EVIDENCE AND VALUE FREEDOM

Elliott Sober

IN THESE POSTPOSITIVIST TIMES, THE SLOGAN "SCIENCE IS VALUE FREE" is frequently rejected disdainfully as a vestige of a bygone age. But, as often happens with babies and their bathwater, there may be something worthwhile in this slogan that we should try to identify and retain. To this end, I'll begin with two absurdities:

(1)   Scientists aren't influenced by their ethical and political values when they do science.

(2)   Scientific inference is independent of values.

I don't know if any philosopher ever believed either of these propositions. Proposition (1) is false for the simple reason that scientists are people, just like the rest of us. Perhaps they often *strive* to leave their ethical and political values at the laboratory door, but who ever thought that all of them have this aim, and that those who do succeed 100% of the time? This, by the way, does not mean that we get to assume that scientific activity can be explained *solely* in terms of the ethical and political values that scientists have. Rather, recognizing the absurdity of (1) should lead us to approach such psychological and sociological questions on a case-by-case basis. Scientists may vary among themselves, and a single scientist

may be more influenced by these values in some contexts than in others. Maybe *some* scientific work proceeds completely independently of these values and *other* parts of science are entirely driven by them, and perhaps a good deal of the real world falls somewhere between these two extremes.[1]

Proposition (2) is absurd because scientific inference is regulated by normative rules. Scientists try to construct *good* tests of their hypotheses, they judge some explanations *good* and others *bad*, and they say that some inferences are *flawed* or *weak* and others are *strong*. The words I have italicized indicate that scientists are immersed in tasks of evaluation. They impose their norms on the ideational entities they construct. However, the obvious falsehood of (2) leaves it open that a restricted version of that proposition might be on the right track:

(3)   The fact that believing a proposition would have good or bad ethical or political consequences is not evidence for or against that proposition's being true.

Is this proposition, or some refinement of it, the kernel of truth in the frequently misstated idea that "science is value free"?

We should not accept proposition (3) just because it "sounds right." After all, the evidence relation often connects facts that seem at first glance to be utterly unrelated. The proposition that the dinosaurs went extinct 65 million years ago because of a meteor hit and the proposition that there now is an iridium layer in certain rocks may appear to have nothing to do with each other. How could the presence of iridium in present-day rocks bear on the question of why the dinosaurs went extinct so long ago? Well, appearances to the contrary, there may well be such a connection (Alvarez and Asaro 1990). Why, then, should we be so sure that the ethical consequences of believing a proposition have no bearing on whether the proposition is true? This is a good question, and in the absence of a good answer, we should not complacently assume that (3) is correct.

Sometimes people believe (3) because they think there are no ethical truths in the first place. If there are no ethical truths, then ethical truths don't provide evidence for anything. Whether or not nonfactualism is correct, I think it fruitful to consider proposition (3) on the assumption that there are ethical facts. In our everyday lives, we treat

ethical statements as if some of them are true. Perhaps this is a mistake, but for present purposes, let's take that practice at face value. If some normative ethical statements *are* true, what is to prevent them from standing in evidential relations with nonethical, scientific propositions?

To investigate this question, let's begin with a useful example of the evidence relation at work—smoke is evidence of fire. This relation holds because the probability of fire if smoke is present exceeds the probability of fire if smoke is absent. The two events ±smoke and ±fire are *correlated*. This is a perfectly objective relation that obtains between smoke and fire; it obtains whether or not anyone believes that it does. Understood in this way, the evidence relation has an interesting property: It is *symmetric*. If smoke is evidence of fire, then fire is evidence of smoke. This means that proposition (3) entails a further claim:

(4)   The fact that a proposition is true is not evidence that believing the proposition would have good or bad ethical or political consequences.

