

# 5

## Prospects for an evolutionary ethics

### 1. TWO KINDS OF QUESTION

Human beings are a product of evolution. This means that the *existence* of our species is explained by the process of descent with modification. But, in addition, it also means that various *phenotypic characteristics* of our species – features of morphology, physiology, and behavior – have evolutionary explanations.

Some see in this simple idea the promise of a new understanding of the problems of ethics. Others grant its truth, but reject its relevance. The first group insists that if evolutionary biology can provide insights into the workings of the human mind, these insights cannot fail to transform our understanding of right and wrong. But for the doubters, the fact that human beings have an evolutionary past has no more significance than the fact that the human body must obey the laws of physics. It is true that our bodies are subject to the law of gravity, but apparently that is of little help in understanding why we think and act as we do.

To assess the relevance of evolutionary ideas to the problems of ethics, we must distinguish two quite different projects. The first is the task of accounting for why people have the ethical thoughts and feelings they do. The second concerns the problem of deciding what the status of those thoughts and feelings is. The difference between these two projects is illustrated by the following pair of questions:

- (1) Why do people have the views they do concerning when it is morally permissible to kill?
- (2) When is killing morally permissible?

I am grateful to James Anderson and Louis Pojman for useful discussion.

### *From a biological point of view*

Problem (1) poses a problem of *explanation*, while (2) engages the task of *justification*. One of the main issues I want to consider is how these questions are related to each other. This is a matter of some intricacy. But at the simplest level, it is important to recognize that there is no automatic connection between the two types of problem.

It is quite obvious that a person can believe the right thing for the wrong reason. Consider my friend Alf, who thinks the earth is round, but does so because he thinks that the earth is perfect and that roundness is the perfect shape. Alf also happens to know various facts that the rest of us recognize as evidence for the earth's being round. But Alf's predicament is that he does not see these facts as having any particular evidential significance. Alf believes *E* and also believes *R*, but he does not believe *R* because he believes *E*. *E* is evidence for *R*; *E* is a reason to believe *R*, but it is not Alf's reason for believing *R*.

This example encapsulates two lessons. First, it is possible to show whether a proposition is justified without saying anything, one way or the other, by way of explaining why someone happens to believe that proposition. *E* justifies *R*, but that says nothing about why Alf believes *R*. Second, and conversely, an explanation of why someone believes a proposition may leave open whether the proposition is in fact well supported by evidence. Alf's strange ideas about perfection explain why he believes *R*. But this quirk of Alf's psychology does not tell us whether *R*, in fact, is a proposition that is well supported by evidence.

So the following two questions are quite different from each other:

- (3) Why does Alf believe that the earth is round?
- (4) Is the proposition that the earth is round strongly supported by evidence?

I hope the parallel between problems (1) and (2) on the one hand and problems (3) and (4) on the other is suggestive. In each pair, it is possible to answer one question without thereby answering the other. In summary, we have the following slogan: *An explanation for why someone believes a proposition may fail to show whether the proposition is justified, and a justification of a proposition may fail to explain why someone believes the proposition.* This simple conclusion is the beginning, not the end, of our discussion of the relationship between such questions as (1) and (2).

## *Prospects for an evolutionary ethics*

### 2. PATTERNS OF EVOLUTIONARY EXPLANATION

Before we can address the question of how the problem of explanation and the problem of justification are related, it is well to get clear on how evolution can contribute to the former task. How might evolutionary theory help explain why we have the ethical thoughts and feelings that we possess?

Although natural selection is just one of the processes that can produce evolutionary change, it is the one of greatest interest to sociobiologists. When natural selection is cited by way of explaining why a trait is currently found in some population, an ancestral population is postulated in which the trait was one of several variants that were represented. These traits are claimed to have differed in their *fitness* – their capacity to help the individuals possessing them to survive and reproduce. Through a process of differential survival and reproduction, the population moved from a mixed condition to the homogeneous condition we now observe. Why do zebras now run fast when a predator approaches? Because, ancestrally, zebras differed in running speed, and the fast zebras survived and reproduced more successfully than the slow ones.

Within this basic outline, it is important to distinguish two quite different patterns of explanation that may be used to elaborate the idea that natural selection explains what we observe. In both instances, the goal is to explain why a certain *pattern of variation* exists.

Why do polar bears have thicker fur than brown bears? The reason is that polar bears live in colder climates than brown bears. In particular, natural selection favored one trait for the organisms in the one habitat and a different trait for the organisms in the other. More specifically still, there are genetic differences between the two species that account for the difference in the thickness of their coats. These genetic differences arose because of selective differences in the bears' environments.

Why do the polar bears who live near the North Pole have thicker fur than the polar bears who live farther south? The reason is that the bears in the first group live in a colder climate than do the bears in the second. In particular, natural selection favored polar bears who possessed *phenotypic plasticity*. The bears evolved a set of genes that permitted them to respond to changes in the ambient temperature. The bears in the two groups differ in coat thickness, but this is not because they are genetically different. Rather, the difference in phenotype is to be explained by appeal to an environmental difference.

