

Two Cornell realisms: moral and scientific

Elliott Sober

© Springer Science+Business Media Dordrecht 2014

Abstract Richard Boyd and Nicholas Sturgeon develop distinctive naturalistic arguments for scientific realism and moral realism. Each defends a realist position by an inference to the best explanation. In this paper, I suggest that these arguments for realism should be reformulated, with the law of likelihood replacing inference to the best explanation. The resulting arguments for realism do not work.

Keywords Abduction · Empiricism · Explanation · Likelihood · Moral realism · Naturalism · Scientific realism

Realism about scientific theories and realism about normative moral propositions each begin with a semantic thesis: there are true statements in the category in question and the true ones are true independently of whether anyone thinks they are true, and also are true independently of whether anyone would come to believe them if they thought about the question in a certain way. There are two declarations of independence here. Believing that p is true doesn't make p true, regardless of whether the believer is God, a society, or a single human being. And even if you'd come to believe that p is true were you to consider the relevant evidence and analyze it in the right way, that fact is not what makes p true. The *Euthyphro* is the inspiration in both instances; for example, even if p is true if and only if some procedure would deliver the verdict that p is true, p isn't true in virtue of this fact about the procedure.

The two realisms I want to discuss here defend the semantic theses just mentioned in an indirect way. They don't directly address the class of *all* scientific

E. Sober (✉)
Philosophy Department, University of Wisconsin-Madison, 5185 Helen C. White Hall,
600 N. Park St., Madison, WI 53706, USA
e-mail: ersober@wisc.edu

theories or the class of *all* moral statements, but rather single out special subsets of each. We have good reason to think that certain scientific theories are true because the hypothesis that they are true provides the best explanation for why those theories accurately predict what we have observed. And we likewise have good reason to think that certain normative moral statements are true because those statements provide the best explanations of some of the (non-normative) observations we have made. Charles Sanders Peirce called this style of inference *abduction*; more recently, it has come to be called *inference to the best explanation* (Harman 1965; Lipton 1991). For example, when general relativity accurately predicted the bending of light that Eddington observed in a solar eclipse, the best explanation for why the theory was predictively accurate in this instance is the hypothesis that the theory is true (or approximately true). Similarly, when we observe what Hitler did in World War II, the explanation for why he behaved as he did includes the statement that he was morally depraved. The guiding idea behind this defense of the two realisms is *naturalistic*.¹ Scientists appeal to observations in justifying the scientific theories they endorse; philosophers should do the same thing in justifying philosophical theories. Inference to the best explanation is a tool for both science and philosophy.

How do these points about general relativity and Hitler connect with the broader semantic theses with which I began? The idea is that if general relativity is true, then some scientific theories are true, and if it is true that Hitler was morally depraved, then some normative moral propositions are true. What is more, we have observational evidence for each of the antecedents in these two conditionals. What is supposed to follow is that we have observational evidence for the claim that some scientific theories are true and also for the claim that some moral propositions are true.

Unfortunately, there is a fly in the ointment. The arguments just described use a principle about confirmation that Hempel (1965) called

The Special Consequence Condition: If O confirms H , and H entails C , then O confirms C .

It is widely recognized in philosophy of science that this principle is false. Here's an example that shows why: You are playing poker and a new hand is about to be dealt. You are sitting to the dealer's left, so the first card dealt will come to you. You are wondering whether that first card will be the Queen of Hearts. The dealer is careless in handling the deck, allowing you to catch a glimpse—you see that the card in question is red. The observation that the card is red confirms the hypothesis that the card will be the Queen of Hearts, and the hypothesis that the card is the Queen of Hearts entails that the card is a Queen; however, the observation that the card is red does not confirm the proposition that the card is a Queen. The Bayesian view of confirmation—that O confirms H precisely when O raises the probability of H —makes all this transparent. Before you see that the card is red, your probability that

¹ Naturalism in this sense takes its inspiration from Quine's (1953, 1960, 1963) insistence that philosophy is "continuous" with science. In Sober (2009), I call this thesis naturalism_p and distinguish it from metaphysical naturalism (the thesis that there are no supernatural entities) and methodological naturalism_s (the thesis that science should not make claims about the existence of supernatural entities).

the card will be the Queen of Hearts is $1/52$. When you observe that the card is red, this raises the probability that the card is the Queen of Hearts to a value of $1/26$. However, your observation does not raise the probability that the card is a Queen; it was $1/13$ before you made your observation and has the same value after. The lesson is clear: watch out for arguments that depend on the special consequence condition!^{2,3} Whether the empirical arguments for scientific realism and moral realism sketched above can be reconfigured so as to avoid this faulty premise is something we will consider in due course.

In what follows, I'll consider Sturgeon's (1984, 1986) argument for moral realism and Putnam's (1975b) and Boyd's (1980, 1983) arguments for scientific realism. All these arguments are naturalistic in the sense I've described and all make use of inference to the best explanation. I'll suggest that each should be reconfigured by replacing inference to the best explanation with the law of likelihood (an idea I'll explain). The question is whether this reformulation allows these empirical arguments for the two realisms to stand.

1 Sturgeon on moral realism

Sturgeon (1984) uses three examples to motivate the idea that moral propositions can help explain nonmoral matters of fact. The two best, I think, are these:

- Hitler started World War Two because he was morally depraved.
- Abolitionism became more popular in the eighteenth and nineteenth centuries in Europe and North America because slavery became morally worse during that period.

Sturgeon hopes that readers will find it unproblematic to regard these two statements as true, but, to his credit, he forthrightly addresses the more general question of why the moral propositions should be regarded as explanatorily relevant to the nonmoral propositions cited. Here is his criterion:

... if a particular assumption is completely irrelevant to the explanation of a certain fact, then that fact would have obtained, and we could have explained it just as well, even if the assumption had been false (p. 223).

Sturgeon's proposal, I take it, provides the following sufficient condition for explanatory relevance:

² I argue in Sober (1993, 2011a) that the indispensability argument for mathematical Platonism due to Quine (1953, 1981) and Putnam (1971) is flawed because it uses the special consequence condition.

³ Carnap (1950b) discerns two concepts of confirmation—the incremental concept (wherein O raises the probability of H) and the absolute concept (wherein $\Pr(H | O)$ is high). My criticism of the special consequence condition pertains to the first; the condition is correct when the second concept is used. This is no help to the realisms under discussion, however, since the absolute concept of confirmation fails to capture the idea of *evidential relevance*. $\Pr(H | O)$ can be high even when O is evidentially irrelevant to H , and it can be low even when O is positively relevant to H . It is unfortunate that Carnap used the term “confirm” to label this second concept.

If a fact F would not have obtained or F would not be just as well explained were assumption A false, then A is explanatorily relevant to F .

This proposition is somewhat circular, since “well explained” appears in the antecedent and “explanatorily relevant” appears in the consequent. However, the principle entails something that involves no such circularity:

(SERP) If a fact F would not have obtained were assumption A false, then A is explanatorily relevant to F .