If (3) is true, so is (4). But surely there are counterexamples to proposition (4). Consider a physician who will give a drug to her patients if she thinks the drug is safe but will withhold the drug if she thinks it is not. Suppose that the drug will provide significant health benefits if it is safe. And suppose further that the physician is a pretty good judge of whether the drug is safe. We then have a causal chain in which earlier links raise the probability of later ones:

the drug is safe → the doctor thinks the drug is safe → the patients receive the drug → good consequences accrue to the patients

In this instance, correlation is *transitive*. The nonethical statement "the drug is safe" is therefore evidence for the ethical statement "good consequences accrue to the patients." An ethical and a nonethical fact are evidentially related, just like smoke and fire. I conclude that (4) and therefore (3) are false (Stephens 2000).

It might be replied that (3) and (4) can be saved from refutation by focusing exclusively on the two propositions under discussion. Consider just the two statements "the drug is safe" and "good consequences accrue to the patients." In the absence of any further information, there is no saying whether these statements are positively evidentially relevant

to each other, negatively relevant, or entirely irrelevant. We assumed in our story that the physician is well meaning and discerning. This assumption was enough to bring the two statements into a positive relation. However, if the doctor were malevolent or a very bad judge of drug safety, the two statements would stand in a relation of negative relevance or be mutually irrelevant. With no further information, the relation of the two statements is indeterminate. This is a probabilistic analog of Duhem's thesis (Sober 1988).

The trouble with this reply is that what is true for the pair of statements about the physician is true for practically any pair of statements.[2] The evidence relation isn't binary; it has at least three places. When one statement is evidence for a second, this usually is due to the mediation of a third, which provides background information.[3] Given this, it is hardly surprising that ethical facts can provide evidence about scientific propositions; they can do this if one's background assumptions include other, ethical claims.

How is this criticism of (3) and (4) related to Hume's famous dictum that an *ought* cannot be inferred from an *is*? Although Hume was talking about deduction, it is natural to generalize his thesis to a claim about evidential relationships that are nondeductive and probabilistic. What one needs to say here is that an *is*-statement and an *ought*-statement are not evidentially relevant to each other, unless one's background assumptions include other *ought*-statements. But because one's background assumptions often *do* include such "bridge principles," evidential reasoning can run from *is* to *ought* and from *ought* to *is*.[4]

How might propositions (3) and (4) be refined? What is needed is a way to separate the pattern exemplified by the physician from a second class of examples, which William James (1897) addressed in his famous discussion of *the will to believe*. James argued that believing in God can provide substantial psychological benefits. Faith in God can help people feel that their lives are meaningful, which may save them from depression and allow them to lead worthwhile lives. It is debatable how general this point is; there are enough happy atheists and depressed theists around to make one suspect that, at least for many people, mental health is independent of theological conviction. But let's restrict our attention to people who fill James's bill. These people would benefit from believing that God exists; I'll go further and say that it is a *good thing* for

	P is true	P is false
S believes P	w	x
S does not believe P	y	z

Table 5.1. Comparing utilities as a function of what is believed and what is the case.

these people to embrace theism and thus save themselves from the slough of despond. However, I still want to claim that the fact that they would benefit from believing in God provides no evidence that God in fact exists. It is examples like this that make proposition (3) sound so plausible. What distinguishes James's theist from the well-meaning physician?

We can separate these cases by considering the two-by-two table above (table 5.1). The entry in a cell represents *utility*—how good or bad the consequences are of being in that situation. In the physician case (where P = "the drug is safe"), there are both "vertical" and "horizontal" effects. The well-being of the patients is affected both by whether the drug is safe $(w > x)$[5] and by what the physician believes $(w > y$ and $z > x)$. In the case of James's believer (where P = "God exists"), however, there are only vertical effects. As far as the individual's psychological well-being is concerned, the only thing that matters is that he or she believes in God $(w > y$ and $x > z)$; whether God actually exists doesn't matter $(w = x$ and $y = z)$.

This, I think, provides the key to revising propositions (3) and (4). Our question—when do the ethical consequences of believing a proposition provide evidence as to whether the proposition is true?—can be represented by using the tools of decision theory. We begin by identifying the expected value of each of the two "acts."[6]

EV[Believe P] = wp + x(1−p)
EV[Don't Believe P] = yp + z(1−p).