There is a truism that neither of these patterns of explanation contra-

### *From a biological point of view*

dicts. Every trait a bear has is a product of the genes it possesses interacting with the environment in which it lives. This remark is a truism because “environment” is defined as a garbage can category; it includes all factors that are not considered “genetic.”

If genetic and environmental causes both play a role in the development of a trait, we still may wish to say which was more “important.” This question quickly leads to nonsense if we apply it to the trait exhibited by a single organism. It is meaningless to say that 3 inches of Smokey’s fur was provided by his genes and 2 inches by his environment. If we wish to assess the relative importance of causes, we must examine a population of bears in which there is variation. And, as we have seen, some patterns of variation will mainly be due to genetic differences, while others will mainly be due to differences in environment.

When sociobiologists attempt to provide evolutionary explanations of human behavioral and psychological characteristics, they are often accused of endorsing the doctrine of “genetic determinism.” However, I hope these two examples show that evolutionary explanations are not automatically committed to the primacy of genetic over environmental explanation. Natural selection can have the result that phenotypic differences are due to genetic differences. But natural selection also can have the consequence that phenotypic differences are explained by differences in environment.

Let us now apply this lesson to problems like (1). Just as coat thickness can vary among bears, so views about killing can vary within and among human societies. Once again, it is useful to think about the pattern of variation that needs to be explained. Contemporary societies differ from each other, and individual societies have changed their views over time. In addition, we must remember that societies are not monolithic entities; if we look carefully *within* a society, we often will discern variation in opinion about when killing is permissible.

It is important not to confuse the code followed in a society with the simple slogans that the society endorses. First, there is the gap between ideals and reality. But, in addition, the slogans often do not begin to capture what the ideals really are. The Sixth Commandment says “Thou shalt not kill,” but very few believers have thought that killing is *always* impermissible.

The ethical code endorsed by a society, or by an individual, is a complex object, one that is often difficult to articulate precisely. Ethical codes are in this respect like languages. You and I know how to speak English, but which of us is able to describe precisely the rules that define grammatical English? If we wish to explain why people have the views

## *Prospects for an evolutionary ethics*

they do about the permissibility of killing, we first must obtain an accurate description of what those views actually are. "Why do people think that killing is wrong?" is a poorly formulated question.

Social scientists and evolutionary biologists have different contributions to make to our understanding of problem (1). In many parts of the world, capital punishment is now much less popular than it was a hundred years ago. Why did this transition occur? In contrast, consider the fact that all (or virtually all) human societies have regarded the killing of one's children as a much more serious matter than the killing of a chicken. Why is this so? Perhaps historians have more to tell us than evolutionists about the first question, but the reverse is true with respect to the second. There need be no battle between these disciplines concerning which holds the key. Ethical beliefs about killing are complex and multifaceted; different beliefs in this complex may fall within the domains of different disciplines.

If recent changes in view as to the permissibility of capital punishment don't have an evolutionary explanation, what does this mean? Of course, it is not to be denied that having ethical opinions requires a reasonably big brain, and our big brain is the product of evolution. What I mean is that the *difference* between people a hundred years ago who favored capital punishment and people now who oppose capital punishment is not to be explained by appeal to evolutionary factors. Again, the problem that matters concerns how one should explain a pattern of variation.

My example about the chicken may suggest that evolutionary considerations can be brought to bear only on characteristics that are invariant across cultures. Many sociobiologists reject this limitation. For example, Alexander (1987) and Wilson (1978) grant that human beings have enormous behavioral plasticity. We have an astonishing ability to modify our behavior in the light of environmental circumstances. Nonetheless, these sociobiologists also maintain that evolution has caused us to produce behaviors that, in the environment at hand, maximize fitness. For them, human behavioral variation is to be explained by the same kind of account I described above for geographical variation in polar bear coat thickness.

An example of this sort is provided by Alexander's (1987) explanation of the kinship system known as the avunculate. In this arrangement, a husband takes care of his sister's children far more than he takes care of his wife's. Alexander's hypothesis was that the avunculate occurs when and where it does as an adaptive response by men to high levels of female promiscuity. If a husband is more likely to be geneti-

### *From a biological point of view*

cally related to his sister's children than to his wife's, he maximizes his inclusive fitness by directing care to his nephews and nieces.

I am not interested here in exploring whether this suggested explanation of the avunculate is empirically plausible. My point has to do with the form of the explanation proposed. Alexander is trying to explain a behavior that is *not* a cultural universal. And his preferred explanation is *not* to say that societies with the avunculate are genetically different from societies with other kinship systems. Rather, Alexander's hypothesis is that there is a universal tendency to behave in fitness-maximizing ways; this universal tendency leads people in different environments to behave differently.

My pair of examples about coat thickness in bears was intended to demonstrate that we should not equate "evolutionary explanation" with "genetic explanation." The evolutionary process can have the result that phenotypic differences are explained by genetic differences, but it also can have the result that phenotypic differences are explained by environmental differences. This flexibility in evolutionary explanation raises a question, however. If evolutionary explanations can encompass both genetic and environmental factors, what would a *nonevolutionary* explanation look like?

A useful example is provided by the rapid decline in birth rates that occurred in Europe in the second half of the nineteenth century. In some areas of Italy, for instance, average family size declined from more than five children to a little more than two (Cavalli-Sforza and Feldman 1981). Why did this demographic transition occur?

