

Introduction: What is the Philosophy of Science?

Christopher Hitchcock

What is the philosophy of science? It is the application of philosophical methods to philosophical problems as they arise in the context of the sciences. That's not a particularly helpful answer as it stands, but at least it allows us to break our original question into parts: What are the methods of philosophy? What are philosophical problems? How do these problems arise within different scientific fields?

0.1 Philosophical Methods

The first question is the most difficult. In the first half of the twentieth century, a prominent school of thought (particularly associated with the Austrian philosopher Ludwig Wittgenstein) held that the philosopher's task was to clarify the meanings of words. The great problems of philosophy, it was thought, were mere confusions resulting from a failure to understand the meanings of the words used to frame those problems. Few philosophers today would subscribe to such an extreme view; nonetheless, the clarification of meanings is still an important part of the philosopher's repertoire. Particularly important is the ability to draw distinctions between different things that a term or phrase might mean, so as to assess more accurately claims involving those terms or phrases. The chapters on genetics by Sahotra Sarkar and Peter Godfrey-Smith (chapters 13 and 14), for example, involve careful analysis of the various things that one might mean by "information."

Perhaps even more fundamentally, philosophy involves the analysis of arguments, often aided by the formal methods and conceptual resources of symbolic logic (and other areas, such as probability theory). Philosophers, when defending a position, will construct arguments in support of that position. In addition, they will examine arguments that have been proposed by opponents. For each such argument, they may ask: What is the structure of the argument? Is it logically valid? If not, would it be valid

if one were to add certain specific premises? Does it employ inferential methods other than those of deductive logic? What are the premises of the argument? Are the premises true? – and so on. Moreover, philosophers will try to anticipate objections to their own arguments, and defend their arguments against these objections before they are even raised. Almost every philosophy paper employs these methods to some extent or other; the two chapters on unobservable entities, by Jarrett Leplin and by André Kukla and Joel Walmsley (chapters 5 and 6), provide particularly clear examples – both chapters examine, criticize, and propose a variety of arguments on both sides of the debate.

Nonetheless, it is almost impossible to isolate any uniquely philosophical methods. In the philosophy of science, especially, there is no clear line where the philosophy ends and the science begins. While few (but still some!) philosophers actually conduct experiments, many philosophers will freely make use of empirical findings to support their positions. Consider chapters 15 and 16, by Peter Carruthers and by James Woodward and Fiona Cowie, for example. These chapters tackle the question “Is the mind a system of modules shaped by natural selection?” This involves traditional philosophical issues, such as the relationship between the mind and the brain; it involves careful analysis of the concept “module”; but it also requires the consideration of empirical results in psychology, as well as theoretical issues in evolutionary biology. Like empirical scientists, philosophers sometimes construct mathematical models of the “phenomena” that they seek to understand. In his chapter on scientific confirmation (chapter 3), Patrick Maher uses probability theory to construct a mathematical relation that, Maher argues, captures important features of the relation between scientific theory and empirical evidence. In general, then, it appears that philosophers are willing to employ almost any tools that can shed light on philosophical problems.

0.2 Philosophical Problems

It is hard to say what makes a problem “philosophical.” There are, nonetheless, certain collections of problems that, over the past two and a half millennia, have come to be seen as paradigmatically philosophical problems. Three central areas of concern are *ethics*, *epistemology*, and *metaphysics*. This is by no means an exhaustive list – a fuller list would have to include aesthetics (the study of art and beauty), logic, social and political philosophy, the philosophies of language, mind, and religion (not to mention the philosophy of science itself), and the history of philosophy. Nonetheless, the core areas of ethics, epistemology, and metaphysics intersect with all these branches of philosophy; understood broadly, these three areas cover much of the field of philosophy.

Ethics deals with issues of right and wrong – both the morality of specific types of behavior and also more fundamental issues concerning the ultimate sources of moral value. Epistemology deals with the nature of knowledge and belief: What is knowledge, and how is it distinguished from mere belief? What are the sources of knowledge? What constitutes *justified* belief? Metaphysics is the most difficult to characterize; roughly, it involves the examination of concepts that play a fundamental role in other areas of philosophy, and in other disciplines. For example, metaphysi-

cal issues fundamental to ethics involve concepts such as the freedom of the will, and the nature of personal identity.

0.2.1 Ethical issues in science

Ethical issues can arise in a number of ways within the scientific context. Most obviously, technical innovation can create new possibilities whose moral status is in need of evaluation. For example, only recently has it become possible to clone large mammals such as sheep. It may soon be technologically possible to clone human beings (at the time of this writing, there are unsubstantiated reports that this has already happened). Many people react in horror at the thought of human cloning; similar reactions met other forms of technologically aided reproduction, such as artificial insemination and *in vitro* fertilization. Just what, if anything, is wrong with creating a genetic copy of a human being? Does this outweigh the possible benefits of cloning as a form of reproductive technology, especially for individuals or couples who have no other option? Obviously, ethical theorists such as Aristotle, Kant, and Mill were not able to anticipate these sorts of issues when developing their moral theories.

Another set of issues arises in connection with the treatment of experimental subjects. Presumably, the sub-atomic particles that are forced to follow very constrictive paths only to be annihilated in a super-collider are not *harmed* in any morally relevant sense. Experiments involving human beings, or even nonhuman animals, are more problematic. For human subjects, a consensus has emerged (although surprisingly recently) that *informed consent* is essential: experimentation upon human subjects is permissible only when the subjects have voluntarily given their consent after being informed of the potential risks and benefits involved. By their very nature, however, experimental treatments are such that the potential risks and benefits are not fully known in advance. Moreover, the notion of consent is much more complex than it appears. Various forms of coercion may affect a person's decision to participate in an experiment. In medicine, there is often a power asymmetry between patient and doctor, and a patient may feel that she has to participate in order to receive the best treatment. In psychology, it is a common practice for professors to require students to participate in experiments to receive course credit. In the case of animal subjects, informed consent is, of course, impossible. The key issues here involve the moral status of animals. Mammals, at least, are quite capable of suffering physical pain as well as some forms of psychological distress. How is this suffering to be weighed against the potential benefits of experimentation for human beings?

Recently, there has been considerable concern about the status of women and minorities in the sciences. There can be little doubt that the scientific profession has discriminated against women as well as members of racial and religious minorities in a number of ways. Perhaps most obviously, there have been considerable barriers preventing women and minorities from pursuing scientific careers (taking an extreme form, for example, in the expulsion of Jews from scientific posts in Nazi Germany). Some have argued that the exclusion of such alternative voices has harmed science by narrowing its vision.

