Introduction

Philosophers have tried to learn from science in two quite different ways. First, the contents of particular scientific theories have thrown light on philosophical problems. For example, philosophers interested in the nature of space and time have found much of interest in relativity theory. As sometimes happens in the history of thought, a problem that begins its life as a problem in philosophy later turns out to receive illumination from a body of scientific results. Philosophy then has to catch up with the news from outside.

A second avenue of influence develops when philosophers contemplate the character of science itself. Here scientific theories serve as examples. Relativity theory says nothing about what scientific knowledge is; its subject is space, time, and motion, not the nature of inquiry. Still, the hypotheses of relativity theory have been an important point of reference for philosophers seeking to understand how the human mind is able to grasp a reality outside itself.

I've called this collection of papers *From a Biological Point of View* because it is biology – particularly the theory of evolution – that has performed, in my own work, the two functions just described. I believe that the contents of biological theories speak to problems that first saw the light of day as problems in philosophy. And I have continually been struck by the fact that biological theories provide philosophically interesting examples of how the enterprise of knowledge proceeds.

The title I've chosen for this volume is an homage to, and is intended to contrast with, W. V. Quine's collection *From a Logical Point of View* (Harvard University Press, 1953). In his work in metaphysics, epistemology, and other areas, Quine constantly reached for logical ideas as tools of clarification. Although Quine developed his views in opposition to logical empiricism and positivism, he shared with those he opposed a practical commitment to the centrality of logic as a tool of

philosophy. I think it is fair to say that in the forty years since Quine's book appeared, logic and philosophy of science (to say nothing of logic and philosophy as a whole) have drifted further apart. Logic has edged closer to mathematics and computer science. And philosophers have become more skeptical that logic provides general techniques for solving philosophical problems. No new overarching paradigm now commands wide acceptance, not do I recommend my own approach as a panacea. Perhaps the problems we call "philosophical" are too diverse for any single method to achieve what logic was once thought to be able to do.

A work of philosophy should attempt to solve a philosophical problem. A work of *meta*philosophy, which is what this Introduction purports to be, discusses the character of philosophy itself. Metaphilosophy is a pursuit that inspires me with feelings of suspicion. Philosophical programs, I believe, should be judged by their fruits; they can't be justified in advance. What is more, philosophers often misrepresent the real nature of their methods when they self-consciously identify themselves with this or that program. With the demise of positivism and the rise of philosophical "naturalism," it has become fashionable to describe philosophy as a project that is continuous with the sciences. The idea that philosophers go in for conceptual analysis whereas scientists do empirical work is supposed to be an untenable dualism that sophisticates are expected to recognize as naive. Yet, in spite of this change in the metaphilosophical landscape, much of philosophy proceeds as it did before. It now is unfashionable to say that one is giving a conceptual analysis of the meaning of terms such as "justice," "knowledge," or "freedom"; instead, it is more acceptable to say that one is trying to describe what justice, knowledge, or freedom *is.* It is worth pondering how much of this shift is substantive, and how much is just window-dressing.

These cranky remarks should not be taken to suggest that I subscribe to some rigid separation of conceptual clarification and empirical insight. I love blurred boundaries as much as the next person. However, I do think that in *practice,* if not in *theory,* philosophy of science (and philosophy as a whole) largely remains a subject driven by its own problems and methods. It is no accident that we philosophers do not do experiments and do not go into the field to make natural observations. If metaphilosophy is what we want, we need a theory of philosophy that is true to the actual practice of that subject.

It should be no embarrassment to philosophy that its conundrums often fail to engage the attention of working scientists. If philosophy is

a discipline in its own right, this is to be expected. Nor should a misplaced scientism lead us to think that philosophy is barred from making normative claims by way of criticizing what scientists do. I once heard a justly respected philosopher say that naturalism means that when philosophy and science conflict, it is almost always the philosopher who is wrong. This is correct, of course, when the philosopher is poorly informed and the scientist has his or her head screwed on straight. But we should not forget that scientists are mortals, and philosophers can learn from experience.

Even if philosophy is a discipline that is driven by its own evolving set of problems, this does not mean that scientific ideas are irrelevant to the solution of philosophical problems. What has happened in the past is doubtless happening right now. Science encroaches little by little on what philosophy views as its own terrain. Given the success of quantum mechanics as an empirical theory, we may be puzzled as to how earlier generations of philosophers could have thought that determinism is *a priori* true. Yet the idea that science right *now* may provide evidence that decides a problem that we *now* regard as philosophical will strike many philosophers as a quixotic dream.