I'll call this “Sturgeon’s explanatory relevance principle.” SERP entails that the first bulleted statement is true if Hitler wouldn’t have started World War Two had he not been morally depraved, and SERP entails that the second statement is true if abolitionism wouldn’t have become more popular at the time and place described had slavery not become more morally abhorrent. I will not quarrel with these counterfactuals but instead will pick some bones with SERP. I see three kinds of counterexample to this principle, which are easier to state and understand if I shift from talking about facts and assumptions to events: (i) cases where F not only causes A , but F is a necessary condition for A ; (ii) cases where F and A are joint effects of a common cause C , where C suffices for A , and C is necessary for F ; (iii) cases where A and F are simultaneous and A supervenes on F .⁴

Can Sturgeon’s principle of explanatory relevance be improved? Harman (1977) does not state a principle explicitly, but an example he describes and the anti-realist conclusion⁵ he draws about it suggest such a principle. The example involves a group of hoodlums who, just for fun, set fire to a cat. Harman says that the hoodlums’ behavior can be explained by their upbringing and their resulting psychological properties (which can be characterized non-normatively). There is *no need*, Harman says, to mention the hoodlums’ moral failings. For Harman, the best explanation of the hoodlums’ behavior omits the very consideration that Sturgeon wants to cite in his explanation of Hitler’s behavior. Harman’s argument is a parsimony argument (Sober 1990b, 2009; Shafer-Landau 2007), but here I want to examine it from the point of view of the idea of explanatory relevance. As I mentioned, Harman does not explicitly state a general criterion, but I think it is natural to take the following principle to be in the background of Harman’s discussion of his example:

(HERP) If C explains E , and if $\Pr(E \mid C \& X) = \Pr(E \mid C)$, then X is explanatorily irrelevant to E .

The idea here is that X is explanatorily irrelevant to E if C is explanatorily relevant to E and adding mention of X to C doesn’t change the probability of E ; in this case,

⁴ Harman (1986, p. 63) suggests a counterexample to SERP that involves epiphenomenalism; I would categorize some cases of epiphenomenalism under case (ii) above. Sturgeon (1986, pp. 74–75) takes issue with Harman on this point.

⁵ Harman (1977) counts himself a cognitivist, but he isn’t a realist in the sense I defined. He thinks that normative moral propositions have truth values, but denies the independence claims I associate with realism.

C screens-off X from E .⁶ For example, the barometer reading is explanatorily irrelevant to the occurrence of the storm because the barometric pressure is relevant, and the barometric pressure screens-off the barometer reading from the storm.⁷ Although HERP seems plausible in light of this and other similar examples, HERP is just as flawed as SERP. Consider a causal chain $X \rightarrow C \rightarrow E$. In many such causal chains, the more proximate cause C screens-off the more distal cause X from the effect E . But surely we don't want to say that distal causes are always explanatorily irrelevant to their effects.⁸ My dialing your telephone number causes your phone to ring, and the phone's ringing causes you to pick up. It would be a mistake to say that my dialing is explanatorily irrelevant to your picking up, even though it is true that the phone's ringing screens-off my dialing from your picking up.⁹

Rather than try to formulate a better principle of explanatory relevance, I suggest that we change the subject. My reason is not just that I think that there are many explanations of a given event and that which are better and which are worse depends on our interests (Sober 1999). In connection with the question of whether a naturalistic defense of moral realism can be made to work, there is the following, additional, motivation: *if an observation O isn't evidence for a hypothesis H , then it hardly matters whether H , if true, would explain O* . So let us jettison inference to the best explanation and take up the following principle about evidence:

Observation O favors hypothesis H_1 over hypothesis H_2 if and only if $\Pr(O \mid H_1) > \Pr(O \mid H_2)$.

This is the principle that Hacking (1965) calls *the law of likelihood*. Notice that it involves probabilities of the form $\Pr(O \mid H)$, not ones of the form $\Pr(H \mid O)$. The question is not which hypothesis has the greater probability of being true, but which confers the greater probability on the observations at hand. The name for this idea, due to Fisher (1922), is misleading; “likelihood” and “probability” are synonyms in English, but not in what is now a standard technical usage.

I won't try to justify the law of likelihood here,¹⁰ but will merely illustrate it via an example due to Salmon (1984). You, an instructor in a philosophy course, have assigned your students the task of writing an essay on a particular topic. Two of the students submit papers that are identical. Not wishing to jump to a hasty conclusion, you consider two hypotheses. The first says that the students plagiarized from the

⁶ This principle is also in the background of Putnam's (1975a) influential argument against reductionism. I discuss this in Sober (1999).

⁷ Salmon (1971) uses “statistical relevance” to characterize explanatory relevance, and cited common cause set-ups like the one about the barometer as evidence favorable for his thesis.

⁸ This example also refutes a principle about causal relevance: if C causes E and if C screens-off X from E , then X is not causally relevant to E .

⁹ Here I ignore the brave new world of caller ID.

¹⁰ Fitelson (2011) has criticized the law of likelihood on Bayesian grounds; I reply in Sober (2011b). There is also a body of frequentist objections to the likelihood principle (which differs subtly from the law of likelihood); I discuss some of these in Sober (2008b).

same source, a file they found on the Internet. The second says that the students worked separately and independently. Although the two hypotheses both allow that what you observe was not impossible, they differ in this respect: the first says that the matching was highly probable (“it was to be expected”), whereas the second says that the matching was very improbable (“it was almost a miracle”). The observed matching of the papers favors the plagiarism hypothesis (which postulates a common cause) over the hypothesis of separate origination; the law of likelihood is the principle that is doing the epistemic work.

What happens to Sturgeon’s two examples when inference to the best explanation is set to one side and the law of likelihood is used instead? The examples about Hitler and slavery get interpreted as follows:

$$\begin{aligned} & \text{Pr}(\text{Hitler started World War Two} \mid \text{Hitler was morally depraved}) > \\ & \text{Pr}(\text{Hitler started World War Two} \mid \text{Hitler was morally sound}).^{11} \end{aligned}$$

$$\begin{aligned} & \text{Pr}(\text{abolitionism became more popular} \mid \text{slavery became morally worse}) > \\ & \text{Pr}(\text{abolitionism became more popular} \mid \text{slavery became morally better}). \end{aligned}$$

I think that both these inequalities are *true*; observations *do* favor some moral hypotheses over others. However, I don’t think that the observations cited are evidence in favor of moral realism. I’ll explain why by concentrating on Hitler.