To provide just one more example of ethical issues concerning science, let us remind ourselves that scientific research costs money. The funding that is necessary to support

scientific research comes from a finite pool, and a decision to fund one research project is inevitably a decision to withhold funds from other projects, both within and outside of science. How are these decisions made? How can we evaluate the financial value of pure research as balanced against health care, education, defense, and other needs? How can we evaluate the financial value of one research project against another quite different one? Should these decisions be made by scientists alone, or should lay citizens participate in these decisions? If the latter, what form would lay participation take?

While these and related ethical issues in science are not taken up by the contributions in this volume, some of them will be covered in future volumes of the Blackwell *Contemporary Debates* series.

0.2.2 Epistemological issues in science

Science is in the business of producing knowledge, so it is not particularly surprising that epistemological problems arise in the scientific context. One of the most fundamental questions concerns the ultimate *source* of knowledge. *Empiricism* holds that all our knowledge of the world derives from sense experience. If you want to know what the world is like, you have to go out and look (or listen, feel, smell, or taste). It is a little tricky to say just what is meant by knowledge of *the world*, but there is an intended contrast with, for example, knowledge of mathematics or logic. Empiricism is most closely associated with three British philosophers of the seventeenth and eighteenth centuries: John Locke, George Berkeley, and David Hume. Locke, especially, held that experience is the ultimate source of all our *ideas*. More modern versions of empiricism hold that experience alone can provide the *justification* for our beliefs about the world: we may perhaps be able to formulate hypotheses without the benefit of sensory input, but only observation can tell us which hypotheses are correct. This form of empiricism is widely espoused by philosophers today.

Empiricism is most often contrasted with rationalism. This view, associated most strongly with the seventeenth-century philosophers René Descartes, Gottfried Leibniz, and Benedict de Spinoza, holds that human reason is the ultimate source of knowledge. Descartes, in particular, held that all knowledge should be constructed on the model of mathematics, deducing conclusions rigorously from basic premises whose truth could not be doubted (such as “I think, therefore I am”). Another alternative to empiricism can be traced to the work of the ancient Greek philosopher Plato (and hence is referred to as “Platonism”). Plato believed that an appropriately trained philosopher could acquire the ability to “see” past the appearances into the reality that lay behind those appearances. The word “see” here is metaphorical, and refers to a special kind of insight, rather than literal vision. This view has recently been revived in an interesting way by the philosopher James Robert Brown. Brown has argued that *thought experiments* can provide us with new knowledge of the world, even though by definition those “experiments” do not involve any new observations. Rather, thought experiments provide us with a kind of direct insight into the nature of things. Brown presents his views in chapter 1. In rebuttal, John Norton argues that thought experiments can be understood in empiricist terms, and calls for a more thorough analysis of the epistemology of thought experiments. The subject of thought experi-

ments, then, provides one small arena in which the larger battles between empiricism, rationalism, and Platonism may be played out.

Despite these disputes, no one would deny that observational evidence plays a prominent (although perhaps not exclusive) role in supporting and undermining scientific hypotheses. How does this happen? In formal logic, there are explicit rules that tell us whether or not a certain conclusion follows from a particular set of premises. These rules are demonstrably truth-preserving: they guarantee that a logically valid inference from true premises will always yield a true conclusion. Let us put aside worries raised by Descartes and others about the reliability of our senses, and assume that the beliefs we form on the basis of direct observation are correct. Are there rules of inference, akin to the rules of deductive logic, that would take us from these observational premises to theoretical conclusions with no risk of error? In general, this is not possible. Any interesting scientific hypothesis has implications whose truth cannot be established by direct observation. This may be because the hypothesis has implications for what goes on at distant locations, or in the future, or at scales too small to be seen by the human eye, or for any number of other reasons. There is thus little hope that we will be able to simply deduce the truth of scientific hypotheses and theories from observations in the way that conclusions can be deduced from their premises in logic. This gloomy conclusion is supported by the history of science, which tells us that even the best-confirmed theories (such as Newtonian gravitational theory) can be undermined by further evidence. Thus, while fields such as mathematics and logic trade in certainties, scientific hypotheses always remain at least partly conjectural.

In light of this situation, some philosophers have attempted to apply the concepts of probability to scientific theories and hypotheses. While it may be impossible to establish a scientific hypothesis with certainty, a hypothesis may be rendered more or less probable in light of evidence. Evidence that increases the probability of a theory is said to support or *confirm* that theory, while evidence that lowers the probability of a theory is said to *undermine* or *disconfirm* it. This way of thinking about the relationship between theory and evidence was pioneered by the eighteenth-century English clergyman Thomas Bayes, and further developed by the great French physicist Pierre Simon de Laplace. The probabilistic approach became very popular in the twentieth century, being championed in different ways by the economist John Maynard Keynes, the English *wunderkind* Frank Ramsey (who died at the age of 26), the Italian statistician Bruno de Finetti, the Austrian (and later American) philosopher Rudolf Carnap, and a host of later writers. One version (or perhaps a collection of interrelated versions) of this approach now goes by the name of “Bayesianism” (after the Reverend Thomas Bayes). The Bayesian position is sketched (and criticized) by Kevin Kelly and Clark Glymour in chapter 4. Patrick Maher, in his contribution to this volume (chapter 3), provides us with a different way of understanding confirmation in probabilistic terms.

A different line of response is most prominently associated with the Austrian (and later British) philosopher Karl Popper, who was knighted for his efforts (one of the perks of being British). This approach denies that it is appropriate to talk about the confirmation of theories by evidence, at least if this is to be understood in terms of epistemic justification. The process whereby scientists subject their theories to empirical

test is not one in which they seek to justify belief in that theory. The scientific method, rather, is one of formulating hypotheses, subjecting them to empirical test, and winnowing out (or at least modifying) those hypotheses that don't fit the results. It is possible that this process will eventually lead us to the truth, or at least to partial truth, but at no point will the empirical data collected to that point provide reason to believe in any of those hypotheses that remain. Kevin Kelly and Clark Glymour, the authors of chapter 4, present their own account of scientific inquiry that shares Popper's skepticism about the idea that data can partially support a general conclusion.

