Quine, W. V. O. 1975. "On Empirically Equivalent Systems of the World." *Erkenntnis* 9: 313–28.
van Fraassen, B. C. 1980. *The Scientific Image*. Oxford.
van Fraassen, B. C. 1985. "Empiricism in the Philosophy of Science." In *Images of Science*, edited by P. M. Churchland and C. A. Hooker. Chicago.
van Fraassen, B. C. 1989. *Laws and Symmetry*. Oxford.

MIDWEST STUDIES IN PHILOSOPHY, XVIII (1993)

Epistemology for Empiricists

ELLIOTT SOBER

In 1963, Paul Feyerabend published a paper called "How to Be a Good Empiricist." Feyerabend took empiricism to be the reigning orthodoxy of his time, but argued that it had lapsed into an anti-empirical dogmatism. Feyerabend tried to show empiricists how to do better. Had Feyerabend not beat me to it, I would have used the title he chose for his paper as the title for the one you now see before you. How one is to be a good empiricist is still a problem worth posing. Not that the ideas I will advance overlap very much with the ones he discussed. Questions about how theoretical terms change meaning when old theories give way to new ones and about whether science is cumulative are important, but they will not be important to me in what follows. And, of course, another difference between my project and Feyerabend's is that empiricism can no longer be called a reigning orthodoxy. Feyerabend's paper appeared during a time in which logical empiricism was giving way to scientific realism. My hunch is that if there is now a dominant ideology, it is realism, not empiricism.

Empiricism is nonetheless still a doctrine worth reckoning with, mainly due to the formulation and defense that Bas van Fraassen (1980, 1985, 1989) has more recently provided. Van Fraassen criticized scientific realism; realists replied in kind. In my opinion, some of these criticisms of van Fraassen's views are fundamentally sound. The question I wish to address is whether anything of value remains of empiricism. Empiricism is not now a reigning orthodoxy, but it still is worth asking whether empiricists can do better.

Empiricism is a thesis about the importance that experience should have in shaping our beliefs about the world. It is a normative claim—a thesis of epistemology, not psychology. The exact formulation the doctrine should receive is a problem to which I will return. For now, I will begin with van Fraassen's statement of the difference between his own version of empiricism—*constructive empiricism*—and scientific realism. According to van Fraassen (1980, 8), realism maintains that

> Science aims to give us, in its theories, a literally true story of what the world is like; and acceptance of a scientific theory involves the belief that it is true.

Constructive empiricism, on the other hand, says that

> Science aims to give us theories which are empirically adequate; and acceptance of a theory involves as belief only that it is empirically adequate. (van Fraassen 1980, 12)

Fundamentally, the conflict concerns the rational attitude we should have towards theories in science.

What does it mean to say that a theory is empirically adequate? The rough idea is that what the theory says about observable objects is true. Many theories talk about both observable and unobservable entities. According to van Fraassen, if such theories pass reasonable standards of adequacy, we should conclude that what they say about observables is true, but we should remain agnostic with respect to what they say about unobservables. The most we can assert, then, is that such theories are empirically adequate; we cannot say that they are true. Because true theories are empirically adequate, but not conversely, constructive empiricists are more circumspect than realists in the way they assign truth values.

Before taking my discussion of empiricism and realism any farther, several points need to be clarified concerning van Fraassen's statement of the problem and the gloss I have given of what he says. I have described both philosophical doctrines as *ought* statements. But van Fraassen's formulations are in the indicative mood—they describe what the aim of science *is*. This suggests that realism and empiricism are not epistemological theses at all (at least not if epistemological claims are normative claims concerning the regulation of belief).

Reflection on how van Fraassen understands theses about the "aim" of science reveals that his indicative formulations in fact express normative claims. His contrast between realism and empiricism is intended to offer two rival theories about what is required for one to play "the game of science." The empiricist says that the most that is *required* in science is belief in the empirical adequacy of theories; the realist says that belief in the truth of theories is sometimes *required* as well. Claims about the aims of science are, for van Fraassen, theses about the norms that constitute scientific activity.[1]

So there is no conflict between the indicative formulations of realism and empiricism that van Fraassen provides and my claim that both are normative theses in epistemology. But there is a further question, this one concerning the *kind* of normative claim that is at issue. Van Fraassen (1985, 252) describes empiricism and realism as views about what is *permissible*, not about what is epistemologically *obligatory* (see also van Fraassen 1989, 171–72). On this reading, van Fraassen's empiricism does not say that the realist is *irrational* in believing various theories true, only that such belief is not rationally required.

According to this formulation, the empiricist says only that one *may* remain agnostic concerning theories about unobservables, not that one *must* do so.

Whether epistemological norms should be thought of as norms of obligation or norms of permission is an interesting question, but it is not a question I will discuss in what follows. My reason for skirting this issue is that the problems I wish to address are not affected by which formulations one considers. For ease of exposition, I will describe empiricism as *requiring* agnosticism, and realism as *requiring* assent, with respect to claims about unobservables.[2] As will become clear shortly, these formulations are only initial points of departure.

One standard criticism of constructive empiricism is that it attributes to the distinction between observables and unobservables an epistemological significance it does not possess. According to van Fraassen, the moons of Jupiter are observable; they are observable because a properly situated astronaut could see them without the aid of instrumentation. In contrast, the AIDS virus is not observable, because detecting its presence and properties requires the use of instruments.

Although moons and viruses differ in the respect just mentioned (let us grant), they are similar in another respect. Neither has *actually* been observed. According to van Fraassen's preferred usage, we earthlings have not in fact *observed* the moons; rather, we have indirectly *detected* their presence by using instrumentation—a telescope. A parallel remark applies to viruses. We have not directly observed them either; rather, we have detected their presence with the help of an instrument—a microscope. The moons are observable, but unobserved; viruses, on the other hand, are both unobserved and unobservable.

Van Fraassen (1980, 16) is right to construe observability as a dispositional property:

> X is observable if there are circumstances which are such that, if X is present to us under those circumstances, then we observe it.

He also is right to say that empirical science will help us understand what makes some objects, but not others, observable.[3] Nonetheless it is entirely unclear why this difference between moons and viruses should matter to us epistemologically. Perhaps there are specific reasons why we should be more circumspect about some particular hypothesis about the AIDS virus than we are about some particular hypothesis concerning Jupiter's moons. But the mere fact that Jupiter's moons are observable while individual viruses are not is epistemically irrelevant.

This argument against van Fraassen has become standard. I think it is fundamentally correct. The question I wish to ask is whether some version of empiricism can survive its corrosive effect. What, if anything, is epistemologically special about observation? To begin answering this question, I want to locate the objection to van Fraassen's version of empiricism within a larger epistemological context. There is a general lesson to be drawn from this criticism that illuminates how the concept of evidence should be understood.

1. ACTUALISM

Consider two diagnostic problems that physicians confront. They want to gather evidence concerning whether a patient has diabetes; they also want to gather evidence about whether a patient has small pox.[4] In each case, a laboratory test is performed, and the outcome of the test is interpreted. For simplicity, let us suppose that each test produces one of two results, which we designate "positive" and "negative." Each test has two error characteristics (a,b and c,d, respectively), representing the chance of a false positive and the chance of a false negative. These are described in table 1.