Reducing family size did not enhance a parent's biological fitness. In fact, exactly the reverse was true. It is important to remember that fitness concerns survival *and* reproductive success. In this instance, forces of cultural change worked in opposition to the direction of natural selection. Cultural change overwhelmed the weaker opposing force of biological evolution. It is the historian, not the evolutionist, who will explain the demographic transition.

In discussing the fur thickness of bears, I mentioned a truism: Every trait a bear has is due to the genes it possesses interacting with the environment in which it lives. This truism applies to human beings with equal force. When a woman in nineteenth-century Italy had only two children, this behavior was the product of a developmental process in which her genes and her environment both participated. However, I hope it is clear that this truism does not tell us what sort of explanation will be relevant to the explanation of *patterns of variation*. The demographic transition may have a quite different type of explanation from

### *Prospects for an evolutionary ethics*

the account that Alexander recommends for understanding the distribution of the avunculate.

The question “Can evolution explain ethics?” has usually elicited two simple responses. The first is “yes, since ethics is a facet of human behavior and human behavior is a product of evolution.” The second is “no, since ethics is a facet of human behavior and human behavior is a product of culture.” Both of these sweeping pronouncements neglect that fact that “ethics” is not the name of a single simple trait, but names something complex and multifaceted. *Ethics* is a “supertrait,” which has many “subtrait” components. The whole must be broken into its parts if we are to make headway on problems of explanation. It is reasonable to expect that some subtraits will have an evolutionary explanation while others will be explained culturally. For each subtrait that we wish to study, we must determine which form of explanation is more plausible.

### 3. THREE META-ETHICAL POSITIONS

In my opinion, it is *obvious* that biology and the social sciences can collectively claim problem (1) as their own. The alternative is to think that human behavior is not susceptible to scientific explanation at all. No argument for this negative thesis has ever been remotely plausible. And progress in biology and the social sciences lends empirical support to the idea that science does not stop where human behavior begins.

Even if the status of the explanatory problem (1) is reasonably clear, the question of justification (2) is not nearly so straightforward. Question (2) concerns a *normative* matter. Does it have an answer, and if it does, how is that answer related to nonnormative matters of fact?

First, some terminology. Let us say that a statement describes something *subjective* if its truth or falsity is settled by what some subject believes; a statement describes something *objective*, on the other hand, if its truth or falsity is independent of what anyone believes. “People believe that the Rockies are in North America” describes something subjective. “The Rockies are in North America,” on the other hand, describes something objective. When people study geography, there is both a subjective and an objective side to this activity; there are the opinions that people have about geography, but in addition, there are objective geographical facts.

Many people now believe that slavery is wrong. This claim about present-day opinion describes a subjective matter. Is there, in addition

### *From a biological point of view*

to this widespread belief, a fact of the matter as to whether slavery really is wrong? *Ethical subjectivism*, as I will use the term, maintains that there are no objective facts in ethics. In ethics, there is opinion and nothing else.

According to subjectivism, neither of the following statements is true:

Killing is always wrong.

Killing is sometimes permissible.

Naively, it might seem that one or the other of these statements must be true. Subjectivists disagree. According to them, no normative ethical statement is true.

*Ethical realism* is a position that conflicts with ethical subjectivism. Realism says that in ethics there are facts as well as opinions. Besides the way a murder makes people feel, there is, in addition, the question of whether the action really is wrong. Realism does not maintain that it is always obvious which actions are right and which are wrong; realists realize that uncertainty and disagreement surround many ethical issues. However, for the realist, there are truths in ethics that are independent of anyone's opinion.<sup>1</sup>

There is a third position that bears mentioning. The position I'll call *ethical conventionalism* asserts that some ethical statements are true, but maintains that they are made true by some person's or persons' belief that they are true. For example, so-called *ethical relativists* say that murder is right or wrong only because various people in a society have come to believe that it is. And advocates of the *divine command theory* say that murder is wrong only because God believes (or says) that it is. Some versions of *existentialism* maintain that what is right or wrong for a person to do is settled by a decision the person makes about what sort of person he or she wishes to be.

This triplet of positions is summarized in Figure 5.1. They constitute the three possible answers that there can be to two sequential yes/no questions. Although it may not be clear which of these three positions is most plausible, it should be clear that they are incompatible. Exactly one of them is correct.<sup>2</sup>

The distinction between subjectivism and its alternatives is important to bear in mind when one considers what evolution can tell us about the status of problems like that posed by (2). Are facts about human biology going to tell us when and why murder is wrong? Or are biological facts going to show that all normative beliefs are untrue? Here we are forced to choose; evolution cannot simultaneously tell us which eth-



## *Prospects for an evolutionary ethics*




Figure 5.1

ical norms are correct *and* that all normative statements are illusions.