As we noted above, one of the reasons why there is a gap between observational evidence and scientific theory is that the latter often makes claims about entities that are *unobservable*. Are scientific claims about unobservable entities *especially* problematic? Questions of this sort are central to millennia-old debates between *realism* and *anti-realism*. Such debates arise in many different areas of philosophy – we will here be concerned with *scientific* realism and anti-realism, as opposed to, say, moral realism and anti-realism – and they can take on a number of different forms. Sometimes the debates are metaphysical in nature: George Berkeley held that nothing can exist without being perceived (although God perceives many things that are not perceived by humans). Sometimes the debates are semantic: members of the Vienna Circle, a school of philosophy centered in Vienna in the 1920s and 1930s, held that statements apparently referring to unobservable entities were to be reconstrued in terms of the empirical consequences of those statements. We will here be concerned with an epistemic form of the realism/anti-realism debate. In chapter 5, Jarrett Leplin argues that observational evidence can (at least sometimes) provide us with grounds for believing in unobservable entities, and in at least some of the assertions that scientific theories make about those entities. More specifically, Leplin argues that theories are particularly worthy of our belief when they successfully make novel predictions. André Kukla and Joel Walmsley (chapter 6) challenge this argument.

0.2.3 Metaphysical issues in science

Three of the most important concepts that appear throughout the sciences are those of law, causation, and explanation. Let us begin with law.

Almost every branch of science has basic principles referred to as “laws.” In physics there are Snell's law, the Boyle–Charles law, the zeroth, first, and second laws of thermodynamics, Newton's laws of motion and gravitation, and so on. In addition, there are several “equations” that are essentially of the same character: Maxwell's equations in electromagnetic theory, Schrödinger's equation in quantum mechanics, and Einstein's field equations in the general theory of relativity. In biology, we have Mendel's laws and the Hardy–Weinberg law; in economics, Gresham's law and the law of supply and demand. The list could easily go on. In general, science seeks not only to discover what particular events take place where and when, but also to reveal the basic principles according to which these events unfold.

Just what makes something a law? According to one account championed by many empiricist writers, a law is a *regularity*. That is, a law is a pattern of the form “Whenever condition A is satisfied, condition B will be satisfied as well.” There may be,

however, any number of regularities of this form that are not laws – all senators for California in 2002 were women, but that is hardly a scientific law. Various proposals have been offered for discriminating true laws from such “accidental generalizations” – for example, laws must be fully general, and not make specific reference to any particular individuals, places or times – but none have become widely accepted. Another problem is that many of the “laws” in science are not universal regularities. There are, for example, certain types of genes (segregation-distorters) that do not obey the genetic “law” of random segregation. Such nonuniversal “laws” are sometimes called *ceteris paribus* laws, laws that hold true other things being equal. The thought is that there exists some specifiable set of conditions, as yet unknown to us, under which the regularity never fails. If these conditions are then built in to the condition A in the formulation “Whenever condition A is satisfied, condition B will be satisfied as well,” then perfect regularity will be restored. John Roberts, in his contribution to this volume (chapter 7), champions the view that laws are regularities while arguing that *ceteris paribus* laws are no laws at all. As a consequence, he holds that all of the so-called “laws” of the social sciences, such as the law of supply and demand, are not true laws.

A very different approach to understanding laws, associated most prominently with the Australian philosopher David Armstrong, takes laws to comprise relationships of “necessitation” that hold between properties, rather than individual entities. James Robert Brown briefly discusses this account in his chapter on thought experiments (chapter 1). Harold Kincaid (chapter 8) argues that at least some laws serve to pick out certain kinds of causal tendency.

The concept of causation is closely related to that of law. According to one view, once widely held, one event A causes another event B just in case B follows A as a matter of law. Obviously, this account of causation will inherit the problems described above concerning the understanding of laws. Moreover, this account of causation will not be very illuminating if laws, in turn, must be understood in terms of causation. Even putting these worries aside, a number of problems remain. Consider, for example, the explosion of the Challenger space shuttle in 1986. One of the causes of this unfortunate event was the freezing of the rubber O-ring used to prevent the leaking of fuel. Are there laws that guarantee that whenever an O-ring freezes (and various other conditions also hold) then a space shuttle will explode? We certainly have found no such laws, and yet we nonetheless believe that the freezing of the O-ring did cause the explosion. Thus we are able to provide evidence for the truth of causal claims, even when that same evidence provides no support for an underlying law. Or suppose that I lick an ice cream cone on a sunny day, after which photons bounce off the cone at a velocity of (roughly) 300,000 kilometers per second. It certainly follows from the laws of physics that anytime I lick an ice cream cone on a sunny day, photons will bounce off the cone at just that speed. However, my licking the ice cream had nothing to do with this – it would have happened regardless of whether I licked the cone, or whether I foolishly watched it melt without ever licking it. So lawful succession appears to be neither necessary nor sufficient for causation.

In response to these problems, a number of alternative approaches to causation have been developed. Both Phil Dowe (chapter 9) and Jonathan Schaffer (chapter 10) canvass some of these alternatives. Dowe himself thinks that A causes B when they

are connected by a *causal process* – a certain kind of physical process that is defined in terms of conservation laws. Jonathan Schaffer, in his chapter, argues that many causes are not so connected to their effects.

The third interrelated concept is that of explanation. At the beginning of the twentieth century, the French physicist Pierre Duhem claimed that physics (and, by extension, science more generally) cannot and should not explain anything. The purpose of physics was to provide a simple and economical system for describing the facts of the physical world. Explanation, by contrast, belonged to the domain of religion, or perhaps philosophy. The scientists of an earlier era, such as Sir Isaac Newton, would not have felt the need to keep science distinct from religion and philosophy; but by 1900 or so, this was seen to be essential to genuine progress in science. This banishment of explanation from science seems to rest on a confusion, however. If we ask “Why did the space shuttle Challenger explode?”, we might mean something like “Why do such horrible things happen to such brave and noble individuals?” That is certainly a question for religion or philosophy, rather than science. But we might instead mean “What were the events leading up to the explosion, and the scientific principles connecting those events with the explosion?” It seems entirely appropriate that science should attempt to answer that sort of question.

