

## Epistemological Similarities and Differences Among the Sciences

A Paper Commissioned by the NAS Board on Science Education (Committee on the Status, Contributions, and Future Directions of Discipline Based Education Research)

Elliott Sober  
Philosophy Department  
University of Wisconsin, Madison  
Madison, WI 53706

Epistemology is the branch of philosophy that studies concepts such as knowledge, evidence, induction, prediction, rationality, justified belief, and probability. Epistemologists study these topics both in the context of scientific practice and in the context of the everyday problems that nonscientists face.

The etymology of “epistemology” often leads philosophers to define epistemology as the theory of knowledge, but the practice of epistemologists shows that they are interested in a wider range of concepts and problems. For example, if you buy a ticket in a lottery, it seems false to say that *you know* that your ticket will lose. But it may be true that *you are justified in believing* that it will lose; after all, the probability of its losing is very high. There is more to epistemology than the topic of knowledge.

Philosophers often draw a distinction between epistemology and metaphysics. Metaphysicians try to describe the basic kinds of entities that exist and to describe how those basic kinds are related to each other. This may sound like what scientists do when they say that quarks, or genes, or tectonic plates exist. Metaphysicians are interested in categories that are broader. For example, all of the above examples fall into the metaphysical category of physical objects. Granted that these and other physical objects exist, are there also entities that are nonphysical? For example, do numbers also exist and are they nonphysical? And when a physical object (e.g., an organism or a ship) persists through time, what makes it true that the temporal stages of that entity are all parts of a single enduring physical object? And how are the objects described in the different sciences related to each other? Are organisms “nothing but” the cells and molecules that are their constituents? This question about reductionism is a metaphysical question.

When I say that epistemologists “study concepts,” the question arises of whether they are attempting to describe how people *actually* understand those concepts, or are attempting to say how people *ought* to understand them. There is disagreement among epistemologists about whether philosophy is purely descriptive or purely normative (or something in between). I myself do not shy away from philosophy’s playing a normative role. Perhaps philosophers of science can help scientists to reason better; if so, we can and should do more than just describe what scientists in fact do. Psychologists describe how people reason; philosophers should attempt to describe how people ought to reason. Logic and psychology are different subjects.

This distinction between normative and descriptive inevitably comes up when philosophers of science discuss whether there is such a thing as “the scientific method.” It is clear that scientists have changed their conception of what scientific inference involves. For example, the statistical practices that now dominate so much of biology and the social sciences were invented in the first part of the 20th century. Nineteenth and 20th century science are dramatically different in this respect. But it is a separate question whether there are inferential rules that all scientists ought to embrace. Perhaps the ideals are eternal, even though the practices often fail to live up to those ideals. Philosophers of science now disagree about whether this idea of timeless rules of inference really makes sense. Those who defend this picture tend to think of rules of inference and methodological principles as things that scientists (and philosophers and statisticians and logicians) discover, not merely invent.

Philosophers of science also disagree about whether the idea of “the scientific method” makes sense when we survey the different branches of contemporary science. Since the 1960s, philosophers of science have paid more and more attention to philosophical questions that arise only in specific scientific disciplines. Before that time, the field was largely dominated by *general* philosophy of science. For example, when Hempel (1965) asked what a scientific explanation is, he assumed that the nature of explanation is the same across the sciences. Since the 1960s, a lot more attention has been paid to philosophy of biology, philosophy of cognitive science, and philosophy of economics. These newer “philosophies of the special sciences” joined the already existing field of philosophy of physics, which has always played an important part in philosophy of science. Philosophers of science now emphasize differences among the sciences more than their predecessors did. Although this shift sometimes reflects a substantive disagreement, often it involves a shift in interest. When one philosopher emphasizes similarities across the sciences and another emphasizes differences, they needn’t be disagreeing. There has been a shift in the field towards splitting and away from lumping.

I’ll begin with an aspect of scientific reasoning about which there is pretty wide agreement. Philosophers of science generally agree that all sciences use deductive reasoning and, moreover, that the rules of deductive reasoning are the same regardless of what the scientific subject matter is. Consider, for example, the following two inferences:

If this organism is a mammal, then it  
is a tetrapod.

This organism is a mammal.

---

This organism is a tetrapod.

If this particle is a proton, then it  
has a positive charge.

This particle is a proton.

---

This particle has positive charge.

In each of these arguments, the statements above the line are the argument’s premises; the statement below is the conclusion. These two inferences have different subject matters—the one from biology, the other from physics—but they also have something in common. Both exemplify the form of inference called *modus ponens*:

If A, then B  
A  
(MP) 

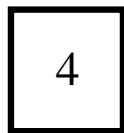
---

B

Here A and B represent any two propositions. Take any two sentences and substitute one of them for A and the other for B in the MP schema. The result will be an argument that has the form *modus ponens*. Logicians use the term “deductive validity” to describe this and other argument forms. Each and every *modus ponens* argument has the following property: if the premises are true, then the conclusion must be true. Just as it doesn’t matter whether the subject matter of the argument is from physics or biology, it also doesn’t matter whether the subject matter comes from a science or from everyday life. The rules of deductive logic are the same. This is the standard view of deductive logic—it describes rules of inference that are not subject-matter specific.

*Modus ponens* will probably strike the reader as an obviously valid form of deductive inference, so this example in which the sciences all use the same rule of inference may seem to be trivial, and so not helpful if the task is to improve scientific practice and science education. My reply is that cognitive psychologists have shown that scientists and nonscientists alike often reason fallaciously when it comes to deductive problems, even ones that are pretty simple (Kahneman, Tversky, and Slovic 1982). A much discussed example is the Wason (1966) selection task. Consider the four cards shown below and the rule

“If a card has an even number on one side, then it says yes on the other.” Which of the cards do you need to turn over if you wish to see whether the rule is true?



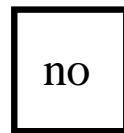
(a)



(b)



(c)



(d)

Most people taking this test mention (a) but many fail to mention (d); cards (b) and (c) are in fact irrelevant to testing the rule, but many subjects mention these. Solving the Wason selection problem involves seeing that the conditional “If something has characteristic X, then it has characteristic Y” is falsified by anything that has X and lacks Y. Cards (a) and (d) are potential falsifiers of the rule stated; this is why it is relevant to check their flipsides. Cards (b) and (c) are not potential falsifiers, which is why there is no point in flipping them over.

The idea that deductive logic describes a set of inference rules that apply to all subject matters has its counterpart in probability and statistics. Although statistics courses are often taught with a focus on specific domains of application (as in undergraduate courses *or biostatistics*), practitioners usually recognize that the procedures they use apply across a much broader range of subject matters. And just as is true with deductive reasoning, there are patterns of mistakes that people often make when reasoning about probability and statistics. Here is an example:

The University of Berkeley some years back found evidence that there was sex discrimination in admission to graduate school (Cartwright 1979). It was observed that

- (1) A smaller percentage of women were admitted to graduate school than men.

But when the University looked more closely at the evidence, they also found that:

- (2) In each department, the percentage of women admitted is the same as the percentage of men.

How can (1) and (2) both be true? The answer is that this is an instance of Simpson’s (1951) paradox. But what does that mean?

Let’s look at a simple, hypothetical example that illustrates the idea. Suppose 1000 women and 1000 men apply to graduate school and that there are just two departments in the university, D1 and D2. In accordance with (1), suppose that 220 women and 380 men are admitted. But how can (2) be true? Suppose that the admissions rates for men and women in the two departments are:

Admission Rates		
	D1	D2
Male	20%	40%
Female	20%	40%

Now suppose that most men apply to D2 and most women apply to D1:

Numbers Applying		
	D1	D2
Male	100	900
Female	900	100

Given the admission rates and the numbers applying, the result is:

Numbers Admitted		
	D1	D2
Male	20	360
Female	180	40

Notice from this last table that 380 men out of the 1000 applying are admitted, while 220 women out of the 1000 applying are. In this example, propositions (1) and (2) are both true.

What this example illustrates is that the following argument is fallacious: “Since each department admits the same percentage of women as men, the University as a whole must admit the same percentage of women as men. This seems clear -- what is true in each part must also be true in the whole.” This line of reasoning will be correct if there is no correlation between an applicant’s sex and the department to which the applicant applies. But, in this example, these two characteristics are not independent. Men tend to apply to one department and women to the other.

Understanding Simpson’s paradox turns out to be important in connection with an important problem in evolutionary biology – the question of how natural selection can cause altruistic characteristics to evolve (Sober 1993). The puzzle begins with two definitional truths in evolutionary theory:

- (3) Altruists have lower fitness than selfish individuals who live in the same group.
- (4) When a population evolves under the control of natural selection, fitter traits increase in frequency and less fit traits decline.

An organism's fitness is its ability to survive and reproduce in its environment. Evolutionary biologists use the term "altruism" to apply to organisms that confer fitness benefits on others at a fitness cost to self. Organisms do not need to think or feel anything to be altruistic; the term is defined solely in terms of fitness consequences of a behavior.

It may seem to follow from propositions (3) and (4) that altruism cannot evolve when the evolutionary process is controlled by natural selection. In fact, this does not follow. To see why, we need to think about Simpson's paradox. Suppose, for example, that there are two groups of organisms that each contain altruists and selfish individuals in the parental generation. Each group contains 100 individuals:

Numbers of Parents		
	Group 1	Group 2
Altruists	90	10
Selfish Individuals	10	90

Suppose that these parents reproduce asexually, with an individual's fitness (its number of offspring) depending on its own phenotype (altruistic or selfish) and also on the type of group it inhabits:

Fitnesses of Parents		
	Group 1	Group 2
Altruists	3	1
Selfish Individuals	4	2

Notice that in each group, altruists are less fit than selfish individuals. However the average fitnesses of the two traits are

$$w_A = 0.9(3) + 0.1(1) = 2.8$$

$$w_S = 0.1(4) + 0.9(2) = 2.2$$

The average fitness of altruism is greater than the average fitness of selfishness. Simpson's paradox strikes again.

Let's suppose that individuals in the first generation reproduce and then die, and that offspring have the same traits as their parents. The following table describes what the next generation will look like:

Numbers of Offspring		
	Group 1	Group 2
Altruists	270	10
Selfish Individuals	40	180

Recall that the parental generation had 50% altruism. In the offspring generation, there are 280 altruists and 220 selfish individuals in total, so the frequency of altruism has gone up. However, the frequency of altruism declines in each group. It drops from 90% to  $270/310 = 87\%$  in Group 1 and from 10% to  $10/190 = 5\%$  in Group 2.

Who would have thought that the same principles of reasoning that are essential to understanding the Berkeley sex discrimination case also are pertinent to understanding the evolution of altruism? The subject matters are so different, how could there be a connecting thread? This result is less surprising when one takes seriously the idea that logic, probability, and statistics describe inference rules that are subject-matter neutral.

The smooth pattern I have described so far – a principle of inference makes sense in one domain and it also applies unproblematically to another – needs to be supplemented by another. Sometimes a principle that seems sensible in one domain turns out to be problematic when applied to another. This makes one realize that the principle needs to be restricted or modified in some way that wasn’t immediately obvious when only the first domain of application was considered.

Consider, for example, Alfred Wegener’s (1924) defense of the hypothesis of continental drift. Wegener noticed that the wiggles in the east coast of South America correspond rather exactly to the wiggles in the west coast of Africa. The pattern is “as if we were to refit the torn pieces of a newspaper by matching their edges and then check whether the lines of print run smoothly across (Wegener 1924, p. 77).” Wegener also noticed that the distribution of geological strata down one coast matches the distribution down the other. And he further observed that the distribution of organisms down the two coasts – both fossilized and extant – also show a detailed correlation. Wegener argued that this systematic matching is strong evidence that the continents had once been in contact and then had drifted apart. Wegener encountered intense opposition from geophysicists, who didn’t see how continents could plough through the ocean floor. It was only with the later development of plate tectonics that geophysicists were able to accept continental drift. Wegener, unfortunately, had argued not just that the continents were once in contact; he also maintained that they slide across the ocean floor. Later work showed that he was half right and half wrong – the continents were in contact, but they sit on plates and a continent and its plate move together (LeGrand 1989).

Wegener’s argument suggests the following principle:

(PCC) If X and Y are correlated, then either X caused Y, or Y caused X, or X and Y trace back to a common cause C.

The philosopher of science Hans Reichenbach (1956) called this the *principle of the common cause*. A shorthand statement of this principle is the claim that correlated events are always causally connected. The example of continental drift, as well as many others, from a variety of sciences, make the PCC sound entirely sensible. Although some early statisticians realized that correlations can exist between events that are not causally connected (Yule 1926), it was not until the flowering of quantum mechanics that it became widely recognized that the PCC may be over-stated. A standard interpretation of quantum mechanics asserts that simultaneous events are sometimes correlated without there being a common cause; if so, the correlation is just a brute fact. This means that the PCC needs to be reformulated. Principles that sound perfectly reasonable in one domain may turn out to be problematic when other applications are considered.

How can this happen? It is easy to over-state a principle, especially if one rests content with considering examples that make the principle seem sensible. Perhaps when we invent principles, our minds tend to gravitate to cases in which the principles seem correct. Maybe your brain child is like your biological child – you want to think well of it! After all, it is yours. A good preventative measure is to look hard for counter-examples. As Popper (1959) emphasized, we should try to refute our theories; it is all too easy to restrict our search so that we look only for favorable evidence. A second preventative measure is to try to derive a proposed principle from principles that are more fundamental. This exercise can help one see that this or that set of assumptions is not sufficient for the target principle to be correct. This negative result may give one pause.

At the start of this essay, I drew a distinction between normative and descriptive issues. Do the sciences differ in the methods they actually use? Ought they to differ in this way? A great deal of normative theorizing about scientific inference (in philosophy, logic, and statistics) assumes that there is one body of inference rules from which all sciences ought to draw. However, this leaves open the question of whether the sciences as actually practiced conform to this normative ideal. It also glosses over the point that there is disagreement about what those allegedly universal normative principles in fact are. These two further issues are exemplified in the debate in statistics, now nearly a century old, between Bayesians and Frequentists. The debate is philosophical. What evaluations can scientists legitimately make of the theories and hypotheses they consider? Bayesians say that it is always appropriate to calculate the probabilities those theories and hypotheses have of being true. Frequentists hold that this goal (often if not always) is impossible to achieve (without introducing objectionably subjective elements) and that other standards of assessment (significance tests, confidence intervals, hypothesis tests, etc.) are appropriate to apply. Scientific disciplines and subdisciplines vary according to how Bayesian or frequentist or pluralistic they are, and this variation among practitioners is reflected in the standards that journal editors impose on the articles they are prepared to publish. Over the past 10-20 years, Bayesian methods have gained in prominence in a variety of disciplines (for example, in phylogenetic inference in evolutionary biology, in earthquake prediction, and in parts of statistical physics), while other disciplines (psychology, medicine, and other areas of physics) remain more resolutely frequentist. And, of course, there are still other areas of science where the data seem to be so unequivocal that scientists often feel that they don't need statistics and probability to reason their way through problems; here Lord Rutherford's opinion that 'if your experiment needs statistics you ought to have done a better experiment' still seems to hold sway. It is unclear to me why this uneven pattern holds across the sciences. Is the explanation that some scientific problems are more amenable to Bayesian solutions

than others? Or is the explanation more inertial and historical – perhaps different disciplines have different traditions owing to the power exerted by a small number of founding figures to shape their discipline’s subsequent development? A mixture of these two explanations is worth considering as well.

As the reader is by now well aware, I am very much in the grip of the idea that all sciences draw on the same set of subject-matter neutral inference principles. But this does not mean that the sciences are all the same. For example, when scientists manipulate objects to do experiments, they need to design their experiments. But many sciences draw on observations that are not the result of experimental manipulation; in such cases, they do not need to design experiments. Granted, but the rules that govern experimental design are supposed to hold whether one is experimenting on bacteria or on metals, and the same is true of the rules that are supposed to govern purely observational studies. And the unification goes even further – when one investigates the rules for experimental design and the rules for purely observational studies, these two problems inevitably trace back to a core of broader considerations (whether it is Bayesian or frequentist).

Another example of a case in which there apparently are methods of inference at work that are specific to a subject matter is furnished by the use of a parsimony criterion in phylogenetic inference problems in evolutionary biology. Consider the fact that all bird species have wings. If these species originated separately, and if the ancestors all lacked wings, then thousands of evolutionary changes (from no wings to wings) would need to have taken place in the separate lineages. On the other hand, if all birds share a common ancestor, then only a single change (from no wings to wings) would need to have occurred to generate the numerous winged bird species we now observe. Phylogeneticists often take this difference in number of evolutionary changes to be an argument in favor of common ancestry, and they use the term “phylogenetic parsimony” to label this principle. Is this type of parsimony uniquely applicable to hypotheses about biological evolution and common ancestry? Apparently not; it also applies to languages and texts. For example, parsimony considerations underwrite the conclusion that the Indo-European languages trace back to a common ancestral language (proto Indo-European) and did not arise separately. And when one investigates the rationale for thinking that phylogenetic parsimony is a good inference principle in evolutionary biology, one quickly is led to broader considerations that unite phylogenetic parsimony with modes of reasoning that do not concern common ancestry or evolution of any kind. Some authors seek to ground phylogenetic parsimony in the broader concept of explanatory power (Farris 1983) while others seek to connect it to the broader concept of likelihood (Sober 2004). It is hard to resist the pull of the idea that methods of inference are subject-matter neutral.

In arguing for the importance of inference rules that are not tied to a specific subject matter, I do not mean to suggest that science education should focus exclusively on those abstract rules, and that all else will magically fall into place after that. Recall my discussion of Simpson’s paradox. The way to see the abstract pattern is not to have it stated in its full and abstract generality (“a correlation in a population can be canceled or reversed within subpopulations”), but to examine different specific examples. I began with sex discrimination at Berkeley and then moved on to the evolution of altruism; after that, there is the challenge of stating the idea of Simpson’s paradox in a form that can be seen to apply these two examples and to many others besides. Abstract rules may be fundamental, but there are good pedagogical reasons not to begin with fundamentals.

I have emphasized the deep epistemological links that connect the sciences. However, it cannot be denied that the sciences have different subject matters, and that science education has those specific subject matters as its main focus. For example, a molecular biologist who does a polymerase chain reaction to produce millions of copies of a given DNA sequence needs to master a specific set of laboratory techniques and a body of theory, both to do the process properly and to make use of the results. This is why the theories and practices that get emphasized in science education are so strongly focused on the specifics of the science being studied. Yet, at the same time, science education needs to prepare future scientists for dealing with subject matters that do not exist now. And scientists trained in one subject matter frequently need to retool so that they can shift to another subject matter. New problems will arise in the future, and many scientists will shift subjects in the course of their careers. Scientists who face these changes have a lot to learn, but they need not start from zero. Reasoning tools carry over. This should be an important part of education in the different scientific disciplines.

But the case for studying logic, probability, statistics, and philosophy of science is stronger. Even when scientists do not change their areas of inquiry, they need to understand the inferential tools that they deploy. Just as understanding a scientific theory or framework involves more than memorizing a list of facts, so too does understanding an inference procedure involve more than memorizing the cookbook recipes that are taught so often under the heading of “methods.”

## References

Cartwright, Nancy (1979): “Causal Laws and Effective Strategies,” *Nous* 13: 419-437.

Farris, J.S. (1983): “The Logical Basis of Phylogenetic Analysis.” In N. Platnick and V. Funk (eds.), *Advances in Cladistics*. New York: Columbia University Press. Reprinted in E. Sober (ed.), *Conceptual Issues in Evolutionary Biology*. Cambridge: MIT Press, 2006.

Hempel, C. (1965): *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press.

Kahneman, J. Tversky, A. and Slovic, P. (1982): *Judgment under Uncertainty — Heuristics and Biases*. New York: Cambridge University Press.

LeGrand, H. (1989): *Drifting Continents and Shifting Theories*. Cambridge: Cambridge University Press.

Popper, K. (1959): *The Logic of Scientific Discovery*. London: Hutchinson.

Reichenbach, H. (1956): *The Direction of Time*. Berkeley: University of California Press.

Simpson, E.H. (1951): “The Interpretation of Interaction in Contingency Tables.” *Journal of the Royal Statistical Society B* 13: 238-241.

Sober, E. (1993): *Philosophy of Biology*. Boulder: Westview Press.

Sober, E. (2004): “The Contest between Likelihood and Parsimony.” *Systematic Biology* 53: 6-16.

Wason, P.C. (1966): "Reasoning". In B.M. Foss (ed.),*New Horizons in Psychology*. Harmondsworth: Penguin.

Wegener, A. (1924): *The Origin of Continents and Oceans*. London: Methuen. Translation of the 3<sup>d</sup> German edition.

Yule, G. U. (1926): "Why Do We Sometimes Get Nonsensical Relations Between Time Series? A Study of Sampling and the Nature of Time Series." *Journal of the Royal Statistical Society* 89: 1-64.