The idea that philosophy is an *a priori* discipline and the idea that it is simply a part of science are both wrong. Philosophy is not a unity; different philosophical problems are structured differently. Nor can one tell in advance how one philosophical problem is related to others, and to matters that arise in other arenas of thought. We should relish the fact that philosophy can be *surprising*. Understanding the nature of a problem is not something we do in advance of trying to solve it.

When philosophy and science are brought into contact, the searchlight can be aimed from philosophy to science, from science to philosophy, or in both directions at once. An essay that is structured in the first way will address a problem that is of interest to working scientists; it will attempt to provide that scientific problem with philosophical illumination. If the essay succeeds, it will be of interest primarily to the scientists themselves. An essay of the second sort will identify a traditionally philosophical problem and will try to show how scientific ideas can be of use in its solution. Here the intended audience will be philosophers – even philosophers who have no antecedent reason to care about the scientific ideas addressed.

The essays in this volume are mainly of type two. They were written for philosophers, not for scientists (though I hope that philosophically inclined scientists may find some of the issues of interest). Although I

believe that type one philosophy of science is important, I feel that my best efforts in that line of work are to be found in my books *The Nature of Selection, Reconstructing the Past,* and *Philosophy of Biology.* In assembling the present collection of essays, I have tried to follow the advice of Ockham's razor by not reprinting essays that I felt would be redundant.

The subjects explored in these essays are diverse. The first four address different features of the human mind that have long interested philosophers. Egoism, solipsism, innate beliefs, and the status of truth-telling and true belief in communication involve problems that are familiar to the philosopher of mind, the epistemologist, and the philosopher of language. What is somewhat less standard is the idea that these issues may be approached from an evolutionary point of view. Rather than ask what behavioral and introspective evidence there is for and against psychological egoism, I explore the question of whether evolution can be expected to have made us psychological egoists. And rather than ask what rational justification people can provide for their belief that a physical world exists external to their states of consciousness, I ask what adaptive advantage organisms obtain by thinking in nonsolipsistic terms. These evolutionary questions do not replace the traditional problems that prompt them. But I believe that they do have considerable philosophical interest in their own right.

Whereas the first four essays defend the bearing of biology on the philosophical problem at hand, the tone of Essay 5 is much more skeptical. Here I consider what evolution may be able to teach us about ethics. The fundamental distinction that needs to be drawn here is that between explaining why we have the ethical thoughts and feelings we do and deciding what status our ethical convictions possess. With respect to the first task, I emphasize the importance of thinking about *patterns of variation* as the fundamental *explanandum.* What needs to be explained is why individuals are similar and different in their ethical beliefs. This emphasis on variation as the proper object of explanation recurs in Essays 10 and 11.

The other main question addressed in Essay 5 concerns whether evolutionary considerations can show us whether our ethical beliefs are true. This is fundamentally an epistemological question – can the evolutionary origins of a belief undermine the claim that it is correct? Can they show, as ethical subjectivism maintains, that no ethical statements are ever true? Here we need to consider the status of *genetic arguments* and the significance of an influential argument that appeals to Ock-

ham's razor – the principle of parsimony – in defense of ethical subjectivism.

Essay 6 addresses the dispute between scientific realists and empiricists. Realists see the discovery of true theories as the proper goal of science; empiricists think that science aims only at the discovery of predictively successful theories. Both philosophies embed particular views about how scientific inference works. In this essay, I propose a synthesis of realism and empiricism that I call *contrastive empiricism.* This position, I believe, retains what is plausible in each traditional doctrine, while avoiding the excesses of each.

Essays 7 and 8 elaborate and implement the epistemology laid out in Essay 6. The idea that we should use the simplicity of a theory to help decide whether the theory is plausible has always been a problem for empiricism. For Ockham's razor seems to say that plausibility is to be judged, in part, on *non*empirical grounds. The same problem arises in connection with the idea that correlated events should be explained by postulating a cause that they share. The challenge to empiricism is to explain the merits of these principles without departing from a reasonable empiricist epistemology.

Although Essays 6, 7, and 8 lay out the fundamentals of the epistemology I wish to defend, they leave considerable work undone. I believe that the essay on the curve-fitting problem that Malcolm Forster and I recently finished lends further credence to the view of simplicity and parsimony that I want to defend (Forster and Sober, 1994). And I have tried to be more precise about the concept of observation, and about the epistemological significance of the distinction between theory and observation, in Sober (1993a). In addition, I explore the bearing of contrastive empiricism on the use of indispensability arguments in philosophy of mathematics in Sober (1993b).

