Parsimony Arguments in Science and Philosophy—A Test Case for Naturalism$_P$

Elliott Sober

University of Wisconsin, Madison

Romanell Lecture delivered before the One Hundred Fifth Annual Eastern Division Meeting of The American Philosophical Association in Philadelphia, Pennsylvania on Monday, December 29, 2008.

Abstract: Parsimony arguments are advanced in both science and philosophy. How are they related? This question is a test case for Naturalism$_p$, which is the thesis that philosophical theories and scientific theories should be evaluated by the same criteria. In this paper, I describe the justifications that attach to two types of parsimony argument in science. In the first, parsimony is a surrogate for likelihood. In the second, parsimony is relevant to estimating how accurately a model will predict new data when fitted to old. I then consider how these two justifications apply to parsimony arguments in philosophy concerning theism and atheism, the mind/body problem, ethical realism, the question of whether mental properties are causally efficacious, and nominalism versus Platonism about numbers.

For many philosophers, the word "naturalism" immediately conjures up a metaphysical and a methodological thesis. Both concern objects that are "in nature," meaning things that exist in space and time; the contrast is with the *super*natural entities that might be thought to exist "outside" of space and time[1]:

> *Metaphysical Naturalism*: The only things that exist are things in nature.
>
> *Methodological Naturalism$_s$*: Science should not postulate the existence of things that are outside of nature.

I put an "s" subscript on the second naturalism to mark the fact that it gives advice to science. These two naturalisms are the ones that get trotted out in discussions of evolutionary theory versus creationism. Evolutionists often say that their theory obeys the strictures of methodological naturalism but is silent on the metaphysical question. They further contend that creationism rejects both these naturalisms; here they are helped by creationists themselves, who often express their belief in a supernatural deity and argue that methodological naturalism_s is a shackle from which science needs to break free. Although it is worth inquiring further into this interpretation of both evolutionary theory and creationism,[2] that is not my topic here. Rather, I am interested in a third naturalism. Like the second, it is methodological, but it is aimed at the practice of philosophy, not of science (hence the "p" subscript that I use to label it):

> *Methodological* $\textit{Naturalism}_p$: Philosophical theories should be evaluated by the same criteria that are used to evaluate scientific theories.

The popularity among philosophers of this form of naturalism owes a lot to Quine's influence. When Quine (1953, 1960, 1963) maintained that philosophy is "continuous" with science, he meant that philosophers should address general questions of ontology in the same way that scientists address more specific questions about what there is.[3]

There are some similarities that link science and philosophy that lend a superficial plausibility to naturalism_p. For example, scientists and philosophers care, or ought to care, about logical consistency. However, there is another context in which this naturalistic thesis is far from obvious. Scientific theories are often evaluated for how parsimonious they are and the same goes for philosophical theories. Taken at face value, this similarity seems to be grist for the naturalist's mill. But is the justification for using parsimony in philosophy really the same as the justification for using parsimony in science? The use of the same word in these two contexts should not lead us to assume that they are.

One way to be a naturalist_p about parsimony is to embrace a kind of nihilism. Perhaps the use of parsimony to adjudicate between scientific theories has no justification whatever and the same is true of its application to philosophical theories. Parsimony arguments in the two fields are on the same footing, namely, none. In this paper I'll describe two kinds of parsimony inference that occur in science for which I think this nihilistic assessment is mistaken. Parsimony isn't always an optional aesthetic frill; there are important types of scientific argument in which parsimony has a demonstrable epistemic relevance. Given this, the question for naturalism_p is how the justification that parsimony has in some scientific contexts bears on the parsimony arguments that philosophers produce.[4]

Parsimony arguments have been advanced in connection with the following philosophical problems:

- Atheism versus theism
- The mind-body identity theory versus dualism
- Epiphenomenalism about the mental versus the claim that mentalistic properties are causally efficacious
- Moral realism versus anti-realism
- Nominalism versus Platonism

Here's an example that concerns ethical realism:

> To explain facts about human thought and behavior, there is no

> need to postulate the existence of ethical facts (evolution and upbringing suffice).
>
> If there is no need to postulate the existence of ethical facts to explain human thought and behavior, then there is good reason to deny that those facts exist.
>
> ---
>
> There is good reason to deny that ethical facts exist.