Many approaches to understanding scientific explanation parallel approaches to causation. The German–American philosopher Carl Hempel, who has done more than anyone to bring the concept of explanation onto center stage in the philosophy of science, held that to explain why some event occurred is to show that it *had* to occur, in light of earlier events and the laws of nature. This is closely related to the “lawful succession” account of causation described above, and it inherits many of the same problems. Wesley Salmon, an American philosopher whose career spanned the second half of the twentieth century, was a leading critic of Hempel’s approach, and argued for a more explicit account of explanation in terms of causation, to be understood in terms of causal processes. Salmon’s view of explanation is thus closely related to (and indeed an ancestor of) Dowe’s account of causation. (Hence it is potentially vulnerable to the sorts of objection raised in Schaffer’s chapter 10.) A third approach, developed in greatest detail by Philip Kitcher, identifies explanation with *unification*. For example, Newton’s gravitational theory can be applied to such diverse phenomena as planetary orbits, the tides, falling bodies on earth, pendula, and so on. In so doing, it shows that these seemingly disparate phenomena are really just aspects of the same phenomenon: gravitation. It is the ability of gravitational theory to unify phenomena in this way that makes it explanatory. While none of the chapters in this volume deals specifically with the problem of analyzing the concept of explanation, the subject of scientific explanation is discussed in a number of them, especially chapters 5, 6, 7, 8, 10, and 11.

0.3 The Sciences

In addition to the problems described above, which arise within science quite generally, there are a number of problems that arise within the context of specific scientific disciplines.

0.3.1 Mathematics

It isn't entirely clear whether mathematics should be regarded as a science at all. On the one hand, mathematics is certainly not an *empirical science*: Mathematicians do not conduct experiments and mathematical knowledge is not gained through observation of the natural world. On the other hand, mathematics is undoubtedly the most precise and rigorous of all disciplines. Moreover, in some areas of science, such as theoretical physics, it is often hard to tell where the mathematics ends and the science begins. A mathematician specializing in differential geometry and a theoretical physicist studying gravitation may well be working on the same sorts of problems (albeit with different notational conventions). Scientists in many disciplines solve equations and prove theorems, often at a very abstract level.

The most fundamental questions in the philosophy of mathematics ask what the *subject matter* of mathematics is, and how we acquire knowledge about it. Let's take a very simple case: arithmetic is about *numbers*. Just what are numbers? They are not "things" in the physical world, like planets, cells, or brains. Nor do we find out about them by observing their behavior. (Of course it may aid our understanding of arithmetic to play with blocks or marbles – if you put two marbles in an empty bag, and then another three, there will be five marbles in the bag. But it would be very odd indeed to consider this an empirical test of the hypothesis that $2 + 3 = 5$.) Of course, the standard method for acquiring knowledge in mathematics is proof: *theorems* are deduced from mathematical *axioms*. What a proof shows then is that the theorem is true *if* the axioms are true. But how do we know whether the axioms are true? We cannot derive them from further axioms, on pain of infinite regress. We cannot assess their truth by observation. One approach to this problem is to claim that mathematical axioms are not true or false in any absolute sense, but only serve to define certain kinds of mathematical system. For example, on this view Euclid's postulates are not assertions about any physical things but, rather, serve to define the abstract notion of a Euclidean geometry. Theorems that are derived from these axioms can only be said to be true in Euclidean geometry; in non-Euclidean geometries, these theorems may well turn out to be false. A mathematical system may be used to model a particular physical system, and it is an empirical matter whether or not the model fits, but this is an independent matter from that of whether the mathematics itself is true.

A different approach is that of mathematical Platonism, often associated with the Austrian mathematician Kurt Gödel of incompleteness theorem fame. (Like many of the great German and Austrian philosophers and mathematicians of the 1930s, he emigrated to America. He never became an American citizen, however, since he believed that there was a logical inconsistency in the American constitution.) According to Platonism, mathematical entities are in some sense *real*: there is an abstract realm in which numbers, sets, scalene triangles, and nondifferentiable functions all live. (This realm is sometimes referred to metaphorically as "Plato's heaven.") We are able to acquire knowledge of this realm by means of a kind of mathematical insight. Mathematical proof then becomes a tool for expanding our knowledge beyond the elementary basis of mathematical propositions that we can "see" to be true. Although this collection has no chapters specifically devoted to the philosophy of mathematics, James Robert Brown defends a Platonist view of mathematics in chapter 1.

0.3.2 Physics

Many philosophers of science have viewed physics as *the science par excellence*. It is certainly true that physics, and astronomy in particular, was the first empirical science to be rendered in a mathematically precise form. Even in the ancient world, it was possible to make very accurate predictions about the locations of the planets and stars. In the seventeenth century, Newton was able to formulate physical laws of unparalleled scope, unrivaled in any other branch of science for almost 200 years (Darwin's theory of evolution by natural selection and Mendeleev's periodic table perhaps being the only close competitors by the year 1900). It would not have been unreasonable, then, for philosophers to predict that all genuine branches of science would ultimately come to look like physics: a few simple laws of vast scope and power. Thus a full understanding of science could be gained simply by understanding the nature of physics.

In the twenty-first century, we have come to learn better. Chemistry and biology have certainly advanced to the stage of scientific maturity, and they look nothing like the model of a scientific system built upon a few simple laws. In fact, much of physics does not even look like this. Nonetheless, physics continues to pose a number of fascinating puzzles of a philosophical nature. The two most fundamental physical theories, both introduced in the early twentieth century, are *quantum mechanics* and *general relativity*. Newtonian physics provides an extremely accurate account of slow, medium-sized objects. It breaks down, however, at the level of very small (or more precisely, very *low energy*) objects such as sub-atomic particles; at the level of objects traveling at near-light velocity; and at the level of very massive objects such as stars. Quantum mechanics describes the behavior of very small objects, special relativity describes the behavior of very fast objects, and general relativity (which includes special relativity) describes very massive objects. All of these theories agree almost exactly with Newtonian mechanics for slow, medium-sized objects. As yet, however, there is no known way of incorporating quantum mechanics and general relativity into one unified theory.

Within quantum mechanics, the most substantial conceptual puzzle concerns the nature of *measurement*. According to the mathematical theory of quantum mechanics, which is extraordinarily accurate in its predictions, there are two different rules describing the behavior of physical systems. The first rule, Schrödinger's equation, describes a continuous and deterministic transition of states. This rule applies to a system unless it is being measured. When a system is measured, a new rule, named after Max Born, takes effect. Born's rule describes a transition that is discontinuous and indeterministic. When a system is measured, it is said to *collapse* into a new state, and the theory provides us only with probabilities for its collapse into one state rather than another. But how does a system "know" that it is being measured? Why can't we treat the original system, together with the measurement apparatus – whatever physical system is used to perform the measurement – as one big system that obeys Schrödinger's equation? And just what is a measurement anyway? It can't just be any old physical interaction, or else any multiple-particle system would be collapsing all the time. The Nobel laureate physicist Eugene Wigner even believed that *human consciousness* is the special ingredient that brings about collapse. Others have maintained

that collapse is just an illusion. In the context of quantum mechanics, then, the concept of *measurement* is a particularly perplexing one.