I am inclined to believe that both of the following likelihood equalities are true:

$$\begin{aligned} & \text{Pr}(\text{Hitler started World War Two} \mid \text{Moral Realism is true}) = \\ & \text{Pr}(\text{Hitler started World War Two} \mid \text{Moral Constructivism is true}). \end{aligned}$$

$$\begin{aligned} & \text{Pr}(\text{Hitler started World War Two} \mid \text{Hitler was morally depraved and Moral} \\ & \text{Realism is true}) = \\ & \text{Pr}(\text{Hitler started World War Two} \mid \text{Hitler was morally depraved and Moral} \\ & \text{Constructivism is true}). \end{aligned}$$

By “moral constructivism,” I mean the thesis that some normative moral propositions are true, and what makes them true is that we would believe them if we deliberated about them in a certain way. Given each of these equalities, the Law of Likelihood concludes that the observation about Hitler fails to provide evidence that favors one hypothesis over another.

When Sturgeon (1984) wrote his article, the main alternative to moral realism that he wanted to consider was moral nihilism, not moral constructivism.

Moral nihilism is the thesis that no normative moral statement is true. What does a likelihood analysis say about realism versus nihilism? First, I think that the following likelihood equality is correct:

$$\begin{aligned} & \text{Pr}(\text{Hitler started World War Two} \mid \text{Moral Realism is true}) = \\ & \text{Pr}(\text{Hitler started World War Two} \mid \text{Moral Nihilism is true}).^{12} \end{aligned}$$

¹¹ I use “morally sound” to represent the opposite of Sturgeon’s “morally depraved.”

¹² I have a worry concerning the three likelihood equalities I just asserted. Maybe there is no saying whether the propositions considered confer the same or different probabilities on the observation cited. That too would undercut the claim that the observation favors one metaethical proposition over another.

Now consider these two probabilities:

Pr(Hitler started World War Two | Hitler was morally depraved and Moral Realism is true)

Pr(Hitler started World War Two | Hitler was morally depraved and Moral Nihilism is true)

Is the first probability larger than the second? It is not, since Hitler's being morally depraved is logically incompatible with moral nihilism. This incompatibility means that the second conditional probability is not defined. The observation cited fails to discriminate between the two conjunctions, not because it judges that they are equally likely, but because one of those "probabilities" doesn't even have a value.

Before I started enumerating all those dispiriting likelihood equalities, I asserted that "Hitler was morally depraved" has a higher likelihood than "Hitler was morally sound," given the observation that Hitler started World War Two. This assertion is incompatible with the thesis that ethical statements lack truth values; if each of these normative ethical statements confers a probability on the observations, then each has a truth value. So my assertion that this inequality is true presupposes that the no truth value thesis is false. Can we say something more—that Hitler's starting World War Two is *evidence* that ethical statements have truth values? Here we must draw back from the precipice. The card's being red favors the hypothesis that the card is the Queen of Hearts over the hypothesis that it is the Queen of Spades, and each of these hypotheses entails that the card is a Queen. However, the observation that the card is red does not favor the hypothesis that the card is a Queen over the hypothesis that it is a Jack. Even if an observation favors one moral statement over another, it does not follow that the observation favors what all of those moral statements entail—that moral statements have truth values.¹³

2 Putnam and Boyd on scientific realism

Just as moral realism needs to be understood in terms of how it contrasts with moral constructivism and nihilism, so scientific realism needs to be understood in terms of how it differs from the alternatives. The debate since the 1960s has not been over whether scientific theories are ever true, or over what makes true theories true. Rather, the dispute has been fundamentally epistemological. Van Fraassen's (1980) formulation is a good place to begin. Realists say that the goal of science is to find

¹³ Philosophers of biology have been interested in "robustness." If each of several competing scientific models entails proposition *R*, then *R* is said to be a "robust theorem." Is robustness evidence of truth? See Sober and Orzack (1993) for discussion.

theories that are true, whereas empiricists say that the goal is to find theories that are empirically adequate (meaning that they are true in what they say about observables).¹⁴ This talk of “goals” does not mean that the debate concerns the psychology of scientists, a question that could be settled by doing an opinion poll (or intensive psychoanalysis) to find out what scientists hope for in their search for theories. Rather, the question concerns what scientific modes of reasoning have the power to deliver. According to Van Fraassen’s empiricism, the evidence that science assembles never obliges one to believe what a theory says about unobservable entities; the most that the evidence could ever require is the belief that the theory is empirically adequate. Realists, on the other hand, think that science is capable of assembling evidence that is so powerful that it compels belief in theories that make claims about unobservables.

Although Van Fraassen’s focus on the attainable goals of science is a good one, I think it is a mistake to construe the dispute between realism and empiricism in terms of the dichotomous concept of rational belief. It is better to think about degrees of confidence (Sober 2008a). I’ll discuss later on how the debate can be emancipated from that restriction.

Just as moral realists have tried to show that there is empirical evidence for thinking that various normative moral propositions are true, so scientific realists have sought to show that the predictive success of a scientific theory is evidence that it is true—both in what it says about observables and in what it says about unobservables. A good place to begin in discussing empirical arguments for scientific realism is Putnam’s (1975b, p. 71) famous remark that “realism is the only philosophy that does not make the success of science a miracle.” This striking statement can hardly be called a fully articulated argument, so let me spell it out a bit. The idea is that when we find that a theory has been successful in making accurate predictions, it is reasonable to demand an explanation for why the theory has been able to do this. Putnam suggests that the only explanation of this remarkable fact is that the theory is true (or approximately true). The most that empiricism can say about the theory is that it is true in what it says about observables. Putnam’s idea is that this is either not an explanation, or it is a poor explanation, of why the theory has made accurate predictions in the tests we have run to date. Explaining why the theory has been empirically accurate to date by saying that the theory is empirically adequate in general is like explaining why this raven is black by saying that all ravens are black; this is hardly a deeply illuminating remark. A similar deficiency can be found in a second alternative to realism—fictionalism. A fictionalist interpretation of a theory says that the theory is false in what it says about unobservables but true in what it says about observables. A fictionalist interpretation of theory *T* entails that *T* has been empirically successful in the predictions we have tested, but that hardly *explains* why the theory has been successful.

¹⁴ It is clear from this formulation that empiricism needs to provide an account of what distinguishes observables from unobservables. I won’t go into that important issue here; see Sober (1990a, 2008a).

Let us now do to this abductive argument for scientific realism what we did earlier to the abductive argument for moral realism; let's substitute the law of likelihood for inference to the best explanation. One reason for changing the subject in this way is something I mentioned before—if the predictive successes of a theory aren't *evidence* that the theory is true, then it hardly matters that the theory, if true, would *explain* why the theory has so far been predictively successful. But there are two additional reasons to embrace this substitution in the present context; opponents of scientific realism often say that the use of inference to the best explanation in arguments for scientific realism is question-begging (Van Fraassen 1980; Fine 1984). And there is Putnam's use of the word "miracle."

Think of the example I discussed earlier concerning the two identical student essays. If the students plagiarized, it is not surprising that the papers matched. But if the students worked separately and independently, it is *almost a miracle* that the papers matched. This use of the *m*-word is, of course, entirely nontheological; it just means that the observed matching would be extremely improbable under the hypothesis of separate and independent origination. The miracle argument for scientific realism wears its connection to the law of likelihood on its sleeve.