By "ethical facts," I mean normative truths about actions—ones that say that they are just or unjust, right or wrong, permissible or impermissible, etc. This parsimony argument, which resembles others that have been deployed in the context of other philosophical problems, has three noteworthy features: (i) the argument concerns explanation, (ii) it invokes a dichotomy between what is needed and what is not, and (iii) it concludes that there is good reason to deny that ethical facts exist, not that we should remain silent about whether such entities exist. We will reflect on these three features in due course.

I am well aware that my survey of this territory will be incomplete in at least two respects. First, I do not pretend that my catalog of cases in which scientific parsimony arguments make sense is exhaustive. Maybe there is more to scientific parsimony than is dreamt of in my philosophy. Second, even if my scientific catalog were exhaustive, that would not settle the question of whether naturalism$_p$ is correct in what it says about parsimony. For if there are parsimony arguments in philosophy that do not measure up to the justified applications of parsimony that are found in science, what should we conclude? The naturalist will say that these philosophical arguments should be consigned to the flames. The anti-naturalist will demur, suggesting that there is more to philosophical parsimony than is dreamt of in science. I will not try to decide who is right here. Rather, my goal is to make a start at understanding what makes parsimony arguments tick. We need to figure out when parsimony arguments should impress us and when they should not. If it turns out that there are irreducibly different types of justified parsimony arguments, and if some of them are not justified at all, so be it.[5]

1. Parsimony as a Surrogate for Likelihood

One of the contexts in which scientific parsimony arguments are justified is illustrated by the Darwinian theory of evolution. This theory has two parts. It says that natural selection has been an important cause of the traits we see in present day organisms *and* that the species we find today trace back to common ancestors. These two elements of the theory are logically independent (Mayr 1982). With respect to common ancestry, Darwin says in the last paragraph of the *Origin* that all of the organisms we see around us trace back to "one or a few" original progenitors, though a few pages earlier he argues, tentatively, that there was one. What evidence is there for this part of the theory? Darwin, like modern evolutionists, sees the answer in *similarity*. It is the similarities that organisms bear to each other that provide evidence of their common ancestry.

To understand the logic of this inference from similarity to common ancestry, consider the problem depicted in Figure 1. We observe that human beings and monkeys have tail bones. Darwin regarded this as evidence that human beings and monkeys trace back to a common ancestor. If the common ancestry (CA) hypothesis were true, it would not be terribly surprising that this trait is found in the two groups; on the other hand, if the two groups had originated separately and independently (SA), it would be a rather surprising coincidence that both have tail bones. This simple argument makes use of a principle that now is called the Law of Likelihood (so named by Hacking 1965):

> *The Law of Likelihood*: Observation O favors H_1 over H_2 if and only if $Pr(O|H_1) > Pr(O|H_2)$.

Here I use "likelihood" in the technical sense that is now standard in statistics, thanks to R.A. Fisher's coinage. Fisher's terminology was unfortunate, since in ordinary English "likelihood" and "probability" are synonyms. In their technical usage, they are not. The probability that a hypothesis H has in the light of observations O is represented by "$Pr(H|O)$" (read this as "the probability of H, given O"). The likelihood that H has, given O, is the probability that H confers on O; it is represented by "$Pr(O|H)$." Notice that the *probabilities* of the two hypotheses play no role in the Law of Likelihood.

Even if we accept the Law of Likelihood, something more needs to be said about its bearing on the question of common ancestry. Why, in the example at hand, should one believe that

> Pr(humans and chimps have tail bones|CA) > Pr(humans and chimps have tail bones|SA)?

This likelihood inequality is not a consequence of the axioms of probability. Rather, it requires substantive assumptions about the evolutionary processes at work in lineages. Here I find inspiration in Reichenbach's (1956) *principle of the common cause*. Although I do not think that Reichenbach's

Figure 1. How does the observation that humans and monkeys both have tail bones bear on the Common Ancestry (CA) and the Separate Ancestry (SA) hypotheses?