Table 1

	Test for Diabetes	
	positive	negative
S has diabetes	1-a	a
S does not have diabetes	b	1-b
	Test for Small Pox	
	positive	negative
S has small pox	1-c	c
S does not have small pox	d	1-d

The entries in these tables denote the probability of a test result, conditional on the patient's situation; they do not describe the probability of a disease, conditional on a test result.

Let us suppose that these tests are *equally reliable*, though each is less than perfectly so. They have identical and nonzero error probabilities: $a = c \neq 0$ and $b = d \neq 0$. This assumption entails that the amount of information delivered by each test will be the same. If a positive result on the first test strongly favors the hypothesis that the patient has diabetes, then a positive result on the second test strongly favors the hypothesis that the patient has small pox. The tests are on an epistemic par.

Now let us consider a third test. This is a test for diabetes *and it is infallible* (table 2).[5]

Table 2

	An Infallible Test for Diabetes	
	positive	negative
S has diabetes	1	0
S does not have diabetes	0	1

This test is infallible because if you have diabetes you are certain to get a positive result and if you do not have diabetes you are certain to get a negative result. If we were to run this infallible test, we would gain more information about whether the patient has diabetes than either of the first two tests could provide. Suppose there is no similarly infallible test for small pox.

Now let's run the first two tests, but not the third. How should the outcomes of these two tests be interpreted? Suppose each delivers a positive outcome. Should we be more confident that the patient has diabetes than we should that the patient has small pox *on the grounds that there exists an infallible test for the former, but none for the latter?* I would say *no*. We should be guided by the evidence we *actually* possess; the fact that we *could* possess greater certainty about diabetes than we *could* about small pox is irrelevant. We could have run the third test, but in point of fact we did not; the mere possibility of running the third test is irrelevant. I will call the idea underlying this judgment the *principle of actualism*. It says that we should form our judgments about hypotheses based on the evidence we actually possess; possible but nonactual evidence does not count.[6]

Actualism is an empty idea until additional epistemic principles are specified. No doubt the modal fact that we *could* be certain about whether a patient has diabetes if we were to run the third test is grounded in some fact about what is *actually* the case. The point is that the mere fact that such a test procedure is available does not help us tell whether a patient has diabetes. Explaining why this is true requires both a substantive theory of evidence and substantive facts about the world.[7]

Actualism is an intuitive principle, on which both realists and empiricists *should* be able to agree. It says nothing about whether our knowledge of electrons and genes is less secure than our knowledge of dogs and tables. However, the point of interest is that van Fraassen's formulation of empiricism violates actualism. Let us return to the comparison of moons and viruses to see why.

We detect some fact about Jupiter's moons by looking through a telescope; we detect some fact about a sample of the AIDS virus by looking through a microscope. These two test problems may differ in many ways that may be relevant to the degree of certainty we are entitled to have about each of the matters under test. However, one thing that is epistemically irrelevant is a certain counterfactual possibility: if we were to get in a space ship and journey to Jupiter, we could see its moons without instrumentation, but no such journey would allow us to see an individual virus without the help of instruments.[8] The reason this difference is irrelevant is given by the principle of actualism.[9]

Empiricism is a thesis about the importance of observation. However, the way to develop the thesis is to show that *what we actually observe* has some special epistemic status. It is not the distinction between *observable* and *unobservable* that is fundamental, but the difference between *observed* and *unobserved*. Empiricists should attempt to show how *actual* observation provides more certainty than other supposed routes to knowledge. Let us bear this lesson in mind; we will return to it later.

2. ACCEPTANCE

I now want to consider a second, and less noticed, property of van Fraassen's formulations of empiricism and realism. It concerns the word *acceptance*.[10] According to van Fraassen (1980), science is in the business of accepting hypotheses; the dispute between realism and empiricism concerns what sorts of hypotheses scientists should accept.

At least since discussion of the lottery paradox (Kyburg 1970), the dichotomy between accepting and not accepting a hypothesis has struck many philosophers of science as too coarse-grained. Instead of deciding between accepting and not accepting, rational agents might be described as assigning *degrees of belief* to various hypotheses, or as deciding how *well supported* various hypotheses are.

A wholesale rejection of the concept of acceptance is not forced on one by the lottery paradox, but it is a plausible and common diagnosis of what the paradox shows. Imagine a fair lottery with a thousand tickets. Exactly one ticket will win and each ticket has the same probability of winning. Consider the statement "ticket #1 will not win." Based on the information given, this statement has a probability of 0.999. If our rule of acceptance is to accept a hypothesis when it is very probable (suppose the threshold for this is specified at 0.95), then we should accept this hypothesis about ticket #1. However, the same line of reasoning applies to each of the tickets. But if we accept *each* such hypothesis, we will have contradicted the starting assumption that some ticket will win.

One might interpret this problem as showing that rules for acceptance must be more complicated than the one described here. Instead of jettisoning the concept of acceptance, maybe we should try to fine-tune it. However, it is hard to see, on reflection, that a cut-off exists that separates what we believe *tout court* from what we do not. Rather, statements tend to grade off from the very plausible, to the moderately plausible, to the moderately implausible, and so on. Let us assume, then, that epistemology should have no truck with the concept of acceptance. How, then, are we to reformulate van Fraassen's opposition between empiricism and realism?[11]

One possibility, to which van Fraassen (1980) alludes, is Bayesianism. Hypotheses are assigned probabilities; as new evidence rolls in, these probabilities are updated by appeal to Bayes's Theorem. When we acquire the evidence E, we assign to the hypothesis H a posterior probability $\mathrm{P}(H/E)$ in accordance with the following formula:

$$\mathrm{P}(H/E) = \mathrm{P}(E/H)\mathrm{P}(H)/\mathrm{P}(E).$$

If we adopt this Bayesian format, empiricism and realism must now be understood as advancing principles about how probabilities are to be assigned and updated.

Van Fraassen's empiricism must delimit an epistemic difference concerning how hypotheses about observables and hypotheses about unobservables should be treated. What might that difference come to in a Bayesian framework? Bayesians say that a hypothesis is confirmed by a piece of evidence when acquiring the evidence makes the hypothesis more probable than it was before. Perhaps empiricism should then be formulated as the claim that *hypotheses about unobservables are not confirmable; if H is about unobservables, then* $\mathrm{P}(H/E) \not> \mathrm{P}(H)$ This proposal articulates empiricist suspicions about hypotheses that talk about things we cannot observe; they are epistemically inaccessible in the precise sense that their probabilities cannot be boosted by observations.

It does not take much to see that this proposal is implausible. Rearranging Bayes's theorem, we obtain the following equality:

$$\mathrm{P}(H/E)/\mathrm{P}(H) = \mathrm{P}(E/H)/\mathrm{P}(E).$$

E confirms H precisely when the left-hand side is greater than unity. However, it is quite obvious that this quantity *can* be greater than unity, even for hypotheses that are about unobservables. For consider the right-hand side. Suppose that H deductively entails E. Then $\mathrm{P}(E/H) = 1$. If E was not certain before the evidence was gathered, then $\mathrm{P}(E) < 1$, which means that $\mathrm{P}(E/H)/\mathrm{P}(E) > 1$. When hypotheses about unobservables deductively entail true predictions that were not known with certainty beforehand, those hypotheses are confirmed.[12] Empiricism cannot deny this.