Putnam's argument does not work when it is framed in terms of the law of likelihood. The observed predictive success of theory *T* does *not* favor a realist interpretation of the theory over other interpretations:

Pr(*T* has been empirically successful so far | *T* is true) =
 Pr(*T* has been empirically successful so far | *T* is empirically adequate) =
 Pr(*T* has been empirically successful so far | *T* is false in what it says about unobservables but true in what it says about observables).

Realism, empiricism, and fictionalism are on a par in this likelihood comparison. Maybe the hypothesis that *T* is true would explain why *T* has been empirically successful, while the other hypotheses about *T* would not. But the law of likelihood does not care about "explanation." If the law of likelihood correctly characterizes how weight of evidence should be assessed, the alleged difference in explanatory power between realism and empiricism does not matter.^{15,16}

Does this mean that science is forever cut off from evaluating theories that describe unobservables? Not at all. Consider two competing theories, *T*₁ and *T*₂, which each make claims about both observable and unobservable entities. The empiricist wants to evaluate these theories by considering how the evidence at hand bears on the empirical adequacy of each. Within the framework of the law of likelihood, the empiricist's question is whether the observations (*O*) discriminate between empiricist readings of the two theories:

¹⁵ Lipton (1991, p. 111) says that his theory of inference to the best explanation helps ground various Bayesian ideas. My own view is that explanation is not a fundamental epistemic category when it comes to understanding how evidence works.

¹⁶ This likelihood criticism of the miracle argument for realism differs from Magnus and Callendar's (2004) point that the argument commits the base rate fallacy (i.e., it ignores prior probabilities).

(E1) $\Pr(O \mid T_1 \text{ is empirically adequate}) > \Pr(O \mid T_2 \text{ is empirically adequate})$.

The realist wants to ask a different question: Do the observations favor what one theory says about both observables and unobservables over what the other says? That is, the realist wants to know whether

(R1) $\Pr(O \mid T_1 \text{ is true}) > \Pr(O \mid T_2 \text{ is true})$.

The point of interest is that the second inequality follows from the first if the following equality is correct:

(*) $\Pr(O \mid T_i \text{ is true}) = \Pr(O \mid T_i \text{ is empirically adequate})$, for $i = 1, 2$.

Empiricists have frequently emphasized that adding the claim that a theory is true to the claim that it is empirically adequate makes no difference to what predictions one can make. They think that empirical adequacy and truth are related by screening-off:

$$\Pr(O \mid T_i \text{ is empirically adequate}) = \Pr(O \mid T_i \text{ is empirically adequate} \ \& \ T_i \text{ is true}).$$

Since truth entails empirical adequacy, a second screening-off idea must be right:

$$\Pr(O \mid T_i \text{ is true}) = \Pr(O \mid T_i \text{ is empirically adequate} \ \& \ T_i \text{ is true}).$$

These two equalities entail (*). The upshot is this: *if science is in a position to solve empiricist discrimination problems like (E1), it also is in a position to solve realist discrimination problems like (R1)*. The empiricist is mistaken in claiming that the evidence can take us only so far (Sober 1990a, 2008a).¹⁷

I mentioned at the start of this section that Van Fraassen (1980) frames his discussion of realism and empiricism in terms of a dichotomous concept of belief and that it would be a step forward if that were replaced with a graded concept of degrees of belief. The discussion of realist and empiricist discrimination problems bears directly on the question of degrees of belief by way of the odds formulation of Bayes' Theorem. I'll state this odds formulation twice, once for realists and once for empiricists:

$$(R2) \quad \frac{\Pr(H1 \text{ is true} \mid O)}{\Pr(H2 \text{ is true} \mid O)} = \frac{\Pr(O \mid H1 \text{ is true})}{\Pr(O \mid H2 \text{ is true})} \times \frac{\Pr(H1 \text{ is true})}{\Pr(H2 \text{ is true})}.$$

$$(E2) \quad \frac{\Pr(H1 \text{ is emp adequate} \mid O)}{\Pr(H2 \text{ is emp adequate} \mid O)} = \frac{\Pr(O \mid H1 \text{ is emp adequate.})}{\Pr(O \mid H2 \text{ is emp adequate.})} \times \frac{\Pr(H1 \text{ is emp adequate.})}{\Pr(H2 \text{ is emp adequate.})}.$$

In each of these equations, a ratio of posterior probabilities equals a likelihood ratio times a ratio of prior probabilities. The ratio of priors compares an agent's degrees of confidence in the two hypotheses before the observation O is made. The ratio of posterior probabilities compares those degrees of confidence after the observation is obtained. These two ratios differ in value precisely when the likelihood ratio isn't equal to one, and the magnitude of the likelihood ratio determines how large a

¹⁷ Boyd (1983, p. 47) says that the miracle argument "... does not address the crucial epistemological claim of the empiricist argument: that since factual knowledge is grounded in experience, it can extend only to observable phenomena." I think the above point about discrimination problems shows why the empiricist claim is mistaken.

change the observation induces.¹⁸ We have seen that the realist's likelihood ratio and the empiricist's likelihood ratio are *identical in value*. This means that the observation O has the same impact on the realist's ratio of probabilities for the two hypotheses that it has on the empiricist's ratio. Not only are we not entirely cut off from getting observational evidence about hypotheses that are about unobservables; in addition, we have exactly the same quantitative evidential input on the realist's question about truth that we have on the empiricist's question about empirical adequacy.^{19,20}

Boyd (1983) is not happy with the miracle argument for scientific realism, though not for the likelihood reasons I have presented. Before I discuss his dissatisfaction, I want to examine how inference to the best explanation figures in one of his own arguments for realism. Boyd (1983, p. 63) says that "the reliability of theory-dependent judgments about projectibility and degree of confirmation can only be satisfactorily explained on the hypothesis that the background theories which determine those judgments are relevantly approximately true." Boyd's *explanandum* differs from Putnam's; whereas Putnam wants to explain why a theory's predictions about observations have been accurate, Boyd wants to explain the "reliability of theory-dependent judgments" of certain kinds. What Boyd has in mind here is the use that scientists make of a background theory B to guide them in deciding which hypotheses are worth testing (this is what he means by "projectibility judgments") and in estimating the degrees to which those hypotheses, once tested, are confirmed by the observations obtained. Suppose B has been a reliable guide in these two respects. Boyd says that a realist reading of B is the only satisfactory explanation of B 's reliability in those two respects. But what if we shift, once again, from *explanation* to *evidence*, substituting the law of likelihood for inference to the best explanation? Let " B was a reliable background theory" be shorthand for the proposition that B was a reliable guide to which theories were worth testing and also to the degree to which those hypotheses, once tested, were confirmed by the observations obtained. The likelihood comparison we need to consider is the following:

$$\Pr(B \text{ was a reliable background theory} \mid B \text{ is true or approximately true}) > \Pr(B \text{ was a reliable background theory} \mid B \text{ is empirically adequate}).$$

¹⁸ A likelihood ratio of 9 induces more of a change than a ratio of 2, but a ratio of 1/9 induces more of a change than a ratio of 1/2. A ratio of 9 is further from unity than a ratio of 1/9, but they induce the same change. A log-scale is the thing to use here.