General relativity raises a host of interesting questions about the nature of space and time. Between 1714 and 1716, Samuel Clarke, a close follower of Sir Isaac Newton, participated in a detailed correspondence with Gottfried Leibniz. (It is believed that Newton himself may have had a hand in drafting Clarke's letters; Clarke's strategy of writing a final reply after Leibniz's death in 1716 certainly smacked of Newton's vindictiveness.) They debated many different issues, including the nature of space and time. According to Newton's theory, *acceleration* has particular sorts of causes and effects. This seems to imply that there is a distinction between true accelerations and merely apparent ones. If I jump out of an airplane (with a parachute I hope!), I will accelerate toward the ground at a little under ten meters per second per second. But from my perspective, it may well seem that it is the ground that is accelerating up to me at that rate! In fact, however, only one of us (me) is being subject to a force sufficient to produce that acceleration. Newton (and hence Clarke) thought that this required the existence of an *absolute* space: one's true motion was one's change in location in absolute space, regardless of what other objects may be doing. Leibniz, by contrast, held that the only true motions were the motions of objects relative to one another. Absolute space was nothing more than a mathematical abstraction used to model the various relative motions. Einstein's special and general theories of relativity add new dimensions to this old debate. On the one hand, general relativity shows that one can formulate the laws of physics relative to *any* frame of reference: it does not matter which objects we think of as moving, and which we think of as being at rest. This would seem to undermine Newton's central reason for believing in an absolute space and time. On the other hand, in the framework of general relativity, matter (or more specifically, energy) interacts with spacetime. (Since, in relativity theory, space and time are intimately bound together, we refer to "spacetime" rather than to space and time separately.) The distribution of mass-energy affects the structure of spacetime, and the structure of spacetime determines how objects will move relative to one another. So in this framework, space and time seem able to causally interact with matter, which certainly suggests that they bear some kind of physical reality.

General relativity also introduces some interesting new physical possibilities, such as the collapse of massive stars into black holes. Perhaps most intriguingly, general relativity seems to be consistent with the existence of "closed causal curves," which would seem to admit the possibility of some kind of time travel. Such a possibility obviously presents serious challenges to our normal understanding of the nature of time. One important spin-off of the general theory of relativity is contemporary cosmology, including the well-confirmed "big bang" hypothesis. Of course, any theory that deals with issues such as the origins and eventual fate of the universe brings in its wake a host of philosophical questions.

One further area of physics that raises interesting philosophical problems is thermodynamics, developed in the first half of the nineteenth century by Clausius, Carnot, Kelvin, and others. This work was given new physical underpinnings by work in statistical mechanics in the second half of the nineteenth century, especially by Maxwell and Boltzmann. Of the three basic laws of thermodynamics, the second is by far the

most philosophically interesting. It says that the entropy of a physical system can increase, but never decrease. Entropy, very informally, is the amount of “disorder” in a system – a more thorough explanation is given in chapter 11 by Huw Price. For example, if milk is poured into coffee, it will very quickly mix with the coffee until the milk is uniformly spread throughout the coffee. Once mixed, however, the milk will never spontaneously separate and jump back out into the pitcher. There are many asymmetries in time: time seems to move toward the future; we remember the past, but not the future; we seem to have some control over the future, but not the past; we would prefer to have our unpleasant experiences in the past, and our pleasant ones in the future; and so on. The second law of thermodynamics presents the promise of an explanation of these phenomena, or at the very least a physical underpinning for the idea that there is a fundamental difference between the past and the future. Unfortunately, the later work of Maxwell and Boltzman raised a number of problems for the second law. First, they showed that it is not, in fact, impossible for a system to decrease in entropy; it is only very unlikely (in some sense of unlikely – see Craig Callender’s chapter 12) that this will happen. More fundamentally, the behavior of thermodynamic systems is ultimately determined by Newton’s laws, which are completely symmetric with respect to the future and the past. There remains a mystery, then, as to how the asymmetry described by the second law can come about.

The two chapters by Price and Callender, chapters 11 and 12, deal with issues at the intersection of cosmology and thermodynamics. Price argues that the recent discovery that the universe was very “smooth” shortly after the big bang helps to explain why entropy increases in time. Price and Callender disagree, however, about whether the initial smoothness itself is something that is in need of explanation.

0.3.3 Biology

Most philosophers of biology have focused their attention on the theory of evolution by natural selection. According to this theory, there is variation within any species: individuals within that species do not all have the same characteristics. Some characteristics will give their bearers advantages in their competition with conspecifics for food, mates, and other resources. Individuals with advantageous characteristics, will, on average, produce more offspring than their rivals. Many of these characteristics can be inherited from one generation to the next. Those characteristics that have all three traits – variability, reproductive advantage, and heritability – will become more plentiful in subsequent generations. The gradual accumulation of such changes over time gives rise to the evolution of diverse life forms, many with highly complex adaptations to their specific environments.

One problem, the “units of selection” problem, concerns the “level” at which these processes occur. Strictly speaking, it is *genes* rather than *characteristics* that are passed on from one generation to another. Perhaps, then, we should also say that it is the genes, rather than individual organisms, that are competing with one another for the opportunity to reproduce themselves. In this picture, championed by the British biologist Richard Dawkins among others, individual organisms are just “vehicles” built by genes in order to aid those genes in reproducing themselves. Is this an appropriate way to describe what is going on? Just what is at stake in saying that it is the

genes, rather than the individual organisms, that are being acted upon by natural selection?

Another issue concerns the extent to which evolution by natural selection can ground teleological claims, claims about the “purpose” or “function” of phenotypic traits (e.g., body morphology or behavioral predispositions). In the case of human artifacts, the function of an object is determined by the intentions of the designer. The function of a screwdriver is to turn screws, for that is what the screwdriver was designed for. I may use a screwdriver to jimmy open a door, or threaten a neighbor’s dog, but these are not the *functions* of a screwdriver; they’re not what the screwdriver is *for*. If (as so-called “intelligent design theorists” believe) organisms were created by an intelligent agent, then it would make perfect sense to talk about the purpose or function of a bird’s feathers, a flower’s petals, a snake’s rattle, and so on: this would just be the use for which the designer intended the characteristic in question. If, however, an organism has evolved naturally, can we sensibly talk about the function of its various characteristics? Some think that we can. The function of a trait is that activity of the trait for which it was naturally selected. The heart is often used as an illustration. The heart does many things: it circulates the blood throughout the body, and it also makes rhythmic sounds. The second effect is beneficial, at least in humans – it allows various heart conditions to be easily diagnosed – but it is because of the first effect, and not the second, that organisms with hearts were able to reproduce themselves successfully in the past. Thus the circulation of blood, rather than the making of rhythmic noises, is the function of the heart.