So the simple fact of the matter is that, within a Bayesian format, the probabilities of hypotheses about unobservables can go up and down, just like the probabilities of hypotheses about observables.[13] But perhaps there is some other asymmetry between observables and unobservables that can be described within a Bayesian framework. Rather than concentrate on *change* in probability, let us focus on the absolute values those probabilities attain. Perhaps hypotheses strictly about unobservables must always be less probable than hypotheses strictly about observables. This is the thesis that *for any evidence* E, $\mathrm{P}(H_1/E) > \mathrm{P}(H_2/E)$ *if* H_1 *is strictly about observables and* H_2 *is strictly about unobservables.*

This also is a proposal that cannot pass muster. Let H_1 be a hypothesis about observables on which we have no evidence whatever; let H_2 be a hypothesis about unobservables on which we have scads of positive evidence. Surely we are right to be more certain about whether someone has the AIDS virus in his blood, based on the results of laboratory tests, than we are about whether a *Brontosaurus* was standing exactly where the White House now is precisely 179 million years ago today.

So we cannot say that *all* statements about observables are more probable than *all* statements about unobservables. However, there is a much attenuated formulation of this asymmetry that cannot be denied. It is the one that van Fraassen (1980, 68–69) notes in the following passage:

> . . . we can distinguish between two epistemic attitudes we can take up toward a theory. We can assert it to be true . . . , and call for belief; or we can simply assert its empirical adequacy. . . . In either case we stick our necks out; empirical adequacy goes far beyond what we can know at any given time. . . . Nevertheless there is a difference; the assertion of empirical adequacy is a great deal weaker than the assertion of truth. . . .

The claim, as I understand it, is this: for any theory T that is about unobservables, and any evidence E, $\mathrm{P}(T \text{ is true}/E) < \mathrm{P}(T \text{ is empirically adequate}/E)$. Indeed, this is a theorem of probability theory: no matter what the evidence is, if one hypothesis entails another but not conversely, the first cannot be more probable than the second. The probability of "T is true" cannot exceed the probability of "T is empirically adequate."

This is true enough, but it is too slender a basis on which to rest the version of empiricism that van Fraassen defends. A realist can grant the mathematical point, but still maintain that the posterior probabilities we assign to theories about unobservables can be driven quite high. If "T is empirically adequate" has a probability of 0.95, then "T is true" cannot be assigned a greater value. But 0.94 is plenty big enough. The simple mathematical constraint on probability assignments is not enough to represent van Fraassen's claim that we should be agnostic as to the truth value of theories about unobservables.

A further reason can be given for being skeptical about this argument of van Fraassen's. Granted, "T is empirically adequate" is logically weaker than "T is true," but there are statements that are weaker still. "T is empirically adequate concerning events that occur on Mondays, Wednesdays, and Fridays" is even weaker, but surely it is arbitrary to say that nothing stronger than Monday-Wednesday-Friday claims about empirical adequacy is defensible. Van Fraassen's *argument from weakness* is defective unless it can be supplemented with an explanation of why reason may venture to a certain point, but no further.[14]

I so far have explored some Bayesian formulations of the epistemology that underlies constructive empiricism. The fact that these are unworkable does not show that no Bayesian approach can be made to work. Indeed, van Fraassen elaborates a more sophisticated proposal in his subsequent book, *Laws and Symmetries*. His idea is to represent the agnosticism that constructive empiricism counsels by using Kyburg's and Levi's ideas about interval probabilities. Van Fraassen's (1989, 193–94) proposal is that an agent is thoroughly agnostic about a hypothesis H precisely when he or she says that the probability of H is between 0 and 1 (inclusive).

Van Fraassen thinks this representation of agnosticism entails that thorough agnostics cannot learn from experience:

> What is the effect of new evidence? If hypothesis H implies E, then the vagueness of H can cover at most the interval $[0, \mathrm{P}(E)]$. So if E then becomes certain, that upper limit disappears. For the most thorough agnostic concerning H is vague on its probability from zero to the probability of its consequences, and remains so when he conditionalizes on any evidence. (van Fraassen 1989, 194)

I want to question the adequacy of this representation of agnosticism as well as its suitability as a vehicle for expressing empiricist views concerning hypotheses about unobservables.

From Bayes's Theorem we know that if $\mathrm{P}(E/H) = 1$, then $\mathrm{P}(H/E) = \mathrm{P}(H)/P(E)$. This is why the prior probability $\mathrm{P}(H)$ must fall in the interval $[0, \mathrm{P}(E)]$. As van Fraassen notes, once one learns that E is true (where E was not certain beforehand), the upper bound on the probability of H increases; the interval associated with $\mathrm{P}(H/E)$ is $[0, 1]$. So learning that E is true does not leave one's attitude to H unchanged. Suppose I learn that smoking may be more dangerous than I had thought. I might say, "I used to think that my chance of cancer was somewhere between 0.001 and 0.01, but now I realize that my chance may be as high as 0.25." Increasing the upper bound reflects learning from experience, not the maintenance of one's previous epistemic state.[15]

If learning that E is true is genuinely to leave one's attitudes to H unchanged, then the interval associated with $\mathrm{P}(H)$ must be the same as the interval associated with $\mathrm{P}(H/E)$. However, when H entails E, this will be true only if $\mathrm{P}(E) = 1$, which in turn is true only if E is entailed by both H and its negation. Observations are rarely related to theories in this way. My conclusion is that the assignment of interval probabilities is not an adequate method for codifying the agnosticism that figures in van Fraassen's constructive empiricism.

If Bayesianism does not provide a hospitable framework for van Fraassen's version of empiricism, it is well to consider other ways of quantifying how evidence bears on hypotheses. Many empirical scientists are reluctant to assign probabilities to the hypotheses they consider; at least such assignments do not figure in their "public" pronouncements about theories. However, this does not stop them from evaluating hypotheses by seeing how those hypotheses probabilify the observations they have made. That is, although scientists rarely discuss probabilities of the form $\mathrm{P}(H/E)$, they often discuss probabilities of the form $\mathrm{P}(E/H)$. The latter sometimes go by the technical term *likelihood*. The likelihood of a hypothesis is the probability it confers on the data, not the probability it possesses in the light of the data.

The likelihood concept is useful when it comes to *comparing* rival hypotheses. One evaluates competing hypotheses by seeing how each probabilifies the data. The basic idea is called the *Likelihood Principle* (Edwards 1972):

$$O \text{ favors } H_1 \text{ over } H_2 \text{ if and only if } \mathrm{P}(O/H_1) > \mathrm{P}(O/H_2).$$

If H_1 says that the data were to be expected while H_2 says that it is almost a miracle that the observations came out as they did, this means that the observations strongly favor H_1 over H_2.

How might the Likelihood Principle be used to carve out a difference between observable and unobservable entities? A scientist will test hypotheses about observables by finding observations that discriminate between them; for example, we might test claims about the dietary habits of different dinosaurs by looking at such things as their fossil remains. If hypotheses about *un*observables were not amenable to this approach, this would mark an important sense in which they are epistemically inaccessible. The idea is that *when* H_1 *and* H_2 *are both about unobservables, every observation* O *is such that* $\mathrm{P}(O/H_1) = \mathrm{P}(O/H_2)$. This expresses the thought that observations cannot distinguish between hypotheses that are about things we cannot observe.

Of course, this suggestion is absurd. As van Fraassen clearly realizes, competing hypotheses about unobservables make different predictions about what we can observe. By seeing which predictions come true, we are able to make judgments about which hypotheses the observations favor.