Unfortunately, some evolutionary ethicists have wanted to have it both ways. For example, Ruse and Wilson (1986) have argued that an evolutionary understanding of why we have the ethical beliefs we do shows that those ethical beliefs cannot be objectively correct. Here is a characteristic passage:

Human beings function better if they are deceived by their genes into thinking that there is a disinterested objective morality binding upon them, which all should obey. We help others because it is “right” to help them and because we know that they are inwardly compelled to reciprocate in equal measure. What Darwinian evolutionary theory shows is that this sense of “right” and the corresponding sense of “wrong,” feelings we take to be above individual desire and in some fashion outside biology, are in fact brought about by ultimately biological processes. (Ruse and Wilson 1986, p. 179)

Ruse and Wilson seem to favor an emotivist account of ethical statements, according to which those statements merely express the feelings of speakers, but never say anything true. Their emotivism is a form of ethical subjectivism as defined above.

At the same time, Ruse and Wilson sometimes suggest that biology provides guidance about which norms we should adopt. Consider the following:

Human mental development has proved to be far richer and more structured and idiosyncratic than previously suspected. The constraints on this development are the sources of our strongest feelings of right and

### *From a biological point of view*

wrong, and they are powerful enough to serve as a foundation for ethical codes. But the articulation of enduring codes will depend upon a more detailed knowledge of the mind and human evolution than we now possess. We suggest that it will prove possible to proceed from a knowledge of the material basis of moral feeling to generally accepted rules of conduct. To do so will be to escape – not a minute too soon – from the debilitating absolute distinction between *is* and *ought*. (Ruse and Wilson 1986, p. 174)

Here biology seems to be in the business of telling us which ethical codes are correct. Biology grounds ethics rather than unmasking it.

Ruse and Wilson sometimes draw a contrast between “internal” and “external” sources of norms. Internal norms are said to be “rooted” in our biology, whereas external ones are divinely given, or hold true independently of details about human biology. They clearly believe that ethics is “internal” and that evolutionary biology shows that this is so.

It is worth noticing that the distinction between subjectivism and realism differs from the internal/external distinction. Realists can easily maintain that an action has the ethical properties it does because of facts about our biology. When you ask your dinner companion to pass the salt and he doesn’t immediately comply, why would it be wrong to stab his hand with your fork? Realists are entitled to answer this question by pointing out that stabbing someone *causes pain*. They are allowed to admit that if stabbing didn’t cause pain, then the action might have quite different ethical properties. Ruse and Wilson maintain that ethics cannot be objective if it is “internal.” No such consequence follows.

#### 4. THE IS / OUGHT GAP

The goal of this paper is not to reach some general assessment of the three meta-ethical positions just described. Rather, the question is whether evolutionary considerations can tell us which is most plausible.

In this section I want to formulate and criticize one argument for subjectivism that I believe has had a great deal of influence. This argument makes no specific reference to evolutionary considerations, but it is in the background of many discussions about evolutionary ethics.

The argument I’ll formulate has its provenance in Hume’s distinction between *is* and *ought*, a distinction to which Ruse and Wilson alluded in the passage quoted before. I will say that an *is-statement* describes what is the case without making any moral judgment about whether this situation is good or bad. An *ought-statement*, on the other hand, makes a moral judgment about the moral characteristics (rightness,

## *Prospects for an evolutionary ethics*

wrongness, etc.) that some action or class of actions has. For example, “thousands of people are killed by handguns every year in the United States” is an *is-statement*; “it is wrong that handguns are unregulated” is an *ought-statement*.

Hume defended the thesis that *ought-statements* cannot be deduced from exclusively *is-statements*. For example, he would regard the following argument as deductively invalid:

Torturing people for fun causes great suffering.

Torturing people for fun is wrong.

The conclusion does not follow deductively from the premisses. However, if we supply an additional premiss, the argument can be made deductively valid:

Torturing people for fun causes great suffering.

It is wrong to cause great suffering.

Torturing people for fun is wrong.

Notice that this second argument, unlike the first, has an *ought-statement* as one of its premisses. Hume’s thesis says that *a deductively valid argument for an ought-conclusion must have at least one ought-premiss*.

The term “naturalistic fallacy” is sometimes applied to any attempt to deduce *ought-statements* from exclusively *is-premisses*. The terminology is a bit misleading, since it was G. E. Moore in *Principia Ethica* (1903) who invented the idea of a “naturalistic fallacy,” and Moore’s idea differs in some respects from the one just described. Unfortunately, most people discussing evolutionary ethics tend to use Moore’s label to name Hume’s insight. I want to keep them separate. The proposition I’ll call *Hume’s thesis* says you can’t deduce an *ought* from an *is*. Hume’s thesis, I believe, is correct.<sup>3</sup>

Hume’s thesis, by itself, does not entail subjectivism. However, it plays a role in the following argument for subjectivism:

*Ought-statements* cannot be deduced validly from exclusively *is-premisses*.

If *ought-statements* cannot be deduced validly from exclusively *is-premisses*, then no *ought-statements* are true.

No *ought-statements* are true.

### *From a biological point of view*

The first premiss is Hume's thesis. The second premiss, which is needed to reach the subjectivist conclusion, is a *reductionist assumption*. It says that for an *ought*-statement to be true, it must reduce to (be deducible from) exclusively *is*-premisses.