¹⁹ I present this Bayesian argument even though I have reservations about Bayesianism (Sober 2008b). My point is addressed to empiricists and realists who think that the two philosophies make recommendations about what we should believe or about what our degrees of belief should be.

²⁰ Van Fraassen (1980, p. 34) notes that it is a consequence of the axioms of probability that $\Pr(H \text{ is true} \mid X)$ cannot exceed $\Pr(H \text{ is empirically adequate} \mid X)$, no matter what X is; this is because truth entails empirical adequacy. Van Fraassen takes this to mean that the empiricist's interpretation of H is "safer" than the realist's. The point is correct, but it hardly shows that we are never rationally obliged to believe that a theory about unobservables is true.

I have three problems with this inequality.

My first two concern the two sorts of reliability that B is observed to exhibit. Let's start with degree of confirmation. What does it mean for B to have been "reliable" in the judgments it delivered about the degree to which different theories were confirmed by observations? I guess this means that $\text{DOC}_B(T, O) \approx \text{DOC}_{\text{true}}(T, O)$ —that using B to assess the degree of confirmation that T receives from O yields close to the same numerical value as using the relevant true background assumption. How can we tell whether this closeness relation obtains? I'm not sure. Maybe the closeness relation will obtain if B is true or approximately true. But the likelihood reformulation of Boyd's argument is supposed to provide *evidence* that B is true or approximately true. This means that the evidence cited must not beg the question; the argument goes wrong if the "evidence" is something we can know only if we already know that B is true or approximately truth. For this reason, I suggest that we delete the part of Boyd's picture of B 's reliability that involves B 's judgments about degree of confirmation.

My second question about the likelihood inequality concerns Boyd's point about projectibility—that B has been a reliable guide to which hypotheses are worth testing. How is this sort of reliability to be understood? I wonder if Boyd is thinking about reliability as follows: When B instructs us to test some small handful of possible hypotheses and we find that one of them is true, B scores a success, but when the hypotheses that B tells us to test all turn out to be false when we test them, B scores a failure. If B gives advice about multiple problems, we can compute how reliable it has been by seeing how often it scores a success. Perhaps this dichotomy between finding that one of the tested hypotheses is true and finding that none of them is should be replaced by something more nuanced. For a Bayesian, testing might result in the discovery that one of the tested hypotheses is probably true, or it might result in none of the tested hypotheses' attaining that status. We then can define B 's successes and failures in those terms.

Here are my qualms about the projectibility part of Boyd's argument: Background theories, by themselves, do not give advice on which hypotheses are worth testing. The most that B , by itself, can tell us is which hypotheses are more probable and which are less. Testing is an action, and rational action requires attention to utilities, not just to probabilities. For example, even if you believe B and B tells you that a given hypothesis is probably false, you still may have good reasons to test that hypothesis. Furthermore, it is a mistake to think that a background theory makes reliable recommendations about which hypotheses to test only when at least one of those hypotheses turns out, after the test, to be true or probably true. A set of hypotheses might be worth testing even if the evidence obtained turns out to be equivocal, and even if the evidence indicates that all of them are false. In both cases, the information obtained might be well worth having. I also think it is a mistake to rely too heavily on one's background beliefs in deciding what to test. To do so is to refuse to test hypotheses that are incompatible with what one already believes. All this leaves me unclear on what it means for an auxiliary assumption to be "reliable" in the recommendations it makes concerning which hypotheses should be tested.

My third objection to the likelihood inequality stated above begins with Boyd's idea that B 's reliability about projectibility (i.e., judgments about which hypotheses were

worth testing) is something we can observe. I take Boyd's idea about this type of reliability to include the thought that B scores a success when it tells you to test H , and H turns out, after the test, to be true or probably true. My question is whether " H is true or probably true" can count as an observation in the context of the problem of comparing a realist and an empiricist reading of B . In particular, can empiricists and realists agree that a test outcome shows that H is true or probably true? Empiricists will get off the bus here if H is even partially about unobservables. So if the likelihood argument is not to be question-begging against empiricism, we must restrict ourselves, in thinking about the hypotheses that B tells us to test, to hypotheses that are strictly about observables. But in that case " B is true" and " B is empirically adequate" make the same recommendations about whether we should test H . So the likelihood inequality (with B 's reliability restricted to Boyd's projectibility idea, the degree of confirmation idea having been set to one side) is false; what is true is an equality.

Let's now consider Boyd's (1983, p. 54) reservations about Putnam's miracle argument. He grants that the argument "has considerable force," but adds that it has a defect: "while [it] provides good reason to think that there must be *something* wrong with the empiricists' argument, it affords us no diagnosis of *what* is wrong with it." A few lines down, Boyd says that the miracle argument provides "a reason to suppose that realism is true, but we are not provided with any epistemology to go with that conclusion." I take Boyd to be saying here, not that the miracle argument is unsound, but that it is incomplete.

Boyd (1983, p. 54) thinks the argument is incomplete because it fails to explain what is wrong with a central empiricist idea, the *evidential equivalence thesis*; this is the claim that observationally equivalent theories are evidentially equivalent. Two observationally equivalent theories can be incompatible; though they agree about all possible observations, they nonetheless can disagree about unobservables. Boyd takes empiricism to be arguing that since no observation can discriminate between empirically equivalent theories that are incompatible, we are forever cut off from making justified claims about which of those theories about unobservables is true.

Boyd argues that the evidential indistinguishability thesis is false because it focuses exclusively on the "direct" testimony of observations and ignores the fact that there are "indirect" sources of evidence. Scientists rightly evaluate a set of empirically equivalent hypotheses by using their background beliefs about theoretical plausibility. Boyd says that these background beliefs are themselves justified by observations; he isn't saying that the judgments about theoretical plausibility that are used to break the tie between empirical equivalent hypotheses are knowable a priori.²¹

As an example (not one that Boyd uses, as far as I know), let's consider Reichenbach's (1928) claim that the following two conjunctions are observationally

²¹ Boyd (1983, p. 67) says that "... considerations of the theoretical plausibility of a proposed theory in the light of the actual (and approximately true) theoretical tradition are evidential considerations: results of such assessments of plausibility constitute evidence for or against proposed theories. Indeed, such considerations are a matter of theory-mediated empirical evidence, since the background theories with respect to which assessments of plausibility are made are themselves empirically tested (again, in a theory-mediated way). Theory-mediated evidence of this sort is no less empirical than more 'direct' experimental evidence."

indistinguishable: “Euclidean geometry & there are universal forces” and “non-Euclidean geometry & there are no universal forces.” For the sake of argument, let’s assume that this is a genuine case of empirical equivalence; any measurement outcome that is predicted by one of these conjunctions is also predicted by the other. Reichenbach, a good empiricist, concluded that the choice between these conjunctions must be a matter of convention. I take it that Boyd would reply by saying that if the theoretical tradition tells us that there are no universal forces, or if it tells us that space is non-Euclidean, then we can use either of these deliverances to cut the Gordian knot. Theories that don’t disagree about any possible observation can differ in their theoretical plausibility and that is enough to show that the empiricist’s evidential indistinguishability thesis is mistaken.