Not only does the Likelihood Principle apply both to hypotheses about observables and to hypotheses about unobservables. In a certain sense, it entails that observables and unobservables are on an epistemic par. If the observations we have made discriminate between "H_1 is empirically adequate" and "H_2 is empirically adequate," those observations must also discriminate between "H_1 is true" and "H_2 is true." Van Fraassen's constructive empiricism says that it is sensible to judge a theory empirically adequate, but metaphysical to say that the theory is true. The Likelihood Principle entails that this is an untenable dualism (Sober 1990a).

I have suggested that empiricism and realism should be formulated in terms of some matter-of-degree concept such as *degree of belief*, not in terms of the absolute concept of *acceptance*. The reason I have given for this suggestion has nothing special to do with realism and empiricism, but simply reflects a widely held interpretation of the lottery paradox.[16] However, there is an additional argument for this reformulation, one that is internal to the dispute between realism and empiricism.

Van Fraassen (1980, 16) agrees with earlier realists (e.g., Maxwell 1962) that the distinction between observable and unobservable entities is a matter of degree. However, unlike some of his realist opponents, van Fraassen insists that the absence of a precise boundary does not obliterate the distinction or its importance. At one end of the continuum there are objects that are clearly observable, while at the other there are objects that are clearly unobservable. In between is a gray area. I agree with *part* of what van Fraassen is saying here; we cannot reject the use of a distinction—in science, in philosophy, or in ordinary life—just on the ground that it is vague.

Nonetheless, there is a further issue, this one specific to how van Fraassen formulates empiricism. If our epistemic attitudes to a proposition are supposed to depend on whether the proposition is about unobservable entities, what attitude should we take to propositions that are about objects that fall in the gray area? Van Fraassen could elect to say nothing about these borderline cases. But surely empiricism can be given a less arbitrary formulation. If observability is

a matter of degree, and if epistemic attitudes are to be keyed to observability, then those epistemic attitudes should themselves be formulated as matters of degree. A more natural formulation of empiricism might say something like this: *how certain we are entitled to be about a given proposition will depend on how observable the objects are that the proposition is about.* This is an internal reason for empiricists to shy away from the dichotomous category of acceptance.

3. SOURCE VERSUS SUBSTANCE

Van Fraassen's constructive empiricism is first and foremost a position in epistemology. Its fundamental claims address the question of how rational acceptance should be understood. Yet, when we try to equip that doctrine with an epistemology in which the all-or-nothing category of *acceptance* is replaced by the more nuanced idea of *degree of evidential support*, we run into trouble. It is entirely unclear how the distinction between observables and unobservables could be grafted onto a theory of the evidence relation that is probabilistic in character.

It is useful at this point to back up one step and try to see *why* van Fraassen's formulation of empiricism goes astray. I think a hint is provided in the following passage:

> . . . for James and Reichenbach, the core doctrine of empiricism is that experience is the sole source of information about the world and that its limits are very strict. . . . I explicate the general limits as follows: *experience can give us information only about what is both observable and actual.* (van Fraassen 1985, 253)

It is important to recognize a gap between the "core doctrine" and van Fraassen's proposed explication. The formulation he attributes to James and Reichenbach concerns the *sources* of evidence that are available to us; in contrast, van Fraassen's own proposal concerns the *content* of the hypotheses to which we have epistemic access. However, the claim about sources and the claim about substance are quite different. Even if evidence comes to us only through observation, that does not mean that our evidence bears only on what we believe about observables.

Perhaps it can be shown that the claim about *sources* leads inevitably to the claim about *substance*. But this is not at all obvious. For example, suppose that for some reason I am unable to look inside the gas tank of my car to see how much gas it contains. Perhaps the only access I have to claims about the contents of the gas tank derives from looking at the gas meter on the dashboard, noticing when the car stops moving (such events are curiously correlated with the needle's pointing to the letter "E"), and so on. My *sources* of information are limited to what I can observe, and I am supposing that I cannot observe what is in the gas tank. But surely it does not follow that the *content* of what

I believe should be limited to claims about what the meter says, about when meter readings are correlated with the car's grinding to a halt, and so on.[17]

It is interesting that immediately after giving the above quoted explication of what empiricism says, van Fraassen elaborates on it as follows:

> We may be rational in our opinions about other matters—Augustine's "faith in things unseen," which he rightly said, pervades even our everyday opinions about everyday matters—but any defense of such opinions must be by appeal to information about matters falling within the limits of the deliverances of experience. (van Fraassen 1985, 253)

This more modest formulation says that our evidence is limited to the deliverances of experience; however, it allows that experience can bear on hypotheses that describe "things unseen."

One salient feature of van Fraassen's approach is that he criticizes earlier formulations of empiricism in which scientific theories are construed nonliterally. The philosophical positions he criticizes include proposed phenomenalistic reductions, in which a scientific theory apparently about electrons is said to be translatable into a theory about experiences. Van Fraassen wants to emancipate empiricism from this semantic formulation. Indeed he does so, but a residue of semantics remains. According to van Fraassen, the epistemic attitude we are supposed to have towards a theory depends on what the theory is about; we may assign truth values to statements strictly about observables, but should remain agnostic with respect to statements that are about things we cannot observe. Perhaps the difficulties facing van Fraassen's formulation can be overcome if we formulate empiricism as a claim about *sources* of information, and not, in the first instance, as a claim concerning the *subject matter* of hypotheses about which we can hope to get evidence.[18]

Even after the claim about sources is separated from the claim about subject matter, the empiricist claim that experience is our sole source of evidence requires refinement. It should not entail that reason plays no role in evaluating hypotheses. Suppose we are comparing various hypotheses that seem *prima facie* to be possible explanations for some body of data; we then discover that one of those hypotheses contains a logical inconsistency. Surely this discovery can be grounds for rejecting the hypothesis. Here it is reason, not sense experience, that drives our decision. In the light of this, what claim can empiricism propose about the privileged status of experience?

This question suggests that one needs to think of stronger and weaker forms of empiricism; these differ with respect to how determining a role they attribute to experience. The most extreme formulation will deny that *a priori* considerations play any role at all in hypothesis evaluation. If deductive logic is relevant to hypothesis evaluation, then this sort of empiricism will claim that the credentials of logic derive entirely from sense experience. A less extreme formulation will grant that deductive logic is *a priori*, but will draw the line there. Further along the continuum is the view that deductive logic and Bayesian coherence are the sole *a priori* constraints on what we ought to believe. If we move a good deal farther in this direction, we come to a point of view in which a slew of supposedly *a priori* principles are thought to constrain which theories we may accept. Kantian theses about synthetic *a priori* knowledge exemplify this type of position; coherent theorizing must assume that every event has a cause, that space is Euclidean, and the like. Here one has moved along the continuum sufficiently far to say that the doctrine at hand is no longer an instance of empiricism.