Here *p* denotes the probability that the proposition P is true (and I assume that acts are independent of states of the world). When does the

fact that EV[Believe P] > EV[Don't Believe P] provide information about the value of p? A little algebra reveals that

EV[Believe P] > EV[Don't believe P] if and only if $p/(1-p) > (z-x)/(w-y)$.

Suppose that the left-hand inequality is true. When will filling in the values for the utilities w, x, y, and z provide a nontrivial lower or upper bound on the value of p? This fails to happen in the case of James's theist because $(w-y)$ is positive while $(z-x)$ is negative. With these values, all that follows is that $p/(1-p)$ must be greater than some negative number; this is entirely uninformative, since no ratio of probabilities can be negative. The case of the physician is different. Here $(z-x)$ and $(w-y)$ are both positive; therefore, their ratio provides a nontrivial lower bound for the value of p.[7] Thus, science and ethics are not always as separate as propositions (3) and (4) suggest; sometimes information about the ethical consequences of believing a proposition *does* provide information about the probability that the proposition is true.

I have represented the ethical consequences of believing a proposition in terms of the expected utility of doing so. This may sound as if it requires a consequentialist ethics, but in fact it does not. A nonconsequentialist also can enter payoffs in the four cells. Furthermore, the argument is general. Utilities may be calculated in accordance with an ethical theory or on some nonethical basis. What we have here is a general format in which judgments about which acts are better, plus information about the utilities of outcomes, can have implications about the probabilities of propositions.[8]

Before phoning the National Science Foundation with the news that ethics can be a source of evidence for scientific claims, we should reflect on the fact that the expected utility of an action is a composite quantity, built up from probabilities and utilities. If one's ethical judgments about which actions one should perform are based on comparing their expected values, then ethical judgments *require* information about probabilities for their formulation. If so, how can ethical judgments be a source of information about probabilities? This point does not rescue propositions (3) and (4) from counterexample, but it does suggest a new way to think about the problem.

There is a special circumstance in which it is possible to decide which of the two actions (believe P, or don't) is better without any information about the probability that P is true. This is the case in which one action *dominates* the other. James's argument is of this type: He contends that belief in God is beneficial, whether or not God in fact exists.[9] However, we have already seen that the inequality of the expected utilities provides no information about the probabilities in James's case. Without dominance, no conclusion about which action is better can be reached unless one *already* has information about the probabilities. What we have here is an instance of the maxim "out of nothing, nothing comes." If comparing the ethical consequences of believing P and of not believing P has implications about the probability of P, this must be because the description of the ethical consequences already has built into it some information about those probabilities. Thus, ethical facts (about expected utilities) and scientific facts (about probabilities) *are* connected, contrary to what (3) and (4) assert. However, the problem is that this connection is *useless*; we can't use ethical information to gain information we don't already have about probabilities.

The situation would be different if we were able to discover which actions are better than which others, when dominance fails, without already having to have information about probabilities. For example, if there were an infallible guru who would simply tell us what to do, and who would reveal the utilities that go with different states of the world, we could use these inputs to obtain new information about probabilities. But in the absence of such an authority, we are left with the conclusion that our access to information about what we should do must be based on information about probabilities (except when there is dominance). When an ethical conclusion requires information about probabilities (as in the case of the physician), that conclusion can't be a source of new information about those probabilities. And when the ethical conclusion can be reached without information about probabilities (as in the case that James describes), the conclusion tells us nothing about the probabilities. This suggests the following dilemma argument:

> The question of whether believing P has better ethical consequences than not believing P either depends for its answer on information about the probability of P, or it does not.

If it does so depend, then we need information about probabilities to answer the ethical question, and so the ethical judgment cannot supply information about probabilities that we don't already have.

If it does not so depend, then the ethical judgment has no implications about the probability of P.