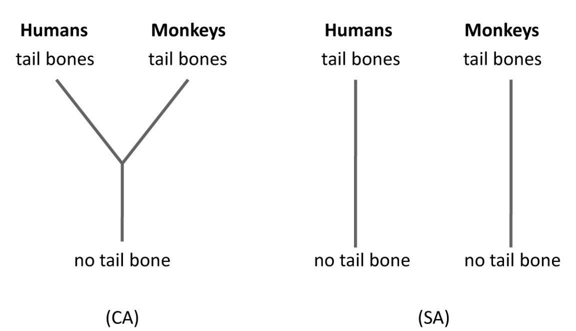

principle is correct, I do think that he described a set of assumptions about common cause hypotheses that can be supplemented by assumptions about how separate cause hypotheses should be understood; these assumptions provide a sufficient condition (described in the Appendix) for the above likelihood inequality.[6] It is this Reichenbachian picture that I think underlies one's intuitive judgment that the tail bone similarity would be more probable if humans and chimps had a common ancestor than it would be if they did not.

Parsimony never got mentioned in the likelihood argument just sketched. Yet, many evolutionists have thought that parsimony is a guide to reasoning about phylogenetic relationships. They would claim that the CA hypothesis provides a more parsimonious explanation of the tail bone similarity than the SA hypothesis does. Consider Figure 1 once more. If we assume that the remote ancestors of human beings and of monkeys lacked tail bones, then the CA hypothesis requires that there was at least one change, from no tail to tail, in the lineages leading to the present. The SA hypothesis, on the other hand, requires that at least two such changes occurred. CA is more parsimonious than SA because 1 is a smaller number than 2. Here we measure parsimony by counting how many changes in character state a genealogy requires to yield the characteristics found at its tips. If the difference between 2 and 1 does not strike you as compelling, think of the 260-some monkey species that are known and consider why they all have tail bones. Isn't the hypothesis of 260 separate originations of this trait far less parsimonious than the hypothesis that the trait originated just once?

Biological questions about common and separate *ancestry* are special cases of a more general problem—the problem of comparing hypotheses about common and separate *causes*. If two students in a philosophy class hand in essays that are word-for-word identical, the professor may consider whether this matching is due to the students' copying from a common source (maybe a file on the Internet) or was due to the students working separately and independently (Salmon 1984). The matching is more probable under the plagiarism hypothesis than it is under the hypothesis of separate origination. And the plagiarism hypothesis provides the more parsimonious explanation of the matching.[7]

If the common and separate ancestry hypotheses can be evaluated in terms of their likelihood and also in terms of their parsimony, how are these two epistemologies connected? The idea that parsimony is justified because of its link with likelihood was proposed by Edwards and Cavalli-Sforza (1964), two students of R.A. Fisher. They proposed a *principle of minimum evolution*. Edwards and Cavalli-Sforza formulated this principle because of a technical problem they encountered when they tried to evaluate the likelihoods of genealogical hypotheses that are more complicated than the toy example depicted in Figure 1.[8] They proposed parsimony as an expedient solution, one that they believed would usually order genealogical hypotheses in the same way that a likelihood evaluation would do. For them, likelihood was the foundation and parsimony was justified in terms of it. Their conjecture, advanced in the context of phylogenetic inference, was that parsimony is a surrogate for likelihood.

A lot has been learned about phylogenetic inference since that early paper. We now know that likelihood and parsimony do not always agree on which hypotheses are better and which are worse. With the wisdom of hindsight, we now can describe some simple examples in which they disagree. Cases of disagreement are the ones to think about with respect to the question of how parsimony and likelihood are related. The thought that parsimony's justification depends on likelihood has its mirror image—that likelihood has authority only to the extent that it reflects parsimony. This is not an abstract possibility, but is the position taken by many cladists.[9]

Here's a simple inference problem. Consider a group of fathers and their children, each father having exactly two. Each father has ten pennies and will give some or all of these pennies to his two children. Jack and Jill are two children in the group. Jack received four pennies from his father and Jill received four pennies from hers; there is a matching in the number of pennies received. The CA hypothesis says that Jack and Jill are siblings. The SA hypothesis says they are not. Does the matching favor CA over SA in the sense of the Law of Likelihood? That depends on the rules that fathers follow in distributing pennies to their children. Here are two possibilities:

- Each father chooses a number between 1 and 5 (with equal probability) and gives each of his two children that number of pennies.
- Each father chooses between the numbers 4 and 5 (with equal probability) and gives one of his children that many pennies and the other a lesser number (choosing among those options with equal probability).