To decide where along this continuum the truth is to be found, one must address the issue of whether the "superempirical" virtues of a theory—its simplicity, elegance, generality, degree of unification—play a role in determining how we should evaluate theories. If the simplicity of a theory has nothing to do with how well supported it is by available observations, then an empiricist should say that simplicity is not a sign of truth—it is irrelevant to assessing the theory's plausibility. Although few present-day realists would endorse the Kantian principles alluded to above, many uphold the relevance of such superempirical virtues. Accordingly, this is a very important issue for the dispute between realism and empiricism.[19]

The intuitive idea that the empiricist must try to explicate is that theory evaluation should be driven by observations, not by *a priori* preconceptions. If we observe many emeralds and find that all have been green, the empiricist has no trouble explaining why we should prefer the hypothesis that all emeralds are green over the hypothesis that only 87 percent of them are. Logic may be *a priori*, but the use of logic in this instance allows the observations to decide which hypothesis is more plausible. A more difficult problem for the empiricist is to explain why we should prefer "all emeralds are green" over the claim that "emeralds are green until the year 2100, but thereafter are blue." Can observations be cited that help one discriminate between these two hypotheses, or is some irreducible appeal to simplicity the only thing that stands between us and skepticism? If the latter, empiricists will be driven to skepticism; the question will then be whether we should follow them there, or view this consequence of their position as a *reductio*.

So the dispute about empiricism should not be understood as a clash between reason and experience. The question is not "should we reason or should we observe?" Rather, the problem concerns what reason dictates as policies of belief change. If hypothesis evaluation should be driven by observations, the empiricist wins; but if legitimate discriminations among self-consistent hypotheses far outrun what is mandated by the observations we have available, then empiricism is mistaken. Again, let me emphasize that this is not a dispute between clearcut dichotomies, but concerns matters of emphasis.

4. THREE DESIDERATA AND A PROPOSAL

I have suggested three *desiderata* concerning how empiricism should be formulated. First, the doctrine should claim that observations are special *sources* of information concerning which theories are true; it should not claim that

observables are the sole *subject matter* about which we can have reasonable opinions. Second, the doctrine should not be expressed in terms of the dichotomous concept of acceptance, but should be formulated in terms of some matter-of-degree concept like evidential support. And finally, the doctrine should conform to the principle of actualism; the significance attributed to the distinction between observable and unobservable entities should attach instead to the distinction between what is observed and what is not.

One natural way to satisfy these *desiderata* is by viewing the Likelihood Principle as the sole vehicle by which competing hypotheses can be evaluated. If one theory is claimed to be more plausible than another, this cannot be due to some irreducible fact about how simple or unified the two theories are; comparisons of plausibility must be based on observations that discriminate among the competitors.

The Likelihood Principle imposes no limitation on the subject matters or vocabularies of competing hypotheses. If two hypotheses about electrons make different predictions about what we can observe, then observations will allow us to say which is more plausible. In this respect, hypotheses about electrons are on the same footing as hypotheses about the moons of Jupiter. If the data we actually possess discriminate among the competitors, we can make a scientific assessment; if not, not.[20]

It follows from this epistemology that certain kinds of skeptical challenges cannot be answered. If the evil demon hypothesis is formulated in such a way that no possible experience can discriminate between it and the "normal" hypotheses we believe, then science is in no position to say which of them is true. Philosophers may talk of *ad hocness*, simplicity, and explanatoriness all they wish, but they should realize that they are not invoking the epistemological standards of science, but ones of their own devising. Science is not in the business of attempting to discriminate between predictively equivalent hypotheses. The scientific method provides no handy solution to such problems; rather, such problems are passed over as unscientific (Sober 1990a).

The name I give to this approach to scientific inference is *contrastive empiricism* (Sober 1990a; 1993). It involves a compromise between realism and empiricism. It draws from realism the idea that hypotheses about unobservables are just as much within the scope of scientific evaluation as hypotheses about observables. From empiricism, it draws the idea that theory evaluation must be driven by observations; appeal to the so-called "superempirical" virtues is idle.[21]

5. OBSERVATION: BETWEEN SCYLLA AND CHARYBDIS

I have already pointed out that van Fraassen (1980) grants that the distinction between observable and unobservable is a matter of degree, though this does not stop him from imposing on it a dichotomous epistemology, in which some statements, but not others, may be regarded as true. However, it is only fair to note that this fixation on dichotomous epistemological categories also is

prominent in the writings of van Fraassen's realist opponents. For example, after emphasizing the highly fallible nature of scientific theorizing, Churchland (1985, 36–37) asks:

> Why, then, am I still a scientific realist? Because these reasons fail to discriminate between the integrity of observables and the integrity of unobservables. If anything is compromised by these considerations, it is the integrity of theories generally. That is, of *cognition* generally. Since our observational concepts are just as theory-laden as any others, and since the integrity of these concepts is just as contingent on the integrity of the theories that embed them, our observational ontology is rendered *exactly as dubious* as our nonobservational ontology.

According to Churchland, observables and unobservables are *in the same boat*, because *cognition in general* is fallible and theory-laden.

It is a startling thesis that our confidence in the existence of dogs and tables is no more secure than our confidence in the existence of oncogenes and leptons. Granted, all cognition is fallible and theory-laden. But is all cognition fallible and theory-laden *to the same degree*? This monism is no less bizarre than van Fraassen's dualism; scientists would no doubt scratch their heads in wonderment at the claim that everything is just as dubious as everything else.

Perhaps Churchland's choice of words in this passage overreaches what he wants to assert. He advances the *positive* claim that statements about observables and statements about unobservables are equally dubious. But maybe all he had in mind was a *negative* thesis; perhaps his point was just to reject the idea that every statement about unobservables is more dubious than every statement about observables. As noted in Section 3, I reject this idea as well. However, much is left unsettled by this very modest point. Granted, there exists at least one statement about unobservables that is more certain than at least one statement about observables. Still, it does not follow that the observable-unobservable distinction is epistemologically irrelevant. There is at least one smoker who has better health than at least one nonsmoker; however, it would be a mistake to conclude from this that the smoking-nonsmoking distinction is irrelevant to health.

So the question that needs to be addressed is whether there is any matter-of-degree epistemic difference that is associated with the distinction between observables and unobservables, given that the latter distinction may itself also be a matter of degree. I think the answer to this question is *yes*. I will argue for this thesis in two steps. First, I will map the distinction between observable and unobservable onto the distinction between *direct* and *indirect*; then I will provide a likelihood representation of the idea that we have more decisive evidence about matters with which we have direct contact than we have concerning matters that are known only by indirection.

I begin with what I regard as a truism: *when something is unobservable, the only way we can find out about it is by finding out about something that is observable*. Electrons are too small to see. So we use instruments. We usually

encounter few problems in seeing what state a measuring device is in; we then infer from the state of the measuring device what state the electrons occupy in our experiment. This intuitive idea may be summarized as follows: *our knowledge of unobservables is necessarily indirect.* Here I use "unobservable" in a wider sense than van Fraassen does—we cannot observe flesh-and-blood dinosaurs because of our temporal location, the moons of Jupiter because of our spatial location, and present-day viruses because of our size. Since we are large earthlings who live in modern times, we can know about each only indirectly.

Things from which we are not cut off in this way can be known without a pathway through unobservables. We can know what state the voltage meter is in—what number the pointer points to—without knowing the state of the electrons in the experiment. And we can know what a fossil looks like without already knowing what the dinosaur was like from which the fossil derives. Knowledge of unobservables is necessarily mediated by knowledge of observables, but knowledge of observables need not pass through knowledge of unobservables. Here is an asymmetry between observables and unobservables that is worth pondering.[22]

The asymmetry I am proposing does not entail that our knowledge of observables is infallible or theory-neutral. Granted, we can be wrong when we judge what state a measuring device is in; and we can tell what state the device is in only by exploiting background information. When I say that unobservables are knowable only indirectly, while observables are knowable directly, I am positing a difference in degree. Maybe it is true that *all* perceptual judgments about physical objects are generated by inferential processes (Marr 1982). *Direct* does not mean totally *non*inferential, but that inferences are present to a lesser extent. Perhaps, when we tell what state an electron occupies by looking at a measuring device, there are *two* layers of inference. First, we figure out what state the device is in; then, on that basis, we infer what state the electron occupies.