If science is the activity of putting questions to nature, then an important kind of scientific question formulates a request for explanation. What presuppositions do those why-questions have? In Essay 9, I argue that the explanatory problems we pose about nature embody particular assumptions about causal structure. The idea that different events should be traced back to a common cause is not just the inferential maxim discussed in Essay 7; it also is a guide to how we organize what we already know when we try to generate explanations.

The nature of causality – what it means for one event to cause another – is a perennial philosophical problem. A related, though distinct, problem concerns the nature of causal importance, the subject of

Essay 10. If an event has several causes, what does it mean to say that one of them is "more important" than another? This question repeatedly arises in biology and the social sciences, and the sciences have well-defined methods for handling them. Here is a metaphysical problem that the practice of biology informs.

Essentialism provides a view of science that many philosophers now associate with scientific realism. Just as science aims to discover unobservable causal mechanisms that account for observable phenomena, so it attempts to discover what the natural kinds are to which individuals entities belong. For a long time, the two standard examples of this essentialist interpretation of the activity of science have been the periodic table of elements in chemistry and the identification of biological species. In Essay 11, I argue that evolutionary theory undermines essentialism as a *global* philosophy of science, and does so for subtle reasons that have not been widely appreciated.

The last essay in this collection is on a problem in the philosophy of time. A law of nature can describe how a system in a given state will probably evolve or it can describe how a system in a given state probably did evolve. I call these two types of law *forward-directed* and *backward-directed,* respectively. It may seem at first glance that laws of nature can be of both sorts. However, I argue in Essay 12 that this appearance is misleading; when the concept of law is made precise in a certain way, one can show that systems that can be expected to evolve cannot have both sorts of laws. I also float a proposal concerning why science seems to opt for forward-directed more than backward-directed laws.

Even though the biological content of some of the essays I've described is obvious, others may seem to have only a scant connection with what might be termed "a biological point of view." After all, what could the principle of the common cause or a problem about time asymmetry have to do with the theory of evolution? Although a fuller answer to this question is contained in the essays themselves, I'll provide some brief comments here.

Evolutionists implicitly use the principle of the common cause when they argue that similarities between species should be explained by the hypothesis that the species have a common ancestor from which the similarities were inherited as homologies. It was this phylogenetic context that encouraged me to think about the justification of the principle of the common cause from a probabilistic point of view.

A similar biological background guided my thinking about the issue of temporal asymmetry discussed in Essay 12. The rules of Mendelian

inheritance provide conditional probabilities that are forward-directed; they describe the probability that an offspring will have a particular genotype, given that the parents have this or that genotype. Those conditional probabilities do not change with time; even when the frequencies of the *A* and *a* genes evolve, the probability that an offspring will be *Aa,* given that its parents are both *Aa,* remains the same.

As natural and familiar as the Mendelian rules are, it is much harder to think about backward-directed probabilities. Why aren't the laws of inheritance stated by specifying the probability that both parents are heterozygotes, given that an offspring is a heterozygote? If such "laws" could be formulated, would they remain true as the population evolves? This simple biological example led me to consider a quite general problem about the character of scientific laws.

I would be lying if I described this collection of essays as a systematic and linear development of a single thesis. What unity the essays possess stems from the fact that I have tried to take seriously the content and the practice of evolutionary biology. The philosophical claims I defend are various; they do not stand or fall together. However, I would like to think that this variety, far from being a defect, reflects favorably on what a biological point of view in philosophy has to offer.

REFERENCES

Forster, M., and Sober, E. (1994). How to Tell When Simpler, More Unified, or Less *ad Hoc* Hypotheses Will Provide More Accurate Predictions. *British Journal for the Philosophy of Science, 45*: 136.

Sober, E. (1993a). "Epistemology for Empiricists." In H. Wettstein (ed.), *Midwest Studies in Philosophy,* volume 18, University of Notre Dame Press, pp. 39–61.

Sober, E. (1993b). Mathematics and Indispensability. *Philosophical Review, 102:* 35–57.

1

Did evolution make us psychological egoists?

1. TWO CONCEPTS

The concept of altruism has led a double life. In ordinary discourse, as well as in psychology and the social sciences, *altruism* refers to behaviors that are produced because people have certain sorts of motives. In evolutionary biology, on the other hand, the concept is applied to behaviors that enhance the fitness of others at expense to self.

A behavior can be altruistic in the evolutionary sense without being an example of psychological altruism. A plant that leeches insecticide into the soil may be an altruist, if the insecticide benefits its neighbors and imposes an energetic cost on the producer. In saying this, I am not attributing a mind to the plant. Evolutionary altruism has to do with the fitness consequences of the behavior, not with the mechanisms inside the plant (mental or otherwise) that cause the plant to behave as it does.