(5)   Judgments about the ethical consequences of believing P cannot supply *new* information about the probability of P.

The conclusion of this argument, proposition (5), is a reasonable successor to the failed propositions (3) and (4).

The argument just presented is reminiscent of Rudner's (1953) well-known argument that the scientist qua scientist makes value judgments.[10] Rudner describes a physician who must decide whether a drug is safe and argues that this decision must be based on considering ethical features of the four possible outcomes depicted in table 5.1, which we have already discussed. Rudner's argument elicited two criticisms. Levi (1967) contended that accepting a proposition and acting on one's belief are distinct and that the former should not be based on ethical values; Jeffrey (1956) maintained that science is not in the business of accepting and rejecting but merely seeks to assign probabilities to hypotheses. My own argument is neutral on Rudner's position. Perhaps deciding what to believe depends on ethical values; perhaps it does not. My point is that no matter how one decides what to believe, one still can consider what the ethical consequences are of that decision. My question was whether this ethical consideration has implications concerning the probabilities of hypotheses. It does in the case of the physician but not in the case of James's theist.

In the example about the physician, and in many other examples of moral deliberation about which action to perform, one's ethical decision depends on matters of scientific fact. In terms of table 5.1, the ubiquitous pattern is that an inequality gets reversed as one moves from the first column to the second. Although one's ethical decision thus depends on a judgment about a matter of scientific fact, it is possible to form a judgment about the scientific facts without having a commitment, one way or the other, on the ethical question. The physician can't decide whether to administer the drug without knowing something about its probability of being safe, but it is perfectly possible to discover whether a drug is safe without having a view, one way or the other, on

whether unsafe drugs should be withheld from patients. Moral ignoramuses can assess the weight of evidence, but scientific ignoramuses cannot make good moral decisions (when those decisions depend, as they almost always do, on scientific matters of fact).

Let's review our progress from propositions (1) and (2) through (3) and (4) and then to (5). Proposition (1) concerns the *behavior of scientists*, whereas (3) and (4) concern the logic of various scientific *concepts*. Proposition (2) ambiguously straddles this distinction; "scientific inference" can be taken to refer to what scientists do or to the formal properties of various types of argument. Proposition (5) addresses the concept of evidence; it does not assert that scientists are immune from political and ethical influence when they decide whether one proposition is evidence for another. More specifically, the claim is not that ethical values (represented by the expected utilities of actions and the utilities of outcomes) *have no implications* about the probabilities of hypotheses, but that ethical inputs are *not needed* to estimate those probabilities. This is why looking to ethics for evidence concerning the truth of scientific hypotheses is to place the cart before the horse.

NOTES

1. I am grateful to Ellery Eells and Dan Hausman for useful comments.

2. The exception arises when one statement deductively entails the other; then they must be positively relevant or of zero relevance, and negative relevance is ruled out.

3. I say "at least" three places because in most scientific contexts, what one can discuss is whether the evidence discriminates between a pair of hypotheses, given a set of background assumptions. See Sober (1994a) for discussion.

4. Does the *ought implies can* principle undermine Hume's thesis? The principle's contrapositive asserts that if it is impossible for an agent to perform an action, then it is false that the agent ought to do so. Here an *is* implies the negation of an *ought*. Hume's thesis can be preserved by insisting that the negation of an *ought*-statement is not itself an *ought*-statement. Similar reasoning is required if one wishes to reconcile Hume's thesis with the fact that philosophers have presented philosophical arguments (of varying quality) for the claim that normative ethical statements lack truth values. These arguments do not contain premises that are normative ethical statements. For example, Harman (1977) and Ruse and Wilson (1986) each present parsimony arguments for the nonexistence of ethical facts; see Sober (2005) and (1994b), respectively, for discussion of each.

5. I take it that it doesn't matter whether the drug is safe if the doctor doesn't believe that it is ($y=z$), because patients won't receive the drug in that situation regardless of whether the drug would be good for them.