There is a second misinterpretation that also should be avoided. I am not saying that every belief that is about an unobservable must involve "more" inference than every belief that is about an observable. I do not know how to compare such quantities, especially when they attach to hypotheses that are about quite different subject matters. Doubtless, it is possible to imagine an inference about electrons that involves three "steps" and an inference about the birds living in Antarctica that involves four. In any case, there is no need to propose a *total ordering* of *all* propositions in the one category and *all* propositions in the other. Rather, my proposal posits a *partial ordering;* it is, so to speak, "chain internal." When we infer something about an unobservable u, based on beliefs we have concerning an observable o, u will be known more indirectly than o is; the inferences grounding our belief about o will form a subpart of the inferences grounding our belief concerning u.

I hope it also is clear that my thesis requires no precise boundary between observable and unobservable, or between what is observed and what is inferred from observation. Van Fraassen says we do not "observe" when we use instrumentation; others have adopted a more liberal terminology. For me, it makes no difference whether or where a line is drawn. Whether we "observe" electrons in a cloud chamber or merely "detect" them, the point is that our belief about the electrons is based on a judgment we form about the state of the cloud chamber screen.

Having argued that the distinction between observable and unobservable maps onto the distinction between direct and indirect, I now need to explain why this fact is epistemologically significant. Why should extra inferences, so to speak, have the effect of weakening the testimony of experience? Even if I decide what state the electron occupies by looking at a measuring device, why can't my experience provide more evidence about the electron than it does about the state of the measuring device?

The reason indirection affects strength of evidence can be understood within the framework of the Likelihood Principle. Consider a causal chain, in which an effect E traces back to a (relatively) direct cause (D) and thence to a more indirect cause (I):

$$I \longrightarrow D \longrightarrow E$$

Let us suppose that this causal chain has the *Markov* property: the intermediate link D "screens off" E from I. To see what this means, imagine that each of the three nodes can be in one of two states, which I will denote as "0" and "1". When E is in state 0, I will express this fact by saying that "$E = 0$," and similarly for other nodes and other states. Screening off means that

$$\mathrm{P}(E = i/D = j) = \mathrm{P}(E = i/D = j \;\&\; I = k) \neq \mathrm{P}(E = i/I = k),$$
for $i, j, k = 0, 1$.

The idea is that if one knows the state of D, one's ability to predict the state of E remains the same, whether or not one also knows the state of I.

Suppose we know that $E = 0$. On this basis, we wish to infer what state D is in and also what state I occupies. In each case, we will use the Likelihood Principle to interpret the data. We will compute the likelihood ratio $\mathrm{P}(E = 0/D = 0)/\mathrm{P}(E = 0/D = 1)$. This will tell us whether and to what degree the observation that $E = 0$ favors $D = 0$ over $D = 1$. Likewise, we will compute $\mathrm{P}(E = 0/I = 0)/\mathrm{P}(E = 0/I = 1)$, the point this time being to see how strongly the observation favors one hypothesis about the indirect cause over the other. In particular, we will be interested in the relationship between these two likelihood ratios. When does the state of E provide more information about the state of D than it does about the state of I?

The answer is *always* (so long as the screening-off relation obtains and the probabilities are strictly between 0 and 1): E tells us more about D than E tells us about I. More specifically, the claim is:

If $P(E = 0/D = 0)/P(E = 0/D = 1) > 1$, then
$P(E = 0/D = 0)/P(E = 0/D = 1) >$
$P(E = 0/I = 0)/P(E = 0/I = 1)$ and
$P(E = 0/D = 0)/P(E = 0/D = 1) >$
$P(E = 0/I = 1)/P(E = 0/I = 0)$.

The antecedent of this conditional means that $E = 0$ favors $D = 0$ over $D = 1$; the consequent means that $E = 0$ favors $D = 0$ over $D = 1$ more than $E = 0$ favors either state of I over the other. The same point holds for the other state that E might occupy, namely $E = 1$; this observation also provides more information about D than it does about I. So no matter which state E happens to occupy, observing that state tells us more about the direct cause than it does about the indirect cause. A proof is provided in the Appendix.

Consider an example of this three-step process. The state of an electron in some apparatus causes a measuring device to go into some particular state; the state of this measuring device causes a visual image to form in the mind of the experimenter. The chain goes from electron to meter reading to sensation. The sensory image is evidence about the state of the meter, in that the sensory image favors one hypothesis about the meter over another; and the state of the meter is evidence about the state of the electron, in that the state of the meter favors one hypothesis about the electron over another. As it happens, the favoring relation is *transitive;* it follows that the sensory state favors one hypothesis about the electron over the other (see the Appendix). But more to the point, the sensory state provides better evidence about the state of the meter than it does about the state of the electron.[23]

Perhaps the belief that *the meter reads 9.3* results from an inference of the same general type as the inference that generates the belief that *the electron is in state F*. Both involve "abduction," which is to say that in both cases data are interpreted by appeal to the Likelihood Principle. However, this *qualitative* parity between the two hypotheses does not entail that they are on a *quantitative* par. And indeed they are not. The hypothesis that is more directly related to sensation is the one about which sensation has more to say. *Experience teaches us the most about matters that are closest to experience.*[24]

Empiricism says that belief revision ought to be driven by observation. In privileging observation, the claim is not that observations are known with certainty nor that they can be known without using an interpretive theory. Nor is it plausible to claim that statements about observables are knowable, whereas other statements are not. These *absolute* categories have no place in a probabilistic epistemology. Observation is special because of its *relative* status. Our knowledge concerning "things unseen" is mediated by our knowledge of the things we see. This asymmetry has epistemological consequences, which empiricism must seek to describe without exaggeration and without the imposition of false dichotomies.

APPENDIX

In the Markov chain from I to D to E, we define the following conditional probabilities:

$P(E = 0/D = 0) = a$ $\quad$ $P(D = 0/I = 0) = c$
$P(E = 0/D = 1) = b$ $\quad$ $P(D = 0/I = 1) = d$,

all of which we assume to be strictly between 0 and 1. Then

$$P(E = 0/D = 0)/P(E = 0/D = 1) > P(E = 0/I = 0)/P(E = 0/I = 1)$$

if and only if

$$a/b > [ca + (1 - c)b]/[da + (1 - d)b].$$

This latter quantity simplifies to

$$(a - b)[ad + b(1 - c)] > 0,$$

which must be true if $a > b$. A symmetrical argument goes through for the observation that $E = 1$.

Another consequence of this format is that the favoring relation is transitive:

If $P(E = 0/D = 0)/P(E = 0/D = 1) > 1$ and
$P(D = 0/I = 0)/P(D = 0/I = 1) > 1$,
then $P(E = 0/I = 0)/P(E = 0/I = 1) > 1$.

In other words: if $a > b$ and $c > d$, then

$$ca + (1 - c)b > da + (1 - d)b.$$

The consequent of this conditional simplifies to $(a - b)(c - d) > 0$.