My reply is that I think empiricists will deny that *justified* background beliefs can do the work that Boyd says they can do. It may be a background belief shared by some or all physicists that there are no universal forces, but, according to Reichenbach, that isn’t because they have observational evidence that universal forces do not exist. The “theoretical tradition” needs to be understood in the right way. For Reichenbach, this involves recognizing how conventional elements enter into scientific theories. Carnap (1950a) concurs. My point here is not to defend empiricism but to suggest that Boyd’s refutation of the empiricist’s evidential equivalence thesis does not work.

As noted above, Boyd thinks that the theoretical considerations that permit one to discriminate between empirically equivalent hypotheses are themselves supported by observational evidence. This raises a question. If observation O suffices to justify one’s believing B , and B discriminates between hypotheses T_1 and T_2 , is it also true that O discriminates between T_1 and T_2 ? If O does discriminate, then T_1 and T_2 are not empirically equivalent. This means that Boyd must be focusing on cases that belong to what I will call

Boyd’s World: (i) O rationally justifies believing B ; (ii) B discriminates between T_1 and T_2 ; (iii) O does not discriminate between T_1 and T_2 .

Without a characterization of what rational belief and discrimination are, it is hard to tell whether the three propositions that characterize Boyd’s World are logically consistent and, if they are, what the circumstances are in which all of them are true. We nonetheless can explore Boyd’s World by tracing the consequences of the following two assumptions—that O justifies believing B only if $\Pr(B \mid O)$ is high, and that X discriminates between T_1 and T_2 if and only if X confers different probabilities on the two theories.²² Given these, the following assumptions about probabilities provide a necessary condition for being in Boyd’s world:

²² Notice that I am here using “discriminates” to mark a difference between the probabilities of hypotheses, not a difference between their likelihoods. Once again, I am talking the Bayesian talk, despite my reservations.

(iv) $\Pr(B \mid O)$ is high; (v) $\Pr(T_1 \mid B) \neq \Pr(T_2 \mid B)$; (vi) $\Pr(T_1 \mid O) = \Pr(T_2 \mid O)$.

Conditions (iv), (v), and (vi) cannot be jointly satisfied if T_1 and T_2 are incompatible and $\Pr(T_1 \mid B \ \& \ O)$ is high. To see why, consider the following consequence of the axioms of probability:

$$\Pr(T_1 \mid O) = \Pr(T_1 \mid B \ \& \ O)\Pr(B \mid O) + \Pr(T_1 \mid \text{not}B \ \& \ O)\Pr(\text{not}B \mid O).$$

The right-hand side of this equation has two addends; their sum exceeds $\frac{1}{2}$ if the first addend does so. Note that the value of $\Pr(B \mid O)$ is addressed by (iv); it is “high.” If $\Pr(T_1 \mid B \ \& \ O)$ is also high, we get the result that $\Pr(T_1 \mid O) > \frac{1}{2}$.²³ If the two theories are incompatible, it follows that $\Pr(T_1 \mid O) \neq \Pr(T_2 \mid O)$, thus violating (vi). The bottom line is this: if O makes B very probable and $B \ \& \ O$ makes T_1 very probable, then O cannot fail to discriminate between T_1 and T_2 when the two theories are incompatible.²⁴

What if $\Pr(T_1 \mid O)$ is not greater than $\frac{1}{2}$? In that case it is *possible* for conditions (iv), (v), and (vi) to all be satisfied, but whether they are will depend on the exact values of the relevant probabilities. Condition (vi) is very demanding.

Naturally, this analysis of Boyd’s idea that empirical background knowledge can discriminate between empirically equivalent theories depends on the explications of justified belief and discrimination that I have used. Even so, I suggest that there are lots of cases that don’t fall in Boyd’s World. Furthermore, it is unclear how examples of empirical equivalence like Reichenbach’s can be analyzed in the way that Boyd suggests.

Boyd’s strategy for addressing the problem of empirically equivalent theories via appeal to empirical background knowledge isn’t the only one that realists adopt. They sometimes maintain that there are a priori considerations that do the trick. For example, realists sometimes suggest that simplicity is a “super-empirical virtue” that theories exhibit and that simplicity has more than a pragmatic significance (see, e.g., Churchland 1985). It isn’t just that we find simple theories beautiful and easy to use; the idea is that the greater simplicity of one theory over another is evidence (not proof!) that the first is true and the second is false. It also is part of this realist picture that just as simplicity can be used to discriminate between empirically nonequivalent hypotheses, it also can be used to discriminate between empirically equivalent hypotheses. If Reichenbach’s conjunctions differ in simplicity, that is evidence (of a nonempirical kind) that favors one of them over the other, or so realists sometimes suggest. Reichenbach (1938) replies that there are two kinds of simplicity, logical and inductive, and that it is the first but not the second that discriminates between empirically equivalent theories, while it is the second but not the first that has evidential significance. Reichenbach’s suggestion that there are two kinds of simplicity can be separated from the verificationist idea that empirically equivalent theories are

²³ If $\Pr(T_1 \mid B \ \& \ O)$ and $\Pr(B \mid O)$ are each larger than 0.72, their product will exceed $\frac{1}{2}$.

²⁴ This point holds whether or not B screens-off O from T_1 .

synonymous. Once separated, I think his basic point is correct. Parsimony arguments come in different varieties—some are epistemically pertinent and some are not, and the pertinent ones differ from each other in terms of what justifies them. In the scientific examples that I have examined in which I can discern why simplicity is epistemically significant, I find that simplicity derives its authority from empirical background assumptions of some sort²⁵; this leads me to think that the so-called principle of parsimony is not a priori and that it offers no resources for discerning an epistemic difference between empirically equivalent theories (Sober 1990c, 1996, 2009).

Boyd thinks the realist needs to argue that scientists have evidence that discriminates between observationally equivalent hypotheses; my inclination is to concede to the empiricist that they do not. Science is in the business of using observations to discriminate between theories that are observationally distinguishable; discrimination problems that can't be settled in this way aren't "scientific," though of course it is sometimes unclear whether a problem has that status and the status of a problem can change as knowledge grows. It is part of the project of science to develop theories in the hope that they will be testable against competitors, though of course it may turn out that a theory, once developed, isn't. This may sound like anti-realism, but the other side of the view I am describing has good realist credentials—we are no more cut-off from knowing about unobservables than we are from knowing about observables.²⁶

3 Concluding comments

The accompanying table summarizes a parallelism (and one small asymmetry) that has emerged in the likelihood assessments of moral realism and scientific realism. Observations discriminate between some normative moral statements and others, just as they discriminate between some scientific theories and others, but both of these results presuppose that the statements in question have truth values.²⁷ This result does not mean that those observations, or any observations, will discriminate among moral realism, moral constructivism, and moral nihilism, nor does it mean that observations will discriminate among scientific realism, scientific empiricism, and scientific fictionalism.