Under the first rule, the fact that Jack and Jill each have four pennies favors the hypothesis that they are sibs; under the second, the matching has the opposite evidential significance. Perhaps it is more parsimonious to think that Jack and Jill received their coins from a common paternal source, but that does not matter. Parsimony isn't central, likelihood is, and which hypothesis is more likely depends on the rules of inheritance (monetary, not genetic) that are in place.

A second example illustrates an opposite parting of the ways between parsimony and likelihood; it is one in which CA is not more parsimonious than SA, though CA still has the higher likelihood. Notice in Figure 1 that CA is more parsimonious than SA only because the roots of the lineages lack tail bones. Let us imagine instead that these remote ancestors had tailbones. The effect of this supposition is to render CA and SA equally parsimonious; neither hypothesis requires a change in character state between root and tips to explain the observations. This change in assumption affects parsimony, but it need not affect the likelihood ordering. The Reichenbachian framework that entails that CA has the higher likelihood if the roots of the genealogies *lack* tail bones has the same consequence if the roots *have* tail bones (see the Appendix). The likelihood ordering does not depend on what the character state of the roots is; the parsimony ordering does.

Philosophical parsimony arguments often concern the existence of some entity—for example, God, ethical facts, and universals—where the more parsimonious hypothesis usually is taken to be the one that denies that the putative entity exists. How does this focus on existence claims apply to the example about tail bones, where we assume that remote ancestors lacked tail bones? The hypothesis of common ancestry entails that a common ancestor exists, while the hypothesis of separate ancestry denies this. If the hypothesis of common ancestry is the one that is more parsimonious, then we have here a case in which the more parsimonious hypothesis is the one that says that something exists while the less parsimonious hypothesis denies this. On the other hand, we might shift our attention from common ancestors to changes in character state. The CA hypothesis says that at least one change (from no tail bone to tail bone) took place while the SA hypothesis says that at least two did. If we count changes rather than common ancestors, now it is the CA hypothesis that postulates fewer entities. So what is the right thing to count—common ancestors or changes? The meaning of the word "parsimony" will not help us here. Parsimony has to do with *less* as opposed to *more*, but the meaning of the term does not reveal how one should answer the question *less of what?* What is needed is a framework that shows why there is an epistemic advantage that hypotheses have in virtue of exhibiting one or another sort of lessness. The Law of Likelihood, coupled with a Reichenbachian understanding of the common cause and separate cause hypotheses, is one such framework. The question of how parsimony should be measured is inseparable from the question of how parsimony inferences should be justified.

2. Likelihood and Model Selection

When it comes to explaining the similarity between two objects, the hypothesis of common cause is often intuitively more parsimonious than the hypothesis of separate causes, and, given Reichenbachian assumptions, the common cause hypothesis has the higher likelihood. Regardless of whether justification flows from likelihood to parsimony, or from parsimony to likelihood, likelihood and parsimony in this case go hand-in-hand. I now want to discuss a kind of parsimony argument that is important in science in which the very opposite is true.

I do so by way of a simple example. Suppose you drive south from Madison and stop at two adjacent fields of corn plants; let m_1 be the mean height in the first field and m_2 the mean height in the second. You sample a number of plants from each field and compute the average height in your two samples; the two sample means are s_1 and s_2. Suppose the difference between these sample means is 4 centimeters. You want to use these data to evaluate two models:

(NULL) $m_1 - m_2 = 0$.

(DIFF) There exists a number d such that $m_1 - m_2 = d$.

I call the first model NULL because it says that there is no difference in the two mean heights. My label for DIFF is a bit of a misnomer, however, since this model does not *require* that the mean heights differ, though it *allows* that they might. DIFF contains a single adjustable parameter (d) whose value might be estimated from the data. The maximum likelihood estimate of d is just the observed difference in the sample means $(s_1 - s_2) = 4$ centimeters.[10] Notice that DIFF has a flexibility that NULL does not possess. The observed difference in sample means might comport very poorly with what NULL says; in contrast, any observed difference in the sample means can be accommodated by DIFF.

