

THE SIMPLEST HYPOTHESIS*

GILBERT HARMAN
Department of Philosophy
Princeton University

Curve Fitting

Suppose you are interested in being able to determine the relationship between two quantities, X and Y. You have reason to believe that Y is a function of X and you have obtained the following repeatable data for the following cases.

$$X = 1, Y = 2$$

$$X = 3, Y = 6$$

$$X = 4, Y = 8$$

$$X = 9, Y = 18.$$

If we let $Y = F(X)$, your data can then be expressed as follows:

$$F(1) = 2$$

$$F(3) = 6$$

* This is a report of work in progress. The preparation of this paper was supported in part by research grants to Princeton University from the James S. McDonnell Foundation and the National Science Foundation (under NSF grant number IST8503968). I am indebted to discussions with Paul Thagard, Michael Ranney and (some years ago) Hugo Margáin.

$$F(4) = 8$$

$$F(9) = 18.$$

Given this data, you would normally take the most reasonable hypothesis to be this:

$$(H1) F(X) = 2X.$$

Why so? It is true that (H1) is compatible with your data as various other hypotheses are not, for example:

$$(H2) F(X) = 3X$$

$$(H3) F(X) = 3 - X.$$

But there are infinitely many hypotheses that are like (H1) in being compatible with your data. For example,

$$(H4) F(X) = 2X + (X - 1)(X - 3)(X - 4)(X - 9).$$

What is it that makes you prefer (H1) to a hypothesis like (H4)?

The natural and immediate answer to this question is that (H1) is simpler and less ad hoc than (H4); and I am not going to dispute that. But this natural and immediate answer suggests two important and related philosophical issues. First, what makes one hypothesis simpler than another? Second, why should we take the simplicity of a hypothesis to be any sort of indication of the truth of the hypothesis?

These are difficult issues. I am going to concentrate on the first and try to say what the simplicity of a hypothesis consists in, but I will have to say something about the second issue as to why simplicity so understood should be used to select among hypotheses that account equally well for the data.

In the example I began with, the first question becomes this: in what way is (H1) simpler than (H4)? One difference between these hypotheses is that (H1)

is expressed in 7 symbols where (H2) uses 28 symbols. So we might consider the proposal that (H1) is simpler than (H4) in that it is shorter and, more generally, that the complexity (or lack of simplicity) of a hypothesis can be measured by the number of symbols used to express that hypothesis (Sober 1975).

However, hypotheses can be expressed in various ways. For any given hypothesis we can arbitrarily introduce notation that would allow us to express the hypothesis with very few symbols. For example, I have been using the symbols "(H1)" and "(H4)" to stand for the hypotheses in question in this particular case, so each can be represented in 4 symbols. I could have used the numerals "1" and "4" instead, in that way representing each hypothesis with one symbol. Clearly, any hypothesis can be expressed using a single symbol in this way. So, if hypotheses differ in simplicity, this difference cannot be measured simply by counting the symbols used to express the hypotheses.

It might be suggested that "(H1)" and "(H4)" are abbreviations of longer expressions and that we must consider how many symbols it takes to express these hypotheses without using abbreviations. But how are we to tell when abbreviations have been used? And how are we to determine what symbols can be used as primitive and undefined? For example, is it permitted to represent 2 times X as "2X" or must we include an operator to represent multiplications as in " $2 \times X$ "? Can we use ordinary mathematical symbols like " \times ", "+", and "-", or are these symbols to be defined in more primitive terms?

One way to answer such questions is to postulate a system of mental representation. In this view, the complexity of a hypothesis for a given person depends on that person's actual mental representation of the hy-

pothesis. (H1) is less complex than (H2) for a given person even though each can be externally represented by a single symbol if that person's mental representation of (H1) uses fewer mental primitives than his or her mental representation of (H4). This is a "psychological" account of simplicity in that the simplicity of a hypothesis is its simplicity for a particular person, given the way in which that person represents the hypothesis in Mentalese.

This leaves the difficult problem of trying to determine what form an inner system of mental representation takes. One test might be to consider how "natural" certain representations are for the person in question. Consider the following two representations:

$$(H1) F(X) = 2 \times X$$

$$(H4) F(X) = 2\$ \times .$$

Here the operator " \times " stands for multiplication and the operator " $\$$ " satisfies the following rule:

$$A\$B = A \times B + (B - 1) \times (B - 3) \times (B - 4) \times (B - 9).$$

The representation of (H1) given seems quite "natural" to me while the representation of (H4) seems quite "unnatural", one that is foreign to my system of mental representation.