Symmetrically, a behavior can be altruistic in the psychological sense without being an example of evolutionary altruism. If I give you a volume of Beethoven piano sonatas (or a package of contraceptives) out of the goodness of my heart, my behavior may be psychologically altruistic. However, the gift giving will not be an example of evolutionary altruism, if the present fails to augment your prospects for survival and reproductive success.

Although the concepts are different, they have a few things in common. Both point to causal explanations of the behaviors so labeled. If I say that a behavior is an example of psychological altruism, I am making a claim about the motives that produced the behavior. If I say

I am grateful to Daniel Batson, Don Moskowitz, and the members of the philosophy departments at University of Pittsburgh, Rutgers University (New Brunswick), and at University of Wisconsin (Madison) for useful discussion.

that a behavior is an example of evolutionary altruism, I am suggesting a certain sort of explanation, which I'll describe in the next section, for why the behavior evolved.[1]

The second common feature of the psychological and the evolutionary concepts is that both have been controversial and even unpopular in much scientific investigation. Psychological egoism, which claims that all of our ultimate motives are selfish, has viewed psychological altruism as a comforting illusion. Egoism has been the dominant position in all major schools of twentieth-century psychology (Batson 1991). And within evolutionary biology, the theory of the selfish gene has been hostile to the idea that evolution produces behaviors that help the group at the expense of the individual (Williams 1966; Dawkins 1976).

My goal in this paper is to clarify these concepts and to further discuss why they are logically independent of each other. Then, having separated them, I will attempt to bring them back into contact. I will explore the question of whether there are evolutionary considerations that help us decide whether we are ever psychological altruists.

2. EVOLUTIONARY ALTRUISM

Altruism has been an important subject for evolutionary theorizing ever since Darwin. I will not describe the history of how this subject has developed, nor will I discuss intricacies that are internal to various theories of current interest.[2] My modest goal in this section is to describe with more care what altruism and selfishness mean in an evolutionary context and to show how each trait is connected with its own picture of how natural selection has operated.

For the most part, Darwin viewed natural selection as a process in which organisms within the same breeding population compete with each other to survive and reproduce. His picture of competition was not the lion versus the lamb, but lions competing with lions and lambs with lambs. In this process, the traits that evolve are the ones that benefit the individual organism. Although Darwinism is sometimes described by saying that characteristics evolve "for the good of the species," this is a major distortion of how Darwin *usually* thought about natural selection.

Usually, but not *always.* There were a small number of occasions on which Darwin took seriously the idea that natural selection involves competition among objects other than individual organisms. One of

the clearest expressions of this alternative occurs in his discussion of human morality in *The Descent of Man.* Here is Darwin's statement of the problem:

> It is extremely doubtful whether the offspring of the more sympathetic and benevolent parents, or of those which were the most faithful to their comrades, would be reared in greater number than the children of selfish and treacherous parents of the same tribe. He who was ready to sacrifice his life, as many a savage has been, rather than betray his comrades, would often leave no offspring to inherit his noble nature. The bravest men, who were always willing to come to the front in war, and who freely risked their lives for others would on average perish in larger numbers than other men. (Darwin 1871, p. 163)

Darwin's point is that if we consider a single tribe that contains both altruistic and selfish individuals, altruists will do *worse* than selfish individuals. If natural selection is the main cause of evolutionary change (as Darwin thought), and if natural selection causes fitter traits to increase in frequency and less fit traits to decline, why hasn't altruism altogether disappeared from human conduct? Here is the answer that Darwin suggests:

> It must not be forgotten that although a high standard of morality gives but a slight or no advantage to each individual man and his children over the other men of the same tribe, yet that an advancement in the standard of morality and an increase in the number of well-endowed men will certainly give an immense advantage to one tribe over another. (Darwin 1871, p. 166).

Although altruistic *individuals* do worse than selfish *individuals* in the same tribe, altruistic *groups* do better than selfish *groups.* Here Darwin was imagining a process of group selection, in which groups compete against each other. This picture of the process of natural selection differs markedly from his more customary formulation, in which organisms within a single population engage in a struggle for existence.

To make sense of the idea of evolutionary altruism, and of the process of group selection that is associated with it, one must be able to think simultaneously about the fitnesses of organisms and the fitnesses of groups of organisms. How are these two levels related? And since altruism is a behavior produced by an individual organism, how do these two kinds of fitness make it possible for altruism to evolve when there is group selection?

Figure 1.1 depicts some of the main conceptual ingredients. It shows that the fitness of an individual depends on two factors. Whether the