DIFF may strike you as a tautology. Well, it is a *near* tautology. The nontautological content comes from DIFF's entailing the existence of the two fields of corn; the tautology is that *if* the two fields exist, there surely is a number that represents the difference in their mean heights. Although DIFF is a near tautology, it can be used to make predictions and these predictions may differ from the ones that issue from NULL. DIFF makes predictions in the following sense. If we replace the adjustable parameter in this model with that parameter's maximum likelihood estimate, the result is a fitted model, which I'll call L(DIFF). L(DIFF) makes a prediction about what new data drawn from the two populations will be like. NULL also makes a prediction about new data, but it doesn't need to be fitted to the old data to do this.

How should we evaluate these two models? In terms of fit-to-data, NULL can't do better than DIFF. The two models will tie if $(s_1 - s_2) = 0$; otherwise, DIFF will fit the data better. And fit-to-data reflects likelihoods. For example, if $(s_1 - s_2) = 4$ centimeters, L(DIFF) will make this observation more probable than the NULL hypothesis will. Another comparison we might make of the two models concerns their probabilities of being true. Since NULL entails DIFF, NULL can't be more probable, no matter what the data are. So in terms of likelihood and in terms of probability, NULL can't be better than DIFF. But what about the model's ability to accurately predict new data when fitted to old? Let us call this property of the model its *predictive accuracy* (Forster and Sober 1994). Predictive accuracy isn't the same as fit-to-old-data, nor is it the same as the model's probability of being true. The Japanese statistician H. Akaike (1973) proved a remarkable theorem about predictive accuracy:

> *Akaike's theorem*: an unbiased estimate of model *M*'s predictive accuracy is $log\{Pr[data|L(M)]\} - k$.

Notice that there are two terms here—the log of the likelihood of *L(M)* and k, which is the number of adjustable parameters that the model contains. Parsimony enters here, since simpler models have fewer adjustable parameters than models that are more complex. Akaike's theorem led to the formulation of the Akaike Information Criterion, which is a method for scoring models. A model's AIC score is the quantity $log\{Pr[data|L(M)]\}$ – k. Since log-likelihoods go up as likelihoods go up, a model's score is improved by its fitting the data better. But it also helps the model if it has a low value for k. In our comparison of NULL and DIFF, DIFF does better in terms of the first consideration, but worse in terms of the second. So which model has the higher AIC score? Since NULL has one fewer adjustable parameter than DIFF, the answer is

AIC(DIFF) > AIC(NULL) iff $log\{Pr[data|L(DIFF)]\} - log\{Pr[data|L(NULL)]\} > 1$.

It is logically inevitable that the difference in the log-likelihoods is greater than or equal to zero. It is not inevitable that the difference is greater than 1; that depends on the data. For DIFF to have the better AIC score, it must fit the data *sufficiently* better than NULL does to overcome the fact that DIFF is more complex. The point of relevance here is that the part of statistics called model selection theory provides a context in which the justification of parsimony is pretty clear. Parsimony is relevant to estimating how predictively accurate a model will be.[11]

Although the difference in parsimony between NULL and DIFF is modest, it is easy enough to tweek the example to make the difference more substantial. Instead of there being one pair of corn plant populations, let there be ten pairs. Consider the model DIFF-10, which has ten adjustable parameters $d_1, d_2, \ldots, d_{10}$, one for each pair, while the model NULL-10 has no adjustable parameters. AIC applies to these two models as follows:

AIC(DIFF-10) > AIC(NULL-10) iff $log\{Pr[data|L(DIFF\text{-}10)]\} - log\{Pr[data|L(NULL\text{-}10)]\} > 10$.

DIFF-10 must fit the data *far* better than NULL-10 does if DIFF-10 is to have the better AIC score. If the two models fits the data about equally well, AIC will say that NULL-10 can be expected to be the more accurate predictor of new data.

