When Science Studies Religion: Six Philosophy Lessons for Science Classes

Massimo Pigliucci

Science & Education
Contributions from History, Philosophy and Sociology of Science and Mathematics

ISSN 0926-7220
Volume 22
Number 1

Sci & Educ (2013) 22:49-67
DOI 10.1007/s11191-011-9355-2

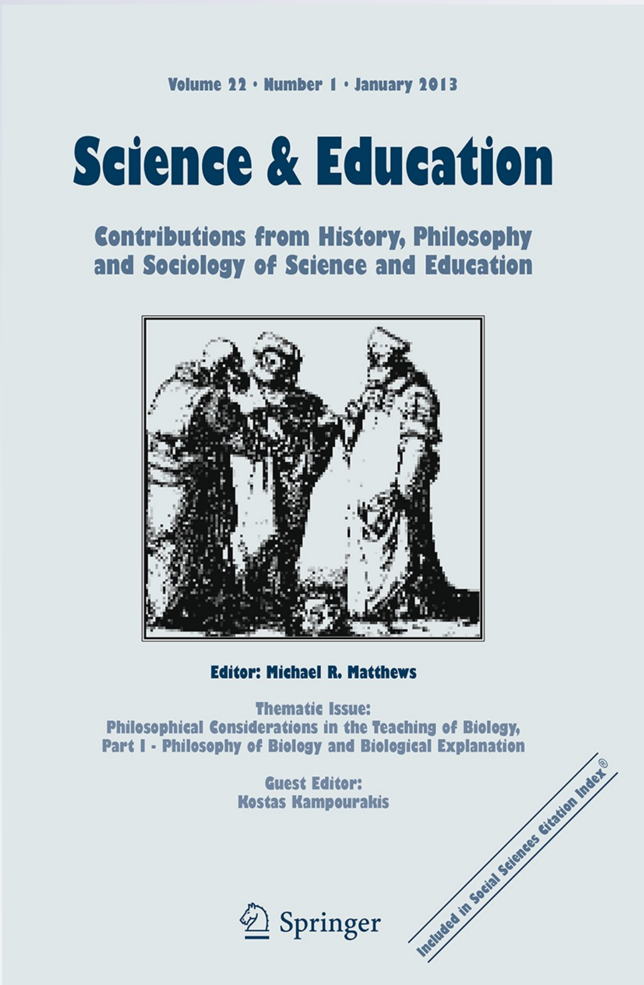

Springer

Sci & Educ (2013) 22:49–67
DOI 10.1007/s11191-011-9355-2

When Science Studies Religion: Six Philosophy Lessons for Science Classes

Massimo Pigliucci

Published online: 17 April 2011

Abstract It is an unfortunate fact of academic life that there is a sharp divide between science and philosophy, with scientists often being openly dismissive of philosophy, and philosophers being equally contemptuous of the naiveté of scientists when it comes to the philosophical underpinnings of their own discipline. In this paper I explore the possibility of reducing the distance between the two sides by introducing science students to some interesting philosophical aspects of research in evolutionary biology, using biological theories of the origin of religion as an example. I show that philosophy is both a discipline in its own right as well as one that has interesting implications for the understanding and practice of science. While the goal is certainly not to turn science students into philosophers, the idea is that both disciplines cannot but benefit from a mutual dialogue that starts as soon as possible, in the classroom.

1 Introduction: Science and Philosophy, a Somewhat Troubled Relationship

"The insights of philosophers have occasionally benefited physicists, but generally in a negative fashion—by protecting them from the preconceptions of other philosophers… Philosophy of science at its best seems to me a pleasing gloss on the history and discoveries of science." So wrote Nobel physicist Steven Weinberg in a chapter of his *Dreams of a Final Theory* provocatively entitled "Against Philosophy" (Weinberg 1994). This is a startling, and yet not particularly uncommon, take on the relationship between science and philosophy, especially popular among scientists. Philosophers, on their part, often retort with equally strong rhetoric. Consider, for instance, Daniel Dennett's famous contention, expressed in *Darwin's Dangerous Idea* (Dennett 1996) that "There is no such thing as philosophy-free science; there is only science whose philosophical baggage is taken on board without examination."

This antagonism, I maintain, does no good to either science or philosophy, and perhaps the best way to eventually overcome it is by exposing science students to the philosophical

M. Pigliucci (✉)
The Graduate Center, City University of New York, New York, NY, USA
e-mail: massimo@platofootnote.org

concepts that are already embedded—often invisibly to the untrained mind—in what they study. In this paper I will use a particular example, biology-based approaches to the understanding of religion, as a case study in how philosophical lessons can seamlessly be introduced into the teaching of scientific methods and hypotheses. I chose a particular class of scientific explanations for the phenomenon of religion because they represent a quintessential example of the complex entanglement of science and humanities—where the very idea of science studying religion is inherently controversial. Adding a philosophical twist to it will show, I think, how such entanglement can actually be productive and conducive to a more sophisticated type of science education.

Of course, there are plenty of other fields of inquiry that could have served the same purpose, especially any subject matter that falls within the general purview of evolutionary psychology, the currently popular approach to the biological study of human behavior. But religion is, I think, a particularly apt choice for a variety of reasons. First, there is a broad and increasingly popular literature on the topic, which makes it a good test study for the issues I am concerned with. Second, the phenomenon itself—religion, that is—is obviously of paramount importance for human societies, and it can therefore hardly fail to stimulate the interest of students. Third, the very same conceptual and philosophical issues that arise in the biological study of religion also characterize the scientific study of pretty much any other complex human behavior, from our mating habits to our propensity for aggression toward outsiders, so that the study of religion can be a useful paradigm to understand both the insights and the pitfalls of a scientific approach to complex human behaviors.

Before proceeding, however, I need to make clear that throughout this paper to study religion scientifically means to study the biological underpinnings of religious feelings and beliefs. This is not at all the same as scientifically "demonstrating" that there is no God, contra to much recent literature taking that approach (e.g., Dawkins 2006). God is far too ill-defined a concept to serve as a scientific hypothesis, so the very idea of scientifically disproving the existence of the supernatural is philosophically hard to defend. Our feelings and beliefs, on the other hand, no matter what the ontological status of their referent (real or imaginary, physical or not), simply have to manifest themselves through our brains, which means that they have to be biologically grounded, and can therefore be studied by the biological sciences. Let us keep this distinction in mind throughout the following, so not to be distracted by unproductive controversies.

2 Lesson One: The Difference Between Proximate and Ultimate Explanations

Religion is, of course, an exceedingly complex phenomenon, which if it is to be understood in naturalistic terms, requires several layers of explanation, calling on disciplines ranging from history to cultural anthropology (on the humanities side) to neurobiology and evolutionary biology (on the science side). In this section I will provide a minimum background concerning the latter two approaches (i.e., including neurobiology), while the rest of the paper will focus mostly on evolutionary biological accounts of the origin and spread of religion.

Religious belief, like any other kind of belief, requires a functional brain. Research in neurobiology, however, has begun to reveal what exactly about the human brain makes people prone to certain aspects of religious belief. This ranges from the actual experience of "mystical" states to the interpretation of patterns of regularity in the world as indications of cosmic design via our natural inclination to project agency on the world itself. For instance, Boyer (2008) convincingly argues that a series of human cognitive traits make people more prone to religious belief. These traits include a tendency to think of gods in anthropomorphic terms (regardless of one's overt religious faith and its precepts); to attribute human-like workings to the minds of gods; to form relationships with imaginary characters (a pre-requisite for forming a relationship with gods); to engage in ritualized behavior, particularly of a "cleansing" nature (often both in the literal sense of washing one's hand or body and in the metaphorical sense of purifying one's soul); and to adopt social norms not because they are intrinsically useful or rational, but solely because they are social.

In terms of actual religious-type experiences, research by Blanke et al. (2002), among others, clearly shows that stimulation of particular areas of the brain—such as the right angular gyrus—similar to what happens during certain epileptic episodes, produces a sense of out of bodyness or of body transformation that are often associated with mystical episodes. Perhaps even more interestingly, it has been known for some time that the same type of experience can be produced by standard religious practices, such as deep prayer or meditation (Newberg et al. 2002). Of course, neither these neurobiological findings nor the above mentioned conclusions from cognitive science show that there is no realm of the supernatural, since even if such a realm existed, human beings presumably would still have to use their biological brains to somehow access it. Nonetheless, it is increasingly hard to avoid the conclusion that the science of the brain is giving us a satisfying naturalistic account of how people believe.

What about *why* we believe? This is where evolutionary biology comes in (e.g., Wilson 2002; Norenzayan and Shariff 2008; Pyysiainen and Hauser 2009; Richerson and Newson 2009; but see Schloss 2009). As we shall see more in detail throughout the rest of this paper, there are fundamentally four types of evolutionary explanations for the origin and spread of religion: what I shall call "classical" adaptive scenarios (i.e., involving individual-level selection); group-selection scenarios (Okasha 2006); by-product (or "Spandrel-type": Gould and Lewontin 1979; Pigliucci and Kaplan 2000) explanations; and memetics. Briefly, the first type of explanation seeks to account for religious belief in terms of a human trait that favors individuals that display it, by augmenting their chances of survival and/or reproduction. The second account uses the same logic—religion as an adaptation—but in this case the objects of selection are groups of humans, not individuals per se. The third explanation understands religion as (at least initially) an accidental by-product of human cognitive traits that evolved for different reasons, with the possibility that religious belief has later been co-opted to become what Gould and Vrba (1982) call an "exaptation." Finally, memetic explanations (whose actual ontological and epistemological status, as we shall see, is very much in doubt) seek to understand religion in terms of hypothetical cultural replicators (memes) analogous to, but mechanistically independent of (and, indeed, in direct competition with), the standard genetic ones (Blackmore 2000).

The contrast between neurobiological-cognitive science and evolutionary explanations of religion introduces the students to a first fundamental philosophical concept: the distinction between proximate and ultimate causes. This distinction was of course first realized by Aristotle. In both his Physics (II, 3) and Metaphysics (V, 2) Aristotle distinguished among four types of causes: material (what something is made of; like the bronze of a statue); formal (what it is to be that something; like the shape of the statue); efficient (the source of the thing; like the sculptor who made the statue); and final (the reason why the thing exists; as in the statue having been sculpted in order to honor a major philosopher) (Aristotle 1998). I have argued that Aristotle's four causes still have a role to play in the understanding of modern biological phenomena (Pigliucci 2003), but the importance of the simpler distinction between proximate and ultimate explanations (respectively Aristotle's material-efficient and final causes) has been stressed ever since Mayr (1961; but see Ariew 2003, who clarifies several problems with Mayr's original treatment, but maintains that the distinction is useful if reframed in terms of dynamic versus statistical explanations in modern biology).

Mayr was concerned with the recurring appearance in biology of so-called teleological concepts, i.e. of suggestions that evolution is driven by internal "vitalistic" mechanisms somehow "aiming" at a particular end result (typically, the "clearly superior" *Homo sapiens*). But as Mayr pointed out, the Aristotelian idea of final cause is perfectly compatible with modern science, once we realize that "final" means whatever mechanism explains why things happen, where "why" does not imply some kind of intelligence lurking behind the process (pace proponents of so-called Intelligent Design: see Pigliucci 2002). Indeed, the Darwinian process of natural selection provides a perfectly naturalistic explanation of "why," for instance, some living organisms have eyes, by operationalizing the concept of function in biological terms, asking what biological structures are "for" (Kaplan and Pigliucci 2001). Accordingly, we can distinguish between teleological explanations based on design etiologies (such as those advanced by ID proponents) and teleological explanations based on consequence etiologies which are absolutely compatible with evolutionary theory (Depew 2008).

To return to our running example, then, neurobiological and evolutionary explanations of religion are not in competition with each other. Rather, they represent respectively proximate and ultimate causes (the latter understood in the modern sense described above). So, for instance, when we say that certain structures in the brain, like the right angular gyrus, are involved in the production of "mystical" experiences, we are considering a type of proximate cause for that phenomenon. When we say that those same experiences are triggered by certain conditions (like deep meditation) we are providing another type of proximate cause. And if we suggest that the reason these experiences happen is that they are a by-product of the evolution of particular brain structures and functions, we are then addressing their ultimate cause.

3 Lesson Two: Doing Science by Multiple Hypotheses

T. C. Chamberlin, then President of the Society of Western Naturalists, published what is now a classic paper in *Science* magazine in 1890, by the title "The method of multiple working hypotheses" (Chamberlin 1890). In it the author expressed concern for the tendency of scientists to begin investigations of a given topic by focusing on one leading explanation of whatever phenomenon is under study. Here is how Chamberlin expressed his worries:

> The moment one has offered an original explanation for a phenomenon which seems satisfactory, that moment affection for his intellectual child springs into existence; and as the explanation grows into a definite theory, his parental affections cluster about his intellectual offspring… There is [then] an unconscious selection and magnifying of phenomena that fall into harmony with the theory and support it, and an unconscious neglect of those that fail of coincidence. The mind lingers with pleasure upon the facts that fall happily into the embrace of the theory, and feels a natural coldness toward those that seem refractory. (Chamberlin 1890, p. 93)

Chamberlin, in other words, was worried about an all too human tendency—affecting scientists, students and lay people alike—to fall in love with a particular explanation because it "feels" right, and to unconsciously engage in one of the most common cognitive biases (and logical fallacies): the selective bias toward remembering "hits" (favorable evidence) and forgetting "misses" (unfavorable evidence), which, incidentally, is one of the prevalent reasons so many people fall for pseudoscientific notions, such as psychic readings.

Not long after Chamberlin's paper, a statistical version of precisely this problematic approach to scientific investigation became codified in the concept of so-called "null hypotheses," which still today are considered a fundamental tool of scientific analysis, despite severe criticisms raised against the practice on methodological (e.g., Berger and Sellke 1987; Cohen 1994; Nix and Barnette 1998) and philosophical grounds (Pigliucci and Kaplan 2006, ch. 10). Yet Chamberlin was in an important sense prescient, as more modern scholarship—again both in statistics and in philosophy of science—has finally began to favor a more formalized version of his "method of multiple working hypotheses," particularly with the onset of Bayesianism (Corfield and Williamson 2001) and of likelihood-based model selection methods (Burnham and Anderson 2002).

In the case study we are examining, we are confronted with four hypotheses concerning the evolution of religion, though it would probably be more appropriate to think of them as four classes of explanations, probably each of which actually contains multiple sub-hypotheses: individual selection, group selection, by-product, and memetics. Let us therefore take a closer look at these four possibilities, an examination that will lay the bases for some of the other philosophy-inspired "lessons" I will highlight in the following sections.

3.1 Individual-Level Selection

Natural selection, the fundamental Darwinian mechanism that is thought by biologists to be the only one responsible for adaptation, has of course been the focus of research even since Darwin (Darwin 1859; Williams 1992). The nature of adaptive explanations requires that one produces hypotheses about what links the characteristic under investigation (in this case, religion) and the survival and/or reproduction (i.e., the fitness) of the entity that is thought to benefit from that characteristic (individual human beings in this case, though the same logic applies to the case of groups, as I shall discuss in a moment). Several possibilities have indeed being proposed so far, which include the idea that religious belief, particularly belief in an afterlife, reduces anxiety and therefore helps managing stress in individuals belonging to a species capable of understanding and reflecting on its own mortality (Wilson 2007). Another scenario sees religious commitment as a behavior attractive to prospective mates because it signals the willingness to pledge significant amounts of resources, a human equivalent of the bowerbird's intricate nest (Wilson 2007).

However, it is relatively easy to level criticisms against these and similar individual-based selection hypotheses. For instance, one could object that other primates seem to find much less costly ways to reduce stress (e.g., bonobo chimpanzees engage in a variety of sexual practices: Manson et al. 1997), and that it isn't clear why a completely arbitrary belief should make someone feel better about dying (unless that belief is the by-product of other human cognitive characteristics, as we will see below). To the religion-as-mark-of-commitment hypothesis it can be reasonably objected that humans could (and in fact do) behave just like other species, where the female is attracted either by purely aesthetic traits (peacock's tail) or by behaviors that not only signal commitment but are specifically addressed to the female (as in the case of metabolically or time expensive courtships). Nonetheless, in terms of teaching the philosophical underpinnings of the method of multiple working hypotheses, both individual selection per se (when considered against the other three classes of explanation), and sub-hypotheses relying on individual selection scenarios (played against each other) do the job nicely.

3.2 Group-Level Selection

The second class of explanations for the evolution of religion relies on the often controversial concept of group selection (Williams 1992; Sober and Wilson 1999; Okasha 2006). Contrary to early discussions of group selection during the 1960s, it is now generally accepted that group-level selection is a coherent concept corresponding to a theoretically possible evolutionary mechanism. The questions concern how widespread the process actually is in nature, and how one may go in order to study it, given that it is often difficult to control experimental conditions in a manner that makes it possible to distinguish quantitatively between the effects of individual and group selection.

From a philosophical point of view additional interesting issues arise when one considers the ontological status of "groups," and in particular how many levels of group selection may in fact operate. For instance, Okasha (2006) proposes cogent reasoning leading to the conclusion that groups of kins, groups of non-kins, as well as species are theoretically legitimate potential targets of selection, but that clades—open-ended groups of species sharing their most recent common ancestor—do not have an ontological status independent from that of species, and cannot therefore be subject to selection, contrary to what sometimes affirmed in the paleontological literature (e.g., Hone et al. 2005).

As in the case of individual-level selective explanations, so for group selection there are several sub-hypotheses that can be considered as alternatives. For instance, religion could have evolved as a group-level adaptation because it augments individuals' propensity for prosocial behavior (something that is in fact empirically substantiated: Norenzayan and Shariff 2008), and presumably therefore increases group persistence—though it should be noted that the latter is not a sufficient condition for group selection to operate, as the group has to somehow "replicate" while inheriting traits that are relevant to its survival and reproduction, in direct analogy to individual-level selection. Another possibility is that religion evolved because it makes it possible for individuals to go as far as sacrificing their life on behalf of the group, a behavior that standard individual selection would not allow to evolve except to help one's close kins (and it has to be noticed that some theoretical biologists and philosophers of science do consider kin selection a form of group, not individual, selective process).

These adaptive scenarios can be criticized in a manner similar to what we have seen above, either by explaining the same behaviors in a non-adaptive manner (by-product hypothesis), or by questioning whether the proposed scenario is sufficient or congruent with the somewhat more stringent requirements that group selection imposes when compared to lower level (individual, gene) selective processes. For instance, again, there are examples of social primates where prosocial behavior evolved by other means, as in reports of other primates putting their life at risk on behalf of non-relatives (Flack and de Waal 2000), clearly without the aid of religious beliefs to prompt them to do so.

3.3 By-Product Hypothesis

Ever since Darwin (1859) biologists have understood that not every biological structure or behavior is the result of adaptive processes. During the shaping of the so-called Modern Synthesis of the 1930s and '40s, still today the standard model of evolutionary theory (Huxley 1942/2010), scientists have acknowledged a plurality of evolutionary mechanisms, including entirely stochastic processes such as genetic drift. More recently, Gould and Lewontin's (1979) and Gould and Vrba's (1982) work has helped usher in a more nuanced view that takes into account that certain characters may appear as a by-product of developmental or other types of constraints always operating on living organisms (spandrels), and that in some cases currently advantageous characters may have began as non-adaptive by-products and later been co-opted by natural selection (exaptations).

In what sense, then, could religion be a by-product of other human cognitive characteristics? Several possibilities have been proposed (Pyysiainen and Hauser 2009), and unlike the cases seen so far, they are usually meant to work in concert (i.e., the by-product hypothesis is made of several sub-components). First, human beings have an ability, shared with several other animal species, to detect patterns of regularity in the external world, an obviously advantageous behavior in terms of survival and reproduction (localization of food, anticipation of threats from predators, etc.). Second, we have a strong propensity to interpret detected patterns in causal terms. Indeed, inferring causation from correlation is such a common type of mistake that the corresponding informal logical fallacy has long been studied by both philosophers and cognitive scientists: it is known as post hoc, ergo propter hoc (after that, therefore because of that). Clearly, inferring causality from correlational evidence is often reasonable, but human beings have been shown to persist in that sort of behavior longer than, say, rats, with a corresponding sharper decrease in returns for their stubborn behavior (Gazzaniga 2003; the idea, of course, is not that rats are aware of the causal inferences they make, but that their behavior is consistent with an assumption of causal inference). Third, human beings have an innate ability and propensity to engage in what philosophers of mind call "agency projection," i.e. in attributing mindfulness not only to other human beings, but often to inanimate and even imaginary objects. Recent research on so-called mirror neurons (Rizzolatti 2005) has began to pinpoint the neurobiological bases (i.e., the proximate causes) of this behavior, while an evolutionary explanation in terms of an enhanced ability to interact socially is the obvious candidate.

The idea, then, is that religion emerged over the course of human history as a natural by-product of these three tendencies: putting together pattern detection, inference of causality, and agency projection one gets the result that gods (imaginary causal agents characterized by mindfulness) are prima facie reasonable explanations for the regularities of complex natural phenomena—from the movement of the planets to the existence of earthquakes, volcanoes and lightning. When pondering this type of explanation, of course, we should consider that the currently widespread idea of a monotheistic religion is in fact a very late comer in cultural evolution, and that most ancient religious beliefs were pantheistic or animistic in nature, consistent with a generic attribution of divinity to the natural world.

3.4 Memetics

We finally come to a brief examination of memetic explanations for religion, though I will come back to this particular scenario below, within the contexts of the ontology of unobservables and of reductionism. The concept of "meme" was famously introduced by Richard Dawkins at the end of *The Selfish Gene*, which popularized ideas about gene-level and kin selection proposed by W. D. Hamilton and G. C. Williams (Dawkins 1976). Dawkins was looking for a generalized version of Darwin's ideas (so-called "universal" Darwinism), and proposed that even non-biological entities would evolve by natural selection, if the proper conditions were met. His example was the evolution of ideas within cultures, and he coined the term "meme" in direct analogy with gene, to indicate a unit of cultural inheritance.

In terms of our current discussion, then, memeticists think of religion as a meme (or, rather, as a set of memes, sometimes referred to as a "memeplex"), which originated by "mutation" of some other (unspecified) idea, and spread by imitation (the standard vehicle of memetic evolution) to become what Dawkins controversially referred to as a "virus of the mind." The human mind is the environment in which memes compete for resources (presumably, human attention and memory) and give rise to ever more efficient—in terms of their own survival and reproduction—new memes. According to Dawkins, and more so to Blackmore (2000), memes can actually be in indirect competition with genes themselves, the two types of "replicators" vying to control us. (If this sounds a bit far fetched and diminutive of the condition of being human, see below, and keep in mind another of Dawkins' famous metaphors, according to which we are just "lumbering robots" controlled by our genes—or memes—for their own sake.)

4 Lesson Three: The Ontological Status of "Unobservables" in Science

Discussions about memetics are an excellent way to introduce students to one of the fundamental debates in philosophy of science: the role of unobservable entities in science and, more broadly the ongoing debate between scientific realists and anti-realists. I cannot do full justice to these areas of philosophical inquiry here, but it is important that science educators consider ways to weave these discussions into the general teaching of science, both to provide students with a more nuanced understanding of the epistemic limits of science itself, and to introduce them to live debates in the philosophy of science.

"Unobservables" and their roles in scientific theory have been under scrutiny ever since the logical positivist attack on metaphysics in the early part of the twentieth century. Of course, unobservables—which are defined as any entity or theoretical construct that is not amenable to direct empirical verification—are very common in science, examples including genes, atoms, quarks, and superstrings, to mention but a few. It is important to consider the fact that what is unobservable today may not be so tomorrow, depending on technical advances in scientific instrumentation. DNA is the clear example here: originally inferred to exist as whatever chemical basis underpinned the biological process of heredity, DNA was first described on indirect chemical grounds; then scientists were able to obtain (still indirect) visual information of its structure via crystallography; and finally we have been able to actually see individual molecules of DNA thanks to high-powered electron microscopy.

Still, it seems likely that at least some theoretical entities that play major roles in science will forever remain unobservable. Quarks, for instance, and a fortiori the even more speculative superstrings, are thought to be much smaller than the wavelength of any kind of radiation we can use to obtain "images," which means that there is no physical way to "see" these objects, even in the enlarged sense of "seeing" allowed by the use of electron microscopes. These true unobservables raise the question of how do we know that they actually correspond to real entities out there in the world, as opposed to being convenient fictional devices that allow scientists to produce coherent theories capable of making empirically verifiable predictions.

Students should not be surprised at the possibility that scientists may introduce theoretical quantities for purely pragmatic reasons—i.e., regardless of their physical existence. The history of science is of course replete with such instances, from pre-Copernican astronomy's idea of epicycles (necessary to account for the actual movement of the planets in the sky) to Einstein's cosmological constant (initially introduced to make general relativity compatible with the idea of a stationary state universe, then regretted by Einstein after the discovery that the universe is in fact expanding, and—somewhat ironically—now being revived because of the further discovery that the universe is not just expanding, but accelerating!).

The debate about the epistemic status of unobservables is part of the broader discussion between realists and anti-realists in philosophy of science. Realists claim that scientific theories aim, and often succeed, to a certain degree of approximation, at describing the world as it really is. The anti-realists counter that we have no way to determine how the world "really is," short of a God's eye-view of things, so that the most that science can claim is that it produces theoretical constructs that "work," i.e., that deliver pragmatic results when applied to certain aspects of the world.

The realism/anti-realism debate is a complex and sophisticated one (Ladyman 2002), but students can be introduced to it by means of a couple of the chief arguments proposed by either side. A major anti-realist point of contention is known as the "pessimistic meta-induction." This line of attack against realism is derived from the history of science, which teaches us that pretty much every major scientific theory that was accepted as true at some point in time eventually turned out to be false and was replaced by a rival theory. Ptolemaic astronomy ruled for centuries, only to fall to the Copernican model. But even Copernicus was wrong in assuming that the orbits of the planets are circular, so the Keplerian version of the theory quickly replaced it. When Einstein's relativity came about, however, even the Keplerian theory had to be modified because scientists now took into account relativistic effects that alter the orbits of planetary bodies in the solar system, and so on. The idea is that if past theories have so far revealed themselves to be false, on what grounds can we ever think that whatever theory currently holds sway within the scientific community is, in fact, true, i.e., provides a realist account of the world?

The obvious counter-move on behalf of realists is colorfully known as the "no miracles" argument, according to which it would be nothing short of miraculous if scientific theories got so much right in terms of empirical predictions, and yet turned out to be wrong in some fundamental sense. Again, though, the history of science presents a problem here: consider Newtonian mechanics, which is still used today to make calculations about planetary orbits and missile launches, even though it has been shown by Einstein to be fundamentally wrong. Newton's absolute conception of space and time is dramatically rejected by Einstein's ideas about curved space–time, yet Newtonian mechanics "works" most of the time because its (wrong) conception of space is good enough for most practical purposes.

But how can a theory that is wrong in the sense of not actually corresponding to physical reality nonetheless provide sufficiently accurate descriptions of, and even predictions about, events in the world? A particularly popular version of realism, structural realism, says that we cannot make claims as to the reality of specific theoretical entities, because new theories may, in fact, do entirely away with those entities, and there is no principled way to distinguish between unobservable entities that do or do not correspond to reality. However, say the structural realists (beginning with John Worrall, who proposed the idea: Worrall 1989), the mathematical structure of scientific theories—where there is one, since not all science is expressed in mathematical formulation—tends to remain stable between theories, thus contradicting the pessimistic meta-induction.

Even shifting attention from unobservables to mathematical descriptions of reality, however, is far from being a knock-down move against anti-realism. Anti-realists will point toward the well known problem of underdetermination of theory by the data, which can be understood by way of a simple example from mathematical analysis. Consider a certain number of data points plotted in an X–Y Cartesian coordinate space. Imagine that the points align themselves on a more or less straight line, so that a simple equation of the type $Y = a + bX$ (where a is the intercept and b the slope of the line) fits the data well enough for practical purposes. We can consider this equation as a structural mathematical model of the physical relationship between X and Y (whatever these two variables may actually indicate in the real world). But of course there is literally an infinite number of equations of the type $Y = a + bX + cX^2 + dX^3 + \ldots$ that will fit the data equally well. In this sense, the available data dramatically underdetermine the equations, so that we cannot tell which equation "truly" describes the data. To point out that some of these equations will be ruled out by additional data does not solve the problem, since the new data set will also have an infinite (though different) set of equations that fit it adequately.

The response of practicing scientists to all of this is that they will likely pick the simplest model compatible with the data (in this case, the original linear equation), invoking a parsimony principle along the lines of Occam's razor. But notice that this choice, while certainly defensible on pragmatic grounds (which are not a problem for anti-realists), relies on a non-empirical criterion (parsimony) and thus introduces significant and somewhat uncomfortable (to the realist) metaphysical elements in scientific theory choice. And so on the debate rages, with interesting and at least initially compelling arguments on either side.

Let us now briefly return to memes and to their status as unobservables within the context of evolutionary theories of religion (or of any other idea, for that matter). There are several serious objections that can be raised to the whole enterprise of memetics, beginning with the observation that it is not easy to make sense of the concept of "fitness" for memes, since we lack any theory of what one might call memetic functional ecology. Without that theory, statements to the effect that the most successful memes are those that spread in a population become entirely tautological: the most successful means are those that spread, and they spread because they are successful. This is a problem that standard Darwinian theory about genes and organisms avoids precisely because we have functional ecological reasons to predict which genes (or organisms) will more likely be successful in any given environmental setting, independently of direct measures of those genes' actual success. For instance, we can predict that a gene causing partial anemia as well as malaria resistance will be successful in environments with high incidence of malaria, but not in environments where malaria is not present (because the anemia excises a cost on individuals carrying that gene). We can then observationally test these predictions and see whether they actually obtain and to what degree.

A more fundamental objection to the memetic approach concerns the ontology of memes themselves: what sort of entities are they? Can we ever observe a meme? There seems to be quite a bit of confusion on these issues. For one thing, apparently anything from as simple a cultural object as a short tune to as complex an entity as "religion" qualifies as a meme. True, memeticists refer to particularly large cultural objects as memeplexes, i.e., complexes made of smaller memes. But this hardly seems to help, since we are still at a loss to say exactly what a meme is, regardless of its alleged size.

Initially a meme was vaguely identified with a given pattern of neuronal firings in the human brain. This, however, is problematic, because it borrows, without argument, a particular—and controversial—view in philosophy of mind about the correspondence between ideas and brain states. For all we know, the same idea can be the result of widely different brain states, just like in everyday life the same idea can be expressed in a variety of forms, written or not. Moreover, some memeticists maintain that inhabiting a human mind is not the only way of survival for memes, since they can reside on, say, computer hard drives, and can be copied not by imitation (from an organism to another), but digitally (Blackmore 2000). This is all well and good, but it is hard to see how memetic theory can retain any coherence if neither the form nor the means of survival and reproduction of its core hypothetical entities is stable or agreed upon by supporters of the theory.

It is definitely the case that, historically, genes—the original analog of memes—were also hypothesized as unobservables well before we discovered the chemical nature and structure of the hereditary material, and certainly well before we had electron microscopes to actually show us real pieces of DNA. But the conceptual confusion in memetics simply does not seem to resemble at all the state of genetics at the beginning of the twentieth century, when that discipline—unlike memetics, whose major academic journal recently folded—was developing a very active research program that led to a rapid succession of discoveries. At the very least, memetics is what Lakatos (1978) famously termed a "degenerate research program," an idea that, even if true, is failing to lead to new insights into how the world works, and should be abandoned on pragmatic grounds. A comprehensively damning assessment of memetics, and one with which I am actually sympathetic, was offered by Jeffrey Scholls (2009):

> It is not entirely clear how it is that positing unseen and undefined entities that infect human minds by unassessed processes involving the entities' own quest for transmission and that cause people to do things that transcend their genetic imperatives is fundamentally different from medieval demonology or, in any case, qualifies as an empirically grounded explanation in terms of natural causes. (Schloss 2009)

5 Lesson Four: The Limits of Reductionism

The preceding discussion of memetics can also be used to introduce in science classes yet another major area of philosophical inquiry: the role of reductionism in our view of the world. Broadly speaking, the issue of reduction in science concerns the extent to which theories, explanations and methods from one branch of science can be accounted for by the theories, explanations and methods of another science (Brigandt and Love 2008). The classic example of reduction is the atomic theory in chemistry, which has been successfully "reduced" to physics, where the theories and methods of physics are considered to be more "fundamental" than those of chemistry, because the latter can be derived from the first, but not the other way around.

Biology and the philosophy of biology have a long history of debates about reductionism, for instance concerning whether classical Mendelian genetics can be reduced to molecular biology (it is doubtful), or whether the principle of natural selection, which is at the center of evolutionary theory, admits of any meaningful reduction in terms of chemistry or physics. (Notice that this is most certainly *not* to say that natural selection somehow violates the laws of chemistry and physics, *pace* creationist claims to the contrary.)

Informally, we can distinguish between a moderate type of reductionism and a more radical one. Radical reductionism rests on the claim that all there is in the universe is made of matter and energy, and that therefore any phenomenon that can be studied by science can ultimately be reduced to fundamental physics. Moderate reductionism, by contrast, aims at reducing our understanding of whatever phenomenon is under investigation to the minimum possible level of organization, without claiming that such level is necessarily the one that falls under the domain of fundamental physics. For instance, let us imagine that we are engineers interested in the structural properties of the Brooklyn bridge. If we took a radical reductionist perspective, we would attempt to describe those properties in terms of quantum mechanics, the most fundamental physical theory devised so far. If we adopted a moderate reductionist stance, however, we would approach the problem at the level of analysis of material science, and perhaps in part of macro-molecular physics, but we would not go any further down, so to speak.

Teachers would have to immediately make a distinction for students when discussions of reductionism are tackled. On the one hand, reductionism can be a claim about what actually exists in the universe, i.e., an ontological position. On the other hand, it can simply be a claim about what we can reasonably know about the universe, i.e. an epistemological position. There is no simple connection between the two types of reductionism introduced above and whether one takes an ontological or epistemological stance in this debate. For instance, one can reasonably be an ontological reductionist (i.e., claiming that all there is in the universe is matter and energy) and yet take an epistemologically moderate reductionist position, simply because, say, a quantum mechanical model of the Brooklyn bridge—while theoretically feasible—would be exceedingly impractical.

However, if someone is a moderate reductionist at the ontological level, because she believes that there are genuine emergent properties (O'Connor and Wong 2009) that are not meaningfully reducible to lower levels of analysis, then one is perforce committed to moderate epistemological reductionism as well. Of course, one can also coherently be a radical reductionist both in the epistemic and the ontological senses, though such as position is somewhat rare in the philosophical literature.

What does this have to do with memetics and the evolution of religion? Recall that the idea of memes arose by direct analogy with genes, and in a very strong sense memetics is a (largely failed, I have argued) attempt at reducing human thought and cultural phenomena to the simple mechanics of (hypothetical, speculative) replicator entities—just in the same way in which genetics has attempted, with only very partial success, to position itself as the fundamental science within biology, to which evolutionary biology, developmental biology and all the rest need eventually be reduced to. Memetics then can be understood and discussed as a project of reduction of cultural evolution (typically the domain of the social sciences) to biology (as explicitly advocated, for instance, by Wilson 1999).

There have been other, more or less successful, research programs aiming at understanding culture using the tools of evolutionary biology. Chief among them are the papers produced by a number of authors on gene-culture co-evolution during the 1980s and '90s, where the mathematical approaches of standard population genetics were deployed to create a formalism capable of describing cultural evolution (Aoki 1986; Feldman and Laland 1996). That approach makes for a nice contrast with memetics, because the gene-culture co-evolution tradition attempts to apply the methods of population genetics while recognizing a different ontological status to cultural processes—resulting into a hybrid type of reductionism. Memetics, instead, attempts reduction to evolutionary theory both in terms of methods and as far as the ontological status of the relevant entities (memetic replicators) is concerned.

The broader context of this sort of analysis of reductionism applied to culture is provided by the question of why there is a widespread feeling that the social sciences themselves have not been doing a sufficiently good job with their own subject matter, so to forcefully invite reduction to biology. After all, reduction of biology to physics in turn has been largely unsuccessful, because biology has produced its own overarching theories and methods (Darwinism and its subsequent variations) that left very little space for radical reductionism to work. The social sciences, and in particular psychology, have gone through a brief period of production of overarching theories—indeed, more or less consciously based on the success of Darwinism—as was the case, for instance, with Freudian psychoanalysis (and, in a different context, for Marxist theories of history). Unlike Darwinism, however, Freudianism (and Marxism in history) has proved to be a degenerate research program, which has left a vacuum that gene-culture co-evolutionary theories and memetics have tried to fill, without any spectacular success so far. It remains to be seen whether the social sciences will soon see their own Darwin appear over the horizon, or whether a new type of moderate reduction will once again attempt to fold them into biology.

6 Lesson Five: Causality and Its Multiple Levels

A good complement to discussions of reductionism is the concept of multi-level causality, which can be introduced naturally while talking about the evolution of religion because one of the major explanations for the latter has been proposed to be group selection, perhaps acting in opposition to standard Darwinian individual-level selection. And of course, once we start talking to students about multiple levels of causality, the door is immediately open to generalize to the broad philosophical issues underpinning the very notion of causality itself.

Group selection has been a controversial topic ever since Williams' (1966) damning critique of early, rather naive, versions of the theory. Williams' showed that group selection simply could not happen if it were conceived as selection somehow acting for "the good of the species" (a still popular misrepresentation of evolutionary theory among the general public). However, later on Williams (1992) himself did acknowledge that natural selection can act on group-level traits, although he pointed out that when group selection works against individual-level selection the latter is more likely to prevail because individuals reproduce much faster (thus generating an overall stronger selective pressure) than groups.

Other authors have explored the possibility of group selection playing a major role in the evolution of human behavioral traits, particularly altruism (Sober and Wilson 1999), and recently Samir Okasha (2006) has produced the most complete and convincing philosophical and mathematical analysis of group selection, thus helping to rehabilitate a process that was considered by Darwin but that still today generates a significant amount of skepticism among many evolutionary biologists.

Ultimately, of course, whether or not group selection contributed significantly to the evolution of religious belief (or of any other human trait) is a question to be settled empirically, if at all possible. However, the several group-level scenarios that I mentioned above (e.g., religion being favored because it triggers prosocial behavior, etc.) can be coupled with standard individual-level explanations (e.g., reduction of anxiety about death) to explore the idea that causal explanations need not be inherently reductionist, and may indeed involve legitimate entities positioned at different hierarchical levels of complexity.

Indeed, Okasha's treatment draws, in part, from Price's (1970) classic approach to the problem, which is particularly suitable to introducing students to the idea of multi-level selection theory and consequently of multi-level causality. Price's equation, a generalization of the idea that the Darwinian insight about natural selection can be formalized mathematically, looks (in one form), like this:

$$\Delta z = \mathrm{cov}(W, Z) + E\,\mathrm{cov}(w, z)$$

where Δz is the change in phenotype from one generation to the next, cov indicates statistical covariance, z and Z are the trait means at two different levels of the hierarchy (e.g., individual and group, respectively), w and W are the corresponding fitnesses at those levels, and E indicates the expected value (of the covariance). This version features two hierarchical levels, the one on the left of the plus sign being what Okasha terms the “collective” level, the one to the right of the plus sign being the “particle” level. Depending on the specific situation, individuals can be collectives (in which case the particles are, say, genes), or particles (in which case the collectives may be groups), and the equation can be expanded to include as many terms as it makes biological sense (i.e., it is not limited to two terms, or levels of selection, in principle).

The central message for a science student here should be that non-exclusive causal processes may take place at different levels, and may interact with each other to yield a much more complex picture of a given phenomenon (in this case, the evolution of religion) than a simple reductionist account would. However, one should not pass the opportunity to also introduce basic philosophical discussions of the very concept of causality, starting naturally with Hume and Kant (De Pierris and Friedman 2008). Briefly, Hume’s take on causality is a rather skeptical one. In a section of his *Enquiry Concerning Human Understanding* entitled “Sceptical Doubts Concerning the Operations of the Understanding,” Hume (Hume 1748/1952) suggests that a cause does not logically entail its effect, so that it is only because of experience that we talk about A causing B, while in fact all we have is a constant conjunction in which A is temporally succeeded (in short order) by B. This skepticism about causality naturally dovetailed with Hume’s broader skepticism about induction (in the sense of generalization from experience), which is at the foundation of any scientific inquiry and yet cannot be logically justified on anything other than inductive (and therefore circular) grounds.

In the *Prolegomena to Any Future Metaphysics* (1783/2004) Kant famously wrote: “I freely admit that it was the remembrance of David Hume which, many years ago, first interrupted my dogmatic slumber and gave my investigations in the field of speculative philosophy a completely different direction.” For Kant, causality was a crucial instance of what he called a category, that is a pure concept of the understanding, and therefore most certainly not simply derived from experience via induction. This particular philosophical dispute is most certainly not settled yet, in part because of scholarly disagreement about what, exactly, Hume’s “attack on metaphysics” consisted of, and whether and on what grounds Kant’s response succeeds. Nonetheless, the philosophy of causality can and should be introduced to science students simply because scientists just take the concept of causality itself as primitive and in no need of further analysis—an excellent example of the different scopes and proper domains of science and philosophy.

7 Lesson Six: The Nature of Historical Science

All the preceding discussion should be understood within the broader context of what kind of science evolutionary biology is, which brings us to the last of the proposed six philosophical lessons for science students. A fundamental, if not always sharp, divide in science separates historical disciplines (e.g., astronomy, paleontology) from experimental ones (fundamental physics, chemistry). Evolutionary biology is positioned somewhere in the middle, as clearly biologists can design and carry out controlled experiments in both the laboratory and the field, and yet such experiments are affected by the historicity of the particular system being studied: while an electron is an electron, regardless of how long it has been around and where it has been before finding its way into a particle accelerator, a population of living organisms will react differently to the very same experimental conditions from any other population of living organisms. This is for the simple reason that their genetic makeup has been indelibly shaped by a long succession of ancestors belonging to the lineage that originated the population being experimented upon.

In the context that we are considering, the student may be asked to ponder how exactly different historical hypotheses about the origin of religious belief may be tested and meaningfully compared to each other. To this effect, an excellent framework is provided by two papers published by Carol Cleland (2001, 2002; see also the more recent Cleland 2011) on the epistemic differences between historical and experimental sciences. Cleland's pivotal idea is that the two types of science are separated by what she calls an asymmetry of overdetermination. Building on previous work by David Lewis, she explains that "the basic idea is that localized present events overdetermine their causes and underdetermine their effects." She elucidates the concept by considering the example of a crime being investigated. Once committed, a crime leaves a number of historical traces, no matter how careful the perpetrator was in erasing as many of them as possible. All it takes for a Sherlock Holmes to figure out what happened was a relatively small number of traces that clearly enough point toward a particular sequence of events. Holmes would then be using a type of induction known as inference to the best explanation to pinpoint the culprit. Conversely, the simple act of not committing the crime obviously instantly erases the possibility of any historical trace to be left around. Few currently available clues overdetermine a past event, while so many futures are possible given a particular current state of things that the latter underdetermines the futures.

Cleland cashes in this asymmetry of overdetermination by arguing that—contrary to popular wisdom (and to the opinion of many practicing scientists)—there is nothing inherently epistemically superior about experimental over historical science. This is because of two consequences of overdetermination. On the one hand, while experimental scientists have the ability to strictly control the conditions of their experiments, it is that very strictness that limits the scope of applicability of their results: as soon as one widens the settings of a given experiment by making the conditions more realistic, different factors begin to interact with each other in complex ways, quickly leading to a large number of possible future outcomes; in other words, predictability is purchased at the expense of generality. On the other hand, while historical traces constantly decay through time, and may disappear forever, the historical scientist often needs only a small amount of them to arrive at a sufficiently accurate reconstruction of what happened—just like Sherlock Holmes in the hypothetical example of the impossibility of a perfect crime.

To make things more concrete, Cleland's account makes sense of some surprising limitations of experimental science, as well as some spectacular successes of historical science. In the first case, it is notable, for instance, that non-equilibrium thermodynamics (say in its applications in the guise of atmospheric physics and climate science) quickly reaches a limit in terms of predictive ability, where complex mathematical models are incapable of generating more than very approximate statistical predictions of the future behavior of complex systems, predictions often accompanied by rapidly expanding margins of error. In the second case, however, we have the successes of paleontologists in determining that an extraterrestrial body of massive proportions hit the earth 65 million years ago, helping to cause the extinction of countless numbers of species, chief among them the dinosaurs. The impact was suspected once geologists discovered a worldwide thin layer of iridium in rocks datable to the K/T (Cretaceous/Tertiary—K is the traditional abbreviation for Cretaceous) boundary. This led to a search for a crater, the remnants of which were eventually identified off the Yucatan peninsula via satellite imagery. From there, geologists could calculate the size and direction of the impact, and therefore make fresh predictions concerning additional historical traces, for instance those left by the tsunamis that must have hit the eastern coast of Mexico as a result of the asteroid crush. Sure enough, those traces were found, leading to even more confidence in the conclusion that "the crime" did indeed take place in the way it had been hypothesized.

The above discussion can be used to introduce students to general features of scientific methodology, but also to the limits of Cleland's own analysis. What type of historical traces could be left by an evolutionary process that led to the origin of religion? The answer to this question is not very different from that which applies to much historical research on the evolution of human behaviors. Kaplan and I (Kaplan 2002; Pigliucci and Kaplan 2006, ch. 7) have discussed this issue within the context of the research program of evolutionary psychology, and elucidating the evolution of religion can easily be framed as a particular goal within that research program. We have pointed out that, on the one hand, it is of course perfectly reasonable to assume that evolutionary processes in general, and natural selection in particular, may have shaped at least in part human behavior—though we are somewhat skeptical of evolutionary psychology's cavalier downplaying of the role of culture in doing the same. On the other hand, however, we have highlighted the epistemic reasons why evolutionary psychology may have a particularly hard time testing adaptive hypotheses concerning human behavior as a matter of practice.

Fundamentally, there are three types of evidence that evolutionary biologists use to test adaptive hypotheses in general: fossil records, phylogenetic comparative analyses, and direct measures of ongoing natural selection (the latter being more a test of current adaptive value than of a historical process of adaptation). Most human behaviors do not leave a fossil record, though there are exceptions (for instance, after the discovery of the first Australopithecus we know that bipedal ambulation—a behavior—preceded the evolution of a large brain). Comparative phylogenetic analyses are effective only when there is a large number of closely related species to compare that are variable for the trait of interest; unfortunately, there are no other existing species of Homo, our closest relatives are low in number of taxa (two species of chimpanzee, two of gorillas), fairly distantly related, and do not show sufficiently homologous behaviors for most traits of interest to evolutionary psychologists (and certainly as far as religion in particular is concerned). Finally, while it is possible in principle to measure ongoing natural selection on human behavioral traits (though technically quite difficult for complex abstract behaviors, such as religious belief), this would not be particularly informative because the relevant physical and social conditions that characterized our species for most of its evolution were likely very different from those we experience now, making any extrapolation from current to past selective regimes particularly suspect (this is not the case with most other species because their physical, and especially their social environment have not, presumably, changed that much over long periods of time).

So while Cleland is surely correct in her general analysis of the differences between experimental and historical sciences, and while it is also true that the epistemic power of the former is often overestimated and that of the latter just as frequently unwisely downplayed, there are serious issues when it comes to testing adaptive hypotheses in general, and those pertinent to human behavior in particular. The student of evolution can only benefit from a more sophisticated understanding of the epistemology of the discipline brought about by careful philosophical analysis.

8 Conclusion: Why Study Philosophy in Science?

Famous Nobel physicist Richard Feynman once quipped that "Philosophy of science is about as useful to scientists as ornithology is to birds." It is hard to imagine a more concise way to encapsulate the cross-disciplinary disdain that keeps the divide between the two cultures (Snow 1993). There are, of course, a couple of quick replies one could offer to Feynman. For example that the very existence of some bird species depends on efforts in conservation biology, and therefore ultimately on ornithology; or, more harshly, that birds cannot appreciate ornithology because they are not sufficiently smart. But the real answer should have emerged from the preceding discussion: philosophy is both an independent discipline in its own right, with its goals independent of those of science; and it is a borderline field which overlaps with science when it comes to conceptual and epistemological issues pertinent to scientific research.

One of philosophy of science's goals is to understand how science itself works, regardless of whether that understanding will or will not impact the scientist's actual practice. Another goal is to analyze the details of some of science's conceptual issues—such as the conditions for the testability of adaptive hypotheses. Here clearly the scientist ought to very much be interested in engaging the philosopher simply in order to achieve a more sophisticated command of problems that the scientist herself rarely has to confront directly, problems ranging from epistemic access to specific hypotheses to the general problem of underdetermination of theories by the data. The goal is not to turn scientists into philosophers, and much less students in science classes into students of philosophy. Rather, the idea is that the two cultures have much to offer that we need to learn to respect in its own right, as well as more to offer to each other than the Richard Feynmans of the world seem to appreciate.

In the particular example of scientific studies of religion, we have appreciated a number of things that may not have been obvious to the student who comes across the scientific literature but has no background in philosophy. First and foremost, that to study religion scientifically is not the same as refuting the existence of gods, for the simple reason that the supernatural is too vague a concept to qualify as a scientific hypothesis. Second, we have appreciated the distinction between proximate and ultimate causation, and therefore understood exactly why neurobiological and evolutionary explanations of religion are in no meaningful sense competing with each other. Third, we have examined the advantages of doing science by multiple hypotheses, which has also led us into an examination of multi-level causality and of the concept of causality itself—a primitive concept in science, but one that has been characterized by a rich philosophical tradition. Fourth, we have used memetics to examine both the limits and intelligent application of reductionism in science and the nature of "unobservables" in scientific theory, which in turn has provided us with an entry into the complex philosophical discussions between realists and anti-realists. Finally, we have began to better appreciate the nature of historical sciences and how they compare to experimental ones, which is particularly relevant to the study of evolutionary biology, a discipline characterized by a rich interaction between historical and experimental methods. I have shown by example how easy it is to introduce philosophical concepts in the study of science, and the reason for this ease is that science itself, as Dennett observed, is imbued with philosophical assumptions that should—at least from time to time—be considered by those who practice science as well as by those who teach it.

References

Aoki, K. (1986). A stochastic model of gene-culture coevolution suggested by the "culture historical hypothesis" for the evolution of adult lactose absorption in humans. *Proceedings of the National Academy of Sciences USA, 83*, 2929–2933.

Ariew, A. (2003). Ernst Mayr's 'ultimate/proximate' distinction reconsidered and reconstructed. *Biology and Philosophy, 18*, 553–565.

Aristotle. (1998). *The metaphysics*. Translated by Hugh Lawson-Tancred. Harmondsworth: Penguin.

Berger, J. O., & Sellke, T. (1987). Testing a point null hypothesis: The irreconcilability of p values and evidence. *Journal of American Statistical Association, 82*, 112–122.

Blackmore, S. J. (2000). *The meme machine*. Oxford: Oxford University Press.

Blanke, O., Ortigue, S., Landis, T., & Seeck, M. (2002). Stimulating illusory own-body perceptions. *Nature, 419*, 269–270.

Boyer, P. (2008). Religion: Bound to believe? *Nature, 455*, 1038–1039.

Brigandt, I., & Love, A. (2008). *Reductionism in biology*. The Stanford encyclopedia of philosophy, http://plato.stanford.edu/archives/fall2008/entries/reduction-biology/.

Burnham, K. P., & Anderson, D. R. (2002). *Model selection and multimodel inference*: *A practical information-theoretic approach*. Berlin: Springer.

Chamberlin, T. C. (1890). The method of multiple working hypotheses. *Science, 15*, 92–96.

Cleland, C. E. (2001). Historical science, experimental science, and the scientific method. *Geology, 29*, 987–990.

Cleland, C. E. (2002). Methodological and epistemic differences between historical science and experimental science. *Philosophy of Science, 69*, 474–496.

Cleland, C. E. (2011). Prediction and explanation in historical natural science. *British Journal for the Philosophy of Science*, advanced online publication, 25 February 2011.

Cohen, J. (1994). The earth is round ($p < .05$). *American Psychologist, 49*, 997–1003.

Corfield, D., & Williamson, J. (2001). *Foundations of Bayesianism*. Berlin: Springer.

Darwin, C. (1859). *On the origin of species by means of natural selection*: *Or, the preservation of favoured races in the struggle for life*. London: John Murray.

Dawkins, R. (1976). *The selfish gene*. Oxford: Oxford University Press.

Dawkins, R. (2006). *The God delusion*. New York: Houghton Mifflin Harcourt.

De Pierris, G., & Friedman, M. (2008). *Kant and Hume on causality*. The Stanford encyclopedia of philosophy, http://plato.stanford.edu/archives/fall2008/entries/kant-hume-causality/.

Dennett, D. (1996). *Darwin's dangerous idea*: *Evolution and the meanings of life*. New York: Simon and Schuster, p. 21.

Depew, D. (2008). Consequence etiology and biological teleology in Aristotle and darwin. *Studies in History and Philosophy of Science, Part C, 39*, 379–390.

Feldman, M. W., & Laland, K. N. (1996). Gene-culture coevolutionary theory. *Trends in Ecology & Evolution, 11*, 453–457.

Flack, J. C., & de Waal, F. B. M. (2000). Any animal whatever: Darwinian building blocks of morality in monkeys and apes. *Journal of Consciousness Studies, 7*, 1–29.

Gazzaniga, M. S. (2003). The split brain revisited. *Scientific American, 287*, 26–31.

Gould, S. J., & Lewontin, R. C. (1979). The spandrels of San Marco and the Panglossian Paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Society of London, B205*, 581–598.

Gould, S. J., & Vrba, E. S. (1982). Exaptation—a missing term in the science of form. *Paleobiology, 8*, 4–15.

Hone, W. D. E., Keesey, T. M., Pisani, D., & Purvis, A. (2005). Macroevolutionary trends in the Dinosauria: Cope's rule. *Journal of Evolutionary Biology, 18*, 587–595.

Hume, D. (1748/1952). *An enquiry concerning human understanding*. Forgotten Books.

Huxley, J. S. (1942/2010). *Evolution*: *The modern synthesis*. Cambridge: MIT Press.

Kant, I. (1783/2004). *Prolegomena to any future metaphysics that will be able to come forward as science*. Cambridge: Cambridge University Press.

Kaplan, J. H. (2002). Historical evidence and human adaptation. *Philosophy of Science, 69*, S294–S304.

Kaplan, J. M., & Pigliucci, M. (2001). Genes 'for' phenotypes: A modern history view. *Biology and Philosophy, 16*, 189–213.

Ladyman, J. (2002). *Understanding philosophy of science*. London: Routledge.

Lakatos, I. (1978). *The methodology of scientific research programmes*: *Philosophical papers* (Vol 1). Cambridge: Cambridge University Press.

Manson, J. H., Perry, S., & Parish, A. R. (1997). Nonconceptive sexual behavior in bonobos and Capuchins. *International Journal of Primatology, 18*, 767–786.

Mayr, E. (1961). Cause and effect in biology. *Science, 134*, 1501–1506.
Newberg, A., D'Aquili, E., & Rause, V. (2002). *Why God won't go away*: *Brain science and the biology of belief*. New York: Random House.
Nix, T. W. & Barnette, J. J. (1998). The data analysis dilemma: Ban or abandon. *A review of null hypothesis significance testing*. *Research in the Schools, 5*, 3–14.
Norenzayan, A., & Shariff, A. F. (2008). The origin and evolution of religious prosociality. *Science, 322*, 58–62.
O'Connor, T., & Wong, H. Y. (2009). *Emergent properties*. The Stanford encyclopedia of philosophy, http://plato.stanford.edu/archives/spr2009/entries/properties-emergent/.
Okasha, S. (2006). *Evolution and the levels of selection*. Oxford: Oxford University Press.
Pigliucci, M. (2002). *Denying evolution*: *Creationism, scientism, and the nature of science*. Sunderland: Sinauer.
Pigliucci, M. (2003). From molecules to phenotypes?-The promise and limits of integrative biology. *Basic and Applied Ecology, 4*, 297–306.
Pigliucci, M., & Kaplan, J. (2000). The fall and rise of Dr. Pangloss: Adaptationism and the Spandrels paper 20 years later. *Trends in Ecology & Evolution, 15*, 55–70.
Pigliucci, M., & Kaplan, J. (2006). *Making sense of evolution*: *The conceptual foundations of evolutionary biology*. Chicago: University of Chicago Press.
Price, G. R. (1970). Selection and covariance. *Nature, 227*, 520–521.
Pyysiainen, I., & Hauser, M. (2009). The origins of religion: Evolved adaptation or by-product? *Trends in Cognitive Science, 14*, 104–109.
Richerson, P. J., & Newson, L. (2009). Is religion adaptive? Yes, no, neutral, but mostly we don't know. In M. J. Murray & J. Schloss (Eds.), *The believing primate*: *Scientific, philosophical and theological perspectives on the evolution of religion* (pp. 100–117). Oxford: Oxford University Press.
Rizzolatti, G. (2005). The mirror neuron system and its function in humans. *Brain Structure & Function, 210*, 419–421.
Schloss, J. (2009). Evolutionary theories of religion. Science unfettered or naturalism run wild? In M. J. Murray & J. Schloss (Eds.), *The believing primate*: *Scientific, philosophical and theological perspectives on the evolution of religion* (pp. 1–25). Oxford: Oxford University Press.
Snow, C. P. (1993). *The two cultures*. Cambridge: Cambridge University Press.
Sober, E., & Wilson, D. S. (1999). *Unto others*: *The evolution and psychology of unselfish behavior*. Cambridge: Harvard University Press.
Weinberg, S. (1994). Against philosophy (pp. 166–190). *In dreams of a final theory*. New York: Vintage.
Williams, G. C. (1966). *Adaptation and natural selection*. Princeton: Princeton University Press.
Williams, G. C. (1992). *Natural selection: Domains, levels, and challenges*. Oxford: Oxford University Press.
Wilson, E. O. (1999). *Consilience*: *The unity of knowledge*. New York: Vintage.
Wilson, D. S. (2002). *Darwin's Cathedral: Evolution, religion, and the nature of society*. University of Chicago Press: Chicago.
Wilson, D. S. (2007). *Evolutionary religious studies* (*ERS*): *A Beginner's guide*. http://evolution.binghamton.edu/religion/resources/guide/.
Worrall, J. (1989). Structural realism: The best of both worlds? *Dialectica, 43*, 99–124.