My doubts about this argument center on the second premiss. Why should the fact that ethics cannot be deduced from purely *is*-propositions show that no ethical statements are true? Why can't ethical statements be true, though irreducible? It is important to remember that Hume's thesis concerns *deductive* arguments. Consider an analogy: Scientific theories about unobservable entities cannot be deduced from premisses that are strictly about observables, but this provides no reason to think that theories about unobservables are always untrue.

My remarks on this argument provide no positive defense of ethical realism. However, I hope they do show that one influential argument behind ethical subjectivism is not as decisive as it might seem.

## 5. GENETIC ARGUMENTS

I now want to consider a second argument for ethical subjectivism. It asserts that ethical beliefs cannot be true because the beliefs we have about right and wrong are merely the product of evolution. An alternative formulation of this idea would be that subjectivism must be true because our ethical views result from the socialization we experience in early life. These two ideas may be combined as follows:

- (G) We believe the ethical statements we do because of our evolution and because of facts about our socialization.

No ethical statement is true.

Philosophers are often quick to dismiss arguments like (G) on the grounds that these arguments are guilty of what has come to be called *the genetic fallacy*. A genetic argument describes the genesis (origin) of a belief and attempts to extract some conclusion about the belief's truth or plausibility.

The dim view that many philosophers take of genetic arguments reflects a standard philosophical distinction between the *context of discovery* and the *context of justification*. This distinction, emphasized by the logician Gottlob Frege, was widely embraced by the positivists.

Hempel (1965) illustrates the point of this distinction by recounting the story of the chemist Kekulé, who worked on the problem of de-

## *Prospects for an evolutionary ethics*

termining the structure of benzene. After a long day at the lab, Kekulé found himself gazing wearily at a fire. He hallucinated a pair of whirling snakes, which grabbed each other's tails and formed a circle. Kekulé, in a flash of creative insight, came up with the idea of the benzene ring.

The fact that Kekulé arrived at the idea of the benzene ring while hallucinating does not settle the question of whether benzene really has that structure. It is for psychologists to describe the context of discovery – the (possibly) idiosyncratic psychological processes that led Kekulé to his insight. After Kekulé came up with his idea, he was able to do experiments and muster evidence. This latter set of considerations concerns the logic of justification.

I agree that one can't *deduce* whether Kekulé's hypothesis was true just from the fact that the idea first occurred to him in a dream. The same holds true of my friend Alf; one can't *deduce* that his belief about the shape of the earth is mistaken just from the fact that he reached this belief because of his weird ideas about perfection. However, it is a mistake to overinterpret this point. I want to suggest that there can be perfectly reasonable genetic arguments. These will be *nondeductive* in form.

Consider my colleague Ben, who walks into his introduction to philosophy class one day with the idea that he will decide how many people are in the room by drawing a slip of paper from an urn. In the urn are a hundred such slips, each with a different number written on it. Ben reaches in the urn, draws a slip that says "78," and announces that he believes that exactly 78 people are present.

Surely it is reasonable to conclude that Ben's belief is probably incorrect. This conclusion is justified because of the process that led Ben to his belief. If so, the following is a perfectly sensible genetic argument:

Ben decided that there are 78 people in the room by drawing the number 78 at random from an urn.

$p$

---

---

It isn't true that there are 78 people in the room.

I have drawn a double line between premiss and conclusion to indicate that the argument is not supposed to be deductively valid. The  $p$  next to the double line represents the probability that the premiss confers on the conclusion. I claim that  $p$  is high in this argument.

It is quite true that one cannot *deduce* that a proposition is untrue just from a description of how someone came to believe it. After all, the fact that Ben drew the number 78 from the urn doesn't absolutely

*From a biological point of view*

rule out the possibility that there are exactly 78 people in the room. Still, given the process by which he formed his belief, it would be something of a miracle if his belief just happened to be correct. In this case, the context of discovery *does* provide evidence as to whether a belief is true. If so, we must be careful not to conflate two quite different formulations of what the genetic fallacy is supposed to involve:

- (5) Conclusions about the truth of a proposition cannot be *deduced validly* from premisses that describe why someone came to believe the proposition.
- (6) Conclusions about the truth of a proposition cannot be *inferred* from premisses that describe why someone came to believe the proposition.

I think that (5) is true but (6) is false. Inference encompasses more than deductive inference. I conclude that argument (G) for ethical subjectivism cannot be dismissed simply with the remark that it commits “the genetic fallacy.”

The genetic argument concerning Ben’s belief is convincing. Why? The reason is that *what caused him to reach the belief had nothing to do with how many students were in the room*. When this *independence relation* obtains, the genetic argument shows that the belief is implausible. In contrast, when a *dependence relation* obtains, the description of the belief’s genesis can lead to the conclusion that the belief is probably correct.

As an example of how a genetic argument can show that someone’s belief is probably true, consider my colleague Cathy, who decided that there are 34 people in her philosophy class by carefully counting the people present. I suggest that the premiss in the following argument confers a high probability on the conclusion:

Cathy carefully counted the people in her class and consequently believed that 34 people were present.

*p* 

---

---

34 people were present in Cathy’s class.

When Cathy did her methodical counting, the thing that caused her to believe that there were 34 people present was *not* independent of how many people actually were there. Because the process of belief formation was influenced in the right way by how many people were actually in the room, we are prepared to grant that a description of the context of *discovery* provides a *justification* of the resulting belief.