A related issue concerns the viability of *adaptationism* as a research strategy in evolutionary biology. This strategy is an inference from the observation that a trait is capable of serving some purpose useful to the organism, to the historical claim that the trait has been naturally selected *because* it served that purpose. This strategy obviously has its limitations: Dr. Pangloss, a character in Voltaire’s novel *Candide*, claimed that the purpose of the human nose was to hold up spectacles. The biologists Stephen Jay Gould and Richard Lewontin have argued that many of an organism’s physical traits may arise as byproducts of developmental constraints on its basic body plan. The human chin is a standard example: there is just no way to get a larynx, an esophagus, a tongue capable of speech, and a jaw strong enough for chewing without getting a chin thrown in for free. (Or at any rate, this is impossible without a major overhaul in human architecture, an overhaul that just wouldn’t be worth the effort from an evolutionary point of view). James Woodward and Fiona Cowie, in chapter 16, criticize the field of *evolutionary psychology* partly on the grounds that it is committed to an implausible form of adaptationism.

Recently, there has been increasing philosophical interest in other areas of biology, such as genetics and molecular biology. One issue concerns the relationship between these two fields. Within classical genetics, as originally formulated by Gregor Mendel in the nineteenth century, the gene is the basic unit of explanation. After the discovery of the structure of DNA by Watson and Crick in 1953, it has become possible to explore the internal structure of genes. This raises the question of whether classical genetics can survive as an autonomous branch of biology, or whether it “reduces” to molecular biology in the sense that all of its basic concepts and principles can be understood in terms of the concepts and principles of molecular biology. One problem

concerns the difficulty of saying just exactly what it means for one branch of science to “reduce” to another. Issues about reduction arise within many other branches of science. Philosophy of chemistry, for example, is a very new field, and one of its most fundamental questions is whether chemistry reduces to physics.

Another issue concerns the extent to which genes determine the phenotypic features of an organism. All researchers agree that both genetic and environmental factors play some role here; and one or the other may play a stronger role in different traits. Nonetheless, some have argued that genes have a distinctive, if not decisive, role to play. This is sometimes expressed in metaphors such as “genetic coding” or “genetic information.” The idea is that genes encode information for phenotypic traits in much the same way that strings of dots and dashes code for words in Morse code. (In genetics, the “words” would not be made up of dots and dashes, but of nucleotides.) In chapters 13 and 14, Sahotra Sarkar and Peter Godfrey-Smith debate the utility of this metaphor.

0.3.4 Psychology

The nature of the human mind has long been a concern of philosophers. René Descartes argued in the seventeenth century that the mind and the body are distinct. In fact, they are made of entirely different substances. This created a problem for Descartes, for he also held that physical matter could only act or be acted upon by contact with other physical matter. How, then is it possible for the physical world to affect the mind, as it does when we form beliefs about the world upon the basis of observation? How is it possible for the mind to affect the physical world, as it does when we form volitions that result in the motions of our bodies? This is the infamous *mind-body* problem. There is, within philosophy, a sub-field referred to as the philosophy of mind that deals with this and other problems about the nature of the mind. When the exploration of these issues makes contact with empirical psychology, we enter the realm of philosophy of psychology and hence, philosophy of science.

Just within the last quarter century, there have been extraordinary advances in neuroscience, due, in part, to technological advances that permit various forms of brain imaging. The empirical exploration of the neural activity that takes place as human subjects engage in various mental tasks would seem to have an obvious bearing on the mind-body problem. If there is an interface between the mind and the body, it is surely in the brain. (Descartes himself believed that it took place in the pineal gland.) It is unclear, however, whether the new understanding of the brain is giving us anything more than “more body” – more knowledge of the physical world – and whether it has got us any closer to making the jump across the mind-body gap. Nonetheless, neuroscience has taught us a great deal about the mind that is of considerable philosophical interest. For example, it has taught us about the role of the emotions in “rational” decision-making, and about the ways in which our conceptions of ourselves may be disrupted. The newly emerging field of *neurophilosophy* explores these connections between philosophy and neuroscience.

Earlier in the twentieth century, a new interdisciplinary field called *cognitive science* emerged out of the disciplines of psychology, philosophy, computer science, and linguistics. This field was driven, in part, by a general admiration for advances made in

computer science. A computer is a physical entity made up of a vast network of transistors, made of doped silicon and embedded in chips, all housed in a container of metal and plastic, and requiring electricity to function. In theory, it may be possible to understand specific operations performed by a computer at this nuts and bolts level. But the operations of a computer can also be understood at a more abstract level, in terms of the *programs* that the computer executes. The guiding idea behind cognitive science was that the mind is to the brain as computer software is to hardware. (The brain is sometimes described as “wetware.”) It should be possible, then, to understand the operations of the mind at a more abstract computational level. This way of thinking about the human mind is closely associated with *functionalism*, the view that a particular mental state (such as the belief that it will rain today) is to be identified in terms of its role within the overall program.

Even more so than in other branches of science, it is often hard to tell where the philosophy of psychology ends, and psychology proper begins. This is, in part, because within academic psychology there is a very strong emphasis on experimentation and data-collecting. Much the same is true within neuroscience. While individual psychologists test and defend hypotheses about particular mental processes, there is a dearth of higher-level theory to unify and explain all of the various empirical findings. To a large extent, philosophers have filled the gap, effectively playing the role of theoretical psychologists.