But this does not meet the objection already considered, since it does not keep there from being various ways to mentally represent any given hypothesis. Just as I can use a single symbol to stand for any given hypothesis while I am talking to someone else, why can't I do the same thing in the system of mental representation I use for thought? If I can, the complexity of a hypothesis for me is not a function of the number of mental symbols I use to express the hypothesis in my inner language of thought.

Indeed, why can't I introduce a mental symbol for the operation "\$" mentioned above so that I can mentally represent (H4) as $F(X) = 2\$X$? That may not be a "natural" representation, but how does that prevent it from being a possible representation?

Now, I suggest that one thing that makes this particular representation seem "unnatural" is that in order to use it I need to unpack it into some other form. If I am to use this representation of (H4) in order to calculate the value of "2\$7", I need first to translate the representation into the equivalent form $2 \times 7 + (7 - 1)(7 - 3)(7 - 4)(7 - 9)$. I then see that this is equivalent to $14 + 6 \times 4 \times 3 \times (-2)$, which is equal to $14 - 6 \times 4 \times 3 \times 2$, which is equal to $14 - 24 \times 6$, which is equal to $14 - 144$, which is equal to -130 .

This calculation is a bit complicated, certainly more complicated than simply multiplying 2×7 . That is what seems to make (H4) a more complicated hypothesis for me. But notice that the complexity in this calculation has little to do with this particular way of representing (H4). However this hypothesis is represented, a relatively complicated calculation is needed in order to determine the value $F(X)$ for particular values of X . The complexity of this hypothesis depending not on how the hypothesis is represented but rather on the calculations needed to use the hypothesis.

I suggest that we generalize this result. Instead of measuring the complexity of a hypothesis by the number of symbols used to represent the hypothesis in English or in Mentalese, measure complexity by the amount of processing required to use the hypothesis in order to connect it with the data—or more generally, the amount of processing that is needed to use the hypothesis in whatever way the hypothesis is to be used. As we have seen, this would count (H1) as simpler than

(H4) on the grounds that fewer calculations are needed to determine $F(X)$ for various X under (H1) than under (H4).

This is still a "psychological" account of simplicity and complexity, but the psychological simplicity or complexity of a hypothesis is taken to depend on the complexity of certain mental processes rather than depending directly on the complexity of particular mental representations. Changing notation to make a hypothesis shorter will not make that hypothesis simpler, unless the new notation allows a psychologically shorter derivation of the data.

Immediately Obvious Steps

The complexity of the connection between a hypothesis and an implication of that hypothesis might be measured by the number of "steps" needed for a person to recognize that connection. For example, the connection between " 2×7 " and "14" is immediate for someone who has mastered the multiplication table for "2". No further calculation is needed to see the connection between " 2×7 " and "14" for such a person. For most people, there is not an equally immediate connection between " $2 \times 7 + (7 - 1)(7 - 3)(7 - 4)(7 - 9)$ " and "-130". Most people can appreciate this connection only by virtue of noting intermediate steps such as those mentioned above, namely, " $14 + 6 \times 4 \times 3 \times (-2)$ ", " $14 - 6 \times 4 \times 3 \times 2$ ", " $14 - 24 \times 6$ ", and " $14 - 144$ ".

The relevant "steps" are "immediately obvious steps". That is, you can recognize that each step follows from a previous step without having to recognize some intermediate steps. I believe that this notion, of an "immediately obvious step" is quite important for the theory

of reasoning. It is important not only for understanding simplicity but for other reasons as well.

For example, I have suggested (Harman 1987) that the concepts of logic might be at least partially explicated in terms of immediately obvious steps of implication between propositions involving these concepts. Simple logical conjunction "A&B" has the property that a conjunction "A&B" implies and is immediately implied by its conjuncts "A" and "B". This is not true of the more complex but logically equivalent concept "%", explicated as follows: "A%B" = "not either not A or not B". If we suppose that the meaning of a concept depends in part on the immediately obvious implications of propositions involving the concept, then we can understand how logically equivalent concepts can differ in meaning.