The argument developed in this Appendix in terms of likelihood ratios is specifically for the case in which nodes each have two possible states. When this is generalized to the case of n states, the natural measure to use is R. A. Fisher's idea of *mutual information*. The mutual information between two random variables X and Y is defined as

$$M(X, Y) = \sum_{x} \sum_{y} P(X = x \,\&\, Y = y) \log[P(X = x \,\&\, Y = y)/P(X = x)P(Y = y)].$$

For the Markov chain from I to D to E, Van Rijsbergen (1983) and Forster (unpublished) have each shown that $M(E, D) > M(E, I)$. See Sober and Barrett (1992) for discussion.

NOTES

I thank Martin Barrett, Ellery Eells, Malcolm Forster, Alan Musgrave, Alan Sidelle, Bas Van Fraassen, and Leora Weitzman for useful comments on earlier drafts of this essay.

1. Van Fraassen does not deny that scientists do other things besides make judgments about empirical adequacy. Sometimes they apply for grants, stab each other in the back, and work long hours. But these are not, for him, *requirements* for one to play the game of science. Scientists who do not participate in these activities may still be "scientific."

Likewise, van Fraassen does not deny that scientists sometimes *ought* to do things besides make judgments about empirical adequacy. Scientists have moral obligations, like the rest of

us. But here again, van Fraassen does not regard the obeying of moral norms as requirements for one to play the game of science. Immoral scientists can still be "scientific."

I thank van Fraassen (personal communication) for helping me clarify the issues involved here, though I do not claim that he will agree with what I say.

2. Readers are invited to replace "requires" with "permits" when constructive empiricism is under discussion to check my claim that the arguments I will advance do not turn on which of these formulations is used.

3. Van Fraassen (i) assigns to observability a crucial role in his version of empiricism, (ii) defines observability in terms of a conditional that presumably must be interpreted nontruth functionally, and (iii) says that elucidating the property of observability is a scientific problem. Nonetheless, van Fraassen (1980, 115–16) also maintains that counterfactual conditionals are not part of the content of any scientific theory, since such conditionals, he says, do not describe objective facts about nature.

4. It won't matter for the purposes of this example whether "S has diabetes" and "S has small pox" are "strictly about observables." If the reader believes that they are not, and feels that this affects the point, he or she should reformulate the argument such that the hypotheses under test concern the number of moons of Jupiter and the number of moons of Mars.

5. Whether the third test is *perfectly* infallible is irrelevant to the point I wish to make; merely imagine that its error probabilities are *lower* than those for the first two tests.

6. Actualism plays a role in a number of epistemic problems—for example, in the optional stopping problem in statistics (Hacking 1965, 107-9) and in the task of explaining why the asymptotic convergence of an inference method is inessential for the method to be justified when applied to finite data (Sober 1988). Eells (1993) formulates a similar thesis: "we have to work with what we have to work with."

7. A scenario can be invented in which the existence of the third test procedure bears on how the results of the first two tests should be interpreted. Merely imagine that Clyde, a mad scientist, would be inclined to invent a procedure of the third kind only if he were inclined to change the error characteristics of the first test. I take it that the mere *possibility* of such scenarios is consistent with the principle of *actualism*.

8. Could we directly perceive the AIDS virus if we were to shrink to the right microscopic size? Even if we could, van Fraassen thinks this would not show that the virus is observable. Van Fraassen's (1980, 17) notion of observability rests on the idea that our location in space is accidental, while our approximate size is essential to us, *qua* human beings. It is hard to resist chiding van Fraassen—an avowed foe of real modality—for these essentialist expostulations. One is reminded of Quine's question about the bicycling mathematician: Is he essentially rational and accidentally bipedal, or accidentally rational and essentially bipedal? Surely empiricism ought not to depend on deciding whether our location is accidental and our size essential, or the other way around.

9. Musgrave (1985, 206) puts the point very well: "it is a curious sort of empiricism which sets aside the weight of *available* evidence on the ground that a casual observer might one day see his mouse or yeti, while the scientist can never see (but can only detect) his electrons."

10. Van Fraassen uses the term "acceptance" so that empiricists and realists have different views about what it means. For empiricists, accepting a theory means thinking that *it is empirically adequate*; for realists, acceptance means believing that *the theory is true*. Notice that in both cases, acceptance involves believing-true some proposition or other (the issue being what the proposition ought to be). I will depart from van Fraassen's usage in what follows. For me, accepting a statement simply means believing that it is true. Realists say that available evidence sometimes warrants accepting theories about unobservables; constructive empiricists deny this. This terminological shift leaves the substantive issues unchanged.

11. This problem is obviously not news to van Fraassen (1980); right after his formulation of realism as a doctrine about acceptance (p. 8), and in his accompanying note 3 (p. 216), van Fraassen says that a realist might adopt a Bayesian framework in which hypotheses are assigned degrees of belief. Van Fraassen remarks that this would be an unusual stance for a realist, but one that is not ruled out by any realist principle. However, the present point is not that realists *could* jettison acceptance, but that both realists and empiricists *should* abandon that category.

I also should mention that van Fraassen takes up the problem of providing a probabilistic representation of his version of empiricism in his book *Laws and Symmetries* (Van Fraassen 1989). I will discuss this proposal in due course.

12. It might be objected, on Duhemian grounds, that hypotheses about unobservables never, *on their own*, entail observational predictions. This point, I reply, can be accommodated. Instead of saying that H entails O, let's say that $H\&A$ entails O (here A provides the requisite "bridge" principles). Then the probability of $H\&A$ is sensitive to the observations, even though $H\&A$ is not strictly about observables.

13. Skyrms (1984) explores a formulation of empiricism in which empirical propositions are the ones whose probabilities can be raised or lowered by observations. It is notable that Skyrms's development of this idea fails to coincide with important elements in van Fraassen's position.

14. Giere (1985, 83) provides a nice version of this response to van Fraassen's argument by locating constructive empiricism in a continuum of positions, ranging from solipsism of the present moment at one extreme to a form of modal-realism at the other.

15. If second-order probabilities were incorporated into this framework, one could go farther. For some second-order distributions (e. g., a uniform distribution), one could say that increasing the upper bound on the first-order probability would entail that the first-order probability has increased. This would further undermine van Fraassen's proposal for how agnosticism should be represented.

16. Using a dichotomous category of acceptance is harmless so long as what one says can be reconciled with the fact that the dichotomy rests, somewhat arbitrarily and contextually, on the matter-of-degree concept of plausibility. In this respect, these epistemic concepts are related to each other the way the concept of *being tall* is related to the concept of *height in inches*. The problem exposed by the lottery paradox is that empiricism and realism should be subjected to a kind of consistency check. Can their pronouncements be represented within an epistemology in which plausibility is construed as a matter of degree?

17. It is worth considering what the word "information" means in the quotation from van Fraassen that is under discussion. I suggest that this concept should be understood within the framework of the Likelihood Principle (see Section 5 and the Appendix). An experience gives us information about a set of competing hypotheses precisely when those hypotheses (perhaps in conjunction with a plausible background theory) confer different probabilities on the experience's occurring. Of course, this explication conflicts with the dictates of constructive empiricism. Van Fraassen needs to explain what alternative explication of "information" he wishes to endorse.

18. The challenge to empiricism, then, is to show that some thesis about substance *follows* from the thesis about sources. I try to defend such an inference in Section 5.