One potential unifying theory within psychology is *evolutionary psychology*. According to this view, the mind consists of a large number of specialized modules. A module performs some one specific task, and does so in relative isolation from what is going on in the rest of the brain. The visual system seems to fit this description: it deals specifically with the interpretation of information gathered by the retina, and does so largely uninfluenced by data from, for example, other sensory systems. More controversially, a host of other tasks are said to be performed by modules. Some evolutionary psychologists maintain, for example, that we are equipped with a “cheater-detection” module, to identify those people who are exploiting the norms of social interaction in order to obtain an unfair advantage. These modules evolved in order to solve particular problems faced by our ancestors in an environment quite different from our own. This evolutionary perspective is thought to provide us with a useful explanatory framework for thinking about various mental traits. For instance, arachnophobia, while a nuisance now, may have been advantageous in an environment in which spider bites posed a genuine risk to life and limb. In chapter 15, Peter Carruthers defends this picture of the human mind, while James Woodward and Fiona Cowie (chapter 16) are highly skeptical of both the methods and the conclusions of evolutionary psychology.

0.3.5 The social sciences

“Social science” is a broad term that encompasses a number of different fields, especially sociology and economics, but also including parts of political science, anthropology, and linguistics. One central question concerns whether the social sciences are genuinely scientific. By and large, the social sciences do not make the sorts of precise predictions that can be cleanly checked against the results of observation. There are

two principal reasons for this. The first is the sheer complexity of the systems studied by social scientists. Particle physics requires a tremendous amount of training and mathematical background, to be sure, but when one is attempting to explain and predict the behavior of single (perhaps even indivisible) particles, it is not so surprising that it is sometimes possible to do so with great precision. Even the simplest organic compounds are substantially more complex than that, a simple living organism considerably more complex than that, and a human brain vastly more complex still; social networks and institutions consisting of large numbers of human beings . . . well, we have come a long, long way from sub-atomic particles. Social scientists study complex systems whose basic units literally have minds of their own.

A second, related reason why prediction is so difficult in the social sciences is that it is so difficult to find or create social systems where all but a handful of factors can be ignored. It is not surprising that astronomy was the first successful predictive science: the objects under study are, for all intents and purposes, affected only by gravity. To be sure, the earth is subject to light pressure from the sun, it has a magnetic field and interacts with charged particles floating around in space, and so on. But these further factors have such a tiny effect on the earth's motion that they can safely be ignored. In experimental physics, it is possible to isolate a small number of significant forces by carefully shielding the experimental system from unwanted sources of influence. By contrast, even when social scientists are able to identify a few of the most important factors that affect the evolution of an economy, social institution, or cultural practice, there are simply too many potentially disrupting forces to be anticipated. Natural disasters, foreign wars, epidemics, political crises, technological advances, and even personal idiosyncrasies can and typically do derail even the best understood social processes. Moreover, in the social sciences, practical and ethical considerations typically prohibit the sort of artificial shielding that takes place in experimental physics. These two factors combine to make precise predictions (as opposed to predictions about general trends) virtually impossible in most areas of social science. (Of course, some physical systems, such as the earth's climate, are also enormously complex, and we are not very good at making predictions about them either. And some areas of the social sciences – for example, microeconomics – do at least sometimes deal with systems that are on a manageable scale.)

A related challenge to the scientific status of the social sciences claims that science aims at the discovery of *laws*, and that there can be no genuine laws of social science. In the third quarter of the twentieth century, especially, it was believed that laws were essential for both explanation and confirmation by evidence. In their contributions to this volume (chapters 7 and 8), John Roberts and Harold Kincaid address the issue of whether there can be social-scientific laws. Kincaid claims that the social sciences can have laws that have much the same character as (at least some of) the laws of physics. Roberts argues that there are no laws of social science, and hence that the social sciences are fundamentally different from physics, although he does not go so far as to deny that the social sciences are scientific.

A different sort of challenge to the status of the social sciences emerges from the claim that the social sciences are largely involved in *interpretation*. Anthropologists, in particular, are often interested in the symbolism involved in certain social rituals, the hidden meaning that lies behind the nominal purpose of the activity. This involves

interpreting a culture's practices, much as a psychoanalyst might interpret a dream in order to uncover its hidden significance. Kincaid addresses this challenge in chapter 8.

The field of economics raises a different sort of worry. That field's basic principles are to a large extent *a priori*. These principles lay down rules for how an individual or firm *ought* to behave, on the assumption that they are interested in maximizing their own well-being. (In the case of firms, well-being is effectively identified with financial profit; in the case of human beings, there is at least some lip-service paid to the idea that money isn't everything.) These rules are then used to determine how such individuals and firms will behave in the marketplace, the prices at which they will be willing to buy and sell goods, the financial risks they will be willing to undertake, and so on. Predictions derived from these principles are then applied to actual economic situations, despite the fact that they are ultimately predicated on purely *a priori* assumptions about the nature of economic agents. This approach raises questions about the epistemology of economics: in particular, it seems to clash with the empiricist doctrine that our knowledge of the world ultimately stems from our experience. Recently, however, there has been an increasing interest in behavioral and experimental economics, which attempt to gather empirical evidence concerning the economic behavior of actual agents; indeed, the 2002 Nobel Memorial Prize in economics was awarded to two pioneers in this area.

In addition to these issues concerning the scientific status of the social sciences, the fields of economics and political science overlap with and influence areas of ethics and social and political philosophy. For example, ethical questions about the most just way of distributing goods within a society cannot be completely divorced from economic questions about the effects of distributing goods in different ways. It may be, for example, that distributing goods on the basis of productivity (rather than simply dividing them equally, for example) can produce a system of incentives for productivity that benefits all, or at least most, members of that society.

0.4 Conclusion

While this brief overview has only scratched the surface, I hope that it provides the reader with some sense of the sorts of issues that comprise the philosophy of science. Moreover, it should help the reader better understand how the eight topics covered by the contributions to this volume fit into the larger picture.

Further reading

All of the chapters in this volume provide references to the most important works on the specific topics covered. Here, I will offer some suggestions for those interested in reading in more detail about some of the other topics covered in this introduction.

There are a number of good anthologies containing some of the most significant contributions to the philosophy of science; among the best are Balashov and Rosenberg (2002), Curd and Cover (1998), Klemke et al. (1998), and Schick (1999).

Westview Press publishes a series of excellent introductory texts, written by leaders in their respective fields. These include *Philosophies of Science: Feminist Theories* (Duran, 1997), *Philosophy of Biology* (Sober, 1999), *Philosophy of Physics* (Sklar, 1992), and *Philosophy of Social Science* (Rosenberg, 1995). Blackwell is the leading publisher of reference works in philosophy: the two most directly relevant are *The Blackwell Guide to the Philosophy of Science* (Machamer and Silberstein, 2002) and *A Companion to the Philosophy of Science* (Newton-Smith, 2000). As with any topic, there is a cornucopia of material on philosophy of science available on the world wide web, much of it unreliable. One outstanding resource is the *Stanford Encyclopedia of Philosophy*, at <http://plato.stanford.edu> (Zalta, 1995–). Still expanding, this online encyclopedia already has a very strong collection of entries in the philosophy of science, all of them written by recognized experts.