6. It might be suggested that believing a proposition is not an action, in the sense that it is not subject to the will. This point is sometimes used against Pascal's wager, but it is an objection that Pascal successfully addressed: He says that if absorbing his argument does not instantly trigger belief, one should go live among religious people so that habits of belief will gradually take hold. Believing a proposition is like other "nonbasic actions": Being president of the United States isn't something one can directly bring about by an act of will, but this does not place it outside the domain of decision theory. For further discussion, see Mougin and Sober (1994).

7. Symmetrically, if not believing P were the better action, this would impose a nontrivial upper bound on the value of the probability.

8. Decision theory from early on has taken an interest in describing the interrelationships of expected utility, utility, and probability. For a brief introduction, see Skyrms (2000, pp. 138–43).

9. Pascal's wager, when the payoffs are finite, does *not* have this property. One needs some information about the probability of God's existing to reach a decision about whether believing is better than not believing. In fact, the theist contemplating Pascal's wager (with finite payoffs) is in the same qualitative situation as the physician deciding whether to believe that the drug is safe. See Mougin and Sober (1994) discussion.

10. The problem that Rudner addresses is the one that James (1897) and W. K. Clifford (1879) debated. It also was central to the debate between the "left" and "right" wings of the Vienna Circle. Neurath argued that evidence does not determine theory choice and that ethical and political values can and should be used to close the gap; Schlick, Carnap, and Reichenbach countered that the intrusion of such values into theory choice is both undesirable and unnecessary: It would compromise the objectivity of science, and scientific inferences can be drawn without taking such values into account. See Howard (2002) for discussion.

REFERENCES

Alvarez, W., and F. Asaro. 1990. "What Caused the Mass Extinction? An Extraterrestrial Impact." *Scientific American*, 263, pp. 78–84.

Clifford, W. K. 1879. "The Ethics of Belief," in *Lectures and Essays*, vol. 2, pp. 177–211. London: Macmillan.

Harman, G. 1977. *The Nature of Morality*. New York: Oxford University Press.

Howard, D. 2002. "Philosophy of Science and Social Responsibility—Some Historical Reflections," in A. Richardson and G. Hardcastle, eds., *Logical Empiricism in North America*. Minneapolis: University of Minnesota Press.

James, W. 1897. "The Will to Believe," in *The Will to Believe and Other Essays in Popular Philosophy*. New York: Longmans Green.

Jeffrey, R. 1956. "Valuation and Acceptance of Scientific Hypotheses." *Philosophy of Science*, 33, pp. 237–46.

Levi, I. 1967. *Gambling with Truth*. Cambridge, MA: MIT Press.

Mougin, G., and E. Sober. 1994. "Betting against Pascal's Wager." *Nous*, 28, pp. 382–95.

Rudner, R. 1953. "The Scientist *Qua* Scientist Makes Value Judgments." *Philosophy of Science*, 20, 1–6.

Ruse, M., and E. Wilson. 1986. "Moral Philosophy as Applied Science." *Philosophy*, 61, pp. 173–92. Reprinted in E. Sober, ed. 1993. *Conceptual Issues in Evolutionary Biology*, 2nd ed. Cambridge, MA: MIT Press.

Skyrms, B. 2000. *Choice and Chance*, 4th ed. Belmont, CA: Wadsworth.

Sober, E. 1988. *Reconstructing the Past—Parsimony, Evolution, and Inference*. Cambridge, MA: MIT Press.

Sober, E. 1994a. "Contrastive Empiricism," in *From a Biological Point of View*. New York: Cambridge University Press.

Sober, E. 1994b. "Prospects for Evolutionary Ethics," in *From a Biological Point of View*. New York: Cambridge University Press.

Sober, E. 2005. *Core Questions in Philosophy*. Upper Saddle River, NJ: Prentice-Hall.

Stephens, C. 2000. "Why Be Rational? Prudence, Rational Belief and Evolution." PhD dissertation, University of Wisconsin, Madison.