²⁵ Here I agree with Boyd's (1991, p. 374) suggestion that "judgments of 'simplicity' are simply special cases of judgments of theoretical plausibility."

²⁶ The question of what to say about empirically equivalent theories has nothing much to do with the distinction between observables and unobservables. Surely "apples exist" is strictly about observables if any statement is, but there are empirically equivalent alternatives to it just as there are to "electrons exist" (Sober 1990a).

²⁷ A Duhemian point applies to these likelihoods; when we talk about $\text{Pr}(\text{observation} \mid \text{theory})$, the probability function typically makes use of various auxiliary assumptions.

Moral Realism	Scientific Realism
$\Pr(\text{Hitler started WW2} \mid \text{Hitler was morally depraved})$ $>$ $\Pr(\text{Hitler started WW2} \mid \text{Hitler was morally sound}).$	$\Pr(\text{Eddington's data} \mid \text{general relativity})$ $>$ $\Pr(\text{Eddington's data} \mid \text{classical physics}).$
$\Pr(\text{Hitler started WW2} \mid \text{moral realism is true})$ $=$ $\Pr(\text{Hitler started WW2} \mid \text{moral constructivism is true})$ $=$ $\Pr(\text{Hitler started WW2} \mid \text{moral nihilism is true}).$	$\Pr(\text{Eddington's data} \mid \text{scientific realism is true})$ $=$ $\Pr(\text{Eddington's data} \mid \text{empiricism is true})$ $=$ $\Pr(\text{Eddington's data} \mid \text{fictionalism is true}).$
$\Pr(\text{Hitler started WW2} \mid \text{Hitler was morally depraved} \ \& \ \text{moral realism is true})$ $=$ $\Pr(\text{Hitler started WW2} \mid \text{Hitler was morally depraved} \ \& \ \text{moral constructivism is true}).$ <p>“Pr(Hitler started WW2 Hitler was morally depraved & moral nihilism is true)” is not defined because the conditioning statement is logically inconsistent.</p>	$\Pr(\text{Eddington's data} \mid \text{general relativity is true})$ $=$ $\Pr(\text{Eddington's data} \mid \text{general relativity is empirically adequate})$ $=$ $\Pr(\text{Eddington's data} \mid \text{general relativity is true in what it says about observables but false in what it says about unobservables}).$

The likelihood equalities described in this table are not a point in favor of anti-realism; rather, the result is a stand-off. However, two of the likelihood comparisons I described earlier (which are not represented in the table) are not so nonpartisan. I think the point about realist and empiricist discrimination problems—the relation of (E1) to (R1) and of (E2) to (R2)—constitutes an important argument in favor of scientific realism.

Sturgeon’s and Boyd’s arguments for realism, as I’ve said, are motivated by a Quinean methodological naturalism about philosophy. Since science uses inference to the best explanation to evaluate scientific theories, philosophers can use that same mode of inference to evaluate philosophical theories. The result, they suggest, is a justification of moral realism and scientific realism. When I substituted the law of likelihood for inference to the best explanation, the results changed. However, I do not interpret that to mean that there is no good argument for moral realism or for scientific realism. The most I can say is that those good arguments, whatever they are, will not come from observational evidence. But even that conclusion may go

too far. Perhaps observations other than the ones I've considered discriminate between the different philosophical theories. For that matter, perhaps the law of likelihood, on which I have leaned so much, needs to be replaced by a better principle about what evidential favoring is, and maybe this replacement will change the conclusions I've drawn about the two realisms.²⁸

In thinking about Boyd's and Sturgeon's naturalism, it is worth remembering a view of philosophical questions that is very different. Carnap (1950a) held that there is a fundamental epistemological difference between internal and external questions. The former are answerable by observational evidence or by mathematical proof; the latter are answerable only pragmatically. We should not seek evidence that there are material objects or that there are numbers, but rather should decide whether adopting these postulates is useful given the larger purposes of inquiry.²⁹

The claim that there are normative ethical truths is entailed by numerous ethical claims about the specifics of right and wrong. Likewise, the claim that there are truths about unobservables is entailed by the specific theories that different sciences have developed. The special consequence condition might lead one to think that evidence for the specifics is evidence for what those specific propositions entail. But evidence does not work like that, as we learn from the Queen of Hearts.

There is one last likelihood wrinkle that I want to mention. It pertains to realist pronouncements that are general ("there are normative ethical truths" and "there are true scientific theories"), not to realist glosses of specific moral claims or specific scientific theories. It is a consequence of the law of likelihood that O favors T_1 over T_2 if and only if $\text{not}O$ would have the opposite evidential significance. This is because

$$\Pr(O \mid T_1) > \Pr(O \mid T_2) \text{ if and only if } \Pr(\text{not}O \mid T_1) < \Pr(\text{not}O \mid T_2).^{30}$$

If Hitler's starting World War Two were evidence favoring moral realism over moral nihilism, then, had Hitler not started World War Two, that would have been evidence favoring nihilism over realism. And if Eddington's data were evidence favoring scientific realism over empiricism, then, had he not obtained those observational results, that would have favored empiricism over realism. Moral realists and scientific realists have never been inclined to think that their philosophies are held hostage by data in this way. This disinclination suggests that realists may need to back away from naturalism; they need to think twice about the idea that the evidence that solves scientific discrimination problems also solves discrimination problems about matters philosophical.

²⁸ I don't think that the law of likelihood is the whole story about the interpretation of evidence; model selection criteria such as the Akaike Information Criterion are also important (Sober 2008b).

²⁹ Putnam (1971, p. 57) criticizes this Carnapian approach to the existence of numbers by siding with Quine, whom Putnam says "has for years stressed both the indispensability of quantification over mathematical entities and the intellectual dishonesty of denying the existence of what one daily presupposes." There is nothing dishonest about admitting that some of one's beliefs cannot be justified by empirical evidence or mathematical proof.

³⁰ Note that this biconditional compares the evidential significance of O and its negation, not the significance of O and one of its contraries.

Acknowledgments I am grateful to Dick Boyd, David Brink, Brandon Conley, Richard Creath, Dan Hausman, Peter Railton, Elena Spitzer, Sharon Street, Nick Sturgeon, Mike Titelbaum, Peter Vranas, and Danielle Wylie for useful comments, and especially to Russ Shafer-Landau for helping me understand some of the ins and outs of moral realism. I gave a talk on this material at the Cornell conference in honor of Boyd and Sturgeon, and then did the same thing at Department Colloquia at MIT, Brown, and University of Wisconsin-Madison, in each case benefitting from the discussion.