It is worth noting that steps that are not immediately obvious at one time can become immediately obvious at some later time. For example, at one time it may not be obvious to you that 4×7 is 28. You may have to note such intermediate steps as these:

$$4 \times 7$$

$$7 + 7 + 7 + 7$$

$$14 + 7 + 7$$

$$21 + 7$$

$$28.$$

At a later time, you can immediately see the equivalence. Your recognition of it does not depend on the recognition of such intermediate steps. According to the view I am defending, when you come to acquire this sort of mathematical skill, hypotheses involving multiplica-

tion become psychologically simpler for you. (And your concept of multiplication undergoes a change.)¹

Goodman's Grue Bleen Problem

I suggest further that the complexity of a hypothesis is relative to what you are interested in. It depends on the complexity of the connections between the hypothesis and things in which you have an interest. In order to illustrate this suggestion, let us look at one aspect of Goodman's famous grue bleen puzzle (Goodman 1965), which is a variant of the curve fitting problem.

We are to suppose that we have examined a variety of emeralds and have determined in each case that the color of the emerald was green at the time of observation. Suppose that it is now January 1, 1990 A.D. and consider the following two hypotheses:

(H5) All emeralds are (always) green

(H6) All emeralds are either observed before the year 2000 A.D. and are (always) green or are not observed before the year 2000 A.D. and are (always) blue.

These hypotheses agree about all emeralds observed before the year 2000 A.D. but disagree about emeralds that do not get observed by then. (H5) implies that such

¹ In supposing that the equivalence between 4×7 and 28 is immediately obvious, representing a "single step", I do not deny that there may be a level of analysis at which the recognition of this equivalence involves several steps. At this deeper level of analysis, the recognition of this equivalence may require more steps and so more time than the recognition of the equivalence between 2×10 and 20, even though both of these equivalences are immediately obvious. I do not at this time know how to give a precise characterization of the relevant level of analysis.

Another complication is that a step of inference might be immediate even though the conclusion of that step is a complex argument involving several steps of implication or explanation. See Harman *et al.* (1987).

emeralds are green. (H6) implies that such emeralds are blue.

Although both hypotheses would account for our evidence, we would prefer (H5) over (H6) on grounds of simplicity. The measure of simplicity I have proposed seems to account for this, since the connection between any item of evidence and (H5) would seem to be immediate whereas the connection between an item of evidence and (H6) would seem to be mediated by several steps. We have

All emeralds are either observed before the year 2000 A.D. and are (always) green or are not observed before the year 2000 A.D. and are (always) blue.

So, this emerald is either observed before the year 2000 A.D. and is (always) green or is not observed before the year 2000 A.D. and is (always) blue.

This emerald is observed before the year 2000 A.D.

So, it is not the case that this emerald is not observed before the year 2000 A.D.

So, it is not the case that this emerald is not observed before the year 2000 A.D. and is (always) blue.

So, this emerald is observed before the year 2000 A.D. and is (always) green.

So, this emerald is green.

But more needs to be said about this case. Following Goodman, let us define a predicate "grue" as follows: x is grue at t if and only if (1) either x is observed before the year 2000 A.D. and is green at t or (2) x is not so observed and is blue at t . Then we can represent this last hypothesis more briefly as follows:

(H6) All emeralds are (always) grue.

Since all emeralds observed so far are observed before the year 2000, all the observed emeralds are grue (at least when observed). And, while it is true that the hypothesis that all emeralds are green offers a simpler account of the fact that observed emeralds have so far been green, it is also true that the hypothesis that all emeralds are grue offers a simpler account of the fact that observed emeralds have so far been grue. Our situation with respect to the two hypotheses may therefore seem perfectly symmetrical and it becomes unclear how we can take (H5) to be any simpler than (H6).

Here there is a temptation to return to counting symbols to measure complexity. Although we can use the term "grue" in order to express (H6) as compactly as (H5) is normally expressed, it might be argued that this compact representation is an abbreviation of our original much more complicated representation of the hypothesis. In this view, "green" expresses a simple concept, whereas "grue" expresses a disjunctive concept.

But this attempt at a solution leads to the difficulties mentioned already. There is no objective way to determine when a representation should be counted an abbreviation. Furthermore, Goodman points out that "green" and "blue" can be seen to express disjunctive concepts via the following definitions (where "bleen" is understood analogously to the way in which "grue" is understood).

x is green at t if and only if (1) either x is observed before the year 2000 A.D. and is grue at t or (2) x is not so observed and is bleen at t .

x is blue at t if and only if (1) either x is observed before the year 2000 A.D. and is bleen at t or (2) x is not so observed and is grue at t .

It might be suggested that "green" is a term that can be applied purely on the basis of observation whereas "grue" is not and that this has something to do with why we prefer (H5) to (H6). However, it is unclear how to make a general distinction between observational terms and other terms. (Is "emerald" an observational term?) Furthermore, it is an accident of this particular example that it involves a term that might count as observational. Many hypotheses in which we are interested do not make use of observational terms in this way and the same problem arises for them as for (H5). That is, we are not primarily concerned with hypotheses of the form, "All A's are B's", where "A" and "B" are observational terms. A distinction between what is directly observable and what is not cannot yield a general solution of Goodman's puzzle.

Goodman himself suggests that the relevant difference between "green" and "grue" lies in our prior use of these terms. In the past, we have projected the predicate "green" from observed to unobserved cases in a way that we have not projected the predicate grue. Our predicate "green" has become entrenched in the history of our past projections as "grue" has not. Goodman argues that we prefer to project such entrenched predicates over unentrenched predicates, other things being equal. In his view, our initial projections are blind. Since there is no past history of projection, no predicates are yet entrenched. We make inferences at this point, even though we are not yet justified in doing so. After we have done this for a while, certain predicates become entrenched and, after that, we can be justified in preferring one hypothesis over another even though they account equally well for the data on grounds of relative entrenchment of predicates.

Goodman's account of this sort of inference or projec-

tion is complicated by other principles and other aspects of his overall philosophical view in a way that makes it hard to assess. There is also a risk of triviality. If the test of a theory of projection is the extent to which it accounts for our past projections of predicates, then Goodman's account may reduce to the claim that we have projected predicates of the sort we have projected.

In any event, I have a somewhat different proposal. I suggest that our interests have a bearing on our inferences. In particular, it is relevant what questions we are interested in answering. Normally, we are interested in what is blue or green, not in what is bleen or grue. (H5) is therefore preferable to (H6), since (H5) gives a simpler route to the sorts of things we are interested in than (H6) does. (H6) provides a simpler route than (H5) does to conclusions about what things are grue or bleen, but we are normally not directly interested in learning such things.

So, I suggest that the simplicity of a hypothesis for a person is determined by the simplicity of the connection between that hypothesis and the data in which that person is interested, as measured by the number of intermediate steps he or she needs to see in order to see the connection.

Since a person will sometimes accept a hypothesis simply because it is the simplest of a group of hypotheses which are otherwise equal in ability to account for the data, this suggestion implies that a person's interests can influence what conclusion the person comes to accept. A natural objection is that this must be a case of irrational wishful thinking. What can we say about that objection?

Wishful Thinking and the Relevance of Interests to Inference

Here it is important to distinguish reasoning that is aimed at what to believe, which (following Aristotle) we can call "theoretical reasoning", from reasoning that is aimed at what to plan to do, which we can call "practical reasoning". Clearly, a person's interests can legitimately help to determine what practical conclusions a person should reach about what to do, so that is an obvious way in which a person's interests are relevant to his or her reasoning. But theoretical reasoning is not practical reasoning and we are now concerned with how a person's interests might affect what theoretical conclusions the person is justified in reaching.

One way in which a person's interests can be relevant to his or her theoretical reasoning is that they help to determine what questions the person has reasons to answer. In that way, a person's interests can legitimately affect which conclusions he or she will draw. At any moment, a vast number of conclusions follow trivially from a given person's beliefs. But the person is not equally justified in drawing each of those conclusions, since at best he or she will only be interested in the truth of a small number of them.

Reasoning is subject to a principle of clutter avoidance. You should not clutter your mind with the trivial consequences of your beliefs, at least if there are certain questions you might be resolving in which you have an interest.

But this is not a general warrant for wishful thinking. Surely, the fact that you want a certain result to be true is not a reason to believe that it is true. Your interests can give you a reason to try to answer a particular question but they are irrelevant to what the answer is (ex-

cept in special cases, for example, in which the question concerns your interests).

This is relevant to my solution to the grue-bleen problem. My solution appeals to a person's interests in order to determine what questions the person wants to answer, not what the answers are. For I have suggested that it is legitimate for someone to accept the simplest account of the data in which he or she is interested, where simplicity is measured by the number of steps needed to get from hypothesis to data. We tend to be interested in whether certain things are blue or green, not in whether they are grue or bleen. We are normally interested in why observed emeralds are green in a way in which we are not so interested in why observed emeralds are grue. It is our interest in answering these questions that leads us to accept (H4) rather than (H5). We do not accept (H4) because you prefer the answers that (H4) gives to other answers. Whether we want emeralds to be green or blue is irrelevant.

But what about the appeal to *simplicity* in choosing (H4) over (H5)? Is that a case of wishful thinking? Again the answer is no.

Simpler hypotheses have pragmatic advantages over more complex hypotheses in that they are easier to use in accounting for data and in making predictions. So, we have a reason to prefer accepting simpler hypotheses to accepting more complex hypotheses. But to accept a simpler hypothesis for this sort of reason is not to engage in wishful thinking.

Wishful thinking involves accepting one hypothesis rather than another because you want the first hypothesis rather than the second to be true. But to prefer the simpler hypothesis for pragmatic reasons is not the same as wanting that hypothesis to be true. A preference for accepting X over accepting Y need not be based

on a prior preference for X's being true over Y's being true.

This still leaves the second of two questions about simplicity mentioned earlier. The first question was "What makes one hypothesis simpler than another". I have made a proposal about that. This leaves the second question, "Why, given that way of measuring simplicity, should you take the simplicity of a hypothesis to be any sort of indication of the truth of the hypothesis?" That is a deep question. I doubt that we have an independent source of information about what is likely to be true over and above the principles of reasoning we actually follow. Since we do (it seems) use simplicity to decide among hypotheses that in other respects are equally satisfactory, when we reflect on particular cases of this sort the simpler hypothesis will seem the most reasonable conclusion, unless we are afflicted with temporary skepticism.

Given a set of hypotheses that all account for the data, we do take the simplicity of a hypothesis as making that hypothesis more likely than others that are less simple. That's what we do, and we have no reason to stop. Perhaps we are justified in continuing to use simplicity in this way in the absence of a serious difficulty with our current practice and the absence of any reasonable alternative.

Similar things can be said about other deep aspects of reasoning, which I will briefly sketch.

Interest, Conservatism, Coherence, and Observation

In Harman (1986) I suggest that the relevant factors in choosing among hypotheses include these: First, as we have just discussed, come your *interests*: it is relevant what questions you are interested in answering:

hypotheses are assessed in terms of how well they answer questions you are interested in answering and a hypothesis will not be chosen if it does not answer such a question. Second, there is *conservatism*: a hypothesis should not involve more of a change in view than is required by the other factors. Third, there is *coherence*: you want to accept the most coherent answer, i.e. the answer that leaves you with the most coherent overall view. Fourth, there is the important role of *observation*: beliefs that are relatively direct consequences of observation tend to be fairly secure.

Coherence has two aspects, negative and positive. Negative coherence is simply the absence of incoherence or inconsistency. Your total view has positive coherence if it coheres or fits together well, minimizing ad hoc assumptions and leaving as little as possible unexplained. Things that promote positive coherence include accepting a general principle that can be used to explain various less general principles, explaining when they work and when they fail; accepting a particular factual claim which, together with principles already accepted, helps to explain several things; accepting instances of general principles already accepted; and accepting things implied by what is already accepted.

Conservatism has a number of aspects too. First, conservatism involves trying to retain as much as possible of your original overall view. Second, it involves minimizing novel additions to that view. Third, it involves avoidance of the sort of fluctuation that would occur if you keep changing your mind, first accepting something, then rejecting it, then accepting it again, and so forth.

The preference for simpler hypotheses over more complex hypotheses that we have been discussing is an in-

stance of conservative bias. You avoid adding complications to your view if you do not have to.

The special weight of observational evidence may reflect the conservatism that rules out unnecessary fluctuation in view. One important feature of observational evidence is that you can always look again. If you give up an observational belief, you will normally be able to get it right back by taking another look. It is pointless to abandon a belief that you can rationally get right back by taking another look or in some other way.

What Justifies the Principle of Reasoning?

In considering how we rely on simplicity, the questions arose whether how the simplicity of a hypothesis could make that hypothesis more likely to be true than other more complex hypotheses that account equally well for the data. If our reason for accepting the simplest hypothesis is that it is easier to use a simpler hypothesis rather than a more complex one, the question arises whether this is not an illicit appeal to wishful thinking.

The same issue arises for the basic principles just discussed. For example, what justifies our conservatism? Given hypotheses that account equally well for the data, how can the fact that we already accept one of them be a reason to think that hypothesis is more likely to be true than any of the others is. If the answer is that it is more convenient to stick with what we have than to switch to something else, the question arises again whether this is not an illicit appeal to wishful thinking.

Again, the fact that one view has more positive coherence than another may make the first view easier to use, but how can that be a legitimate reason to appeal to such a consideration in deciding what to believe?

In reply, I must say, first, that a pragmatic justifica-

tion of this use of simplicity, conservatism, and positive coherence is not a case of wishful thinking. To repeat: there is a difference between accepting X because you have reasons to want X to be true, which is wishful thinking, and accepting X because you have reasons to want to accept X, which need not be an instance of wishful thinking.

Second, consider the skeptical question: "How does such a pragmatic justification show that a simpler or more conservative or more positively coherent hypothesis is more likely to be true?" What does "more likely" mean in this question? If it means "more believable" or "more inferable", then the pragmatic justification is relevant by indicating a practical reason to believe or infer the simpler hypothesis. (For there is no relevant difference in this connection between believing X and believing that X is true.)

Much more might be said about the skeptical issue, but let me leave further discussion of that issue for another occasion and allow me to conclude by mentioning a particular application of my suggested criterion of simplicity.

Parasitical Theories

Let us say that one theory is "parasitic" on a second if the first theory derives its explanations of data by modifying the explanations provided by the second theory. The parasitical explanations include the original explanations and will therefore be more complex than the original explanations. According to my account of simplicity, this added complexity should count against such parasitical explanations.

I am quite familiar with an example from linguistics. It turns out that there is a way to produce phrase

structure grammars that are parasitical on transformational grammars (Harman 1963). But this sort of phrase structure grammar was not taken seriously until phrase structure grammars were able to provide nonparasitical explanations of interesting facts about language (Gazdar *et al.* 1985).

There are also "philosophical" examples, such as the theory that you are a brain in a vat receiving the sorts of perceptual data you would receive if in fact you were wandering around the "external world" in the way you think you are. It is clear that the explanations of data that are offered by this "theory" are parasitic on your ordinary explanations of the data. So, my account of simplicity yields a reason to prefer the ordinary hypothesis over the brain in a vat hypothesis. The calculations are simpler on the ordinary hypothesis.

Finally, consider the "instrumentalist" view of theories which holds that a scientific theory should not be interpreted as claiming literal truth but should be understood to be saying only that it is "as if" the theory were true, that is, the theory is "observationally adequate" (van Fraassen 1980, Churchland and Hooker 1985). Clearly, a theory interpreted in this instrumental way is parasitic on the same theory interpreted realistically. So, by my criterion of simplicity, this sort of instrumentalism is less simple than realism and so should not be accepted!

BIBLIOGRAPHY

- Churchland, Paul and Clifford A. Hooker (1985), *Images of Science* (Chicago, University of Chicago Press).
- Gazdar, G., E. Klein, G. Pulum & I. Sag (1985), *Generalized Phrase Structure Grammar* (Cambridge, Massachusetts, Harvard University Press).

- Goodman, Nelson (1965), *Fact, Fiction, and Forecast* (2nd. edition, Indianapolis, Bobbs-Merrill).
- Harman, Gilbert (1963), "Generative Grammars without Transformation Rules: A Defense of Phrase Structure", *Language*, 39, pp. 597-616.
- (1986), *Change in View: Principles of Reasoning* (Cambridge, Massachusetts, MIT/Bradford Books).
- (1987), "(Nonsolipsistic) Conceptual Role Semantics", *New Directions in Semantics*, Ernest LePore (ed.) (London, Academic Press), pp. 55-81.
- , Marie A. Bienkowski, Ken Salem and Ian Pratt (1987), "Measuring Change and Coherence in Evaluating Potential Change in View", *Ninth Annual Conference of the Cognitive Science Society* (Hillsdale, N. J. Erlbaum), pp. 203-209.
- Sober, Elliot (1975), *Simplicity* (London, Oxford University Press).
- Van Fraassen, Bas (1980), *The Scientific Image* (Oxford, Oxford University Press).