Let us turn now to the argument (G) for ethical subjectivism stated

## *Prospects for an evolutionary ethics*

before. As the comparison of Ben and Cathy shows, the argument for subjectivism is incomplete. We need to add some premiss about how the process by which we arrive at our moral beliefs is related to which moral beliefs (if any) are true. The argument requires something like the following thesis:

- (A) The processes that determine what moral beliefs people have are entirely independent of which moral statements (if any) are true.

This proposition, if correct, would support the following conclusion: *The moral beliefs we currently have are probably untrue.*

The first thing to notice about this conclusion is that it does *not* say that ethical subjectivism is correct. It says that our *current* moral beliefs are probably untrue, not that *all* ethical statements are untrue. Here we have an important difference between (G) and the quite legitimate genetic arguments about Ben and Cathy. It is clear that a genetic argument might support the thesis that the ethical statements we happen to believe are untrue. I do not see how it can show that no ethical statements are true.

The next thing to notice about argument (G) concerns assumption (A). To decide whether (A) is correct, we would need to describe (i) the processes that lead people to arrive at their ethical beliefs and (ii) the facts about the world, if any, that make ethical beliefs true or false. We then would have to show that (i) and (ii) are entirely independent of each other, as (A) asserts.

Argument (G) provides a very brief answer to (i) – it cites “evolution” and “socialization.” However, with respect to problem (ii), the argument says nothing at all. Of course, if subjectivism were correct, there would be no such thing as ethical facts. But to *assume* that subjectivism is true in the context of this argument would be question-begging.

Because (G) says only a little about (i) and nothing at all about (ii), I suggest that it is impossible to tell from this argument whether (A) is correct. A large number of our beliefs stem either from evolution or from socialization. Mathematical beliefs are of this sort, but that doesn’t show that no mathematical statements are true (Kitcher 1993). I conclude that (G) is a weak argument for ethical subjectivism.<sup>4</sup>

Perhaps many of our current ethical beliefs *are* confused. I am inclined to think that morality is one of the last frontiers that human knowledge can aspire to cross. Even harder than the problems of natural science is the question of how we ought to lead our lives. This ques-

### *From a biological point of view*

tion is harder for us to come to grips with because it is clouded with self-deception. Powerful impulses hinder us from staring moral issues squarely in the face. No wonder it has taken humanity so long to traverse so modest a distance. Moral beliefs generated by superstition and prejudice probably *are* untrue. Moral beliefs with this sort of pedigree deserve to be undermined by genetic arguments. However, from this critique of some elements of existing morality, one cannot conclude that subjectivism about ethics is correct.

There is a somewhat different formulation of the genetic argument for ethical subjectivism that is worth considering. Harman (1977) suggests that we do not need to postulate the existence of ethical facts to explain why we have the ethical thoughts and feelings that we do. Psychological and evolutionary considerations suffice to do the explaining. Harman draws the conclusion that it is reasonable to deny the existence of ethical facts. His argument implicitly appeals to *Ockham's razor* (the *principle of parsimony*), which says that we should deny the existence of entities and processes that are not needed to explain anything.<sup>5</sup> We may schematize this argument as follows:

We do not need to postulate the existence of ethical facts to explain why people have the ethical beliefs they do.

It is reasonable to postulate the existence of ethical facts only if that postulate is needed to explain why people have the ethical beliefs they do.

---

---

There are no ethical facts.

As before, I draw a double line between premisses and conclusion to indicate that the argument is supposed to be nondeductive in character. Ethical subjectivism is here recommended on the ground that it is more *parsimonious* than alternative theories.<sup>6</sup>

I think the first premiss of this argument is correct. Psychological and biological facts about the human mind suffice to explain why we have the ethical beliefs we do. For example, we can explain why someone believes that capital punishment is wrong without committing ourselves, one way or the other, on the question of whether capital punishment really is wrong.<sup>7</sup>

Nonetheless, I think that the second premiss is radically implausible. For consider an analogy. Imagine that you are attending a class in statistics in which the professor repeatedly advances claims concerning how people should reason when they face inference problems of various sorts. After several weeks, a student stands up and claims that the



## *Prospects for an evolutionary ethics*

professor's normative remarks cannot be correct, on the grounds that they do not describe or explain how people actually think and behave. The professor, I take it, would be right to reply that statistics is not the same as psychology. Maybe human beings reason in normatively *incorrect* ways. The goal of statistics is not to describe or explain behavior, but to change it.

Precisely the same remarks apply to ethics. Ethics is not psychology. The point of normative ethical statements is not to *describe* why we believe and act as we do, but to *guide* our thought and behavior. We should not endorse ethical subjectivism simply because psychological explanations don't require ethical premisses.<sup>8</sup>

### 6. A NONDEDUCTIVE ANALOG OF HUME'S THESIS

Hume said you can't deduce an *ought*-conclusion from purely *is*-premisses. This thesis, I have emphasized, leaves open whether purely *is*-premisses provide *nondeductive* evidence for the truth of *ought*-conclusions.