References

- Boyd, R. (1980). Scientific realism and naturalistic epistemology. In *PSA: Proceedings of the biennial meeting of the Philosophy of Science Association* (pp. 613–662).
- Boyd, R. (1983). The current status of scientific realism. *Erkenntnis*, 19, 45–90.
- Boyd, R. (1991). Observations, explanatory power, and simplicity—Towards a non-Humean account. In R. Boyd, P. Gasper, & J. Trout (Eds.), *The philosophy of science* (pp. 349–377). Cambridge, MA: MIT Press.
- Carnap, R. (1950a). Empiricism, semantics, and ontology. *Revue Internationale de Philosophie*, 4, 20–40. (Reprinted in *Meaning and necessity*. Chicago: University of Chicago Press, enlarged edition 1956.)
- Carnap, R. (1950b). *Logical foundations of probability*. Chicago: University of Chicago Press.
- Churchland, P. (1985). The ontological status of observables—In praise of the superempirical virtues. In P. Churchland & C. Hooker (Eds.), *Images of science—Essays on realism and empiricism* (pp. 35–47). Chicago, IL: University of Chicago Press.
- Fine, A. (1984). The natural ontological attitude. In J. Leplin (Ed.), *Scientific realism* (pp. 83–107). Berkeley: University of California Press.
- Fisher, R. (1922). On the mathematical foundations of theoretical statistics. *Philosophical Transactions of the Royal Society of London (Series A)*, 222, 309–368.
- Fitelson, B. (2011). Favoring, likelihoodism, and Bayesianism. *Philosophy and Phenomenological Research*, 83, 666–672.
- Hacking, I. (1965). *The logic of statistical inference*. Cambridge, MA: Cambridge University Press.
- Harman, G. (1965). The inference to the best explanation. *Philosophical Review*, 74, 88–95.
- Harman, G. (1977). *The nature of morality*. Oxford: Oxford University Press.
- Harman, G. (1986). Moral explanations of natural facts—Can moral claims be tested against moral reality? *Southern Journal of Philosophy Supplement*, 24, 57–68.
- Hempel, C. (1965). Studies in the logic of confirmation. In C. Hempel (Ed.), *Aspects of scientific explanation and other essays in the philosophy of science*. New York: Free Press.
- Lipton, P. (1991). *Inference to the best explanation*. London: Routledge. (Expanded 2nd edition, 2004.)
- Magnus, P., & Callender, C. (2004). Realist ennuui and the base rate fallacy. *Philosophy of Science*, 71, 320–338.
- Putnam, H. (1971). *Philosophy of logic*. New York: Harper and Row. (Reprinted in *Mathematics, matter, and method*. Cambridge, MA: Cambridge University Press, 1975.)
- Putnam, H. (1975a). Philosophy and our mental life. In H. Putnam (Ed.), *Mathematics, matter, and method: Philosophical papers* (Vol. I, pp. 291–303). Cambridge, MA: Cambridge University Press.
- Putnam, H. (1975b). What is mathematical truth? In H. Putnam (Ed.), *Mathematics, matter, and method: Philosophical papers* (Vol. I, pp. 60–78). Cambridge, MA: Cambridge University Press.
- Quine, W. (1953). Two dogmas of empiricism. In W. Quine (Ed.), *From a logical point of view* (pp. 20–46). Cambridge: Harvard University Press.
- Quine, W. (1960). *Word and object*. Cambridge, MA: MIT Press.
- Quine, W. (1963). Carnap and logical truth. In P. Schilpp (Ed.), *The philosophy of Rudolf Carnap*. LaSalle, IL: Open Court.
- Quine, W. (1981). Success and the limits of mathematization. In W. Quine (Ed.), *Theories and things* (pp. 148–155). Cambridge: Harvard University Press.
- Reichenbach, H. (1928). *Philosophie der raum-zeit-lehre*. English translation by M. Reichenbach as *Philosophy of space and time*. New York: Dover, 1958.
- Reichenbach, H. (1938). *Experience and prediction*. Chicago: University of Chicago Press.
- Salmon, W. (1971). Statistical explanation. In W. Salmon (Ed.), *Statistical explanation and statistical relevance* (pp. 29–88). Pittsburgh: University of Pittsburgh Press.

- Salmon, W. (1984). *Scientific explanation and the causal structure of the world*. Princeton: Princeton University Press.
- Shafer-Landau, R. (2007). Moral and theological realism—The explanatory argument. *Journal of Moral Philosophy*, 4(3), 311–329.
- Sober, E. (1990a). Contrastive empiricism. In W. Savage (Ed.), *Minnesota studies in the philosophy of science—Scientific theories* (Vol. 14, pp. 392–412). Minneapolis: University of Minnesota Press.
- Sober, E. (1990b). *Core questions in philosophy—A text with readings*. New York: Macmillan; 6th edition, Perseus Publishing, 2012.
- Sober, E. (1990c). Let's razor Ockham's razor. In D. Knowles (Ed.), *Explanation and its limits* (pp. 73–94). Cambridge, MA: Cambridge University Press.
- Sober, E. (1993). Mathematics and indispensability. *Philosophical Review*, 102, 35–58.
- Sober, E. (1996). Parsimony and predictive equivalence. *Erkenntnis*, 44, 167–197.
- Sober, E. (1999). The multiple realizability argument against reductionism. *Philosophy of Science*, 66, 542–564.
- Sober, E. (2008a). Empiricism. In S. Psillos & M. Curd (Eds.), *The Routledge companion to philosophy of science* (pp. 129–138). New York: Routledge.
- Sober, E. (2008b). *Evidence and evolution—The logic behind the science*. Cambridge, MA: Cambridge University Press.
- Sober, E. (2009). Parsimony arguments in science and philosophy—A test case for naturalism. *Proceedings and Addresses of the American Philosophical Association*, 83, 117–155.
- Sober, E. (2011a). Evolution without naturalism. In J. Kvanvig (Ed.), *Oxford studies in philosophy of religion* (pp. 187–221). Oxford: Oxford University Press.
- Sober, E. (2011b). Précis of evidence and evolution and responses to commentaries by Branden Fitelson, Roger Sansom, and Sahotra Sarkar. *Philosophy and Phenomenological Research*, 83, 661–665, 692–704.
- Sober, E., & Orzack, S. (1993). A critical examination of Richard Levins' 'The strategy of model building in population biology'. *Quarterly Review of Biology*, 68, 533–546.
- Sturgeon, N. (1984). Moral explanations. In D. Copp & D. Zimmerman (Eds.), *Morality, reason, and truth* (pp. 49–78). Totowa, NJ: Rowman and Allanheld.
- Sturgeon, N. (1986). Harman on moral explanation of natural facts. *Southern Journal of Philosophy Supplement*, 24, 69–78.
- Van Fraassen, B. (1980). *The scientific image*. Oxford: Oxford University Press.