19. Sober (1988; 1990b), and Forster and Sober (forthcoming) develop proposals for bringing simplicity considerations within the purview of a broadly empiricist framework.

20. When prior probabilities are based on a model of a chance process, there is no problem in using them as well as likelihoods to evaluate the overall plausibility of hypotheses. In addition, it is not to be denied that background assumptions shape our evaluation of hypotheses in ways that are additional to the testimony of the data presently at hand. These background assumptions trace their credibility back to likelihood considerations, relative to earlier data sets.

21. If a connection between simplicity and likelihood could be established (in which case simplicity would not be a *super*empirical virtue of theories), then contrastive empiricism would grant that simpler theories are more plausible. There may be *local* contexts in which

this connection can be established. For example, in Sober (1988) I explore the extent to which the principle of parsimony used in phylogenetic inference can be understood within a likelihood framework. And Forster and Sober (forthcoming) provide a framework of considerable generality within which the simplicity of a curve is linked with its predictive accuracy. However, I doubt that there is any global and presuppositionless connection between simplicity and evidential warrant (Sober 1990b).

22. This asymmetry is important in understanding the role that observations play in resolving theoretical disputes. However theory-laden the observations may be, it usually is possible to reach agreement about what the observations are without already agreeing about which theory under test is true (Sober 1993).

23. Although I described the chain from I to D to E as a *causal* chain, the causal idea is quite irrelevant to the above result. What is essential is that the chain have the Markov property. See Sober and Barrett (1992) for further discussion.

When the likelihood ratios discussed above have values different from unity, the states of the different nodes are *correlated*. The likelihood conception of how observation or detection ought to be understood accords well with some of the main ideas behind the treatment in Wilson (1985).

24. To clarify what this Likelihood Result entails, it is useful for a moment to lapse into a Bayesian idiom. If $\mathrm{P}(H)$ is the prior probability of H and if $\mathrm{P}(H/E)$ is the posterior, then $\mathrm{P}(H/E)/\mathrm{P}(H)$ measures the degree to which E confirms H. It is quite possible that H should be quite certain and that E should confirm it only slightly or not at all; simply imagine that $\mathrm{P}(H/E) = \mathrm{P}(H) = 0.95$. Likewise, it is possible that E confirms H a lot, even though H is quite improbable in the light of H; simply imagine that $\mathrm{P}(H) = 0.000001$ and that $\mathrm{P}(H/E) = 0.001$. In short, degree of confirmation is a *diachronic* concept, measuring *changes* in probability, whereas posterior probability is, so to speak, *synchronic*.

Now let us return to the chain from I to D to E and consider the posterior probabilities that propositions about D and I have in the light of the observation that $E = 0$. By Bayes's Theorem we can deduce

$$\frac{\mathrm{P}(I = 0/E = 0)}{\mathrm{P}(I = 1/E = 0)} = \left[\frac{\mathrm{P}(E = 0/I = 0)}{\mathrm{P}(E = 0/I = 1)}\right]\left[\frac{\mathrm{P}(I = 0)}{\mathrm{P}(I = 1)}\right]$$

and

$$\frac{\mathrm{P}(D = 0/E = 0)}{\mathrm{P}(D = 1/E = 0)} = \left[\frac{\mathrm{P}(E = 0/D = 0)}{\mathrm{P}(E = 0/D = 1)}\right]\left[\frac{\mathrm{P}(D = 0)}{\mathrm{P}(D = 1)}\right].$$

The likelihood result developed in the text describes a relationship between the first ratio on the right side of the first equation and the first ratio on the right side of the second. However, this result does not tell us how to evaluate the two left-hand ratios. The result says nothing about how the priors are related, and so does not determine the relation of the posterior probabilities (Sober and Barrett 1992).

This is why the proper motto is that *experience teaches us the most about matters that are closest to experience*, not that *experience allows us to be most certain about matters that are closest to experience*. This latter thesis would be obtainable from the former one, once the likelihood result is supplemented with a thesis about priors. However, I am not Bayesian enough to defend any such further thesis.

REFERENCES

Churchland, P. 1985. "The Ontological Status of Observables: In Praise of the Superempirical Virtues." In *Images of Science: Essays on Realism and Empiricism*, edited by P. Churchland and C. Hooker, 35–47. Chicago.

Edwards, A. 1972. *Likelihood*. Cambridge.

Eells, E. 1993. "Probability, Inference, and Decision." In *Foundations of Philosophy of Science*, edited by J. Fetzer, 192–208. New York.

Feyerabend, P. 1963. "How to be a Good Empiricist—a Plea for Tolerance in Matters Epistemological." In *The Delaware Seminar II*, edited by B. Baumrin, 3–39. Newark, Delaware.

Forster, M. Unpublished. "In Defence of a Causal Theory of Mental Representation."

Forster, M. and E. Sober. Forthcoming. "How to Tell When Simpler, More Unified, or Less *Ad Hoc* Theories will Provide More Accurate Predictions."

Giere, R. 1985. "Constructive Realism." In *Images of Science: Essays on Realism and Empiricism*, edited by P. Churchland and C. Hooker, 75–98. Chicago.

Hacking, I. 1965. *The Logic of Statistical Inference*. Cambridge.

Kyburg, H. 1970. "Conjunctivitis." In *Epistemology and Inference*. Minneapolis.

Marr, D. 1982. *Vision*. San Francisco.

Maxwell, G. 1962. "The Ontological Status of Theoretical Entities." In *Minnesota Studies in the Philosophy of Science*, edited by H. Feigl and G. Maxwell, 3–27. Minneapolis.

Musgrave, A. 1985. "Realism vs. Constructive Empiricism." In *Images of Science: Essays on Realism and Empiricism*, edited by P. Churchland and C. Hooker, 197–221. Chicago.

Skyrms, B. 1984. *Pragmatics and Empiricism*. New Haven.

Sober, E. 1988. *Reconstructing the Past: Parsimony, Evolution and Inference*. Cambridge, Mass.

Sober, E. 1990a. "Contrastive Empiricism." In *Scientific Theories: Minnesota Studies in the Philosophy of Science*, vol. 14, edited by W. Savage, 392–412. Minneapolis.

Sober, E. 1990b. "Let's Razor Ockham's Razor." In *Explanation and Its Limits*. Royal Institute of Philosophy Supplementary Volume 27, edited by D. Knowles, 73–94, Cambridge.

Sober, E. 1993. "Mathematics and Indispensability." *Philosophical Review*.

Sober, E. and M. Barrett. 1992. "Conjunctive Forks and Temporally Asymmetric Inference." *Australasian Journal of Philosophy* 70: 1–23.

van Fraassen, Bas. 1980. The *Scientific Image*. Oxford.

van Fraassen, Bas. 1985. "Empiricism and the Philosophy of Science." In *Images of Science: Essays on Realism and Empiricism*, edited by P. Churchland and C. Hooker, 245–308. Chicago.

van Fraassen, Bas. 1989. *Laws and Symmetry*. Oxford.

Van Rijsbergen, K. 1983. "A Discrimination Gain Hypothesis." In *Proceedings of the 6th International ACM Conference on Information Retrieval*, 101–104, London.

Wilson, M. 1985. "What Can Theory Tell Us About Observation?" In *Images of Science: Essays on Realism and Empiricism*. edited by P. Churchland and C. Hooker, 222–44. Chicago.