On the face of it, the nondeductive connection of *ought* and *is* may seem obvious. Surely the first of the following two statements provides *some* evidence that the second is true:

- (7) Action *X* will produce more pleasure and less pain than will action *Y*.
- (8) You should perform action *X* rather than action *Y*.

I agree that (7) is evidence for (8), but I suggest that the two are connected in this way only because of a background assumption that we find so obvious that it perhaps escapes our notice – that pleasure is usually good and pain is usually bad. Without some such assumption, the evidential connection of (7) to (8) is severed.

In the light of examples like (7) and (8), I propose a generalization of Hume's thesis: *Purely is-premisses cannot, by themselves, provide nondeductive support for an ought-conclusion.* In my opinion, Hume's thesis about deduction also applies to nondeductive relations.<sup>9</sup>

My discussion of genetic arguments may seem to undermine this general thesis. After all, I claimed that facts about how people form their ethical beliefs can provide evidence concerning whether those beliefs are true. If the relevant facts about the process of belief formation were describable by purely *is*-propositions, then we could see genetic arguments as forging a nondeductive connection between *is* and *ought*. Just as Ben's belief about the attendance at his lecture is probably *false*, and

### *From a biological point of view*

Cathy's belief about the size of her class is probably *true*, given the procedures that each followed in forming their beliefs, so we can examine how people form their ethical convictions and draw conclusions concerning whether those ethical beliefs are probably *true*.

But there is a hitch. Descriptions of the process of belief formation cannot provide information about whether the beliefs are true unless we make assumptions about the nature of those propositions and the connections they bear to the process of belief formation.

The nature of these assumptions becomes clearer if we compare the story about Alf with the one told about Ben. When I pointed out that Alf thinks that the earth is round because he believes that the earth is perfect and that roundness is the perfect shape, this did nothing to undermine our confidence that the earth is round. We simply concluded that Alf believes the right thing for the wrong reason. In contrast, the fact that Ben formed his belief about the number of students in his classroom by drawing a ball from an urn did lead us to conclude that his belief was probably untrue. Why did we react differently to these two examples?

The reason is that we have independent reason to think that the earth is round, but none of us has independent knowledge of the number of students in Ben's class. In the former case, we begin by thinking that the earth is round, and what we learn about Alf's peculiar thought processes does nothing to undermine our confidence. In the latter case, we begin by thinking that various enrollment figures are about equally probable, and the information about Ben's draw from the urn does nothing to modify that picture. So we end as we began – by thinking that an enrollment of exactly 78 students has a low probability.

A genetic argument draws a conclusion about whether the members of some class of statements *S* are true. It draws this conclusion by describing the procedures that people follow in forming their opinions concerning the propositions in *S*. My suggestion is that you can't get something from nothing: an argument concerning the status of *S* must include some premisses about the status of the members of *S*. In particular, if you make no assumptions at all about the status of *ought*-statements, no conclusion can be drawn by a genetic argument concerning the status of those *ought*-statements. This was the point of emphasizing the role of consideration (ii) in genetic arguments. I conclude that genetic arguments, if they are to provide evidence that an *ought*-statement is untrue, must make assumptions about normative matters.

## *Prospects for an evolutionary ethics*

My generalization of Hume's thesis, if correct, has interesting implications concerning what evolutionary biology, and science in general, can teach us about normative issues. Let us begin with one clear format in which scientific matters can help us decide what is right and what is wrong.

Consider the relationship between statements (7) and (8). According to the thesis we are considering, (7) provides evidence for the truth of (8) only if further normative assumptions are provided. Hedonistic utilitarianism is one way to bridge the gap. This theory says that one action is morally preferable to another if the first produces more pleasure and less pain than the second. Hedonistic utilitarianism provides the sort of ethical principle that allows (7) to furnish evidence that (8) is correct.

Proposition (7) is the sort of claim that the sciences (in particular, psychology) may be in a position to establish. So, if hedonistic utilitarianism were true, the sciences would provide evidence concerning whether *ought*-statements like (8) are correct. The thing to notice about this dialectic is that the sciences do not provide evidence about (8) *on their own*. Science, on its own, does not generate ethical conclusions; rather, a more accurate formulation would be that science, *when conjoined with ethical assumptions*, can generate ethical conclusions. Hedonistic utilitarianism entails that psychological facts like (7) have ethical implications. But hedonistic utilitarianism is not something that science shows us is correct.<sup>16</sup>

My suspicion is that evolutionary ethics will always find itself in this situation. It may turn out that evolutionary findings do sometimes help us answer normative questions, although the proof of this pudding will be entirely in the eating. Just as hedonistic utilitarianism makes it possible for psychologists to provide information that helps decide what is right and what is wrong, this and other ethical theories may provide a similar opening for evolutionary biologists. This cannot be ruled out in advance. However, evolutionary findings will be able to achieve this result only when they are informed by ethical ideas that are not themselves supplied by evolutionary theory. Evolutionary theory cannot, all by itself, tell us whether there are any ethical facts. Nor, if ethical facts exist, can evolutionary theory tell us, all by itself, what some of those facts are. For better or worse, ethics will retain a certain degree of autonomy from the natural sciences. This doesn't mean that they are mutually irrelevant, of course. But it does mean that evolutionary ethicists who try to do too much will end up doing too little.

## *From a biological point of view*

### NOTES

1. Some philosophers whom I would want to call ethical realists refuse to apply the terms “true” and “false” to normative statements, but prefer terms like “valid” or “correct.” I happen to think that use of the term “true” in this context is fairly unproblematic; if murder is sometimes permissible, then it is true that murder is sometimes permissible. In any event, my use of the term “true” is something of an expository convenience. Realists maintain, for example, that murder is either sometimes permissible or that it never is, and that the permissibility of murder is not settled by anyone’s ethical opinion or say-so. This remark describes an *instance* of the realist position. To state it in full generality, some term such as “true” or “correct” is useful, if not outright necessary.
2. I’ve omitted *ethical skepticism* as an option in this classification of positions. “Don’t know” is not an answer that is incompatible with “yes” or “no.”
3. Hume’s thesis, as I understood it, does not deny that there are terms in natural languages that have both normative and descriptive content. Arguably, to say that someone is *rude* or *cruel* is to advance both a descriptive and a normative claim. It would be enough for my purposes if such claims can be construed as conjunctions, one conjunct of which is purely descriptive while the other is normative.
4. It is useful to represent genetic arguments in a Bayesian format. Consider the case of Cathy. Let  $C$  be the proposition that there are 34 students in her classroom. Let  $B$  be the proposition that Cathy believes  $C$ . The process by which Cathy formed her belief will have implications about the values of  $\Pr(B/C)$  and  $\Pr(B/\neg C)$ . Bayes’ theorem tells us what must be true for facts about the genesis of the belief to have implications about whether the belief is probably correct:  
$$\Pr(C/B) = \Pr(B/C)\Pr(C) / [\Pr(B/C)\Pr(C) + \Pr(B/\neg C)\Pr(\neg C)].$$

If  $\Pr(B/C)$  is high and  $\Pr(B/\neg C)$  is low, then further information about the value of  $\Pr(C)$  will settle whether  $\Pr(C/B)$  is high or low. Notice that the process *alone* (as described by the conditional probabilities  $\Pr[B/C]$  and  $\Pr[B/\neg C]$ ) does not fix a value for  $\Pr(C/B)$ ; additional information about  $\Pr(C)$  is needed.
5. For discussion of the status of Ockham’s razor in scientific inference, see Sober (1988) and Essay 7 in this volume.
6. Ruse and Wilson (1986, pp. 186–87) advance something very much like this parsimony argument when they say that “the evolutionary explanation makes the objective morality redundant, for even if external ethical premisses did not exist, we would go on thinking about right and wrong in the way that we do. And surely, redundancy is the last predicate that an objective morality can possess.”
7. In this respect, I am in agreement with Harman when he points out that

## *Prospects for an evolutionary ethics*

there is a disanalogy between ethical beliefs and many perceptual beliefs. Part of the explanation of why you now believe that there is a book in front of you is that there is a book in front of you.

8. It may be replied that statistics and ethics differ in the following respect: Normative statistical claims are grounded in mathematical truths, whereas nothing comparable can be said for ethical norms. I won't here try to evaluate whether this claimed disanalogy is correct. But note that it is irrelevant to the criticism I've made of Harman's argument: Subjectivism about normative claims (whether they are ethical or statistical in character) isn't well supported by the fact that those claims don't describe or explain how people actually behave.
9. The proposed generalization of Hume's thesis parallels a familiar point about confirmation in nonethical context: An observation confirms or disconfirms a theory only in the light of background assumptions. See Sober (1988) for discussion.
10. Kitcher (1993) is careful to distinguish the following two projects:
  - (B) Sociobiology can teach us facts about human beings that, in conjunction with moral principles that we already accept, can be used to derive normative principles that we have not yet appreciated.
  - (D) Sociobiology can lead us to revise our system of ethical principles, not simply by leading us to accept new derivative statements – as in (B) – but by teaching us new fundamental normative principles.I agree with Kitcher that (B) is far more plausible than (D).

## REFERENCES

- Alexander, R. (1987). *The Biology of Moral Systems*. Aldine DeGruyter.
- Cavalli-Sforza, L., and Feldman, M. (1981). *Cultural Transmission and Evolution*. Princeton University Press.
- Harman, G. (1977). *The Nature of Morality*. Oxford University Press.
- Hempel, C. (1965). *Aspects of Scientific Explanation*. Free Press.
- Kitcher, P. (1993). "Four Ways to 'Biologize' Ethics." In K. Bayertz (ed.), *Evolution und Ethik*. Reclam 1993. Reprinted in E. Sober (ed.), *Conceptual Issues in Evolutionary Biology*. MIT Press, 1993, 2nd edition.
- Moore, G. (1903). *Principia Ethica*. Cambridge University Press.
- Ruse, M., and Wilson, E. (1986). "Moral Philosophy As Applied Science." *Philosophy* 61: 173–92. Reprinted in E. Sober (ed.), *Conceptual Issues in Evolutionary Biology*. MIT Press, 1993, 2nd edition.
- Sober, E. (1988). *Reconstructing the Past: Parsimony, Evolution and Inference*. MIT Press.
- Wilson, E. (1978). *On Human Nature*. Harvard University Press.