matter of *constraints* on the phenomena, rather than detailed predictions—the

[21] See [Shapere, 1974].

resulting inferences can tell us not what will happen but instead such things as how some of the factors and processes involved interact, what would happen if other processes and factors did not intervene, what sorts of processes may be involved in one case or another and (in many cases) how we can draw on various kinds of evidence to reconstruct what happened in particular cases.

BIBLIOGRAPHY

[Albritton, 1980] C. Albritton. *The Abyss of Time* San Francisco: Freeman Cooper & Co., 1980.

[Aristotle, 350 BCE] Aristotle, *Posterior Analytics*, Translated by G. R. G. Mure. `http://classics.mit.edu/Aristotle/posterior.1.i.html`.

[Blount *et al.*, 2008] Z. D. Blount, C. Z. Borland, and R. E. Lenski. Historical contingency and the evolution of a key innovation in an experimental population of *Escherichia coli*. *Proceedings of the National Academy of Sciences USA* 105: 7899–7906, 2008.

[Bonevac, 1982] D. Bonevac. *Reduction in the Abstract Sciences*. Indianapolis, IN: Hackett, 1982.

[Brown, 2004] B. Brown. The Pragmatics of Empirical Adequacy. *Australasian Journal of Philosophy* 82: 242–263, 2004.

[Buffon, 1756] Buffon, *Histoire Naturelle, Générale et Particulière, avec la description du cabinet du roi*. Tome Sixième. `www.buffon.cnrs.fr`.

[Burchfield, 1975] J. D. Burchfield. *Lord Kelvin and the Age of the Earth*. London: MacMillan, 1975.

[Cartwright, 1999] N. Cartwright. *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press, 1999.

[Chamberlin, 1899] T. C. Chamberlin. Lord Kelvin's Address on the age of the earth as an abode fitted for life (Part Two). *Science* 10: 11–18, 1899.

[Drake, 1970] S. Drake. *Galileo Studies*. Ann Arbor: University of Michigan Press, 1970.

[Earman and Norton, 1998] J. Earman, and J. D. Norton. Comments on Laraudogoitia's "Classical Particle Dynamics, Indeterminism and a Supertask". *British Journal for the Philosophy of Science* 49: 123–133, 1998.

[Frisch, 1995] U. Frisch. *Turbulence: The Legacy of A. N. Kolmogorov*. Cambridge: Cambridge University Press, 1995.

[Hallam, 1983] A. Hallam. *Great Geological Controversies*. Oxford: Oxford University Press, 1983.

[Kyburg, 1983] H. Kyburg. Prophecies and Pretensions. In H. Kyburg, *Epistemology and Inference*, pp. 3–17. Minneapolis: University of Minnesota Press, 1983.

[Lampo and Leo, 1998] M. Lampo and G. A. de Leo. The invasion ecology of the toad, *Bufo marinus*: From South America to Australia. *Ecological Applications* 8(2): 388–396, 1998.

[Lyell, 1830] C. Lyell. *Principles of Geology*, Volume 1. London: John Murray, 1830.

[Mathieu and Scott, 2000] J. Mathieu and J. Scott. *An Introduction to Turbulent Flow*. Cambridge: Cambridge University Press, 2000.

[Moin and Kim, 1997] P. Moin and J. Kim. Tackling Turbulence with Supercomputers. *Scientific American* 276(1): 62–68, 1997.

[Morgan and Morrison, 1999] M. S. Morgan and M. Morrison. *Models as Mediators*. Cambridge: Cambridge University Press, 1999.

[Norton, 1993] J. D. Norton. The determination of theory by evidence: The case for quantum discontinuity, 1900–1915. *Synthese* 97(1): 1–31, 1993.

[Norton, 2003] J. D. Norton. A Material Theory of Induction. *Philosophy of Science* 70: 647–70, 2003.

[Oreskes, 2003] N. Oreskes. The Role of Quantitative Models in Science. In Charles Draper William Canham, Charles D. Canham, Jonathan Cole and William K. Lauenroth. *Models in Ecosystem Science*, Ch. 2, pp. 13–32. Princeton: Princeton University Press, 2003.

[Rudwidk, 2007] M. J. S. Rudwick. *Bursting the Limits of Time: The Reconstruction of Geohistory in the Age of Revolution*. Chicago: University of Chicago Press, 2007.

[Sarkar, 2005] S. Sarkar. "Ecology" Stanford Encyclopedia of Philosophy, `http://plato.stanford.edu/entries/ecology/`.

[Sellars, 1963] W. Sellars. Philosophy and the Scientific Image of Man in W. Sellars, *Science, Perception and Reality*, pp. 1–40. London: Routledge and Kegan Paul, 1963.

[Shapere, 1974] D. Shapere. Scientific theories and their domains. In Suppe, F., *The Structure of Scientific Theories*, pp. 518–565. Chicago: University of Illinois Press, 1974.

[Twain, 1883] M. Twain. *Life on the Mississippi*. Boston: James R. Osgood and Company, 1883.

[Wimsatt, 1987] W. C. Wimsatt. False Models as Means to Truer Theories. In Nitecki, Matthew H. and Hofman, Antoni (eds.), *Neutral Models in Biology*, pp. 23–55. Oxford: Oxford University Press, 1987.