Cohen and Wellman (forthcoming) is another volume in the Blackwell *Contemporary Debates* series, dealing with issues in applied ethics. This volume will contain pairs of chapters on cloning and on animal rights; the paired essays will be in much the same format as the contributions to this volume. Munson (1999) is an excellent collection of some of the most important works in bioethics, with extensive introductions and discussions by the editor. This anthology is now in its sixth edition and is updated regularly. (The fourth edition (1991) is now available in mass market paperback form.) Blackwell has several reference works in bioethics, including Kuhse and Singer (2001), and Frey and Wellman (2003). Kitcher (2001) is an engaging discussion of the place of science in a democratic society, including an evaluation various strategies for the funding of scientific research.

Cahn (2002) is an excellent anthology of the works of the great philosophers. It contains, *inter alia*, excerpts from the works of Plato, Descartes, Spinoza, Leibniz, Locke, Berkeley, and Hume, including some of their most important works in epistemology. It is currently in its sixth edition, but earlier editions will have these works as well. Salmon (1990) provides an excellent overview of work on scientific explanation in the twentieth century.

Shapiro (2000) is a good overview of central issues in the philosophy of mathematics. Albert (1992) is an engagingly written book that explains the basic structure of quantum mechanics as painlessly as possible, and provides a critical appraisal of various attempts to solve the measurement problem. Sklar (1974), almost encyclopedic in scope, provides an excellent overview of problems concerning the nature of space and time. For those with a background in general relativity, Earman (1995) provides a rigorous discussion of philosophical issues in cosmology.

Sober (1994) contains an excellent selection of important papers in the philosophy of biology, including papers on the units of selection, functions, adaptionism and the reduction of genetics to molecular biology. Sterelny and Griffiths (1999) provides a good introduction to some of the major issues in the philosophy of biology.

Cahn (2002) includes many of the classic discussions of the mind-body problem, including Descartes's *Meditations*. Churchland (1996) is a highly readable discussion of philosophical issues involving neuroscience; Bechtel et al. (2001) collects a number of important papers on the subject. Clark (2001) is a nice introduction to philosophical issues in cognitive science, while Haugeland (1997) contains a number of classic papers on the subject. Blackwell also publishes *A Companion to Cognitive Science*

(Bechtel and Graham, 1998). For the philosophy of social science, Martin and McIntyre (1994) contains some of the most important works.

Bibliography

- Albert, D. 1992: *Quantum Mechanics and Experience*. Cambridge, MA: Harvard University Press.
- Balashov, Y. and Rosenberg, A. (eds.) 2002: *Philosophy of Science: Contemporary Readings*. London and New York: Routledge.
- Bechtel, W. and Graham, G. (eds.) 1998: *A Companion to Cognitive Science*. Oxford: Blackwell.
- Bechtel, W., Mandik, P., Mundale, J., and Stufflebeam, R. S. (eds.) 2001: *Philosophy and the Neurosciences: A Reader*. Oxford: Blackwell.
- Cahn, S. (ed.) 2002: *Classics of Western Philosophy*, 6th edn. Indianapolis: Hackett.
- Churchland, P. 1996: *The Engine of Reason, the Seat of the Soul: A Philosophical Journey into the Brain*. Cambridge, MA: The MIT Press.
- Clark, A. 2001: *Mindware: An Introduction to the Philosophy of Cognitive Science*. New York: Oxford University Press.
- Cohen, A. and Wellman, C. (eds.) forthcoming: *Contemporary Debates in Applied Ethics*. Oxford: Oxford University Press.
- Curd, M. and Cover, J. A. (eds.) 1998: *Philosophy of Science: The Central Issues*. London and New York: Norton.
- Duran, J. 1997: *Philosophies of Science: Feminist Theories*. Boulder, CO: Westview Press.
- Earman, J. 1995: *Bangs, Crunches, Whimpers, and Shrieks*. Oxford: Oxford University Press.
- Frey, R. and Wellman, C. (eds.) 2003: *A Companion to Applied Ethics*. Oxford: Blackwell.
- Haugeland, J. 1997: *Mind Design II: Philosophy, Psychology, and Artificial Intelligence*. Cambridge, MA: The MIT Press.
- Kitcher, P. 2001: *Science, Truth, and Democracy*. Oxford: Oxford University Press.
- Klemke, E. D., Hollinger, R., Rudge, D. W., and Kline, D. (eds.) 1998: *Introductory Readings in the Philosophy of Science*. New York: Prometheus.
- Kuhse, H. and Singer, P. (eds.) 2001: *A Companion to Bioethics*. Oxford: Blackwell.
- Machamer, P. and Silberstein, M. (eds.) 2002: *Blackwell Guide to the Philosophy of Science*. Oxford: Blackwell.
- Martin, M. and McIntyre, L. (eds.) 1994: *Readings in the Philosophy of Social Science*. Cambridge, MA: The MIT Press.
- Munson, R. (ed.) 1999: *Intervention and Reflection*, 6th edn. New York: Wadsworth.
- Newton-Smith, W. (ed.) 2000: *A Companion to the Philosophy of Science*. Oxford: Blackwell.
- Rosenberg, A. 1995: *Philosophy of Social Science*. Boulder, CO: Westview Press.
- Salmon, W. 1990: *Four Decades of Scientific Explanation*. Minneapolis: University of Minnesota Press.
- Shapiro, S. 2000: *Thinking about Mathematics*. Oxford: Oxford University Press.
- Schick, T. (ed.) 1999: *Readings in the Philosophy of Science*. Mountain View, CA: Mayfield.
- Sklar, L. 1992: *Philosophy of Physics*. Boulder, CO: Westview Press.
- 1974: *Space, Time, and Spacetime*. Los Angeles and Berkeley: University of California Press.
- Sober, E. 1999: *Philosophy of Biology*, 2nd edn. Boulder, CO: Westview Press.
- (ed.) 1994: *Conceptual Issues in Evolutionary Biology*, 2nd edn. Cambridge, MA: The MIT Press.
- Sterelny, K. and Griffiths, P. E. 1999: *Sex and Death: An Introduction to Philosophy of Biology*. Chicago: The University of Chicago Press.
- Zalta, E. 1995–: *Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu>