[Published in *Philosophical Quarterly* 59 (2009): 216-36. Official version: https://academic.oup.com/pq/article/59/235/216/1541665.]

When Is Parsimony a Virtue?

by Michael Huemer

ABSTRACT: Parsimony is a virtue of empirical theories. Is it also a virtue of philosophical theories? I review four contemporary accounts of the virtue of parsimony in empirical theorizing and consider how each might apply to two prominent appeals to parsimony in the philosophical literature: those made on behalf of physicalism and on behalf of nominalism. None of the accounts of the virtue of parsimony extends naturally to either of these philosophical cases. This suggests that, in typical philosophical contexts, ontological parsimony has no evidential value.
KEYWORDS: simplicity, parsimony, Occam's Razor, Ockham's Razor, metaphilosophy

1. The Use of Parsimony in Philosophy

According to the Principle of Ontological Parsimony, a theory that entails the existence of fewer entities or kinds of entity is better than one that entails the existence of more entities or kinds of entity, other things being equal.[1] This principle enjoys widespread acceptance in scientific and ordinary empirical reasoning. In philosophy, however, its applications are more controversial.[2] The principle has been used to defend physicalism in the philosophy of mind, which is said to be *prima facie* superior to dualism for its positing of fewer kinds of substances or properties.[3] It has been called upon by nominalists, who commonly assume,

[1]"Parsimony" or "ontological simplicity" consists in limiting the number of (kinds of) things posited by a theory, taking "thing" in the broadest sense. I use "Ockham's Razor" and "the Principle of Parsimony" interchangeably to refer to the principle that parsimony is a theoretical virtue. Pace Lewis (1973, p. 87), the number of tokens of a kind is relevant in addition to the number of kinds posited (see Nolan 1997).

[2]Cornman (1966, pp. 209-10) and Dieterle (2001, pp. 52-4) question the application of Ockham's Razor to philosophy. Hales (1997) disputes Ockham's Razor in general.

[3]Churchland 1988, p. 18; Smart 1959, p. 156 (appealing to the larger number of laws required by dualism); Melnyk 2003, pp. 244-51. Smart (1984) later thought better of his appeal to simplicity.

following Quine, that unless universals are *needed*—for science or for the understanding of everyday knowledge—realism should be rejected by default.[4] It has even been invoked in defense of David Lewis' modal realism, on the ground that the recognition of possibilia reduces the number of fundamental (unanalyzable) kinds we must recognize in metaphysics.[5]

These examples illustrate the prominence of appeals to simplicity in contemporary philosophical methodology. The appeal is usually made, as in the above cases, without discussion of the reasons for favoring simple theories. At best, authors rely on an analogy with science or other empirical reasoning to motivate their preference for simple philosophical theories. Are these appeals well-founded? Is the analogy to empirical reasoning fair?

To address this, we must first have some idea of why we value parsimony in empirical reasoning. Once we have a handle on that, we can consider whether the same rationale applies in typical philosophical contexts. If it does not, we shall have reason to question the use of Ockham's Razor in philosophy.

2. The Value of Parsimony in Empirical Reasoning

I assume that when parsimony is appealed to in philosophy, it is typically taken to be an epistemic, rather than merely a pragmatic virtue.[6] Therefore, I shall focus on epistemic accounts of the value of parsimony, accounts on which the preference for simpler theories can be justified in terms of its tendency to produce true beliefs while avoiding false ones.

2.1. The Empiricist Account

According to the empiricist account of parsimony, our preference for simpler theories is justified by a general empirical argument. Scientists, relying on the

---

[4]Quine 1961, p. 16. See also p. 18, where he characterizes both mathematics and Platonic ontology as myths, and p. 10, where he decries the lack of explanatory power of universals. However, Quine (1980; 1964, p. 243n) states that he always embraced abstracta, particularly classes, finding them necessary for scientific and everyday discourse.

[5]Lewis 1986, p. 4. For example, Lewis argues that once we accept possible worlds and sets, we need not accept propositions as a distinct basic kind, since propositions can be reduced to sets of possible worlds (1983, pp. 53-5). See also Lewis 1973, p. 87, explaining that modal realism does not offend against parsimony since it does not introduce any new *kinds* of thing, the other possible worlds being the same kind of thing as the actual world.

[6]For pragmatic theories, see Quine 1963; Walsh 1979; Harman 1994.

criterion of simplicity, have been highly successful at identifying true theories. This is suggested, perhaps among other things, by the predictive accuracy of scientific theories. The best explanation for the success of science is that its methodology is truth-conducive, in the sense that it tends to lead one to the truth. Because Ockham's Razor is a prominent and essential element in the scientific method, Ockham's Razor is probably correct. If Ockham's Razor were not correct, scientists would probably be much less successful than they are. In short:

1. Science has been highly successful in identifying truths.
2. The best explanation for this is that its methodology is truth-conducive.
3. Therefore, probably scientific methodology is truth-conducive.
4. The appeal to simplicity is a central part of scientific methodology.
5. Therefore, probably simplicity is a genuine mark of truth.[7]

This argument may raise concerns about circularity.[8] First, one might ask how we know premise (1), that science has succeeded in attaining truth. If our knowledge of (1) depends in part on our use of the criterion of simplicity (for example, if we know it because (1) is *the simplest explanation* for the past predictive accuracy of science) then we have a circular justification of the simplicity criterion. Second, one might wonder how the inference from (1) and (2) to (3) is to be justified. If it depends on an appeal simplicity (perhaps because (2) really means something like "the *simplest* explanation for the success of science is that its methodology is truth-conducive")—then again we are left with a circular justification of simplicity.

Perhaps circularity could be avoided by distinguishing between ontological parsimony and other forms of simplicity, using the value of the latter to defend that of the former. Or perhaps the circularity objection could be defused by appeal to an externalist epistemology.[9] I shall not attempt to resolve these questions here. We should take note before moving on, however, that the empiricist argument, whether persuasive or not, makes no effort at explaining *why* simplicity is a virtue. Thus, even if it convinces skeptics of the value of simplicity, we would still have need of the sort of explanatory theories considered in the following three subsections.

---

[7]This argument was suggested to me by Jonathan Weinberg (informal communication). Nolan (1997, pp. 331-2) makes a similar suggestion.

[8]Swinburne 1997, p. 47.

[9]Melnyk 2003, pp. 249-50.

2.2. The Boundary Asymmetry Account

The boundary asymmetry account starts from the observation that there is a lower bound but no upper bound to the degree of complexity a theory can have. That is, for any given phenomenon, there is a simplest theory (allowing ties for simplest) but no most complex theory of the phenomenon—however complex a theory is, it is always possible to devise a more complicated one. This is most easily seen if we take a theory's complexity to be measured by the number of entities that it posits: one cannot posit fewer than zero entities, but for any number *n*, one could posit more than *n* entities. Similar points hold for other measures of complexity, such as the number of parameters in an equation.

The set of possible degrees of complexity, then, is unbounded in one direction. For any set of possibilities that is ordered and unbounded in the upward direction, any normalizable probability distribution over the possibilities must generally assign decreasing probabilities to later members of the set.[10] Thus, suppose that the degree of complexity of a theory can be measured with integers 1, 2, 3, and so on, with 1 being the measure of the simplest possible theory. Let C_1 be the proposition that the true theory has degree of complexity 1, C_2 be the proposition that the true theory has degree of complexity 2, and so on. Since the C_i are mutually exclusive and jointly exhaustive, their probabilities must sum to 1:

$$P(C_1) + P(C_2) + \ldots = 1$$

Given that there are infinitely many terms, this is possible only if $P(C_i)$ decreases as i increases. To take a simple case, we get the desired result if $P(C_i) = 1/2^i$ for all $i \in \{1, 2, \ldots\}$.

Alternately, if we suppose that degree of complexity is a continuous variable rather than discrete, the probability density function over the possible degrees of complexity should integrate to 1 over the interval from 0 to ∞, which requires that the probability density approach 0 as degree of complexity increases without bound.

---

[10]See below in this section for the more precise statement of this claim. Jeffreys (1973, pp. 38-9; 1961, pp. 45-7) advances this account of the value of simplicity. Decreasing probability (density) with increasing complexity is of course necessary but not sufficient for normalizability.

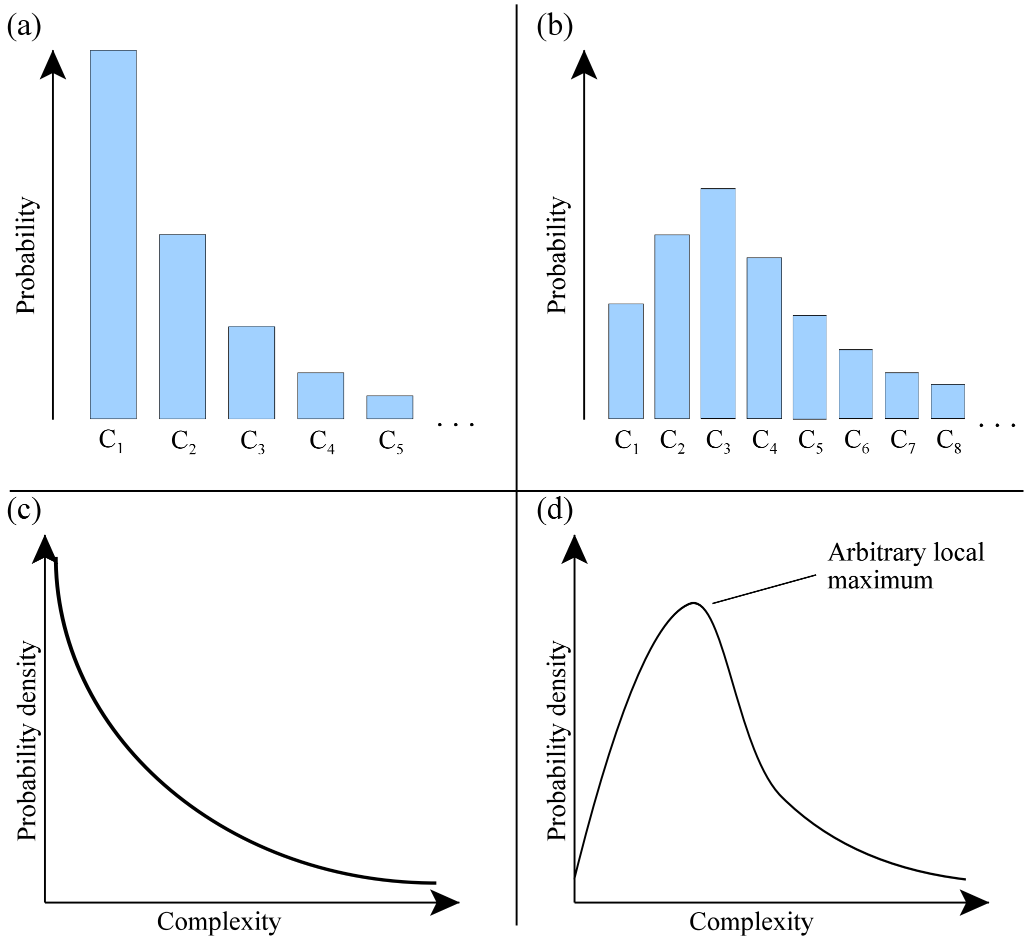

Figure 1. Some normalizable probability distributions over infinite sets of possibilities. (a) Monotonically decreasing discrete probabilities. (b) Non-monotonic discrete probabilities. (c) Monotonically decreasing probability density. (d) Non-monotonic probability density.

This does not show that for any two degrees of complexity, the greater degree of complexity is less likely to be realized. The mathematical argument shows at most that probability (density) must approach zero as a limit, as complexity approaches infinity. Thus, the distributions depicted in figure 1 are all normalizable. Nevertheless, once we see the necessity of letting probability decrease with complexity in the limit, we may feel it more natural to suppose that probability decreases *monotonically* with increases in complexity (figure 1a, 1c). Otherwise, we should find ourselves with one or more seemingly arbitrary local maxima—we should wonder, that is, what was special about certain degrees of complexity such that probability (density) should peak at those points. In figure 1b, for example, we should wonder what is special about C_3.

A monotonically decreasing probability distribution is required if we want to maintain that a simpler theory is *always* more initially probable than a competing, more complex theory. If we are content to maintain only that initial probability approaches zero as degree of complexity approaches infinity, then any coherent probability distribution suffices.

2.3. The Numerousness Account

Suppose one is skeptical of the idea that nature is generally more likely to be simple than to be complex. One could still maintain that simple theories are generally more probable than complex theories, provided that complex theories are *more numerous* than simple ones. To illustrate, suppose that the possible theories of some phenomenon can be divided into two classes, the simple and the complex, and that each class is equally likely to contain the true theory of the given phenomenon. Then simple theories are on average more probable than complex theories, if and only if there are more complex theories than simple ones (figure 2). The point generalizes: if each of n degrees of complexity is equally likely to be realized, then simpler theories are on average more probable than complex theories, if and only if, for different degrees of complexity, on average there are more theories of the higher degree than of the lower.

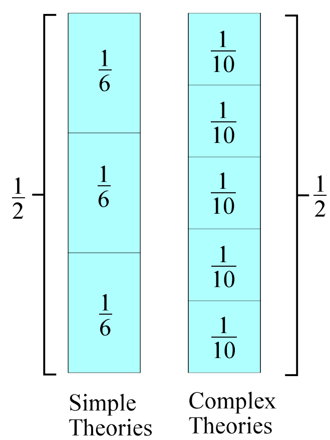

Figure 2. A probability distribution in which each simple theory is more probable than each complex theory, though the truth is no more likely to be simple than to be complex.

There is some reason for thinking that ontologically complex theories are in fact more numerous than ontologically simple theories. The positing of new entities generally allows multiple theories concerning the nature of those entities; consequently, the more entities one posits, the more theories one can construct

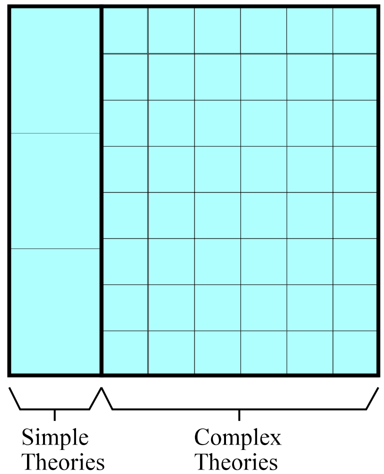

Figure 3. A probability distribution in which each simple theory is more probable than each complex theory, though the truth is more likely to be complex than to be simple.

about those entities.

The same qualitatve result is compatible with weaker assumptions; each degree of complexity need not have the same probability of being realized. We might, for example, consider it more likely that the world is complex than simple, but find that there are so many more complex hypotheses than simple hypotheses that each simple hypothesis is still more probable than each complex hypothesis (figure 3). It is plausible to maintain that not only are there more complex hypotheses than simple ones, but there are *vastly* more complex hypotheses than simple ones. This point shows why—contrary to what is often assumed—those who hold that simple theories are typically more likely to be true than complex ones are not committed to the claim that nature is probably simple.

2.4. The Likelihood Account

The likelihood account, in my own view the most promising account, seeks to show that simpler theories tend to be *better supported* by data that they fit than are more complex theories that fit the data equally well. The essential insight is that typically a simple theory can accommodate fewer possible sets of observations than a complex theory can—the simple theory makes more specific predictions. The realization of its predictions is consequently more impressive than the realization of the relatively weak predictions of a complex theory.[11]

More precisely, the likelihood account directs us to look at *models*, where these are theories or forms of theory that leave the values of some parameters unspecified. One can obtain a *specific* theory by specifying the values of the parameters in a model. For example, suppose we seek to determine the relationship between two variables x and y. Suppose we are considering just two possibilities,

LIN The relationship between x and y is linear, i.e., of the form $y = A + Bx$.

PAR The relationship between x and y is parabolic, i.e., of the form $y = A + Bx + Cx^2$ ($C \neq 0$).

LIN and PAR are models. A specific theory can be obtained by choosing values of A, B, and C. For instance, if one specifies the values of these parameters in PAR as 1, 7, and -9, respectively, then one obtains the specific theory that $y = 1 + 7x - 9x^2$.

If we have just three data points, PAR is guaranteed to fit the data to any desired degree of accuracy: for any three points in a plane, there exists a parabola passing arbitrarily close to them.[12] LIN, however, is not guaranteed to fit the data; there are infinitely many more triples of points that are non-colinear than triples that are colinear.[13] For this reason, if LIN is false, the probability that LIN would

[11] Rosenkrantz 1977, chapter 5; Schwarz 1978. See Jefferys and Berger (1992) for a less technical exposition. Forster (1995) shows that the Bayesian likelihood approach is open to the sort of objections concerning prior probability distributions that apply to Bayesianism in general. Forster and Sober (1994) defend a related approach, based on Akaike rather than Bayes. Popper (1968, pp. 140-42) gives a related account couched entirely in terms of falsifiability. Though I do not discuss the Akaike or Popper approaches separately, my remarks in section 4.4 below may be applied to them as well, insofar as the relevant question is the number of adjustable parameters in each of the philosophical theories discussed in 4.4.

[12] Any three points satisfy some equation of the form $y = A + Bx + Cx^2$. If we require that $C \neq 0$, then (only) triples of points that are colinear fail to satisfy such an equation; nevertheless, equations of the form $y = A + Bx + Cx^2$ ($C \neq 0$) will come arbitrarily *close* to fitting even colinear triples of points. I have stipulated that $C \neq 0$ in my statement of PAR so that LIN and PAR will be incompatible alternatives.

[13] The set of colinear triples of points is a measure-zero subset of the set of all triples of points. To see why, assume that y is a function of x and consider three points, (x_1, y_1), (x_2, y_2), (x_3, y_3). Imagine that the values of x_1, y_1, x_2 ,y_2, and x_3 are all fixed. Then, of the continuum many possible values of y_3, only one makes the three points colinear (namely, $y_3 = [x_3y_2 - x_3y_1 + x_2y_1 + x_1y_2] / [x_2 - x_1]$).

accommodate the data perfectly is 0. Allowing for experimental error, we can allow a nonzero but small probability that LIN would accomodate the data, within the range set by experimental error, if LIN were false. A similar point holds more generally, for larger data sets and equations with more adjustable parameters: the more parameters an equation has, the wider the range of possible collections of points that can be made to fit, to a given degree of accuracy, an equation of that form.

What does this have to do with ontological parsimony? Introducing additional entities into a theory has an effect similar to that of introducing additional adjustable parameters into an equation: suppositions about each of the additional entities can be adjusted to accommodate the data. For example, when Leverrier hypothesized the existence of the planet Neptune to account for observed anomalies in the orbit of Uranus, he had at least two parameters to work with: the mass and orbit of the new planet. The values of these parameters could be adjusted to best accommodate the data about Uranus' orbit. In contrast, had Leverrier hypothesized 83 new planets, he would have had 166 adjustable parameters to work with, enabling the accommodation of a far greater range of possible data. Similarly, if a detective supposes that a crime was committed by a single individual, there is some number of "adjustable parameters" pertaining to that individual: his motivations, beliefs, abilities, whereabouts, and so on. If two individuals are taken to be involved in the crime, twice as many parameters are then available.

Why does the fact that simpler models can typically accommodate a smaller range of data mean that simpler theories are typically *better*? Suppose we have a simple model, S, and a complex model, C, both of which are compatible with some evidence E. Bayes' Theorem tells us:

$$P(S \mid E) = \frac{P(S) \times P(E \mid S)}{P(E)} \qquad P(C \mid E) = \frac{P(C) \times P(E \mid C)}{P(E)}$$

Since we are interested in comparing S with C, we consider the ratio of their probabilities given the evidence:

$$\frac{P(S \mid E)}{P(C \mid E)} = \frac{P(S) \times P(E \mid S)}{P(C) \times P(E \mid C)}$$

S is favored just in case this ratio is greater than 1. Two factors determine this: the ratio of the prior probabilities P(S) and P(C), and the ratio of the "likelihoods," P(E|S) and P(E|C). We have discussed above accounts on which S has the higher prior. The likelihood account argues that S typically has the higher likelihood, P(E|S). Since S is compatible with a smaller range of data, it assigns a higher average probability (or probability density) to those possible sets of data that it allows. C spreads its probability over a larger range of possibilities, consequently assigning a lower probability (density), on average, to the possibilities that it allows (figure 4). For example, if S is compatible with and neutral between possible items of evidence E_1 and E_2, while C is compatible with and neutral among E_1, E_2, E_3, and E_4 (where the E_i are mutually exclusive), then $P(E_1|S) = ½$, whereas $P(E_1|C) = ¼$. S takes a greater risk, since it would be refuted by E_3 or E_4, but if E_1 or E_2 is observed, S is supported twice as strongly as C. Thus, even if a simple model and a complex model have equal prior probabilities, the simple model is usually more probable in the light of data that both models fit.

In several respects, the likelihood account provides at most a qualified defense of the virtue of parsimony: First, it does not suggest that a simple model has any *a priori* advantage over a more complex model. Second, the account only suggests that

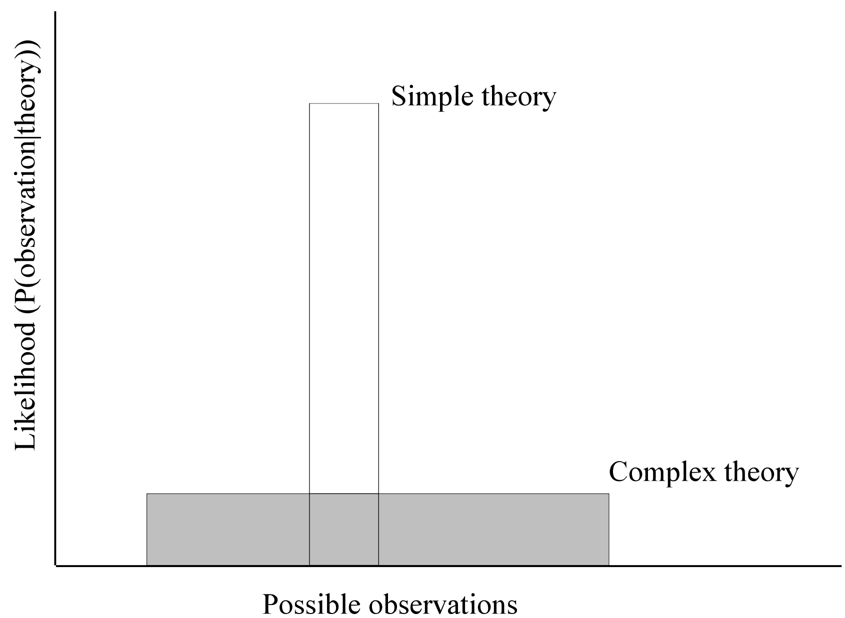

Figure 4. The base of each rectangle represents the range of observations compatible with each theory; the height represents likelihood. The area of each rectangle must be 1. Since S's rectangle has a narrower base than C's, S's rectangle must be taller, indicating a greater likelihood.

parsimony is an epistemic virtue in those cases in which the more parsimonious theory has fewer adjustable parameters. This is typically true, but as we shall see in section 4.4 below, there are cases in which a parsimonious theory does *not* have fewer adjustable parameters than a competing, more complex theory. Third, although it is typically true that a model with more adjustable parameters fits a wider range of possible evidence than a model with fewer parameters, that need not *always* be the case. When it is not the case, the likelihood account again offers no reason to prefer the simpler theory. Finally, even when the simpler of two theories fits a narrower range of data than the more complex theory, the simpler theory need not have a higher likelihood in relation to *every* possible datum that both accommodate; rather, the simpler theory must have a higher *average* likelihood within the range of data that it accommodates than the complex theory has within the range that the complex theory accommodates.

A possible objection to the likelihood account will help us understand the account better. The likelihood account directs our attention to models in which the values of parameters are left unspecified. Why not consider, instead, specific theories in which all parameters are assigned determinate values enabling the theories to fit the data?[14] Such specific theories will generally have likelihoods close to 1. To illustrate, we can convert the models LIN and PAR discussed above into specific theories, such as:

LIN_0 $y = 7 - 2x$
PAR_0 $y = 7 - 2x + x^2$

Each of these hypotheses predicts a specific value of y for each value of x.[15] As a result, if we leave aside experimental error, each hypothesis has a likelihood of either 1 or 0 in relation to any set of data points. Taking account of experimental error, each hypothesis has a likelihood close to 1 for data that it approximately fits. It therefore seems that the likelihood account could not explain why one specific theory has a significant advantage over the other when both theories roughly fit the data.

This is correct; however, when we exchange models for specific theories in this

[14]I thank a referee for this journal for raising this question.

[15]As written, the specific theories would each have prior probabilities of zero, since each is a single point in a space of continuum many possibilities, and thus each would be unconfirmable. To avoid this, we must regard each "specific theory" as really claiming only that the stated parameter values are *approximately* correct.

way, we increase the likelihoods of the resulting theories at the price of lowering their prior probabilities. It is much more probable, initially, that LIN should be true than that LIN_0 should be true, since LIN_0 is just one of the many possible specifications of LIN. Similarly, PAR_0 has a much lower prior probability than PAR. In addition, PAR_0's prior probability is even lower than that of LIN_0, because PAR_0 makes claims about the values of three distinct parameters, whereas LIN_0 only makes claims about the values of two parameters. Other things being equal, it is more likely that two variables take on a specified pair of values than that three variables all take on specified values.[16]

Thus, we may speak either of simpler theories tending to have higher likelihoods, or of their tending to have higher prior probabilities, depending on whether we have in mind models or specific theories. I shall speak in terms of likelihoods and models, because I believe the "theories" we are interested in are usually models, that is, they are individuated in such a way that a range of values for each of their parameters is consistent with the theory. For instance, a revision to the accepted value of the gravitational constant would not generally be taken as a rejection of "Newton's theory of gravity."

3. A Case Study in Parsimony

It is worth looking at a simple, everyday example of parsimony to illustrate some of the ideas of the previous section, as well as to support the contention that the Principle of Parsimony has a probabilistic basis. Imagine that you are sitting at home when your computer and your lamp—the only electrical devices you have on at the moment—shut off simultaneously. Suppose you consider just two accounts of what happened:

H_1: There was a power failure.
H_2: The light bulb burned out and the computer crashed.

---

[16]Again, we must understand the parameter values as approximate, in which case PAR′ picks out a small cube in a three-dimensional space of possibilities, while LIN′ picks out a small square in a two-dimensional space of possibilities. My claim is that the probability measure assigned to the cube should typically be smaller in relation to the 3D space than the measure assigned to the square is in relation to the 2D space. I assume that each parameter in each theory is specified with the same degree of precision, and that each specific theory is located in a region of the possibility space where the probability densities are roughly average (neither especially high nor especially low) for the model that the theory belongs to. Further detail as to the assignment of prior probabilities exceeds the scope of this paper.

H_1 and H_2 can each account for the observed data, but H_1 is the simpler hypothesis, insofar as it postulates a single cause while H_2 postulates two independent causes. H_1 is also, intuitively, the better supported by the evidence. You could of course gather more evidence to settle the matter, but before gathering such evidence, if you had to bet on which hypothesis was correct, you would bet on H_1.[17]

Of the accounts discussed in section 2, the likelihood account handles this case most naturally, revealing one reason why I favor that account. The likelihood account would ask us first to identify the competing models and their parameters. We can think of H_1 as a model in which a single power failure is posited, with the time of the power failure as an adjustable parameter. H_2 is a model in which a burning out of a light bulb and a computer crash are posited, with the times of these events as *two* adjustable parameters. H_1 makes a much sharper prediction than H_2: H_1 can only accommodate experiences in which no time delay is noticed between the shutting off of the lamp and the shutting off of the computer. H_2 can accommodate a wider range of experiences, in which the shutting off of the light and the shutting off of the computer bear any temporal relation to one another. Thus, if E is the observed information as to when the two electrical devices failed, H_1 has a much higher likelihood in relation to E than H_2 does. H_1 is therefore much better supported by E. To state the point more intuitively: if the light bulb were to burn out and the computer crash, they could do so at many different times; H_2 implausibly asks us to consider their simultaneous failure mere coincidence. In contrast, if H_1 were true, the failures would *have* to be simultaneous.

We can test the claim that the preference for simplicity is founded on the likelihood argument by modifying the example in such a way that H_1 no longer enjoys a likelihood advantage over H_2. Thus, consider the following case:

> You go on leave for six months. As you're rushing to the airplane bound for Hawaii, you remember that you left your computer and your lamp on at home. You decide not to worry about it. When you return home six months later, you notice that neither the computer nor the lamp is on.

Once again, you consider two explanatory hypotheses:

H_1: There was a power failure.
H_2: The light bulb burned out and the computer crashed.

The models are the same as in the earlier case; only the evidence has changed. Now

[17] I owe this illustration to Tim Maudlin (in conversation).

the available evidence locates the times at which the two appliance failures occurred only within a six-month interval. Again, both models can account for the evidence. Presumably, H_1 is still the simpler. But this time, the evidence favors H_2 over H_1. What is the difference?

H_1 still has a single adjustable parameter, while H_2 has two adjustable parameters. But, given the relatively short time that power failures usually last, H_1's adjustable parameter, the time at which the current power failure began, would have to be set to within a few hours before the time at which you arrive home. The probability of the evidence on H_1 is therefore low (more simply, background knowledge indicates that the percentage of time during which power failures are in effect at your home is small). In contrast, we can tolerate a wide range of values for H_2's parameters, the times of the two appliance failures: as long as each of these parameters is set to some time within the preceding six months, it would be expected that neither the computer nor the lamp would be on when you arrived home. Therefore, the evidence has a relatively high probability on H_2.

I take it that intuitively, the greater simplicity of H_1 is no longer a factor in deciding which hypothesis to accept. The point is not that H_1's simplicity is outweighed by some other theoretical virtue exhibited by H_2. Rather, talk about simplicity just seems beside the point in this case: we need only consider whether it is more probable that the lamp and the computer should fail during a six month period, or that a power failure should have occurred shortly before you arrived home. Once we notice that the former is more likely, we do not need to talk about simplicity; we do not, for example, need to give an extra boost in credence to the latter hypothesis just because it is simpler. To do so would seem quite arbitrary.

The appliance failure case shows that intuitively, simplicity is sometimes irrelevant, even in comparing two empirical hypotheses. The likelihood account correctly explains such cases.

4. The Failure of Philosophical Appeals to Parsimony

Returning to the question raised in section 1, do philosophical appeals to parsimony have the sort of probative value that appeals to parsimony typically have in other areas? In my view, this is not intuitively obvious. Therefore, a theoretically-based answer is called for. In order to know precisely *when* parsimony is a virtue, we need to know something about *why* parsimony is a virtue. The best way to address our question is thus to look at the most plausible accounts of the virtue of parsimony and to consider whether any of them can be applied to typical philosophical cases. If none of these accounts applies to typical philosophical cases, it is reasonable to presume, until proven otherwise, that appeals to parsimony in philosophy do not have the same probative value as appeals to parsimony in other domains.

In this section, I consider what each of the accounts of parsimony reviewed in section 2 would say about two representative philosophical appeals to parsimony, namely, those made on behalf of physicalism over dualism, and on behalf of nominalism over realism.

4.1. The Empiricist Account

According to the empiricist account of the Principle of Parsimony, the excellent track record of science is evidence that its methodology, including the use of Ockham's Razor, is reliable.

No persuasive case for the use of Ockham's Razor in philosophy can be made along these lines. The kind of evidence we have that science has been highly successful in identifying truths, we do not have in favor of philosophy's efforts to attain the truth. The sort of consensus on substantial bodies of theory that we find in most sciences, particularly the natural sciences, is lacking in philosophy. Nor has philosophy produced the impressive sort of technology or successful predictions that modern science has.

One might argue that once we have established that simplicity is a sign of truth for scientific theories, we can extend the conclusion by induction to other fields of inquiry, such as philosophy. The strength of this inductive inference would depend upon how similar philosophy is to science. Though science and philosophy are both modes of rational inquiry, they seem on their face to have important differences. The apparently greater epistemic success of natural science itself suggests that there may be important epistemological differences between the two. One of these differences seems to be that science relies on experimentation to a greater degree than does philosophy. Philosophical inquiry, on the other hand, relies more heavily on thought experiments, conceptual arguments, and appeals to intuition. Philosophical questions are often more general and fundamental than scientific questions. Philosophical theses are commonly advanced as metaphysically necessary truths, while most scientific hypotheses are contingent. Scientific disputes, unlike philosophical disputes, typically concern either the actual laws of nature or the actual, contingent arrangement and features of concrete particulars in space and time.

Are these differences relevant to the value of simplicity? Given only the empiricist argument, we cannot say. Presumably there is some reason why simplicity is truth-indicative, in the sense that simpler hypotheses tend more often to be true, other things being equal. But the empiricist argument for the value of simplicity does not attempt to identify this reason. In the absence of such identification, the assumption that this reason applies equally to philosophy as it does to science is something of a leap of faith.

I am not at this point arguing that philosophy is disanalogous from science. Perhaps they are analogous; perhaps not. Rather, I am pressing the need for further inquiry into the value of simplicity. We will not know whether the analogy is fair until we have at least some understanding of why simplicity is truth-conducive in science. For this reason, I proceed to the more explanatory views about the value of simplicity.

4.2. The Boundary Asymmetry Account

The boundary asymmetry account argues that there is a lower limit but no upper limit to the degree of complexity a theory can have; therefore, we must assign decreasing prior probabilities to the realization of greater degrees of complexity.

This argument is ineffective in the philosophical contexts we have considered. First, I shall assume, setting aside eliminativism, that the physicalist holds that mental states reduce to or globally, metaphysically supervene on physical states.[18] The dualist holds that mental states are irreducible and do not globally, metaphysically supervene on physical states. The physicalist and the dualist agree both that mental states exist and that physical states exist; they disagree over the metaphysical relationship between the mental and the physical.[19]

On this understanding, the boundary asymmetry argument does not apply to the case of physicalism versus dualism. The boundary asymmetry argument applies only to cases in which the relevant alternatives comprise an infinite set of possible theories that is ordered and unbounded in one direction. The argument is thought to apply to the postulation of entities because there is a lower bound but no upper bound to the number of entities one can postulate. The argument does not apply to disputes about the nature of a given class of entities, or about the relationship between two given kinds of entity, when there are only finitely many contrasting alternatives. In particular, it does not apply to the physicalism/dualism debate, because physicalism and dualism are not naturally viewed as successive steps in some infinite hierarchy of theories. The fundamental question for physicalists and dualists is not, "How many kinds of state there?" The fundamental question is,

---

[18]A more subtle statement of physicalism is that any *minimal* physical duplicate of our world is a duplicate of our world *simpliciter* (Lewis 1983; Jackson 1994). However, the precise formulation of physicalism does not affect the point made in the text, provided that the physicalist accepts that our world contains mental phenomena. The eliminative materialist can escape the argument posed in the text.

[19]Poland 1994, pp. 26-7.

given the existence of mental states and of physical states, what is their relationship to each other? Physicalism and dualism are most naturally viewed as two of the four salient answers to this question.[20]

The same point applies to the realism/nominalism dispute. In this case, the realist posits a kind of entity that the nominalist does not recognize. But the basic issue between nominalists and realists is not, "How many metaphysical categories are there?" The basic issues are more in the neighborhood of "What is the relationship between things and their properties?" or, "Does the fact that things have characteristics or that things resemble each other imply that 'characteristics' *exist*?" Realism and nominalism are most naturally viewed simply as the two possible answers to the latter question.

When we consider Leverrier's introduction of the planet Neptune to explain the anomalies in the orbit of Uranus, we can easily see how Leverrier's theory belongs to an infinite hierarchy of theories, in which each succeeding theory postulates *more of the same*—one could postulate two new planets, three new planets, and so on. But when we consider mind-body dualism or realism about universals, we cannot easily see what would be meant by another theory that postulates more of the same. While one could postulate additional metaphysical categories (for example, sets, possible worlds, propositions), these would be entirely different kinds of things whose postulation was unrelated to the original question about which dualists and physicalists, or realists and nominalists, disputed; nor, in any case, is it easy to envision an infinite hierarchy of such postulates.

4.3. The Numerousness Account

Next, consider the argument that complex theories are typically more numerous than simple theories. This account seeks to motivate lower prior probabilities for more complex theories by first grouping theories according to degree of complexity, then assigning to each equivalence class of equally-complex theories an equal (or at least not wildly different) prior probability.

The difficulty in applying this approach to dualism and physicalism is that it is unclear what classes of theories dualism and physicalism respectively are to be assigned to, such that the members of the first class are more numerous than those of the second class. In a typical case in which some number of unobserved entities are postulated, for example, one might assign a theory to the equivalence class of all

---

[20]The other alternatives are idealism and neutral monism. My arguments here and in the following two subsections vindicate Sober's (1981, p. 146) contention that reduction *per se* does not represent a genuine application of Ockham's Razor, since the reductionist does not deny the existence of anything that the nonreductionist accepts.

theories that posit the same number of unobserved entities (treating the number of postulated unobserved entities as the measure of complexity). When it comes to the mind/body problem, one might try assigning dualism to the class of all theories that recognize exactly *two* distinct kinds of states, and physicalism to the class of all theories that recognize exactly *one* distinct kind of state. However, on this showing, physicalism belongs to the larger class, since the class of monistic theories has at least three members (physicalism, idealism, and neutral monism), while the class of dualistic theories of the mind/body relation has only one member (dualism). I am not, however, arguing that dualism should be regarded as more *a priori* probable than physicalism on these grounds; rather, the task of partitioning the space of theories in this case seems too arbitrary to ground any determinate probability judgements.

It is equally difficult to say what equivalence classes of theories realism and nominalism are to be assigned to. One might try assigning nominalism to the class of "monistic" theories concerning the problem of universals, and realism to the class of dualistic theories. This would fail to deliver the result required by the nominalistic appeal to simplicity, since again there are not more dualistic theories than monistic theories on this issue.

4.4. The Likelihood Account

On the likelihood account, complex models typically have lower likelihoods relative to a given set of data, because complex models have more parameters that can be adjusted to accommodate the data.

To apply this account to the dualism/physicalism debate, we must identify the data to be accommodated and the adjustable parameters that dualists and physicalists respectively have to work with. The adjustable parameters need not be quantities in an equation; they may simply be places where any of a number of specific assumptions could be made by an advocate of the theory, where these assumptions can be adjusted to accommodate the evidence that we have or will have. The most important evidence to accommodate consists in the correlations between mental states and brain events. The dualist's adjustable parameters are assumptions about psychophysical laws—the dualist can adjust assumptions about the laws that underwrite causal connections between mental and physical phenomena, to account for whatever correlations are discovered. Dualism as such has enormous leeway in this regard—any psychophysical laws are consistent with the theory, making dualism virtually guaranteed to be consistent with the data.

This initially seems to favor physicalism, until we notice that physicalism too has an enormous amount of leeway. The physicalist's adjustable parameters are assumptions about psychophysical identities or supervenience relations. The physicalist can adjust assumptions about what sort of physical or functional states mental states *supervene on*, to accommodate the observed psychophysical correlations. Any set of supervenience relations is compatible with physicalism as such, making physicalism, too, virtually guaranteed to be consistent with the data.

One objection to this line of thought is that physicalism is incompatible with at least one conceivable sort of evidence: suppose brain scientists find that certain events in the brain cannot be accounted for in terms of standard physical and chemical causes, but could be accounted for by the hypothesis of nonphysical, mental causes. This is a conceivable empirical development that would seemingly force abandonment of physicalism.

There are two replies to this objection. First, in the imagined scenario, physicalism would not be refuted, for physicalists do not restrict the concept of physical properties to those presently known to physics and chemistry.[21] Hence, in the event that science discovered brain events that could not be causally explained by hitherto recognized physical causes, the physicalist could postulate a previously unknown type of physical cause explaining those events. The point here is not that *any* causal factor discovered by science would automatically count as "physical." The point is that we have no way of ruling out, at any given point in time, the existence of physical causes not yet discovered. Granted, the evidence might render this hypothesis implausible, so that physicalism might be, if not conclusively refuted, at least to some degree disconfirmed. The same, however, is true of dualism: if scientists should discover complete physico-chemical explanations for all processes in the brain and all human behavior, dualism would be rendered implausible, though not conclusively refuted.

Second, even if we overlook the preceding reply, the objection fails because the evidence that would favor physicalism has not in fact been acquired. The essence of the likelihood account is that simpler theories are typically *better supported by data that they accommodate* than complex theories are. This does not confer any advantage, even *pro tanto*, on a simpler theory if data that it accommodates has not actually been gathered—whether because incompatible data has been gathered or simply because the relevant observations have not been carried out. In the present case, the evidence that would favor physicalism has not been gathered, because scientists have not discovered a complete causal explanation of all changes in the brain and all human behavior; that is, it is not yet known whether the observable facts will fall within the range allowed by the theory. The alleged simplicity of physicalism therefore does not as yet provide evidence for the truth of physicalism.

Now consider a further objection designed to show that physicalism enjoys a

[21]Poland 1994, pp. 41, 118-19.

likelihood advantage over dualism. Perhaps physicalism could be refuted by the failure of concerted efforts to identify robust type-type correlations between physical and mental states. By contrast, dualism is refuted neither by the failure to find such correlations nor by the finding of such correlations, for the dualist can always explain any correlations by postulating appropriate psychophysical laws.

Again, there are two replies. First, the imagined evidence would not refute physicalism, for most physicalists already embrace the thesis of multiple realizability; thus, they are not committed to the existence of type-type correlations between mental and physical states. Contemporary physicalism is committed only to the supervenience of mental states on physical states, that is, the thesis that any two individuals in the same total physical state must be in the same mental state.[22] This supervenience thesis could not be experimentally falsified, as it is impossible to verify that two individuals are in *exactly* the same physical state, and it is consistent with supervenience that a very small physical difference should correlate with a large mental difference. Even if we could somehow verify that two individuals had the same known physical properties, it would always be open to the physicalist to hypothesize one or more hitherto undiscovered, theoretical physical properties to account for the psychological difference. Furthermore, if we weaken the physicalist thesis from a strong supervenience to a *global* supervenience thesis, as required by popular externalist theories of mental content, the thesis becomes unfalsifiable for the added reason that we cannot observe two different, complete worlds.

Second, the evidence that the objection would identify as supporting physicalism has in any case not been gathered. Even if we accept that physicalism would be refuted by the failure to find type-type correlations between mental and physical states, this point does not favor physicalism over dualism until such type-type correlations are found. Again, on the likelihood account, simpler hypotheses are more probable, not *a priori*, but in the light of evidence that they accommodate—or more precisely, in the light of the confirmed falsity of propositions that they could not have accommodated but that their more complex rivals could have accommodated.

Simplicity considerations are equally impotent to support nominalism over realism. To apply the likelihood account of simplicity to the nominalism-realism debate, we must identify both the data to be accommodated, and the parameters that nominalists and realists can adjust to accommodate that data.

What are the data to be accommodated? We might view the data as consisting in linguistic and metaphysical intuitions, perhaps together with common

---

[22]Horgan 1982; Lewis 1983, pp. 361-4.

observations. Space considerations prevent the discussion of all intuitions that could be brought to bear on the nominalism/realism debate. There are, however, two broad classes of evidence that seem to provide the most important considerations invoked in that debate.

First, perhaps such facts as that a particular object is red and that two particular objects are similar to one another, are among the facts to be explained by a solution to the problem of universals. In this regard, the realist has sufficient adjustable parameters to accommodate any possible data: whatever characteristics objects are found to have, and whatever objects are found to be similar to each other, the realist can recognize appropriate universals to explain the discoveries. In this case, the realist's adjustable parameters are his assumptions about what universals exist.

What about the nominalist? The nominalist who wishes to appeal to Ockham's Razor faces a dilemma. Either nominalism can accommodate the facts that some things are red and that some things resemble each other, or it cannot. If it can, then nominalism, like realism, is consistent with all possible data of the sort under consideration; it therefore enjoys no advantage in terms of likelihood in relation to the data. I assume here that if nominalism can accommodate some things' being red, then it can accommodate any particulars' having any properties. On the other hand, if nominalism *cannot* accommodate the datum that some things are red or that some things resemble each other, then nominalism is refuted by the data and so receives no confirmation boost due to likelihood considerations. Again, on the likelihood account, simplicity matters only given that a theory accommodates the data.

The second main class of evidence to be accounted for consists in our intuitions regarding such sentences as "Yellow is a color" and "Some colors go together better than others"—that is, sentences that appear to name or quantify over universals. Realism can account for the truth or falsity of any of these sentences. What of nominalism? There are three views we might take as to what nominalism predicts:

a. Nominalism predicts that such sentences should seem false or nonsensical.[23]
b. Nominalism predicts that for any such sentence that seems true, we will be able to find a paraphrase that does not name or quantify over universals, such that the latter sentence will appear an adequate substitute for the former.[24]
c. Nominalism predicts neither that such sentences should seem false or nonsensical, nor that we should be able to find adequate nominalistic paraphrases for such sentences.

---

[23]Quine's (1961) criterion of ontological commitment suggests this.

[24]Quine (1964, pp. 242-3, 245) seems to suggest this.

If we take attitude (a), then we must conclude that nominalism conflicts with the data and thus enjoys no likelihood advantage over realism. If we take attitude (b), then we must again conclude that nominalism conflicts with the data, or at best that nominalism is not supported by the data, since there in fact exist statements that appear to quantify over universals and for which no adequate paraphrases have been devised.[25] At best, the nominalist may speculate that adequate paraphrases of these sentences will be discovered in the future. This would not enable the nominalist to claim to have accommodated the data. Finally, if we take attitude (c), then nominalism cannot be refuted by data of the kind in question. In this case, again, nominalism enjoys no likelihood advantage over realism, since nominalism is consistent with the same range of data as realism. This is not to suggest that attitude (c) is reasonable; rather, the lesson to be learned here is that the *most* the nominalist could claim is that his theory is not refuted by the data and is consistent with the same range of data as realism; the nominalist cannot credibly claim to have discovered that the data fall within a narrow range of values consistent with his theory.

5. Conclusion

The attempt to apply any of four contemporary accounts of the value of parsimony to either of two prominent appeals to parsimony in the philosophical literature yields discouraging results, suggesting that many philosophers' "taste for desert landscapes" is indeed more aesthetic than epistemic in motivation.[26] This inference is non-demonstrative—it is possible that parsimony is valuable for a reason we have not discussed, or that parsimony is valuable in the cases of many other philosophical theories despite its lack of value in the cases of physicalism and nominalism. Those wishing to question my conclusion may therefore do so either by identifying a different account of the value of parsimony that is more favorable to philosophical uses of parsimony than the accounts considered above, or by identifying further examples of parsimony reasoning in philosophy that are vindicated by one of the accounts of parsimony already considered.

If the virtue of parsimony were simply an intuition for which no further explanation could be given, then it might be reasonable for the friend of

---

[25]Armstrong (1978, pp. 60-61) adduces the example, "Red is a color." Quine (1980) concedes such examples as "Some zoological species are cross-fertile" and "Some critics admire nobody but one another," though in at least the latter case he seems to have in mind sets rather than universals in the traditional sense.

[26]The "desert landscapes" remark is from Quine (1961, p. 4).

philosophical parsimony to ask us to take it on trust that the use of Ockham's Razor in philosophy is analogous to its use in science. But this is not reasonable after several leading explanations for the virtue of parsimony have been examined and all found to undermine the analogy. At that point, we are within our rights to demand that the advocate of philosophical parsimony explain why *he* thinks parsimony is a virtue in empirical reasoning, such that this explanation would apply equally to most philosophical cases. Until he does so, it is reasonable to reject the typical appeals to ontological simplicity in philosophy.

References

Armstrong, David. 1978. *Nominalism and Realism*. Cambridge: Cambridge University Press.

Cornman, James. 1966. *Metaphysics, Reference, and Language*. New Haven, Conn.: Yale University Press.

Churchland, Paul. 1988. *Matter and Consciousness*, revised ed. Cambridge, Mass.: MIT Press.

Dieterle, J.M. 2001. "Ockham's Razor, Encounterability, and Ontological Naturalism," *Erkenntnis* 55: 51-72.

Forster, Malcolm. 1995. "Bayes and Bust: Simplicity as a Problem for a Probabilist's Approach to Confirmation," *British Journal for the Philosophy of Science* 46: 399-424.

Forster, Malcolm and Elliot Sober. 1994. "How to Tell When Simpler, More Unified, or Less *Ad Hoc* Theories Will Provide More Accurate Predictions," *British Journal for the Philosophy of Science* 45: 1-35.

Harman, Gilbert. 1994. "Simplicity as a Pragmatic Criterion for Deciding What Hypotheses to Take Seriously," pp. 153-71 in *Grue! The New Riddle of Induction*, ed. Douglas Stalker. Chicago, Ill.: Open Court.

Hales, Steven. 1997. "Ockham's Disposable Razor," pp. 356-60 in *The Role of Pragmatics in Contemporary Philosophy*, ed. Paul Weingartner, Gerhard Schurz, and Georg Dorn. Vienna: Holder-Pichler-Tempsky.

Horgan, Terence. 1982. "Supervenience and Microphysics," *Pacific Philosophical Quarterly* 63: 29-43.

Jackson, Frank. 1994. "Armchair Metaphysics," pp. 23-42 in John Hawthorne and Michaelis Michael (eds), *Philosophy in Mind*. Dordrecht, The Netherlands: Kluwer.

Jefferys, William H. and James O. Berger. 1992. "Ockham's Razor and Bayesian Analysis," *American Scientist* 80: 64-72.

Jeffreys, Harold. 1961. *Theory of Probability*, 3rd ed. Oxford: Clarendon.

Jeffreys, Harold. 1973. *Scientific Inference*, 3rd ed. Cambridge: Cambridge University Press.

Lewis, David. 1973. *Counterfactuals*. Cambridge, Mass.: Harvard University Press.

Lewis, David. 1983. "New Work for a Theory of Universals," *Australasian Journal of Philosophy* 61: 343-377.

Lewis, David. 1986. *On the Plurality of Worlds*. Oxford: Blackwell.

Melnyk, Andrew. 2003. *A Physicalist Manifesto*. Cambridge: Cambridge University Press.

Nolan, Daniel. 1997. "Quantitative Parsimony," *British Journal for the Philosophy of Science* 48: 329-43.

Poland, Jeffrey. 1994. *Physicalism: The Philosophical Foundations*. Oxford: Clarendon.

Popper, Karl. 1968. *The Logic of Scientific Discovery*. London: Hutchinson & Co.

Quine, W.V. 1961. "On What There Is," pp. 1-19 in *From a Logical Point of View*, 2nd ed. New York: Harper & Row.

Quine, W.V. 1963. "On Simple Theories of a Complex World," *Synthese* 15: 103-6.

Quine, W.V. 1964. *Word and Object*. Cambridge, Mass.: MIT Press.

Quine, W.V. 1980. "Soft Impeachment Disowned," *Pacific Philosophical Quarterly* 61 450-1.

Rosenkrantz, Roger. 1977. *Inference, Method and Decision*. Dordrecht, Holland: D. Reidel.

Schwarz, Gideon. 1978. "Estimating the Dimension of a Model," *Annals of Statistics* 6: 461-4.

Smart, J.J.C. 1959. "Sensations and Brain Processes," *Philosophical Review* 68: 141-56.

Smart, J.J.C. 1984. "Ockham's Razor," pp. 118-28 in *Principles of Philosophical Reasoning*, ed. James H. Fetzer. Totowa, N.J.: Rowman & Allanheld.

Sober, Elliot. 1981. "The Principle of Parsimony," *British Journal for the Philosophy of Science* 32: 145-56.

Swinburne, Richard. 1997. *Simplicity As Evidence of Truth*. Milwaukee, Wis.: Marquette University Press.

Walsh, Dorothy. 1979. "Occam's Razor: A Principle of Intellectual Elegance," *American Philosophical Quarterly* 16: 241-4.