Popper (1959) maintained that simpler theories are less probable than their more complex competitors. Although this is not always true, it is true for nested models. NULL entails DIFF, and this means that NULL cannot have the higher probability, no matter what data we have. Popper tried to make a virtue of this fact, praising simpler theories for their greater falsifiability. Unfortunately, falsifiability is the wrong tool to use in assessing theories that confer probabilities on possible data sets without entailing any of them. It also isn't clear how higher falsifiability can be an epistemic, as opposed to a pragmatic, virtue of theories. A theory that is more falsifiable is easier to prove false if indeed it is; more falsifiable theories are in this sense easier to test. However, what has that to do with the way the world is? Since the entailment relation linking NULL to DIFF prevents us from saying that the simpler theory has the higher probability of being true, it is tempting to conclude that simplicity has no epistemic relevance at all. It is at this point that AIC is important. Akaike identified a new goal for models—predictive accuracy—and it turns out that simplicity *is* epistemically relevant—not to showing which theories are more probably true, but to estimating predictive accuracy. Popper was on to something important; Akaike solved a problem that Popper identified but could not solve within his own system.

3. How the two concepts of parsimony differ

I have described two contexts in which the justification of parsimony in scientific reasoning is pretty clear. In the first, parsimony is a surrogate for likelihood. In the second, parsimony is relevant to estimating a model's predictive accuracy. These two roles for parsimony are dramatically different. In the first, more parsimonious theories confer *higher* probabilities on the data at hand. In the second, more parsimonious models, when fitted to the data, usually confer *lower* probabilities on the data at hand, but their greater parsimoniousness can provide a reason for expecting them to do a better job of predicting new data.

The example described in the previous section concerning a single pair of corn plant populations provides a useful way to think about the difference between likelihood considerations and model selection criteria like AIC. We assumed that the difference in the sample means is 4 centimeters. If we were to apply the Law of Likelihood to NULL and DIFF, we would have to ask whether

$$\Pr[(s_1 - s_2) = 4|\text{NULL}] > \Pr[(s_1 - s_2) = 4|\text{DIFF}].$$

How are we to evaluate the likelihood on the right-hand side of this inequality? As noted earlier, DIFF is a near tautology. What is the probability that the sample means will differ by 4 centimeters if DIFF is true? DIFF isn't "specific" enough to answer this question. But no such question needs to be answered if we use AIC. To apply AIC, we don't need to contemplate the value of $\Pr[(s_1 - s_2) = 4|\text{DIFF}]$; rather, we must evaluate $\Pr[(s_1 - s_2) = 4|\text{L(DIFF)}]$. This is a more tractable problem. To apply AIC, we need to consider the likelihoods of *fitted* models, not the likelihoods of models that contain adjustable parameters. And the likelihood of the fitted model is just one piece of the puzzle; the other is parsimony (as measured by the number of adjustable parameters), a quantity that goes unmentioned in the Law of Likelihood.

Although the Law of Likelihood does not impose any restrictions on which hypotheses can be "competitors," it is entirely natural to require that competing hypotheses be incompatible. If a coin lands heads 11 times in 20 tosses, it makes perfect sense to compare the hypothesis that the coin is fair ($p = 0.5$) with the hypothesis that the coin is highly biased in favor of heads ($p = 0.9$). It is decidedly odd to compare the first of these hypotheses with the hypothesis that the coin's probability of landing heads is somewhere between 0.4 and 0.6. Yet, in explaining the ABC's of AIC, I considered the models NULL and DIFF. These models are nested; they are related by entailment. Shouldn't proper competitors be incompatible with each other? Well, if the goal is to find hypotheses that are true, maybe they should be. But the goal of AIC is to estimate predictive accuracy. Since nested models can differ in their predictive accuracies, there is nothing amiss in treating them as competitors.

Not that competing models *must* be nested. For example, instead of comparing NULL and DIFF, we could have compared the NULL hypothesis with

(DIFF*) There exists a number $d \neq 0$ such that $(m_1 - m_2) = d$.

This model is incompatible with NULL. I noted earlier that NULL can't fit the data better than DIFF; at best, they tie. However, NULL *can* fit the data better than DIFF*, but the difference will be arbitrary small. If the observed difference in the sample means is exactly 0 centimeters, NULL fits this observation perfectly, whereas L(DIFF*) will say that the difference in the population means is 0.000000000000001. NULL is now just a *tiny* bit more likely than L(DIFF*). However, the question remains of whether NULL or DIFF* will be more accurate predictors, and that isn't settled by the likelihoods.

The Law of Likelihood describes how evidence should be interpreted when the question is whether H_1 or H_2 is true. In contrast, it would be a mistake to describe AIC as providing advice about which of the competing models is true. It is not surprising that false models occasionally make more accurate predictions than true ones. However, it *is* surprising that model selection criteria like AIC sometimes favor models known to be false over models known to be true. With respect to our two populations of corn plants, I think we know that NULL is false. I am as certain as I am about almost anything that the two populations do not have exactly the same mean height—identical to a thousand decimal places and beyond. DIFF, on the other hand, is a near tautology, and so is DIFF*; I am quite sure that both of them are true. If the goal were separating true models from false ones, there would be no need to look at the data. But there is, if the goal is predictive accuracy.

4. What the two concepts of parsimony have in common

In spite of the differences just noted that separate the two parsimony paradigms, they do have some common elements. When parsimony is a surrogate for likelihood, and when parsimony is part of what matters in model selection, parsimony is not an arbitrary aesthetic frill. It has an objective epistemic status. However, in neither case is *explanation* the fundamental consideration. Although it is harmless to describe CA and SA, and NULL and DIFF, as rival "explanations," the real question isn't whether one of these is "more explanatory" than its competitor; rather, what is fundamental is something else—likelihoods for the first pair of hypotheses, estimates of predictive accuracy for the second. Another similarity is that both the likelihood framework and that of AIC are *contrastive*; the frameworks are comparative, concerning better and worse. In neither case do we examine a single hypothesis and decide whether that hypothesis is "needed" to do something (e.g., explain the data). The two frameworks also have in common the fact that they do not discriminate between hypotheses that are empirically equivalent—i.e., that make the same predictions for all possible data. The Law of Likelihood says that evidence favors one hypothesis over another only when the two confer *different* probabilities on the observations. And AIC should be applied to a pair of models only when it is possible for them to make different predictions about new data when fitted to old. Finally, there is the fact that the justification for using parsimony as a surrogate for likelihood and the justification for using it in model selection each rest on empirical assumptions. My example about fathers distributing pennies was intended to demonstrate this point in connection with likelihood. I did not discuss the empirical assumptions behind Akaike's theorem, but the theorem is, as its name suggests, something that is derived from assumptions; these assumptions are often entirely plausible, but they aren't *a priori* in character (Forster and Sober 1994; Sober 2008a).

5. A parsimony principle not covered by the two justifications

When philosophers think about parsimony, they often think of it as a guide concerning what one should *believe*. Ockham's razor, according to the old slogan, is the admonition to not postulate entities beyond necessity. In fact, there are two razors that need to be distinguished.

> *Razor of denial*: If your evidence does not discriminate between "*X* exists" and "*X* does not exist," you should deny the former and affirm the latter.
>
> *Razor of silence*: If your evidence does not discriminate between "*X* exists" and "*X* does not exist," you should suspend judgment about both.

I take it that the "should" in both razors indicates what is *rational* (in the sense of required by reason). Since each says that lack of discriminating evidence suffices to settle what your state of belief should be, both razors deny a role to prudential considerations of the kind that figure in Pascal's wager. Since neither razor says what one should do when the evidence *does* discriminate, neither is committed to a full-blown evidentialism according to which one's beliefs should *always* be guided by the evidence and nothing else. Still, in those cases in which the evidence does *not* discriminate, it is the razor of silence that is the standard bearer for evidentialism. The two razors disagree on whether it is rational to believe that *X* does *not* exist when the evidence fails to discriminate.

Neither razor can be justified in terms of the two parsimony paradigms described earlier. Parsimony in the context of AIC is relevant to predictive accuracy of models, not their truth. And when parsimony is a surrogate for the *likelihoods* of "*X* exists" and "*X* does not exist," it says nothing, by itself, about the *probabilities* that either of these hypotheses possess. The odds version of Bayes' theorem clarifies why this is so:

$$\frac{\Pr(X \text{ does not exist} \mid E)}{\Pr(X \text{ exists} \mid E)} = \frac{\Pr(E \mid X \text{ does not exist})}{\Pr(E \mid X \text{ exists})} \times \frac{\Pr(X \text{ does not exist})}{\Pr(X \text{ exists})}$$

This equation says that the ratio of the posterior probabilities equals the likelihood ratio times the ratio of the prior probabilities. If "*X* does not exist" has the higher likelihood, nothing follows as to whether it also has the higher posterior probability; everything depends on what the prior probabilities are. The two razors therefore raise a new question; if either is justified, we should supplement the two parsimony paradigms already canvassed with a third.

The razors give advice about what to believe, where belief is understood as a dichotomy (you either believe a proposition or you don't) or a trichotomy (you either believe it, disbelieve it, or suspend judgment). Bayesianism, however, is about *degrees* of belief. I will connect the razors to Bayesianism by interpreting the razor of denial as saying that you should have a high degree of belief in "*X* does not exist" when the evidence fails to discriminate. How Bayesianism is related to the razor of silence is something I'll discuss in a moment. Both razors provide advice about what to do when the evidence does not discriminate between the competing hypotheses, but "failing to discriminate" has two interpretations. Let us consider them in turn.

If failing to discriminate means that the likelihoods are equal, Bayesianism interprets each razor as providing advice about the assignment of priors. The razor of silence says that we should suspend judgment about whether *X* exists when the likelihoods are equal. This is correct precisely when "*X* exists" and "*X* does not exist" have prior probabilities that are the same (or approximately so). The razor of denial says that we should believe that *X* does not exist when the likelihoods are equal. This is correct precisely when the hypothesis of nonexistence should be assigned a high prior probability. I see no justification for either of these global recommendations. As long as an assignment of prior probabilities is consistent with the probability calculus, the only justification there is for one such assignment as opposed to another must come from *empirical* information. This may involve a well-confirmed empirical theory or observational data. Logic alone tells us nothing about this, and there is no stand-alone epistemological principle that does so, either. True, there are circumstances in which priors should be taken to be equal and circumstances in which the hypothesis of nonexistence should be assigned a high prior, but there also are circumstances in which the existence claim deserves to have the higher prior. And most importantly, there are lots of cases in which there is no *should* at work at all. This, I believe, is the situation that frequently arises in connection with scientific theories. What is the prior probability of Darwin's theory of evolution or of Einstein's special theory of relativity? Different agents may have different subjective degrees of belief about this. But I see no hope of saying what the prior probabilities should be for either of these theories. Of course, each theory makes observational predictions, and observations may favor those theories over alternatives, but that is a point about likelihoods, not about priors.

Although Bayesianism lends little support to either razor if "failing to discriminate" means equality of likelihoods, the phrase has a second interpretation. It might mean that we have no clue as to what probability the hypothesis that *X* exists confers on any observation, even when we take our background knowledge into account. This is the situation that arises when the hypothesis is *untestable*. It is hard to see what could justify the razor of denial in this context. If we can make no judgment about the value of the likelihood, then we can't say anything about the value of the posterior probability, either. Philosophers now have mostly walked away from verificationism, which is the thesis that untestable sentences are neither true nor false. In contrast, the razor of denial remains a popular tool in philosophy, especially in connection with existence postulates that are taken to be untestable. In this application, the razor of denial is a cousin of verificationism. Both disparage statements that are untestable. And both go too far.

How does the razor of silence fare when the likelihoods of "*X* exists" and "*X* does not exist" are inscrutable? If suspending judgment about the existence claim means that we should assign it a posterior probability of about ½, then this razor is no better than the razor of denial. But if it merely means that we should not assign any probability at all (high, low, or middling), then it makes sense, if a modest sort of evidentialism is right. If weight of evidence is reflected in the likelihood ratio and nothing else, and the likelihoods are unknown, then it is impossible to tell how strongly one should believe the proposition "*X* exists" as opposed to its negation. Given this modest evidentialism, the razor of silence is correct when the propositions in question are untestable.[12]

My treatment of parsimony in terms of the Law of Likelihood and in terms of model selection may seem like *Hamlet* without the Prince, as neither addresses parsimony as a guide concerning what to *believe*. Guidance about belief, I take it, should be understood within a Bayesian framework. But this means that Ockham's razor must be brought up to date. Bayes's theorem makes no overt mention of parsimony, so if parsimony is to have a Bayesian significance, it must do so by being reflected in priors or in likelihoods. The razor of denial and the razor of silence, as I understand them, are meant to apply when the likelihoods "fail to discriminate" between "*X* exists" and "*X* does not exist." This has the two interpretations we have canvassed. It turns out that neither razor makes sense unless it is restricted. Are these restrictions so severe that they render the razors too dull to be of any use? I think that is often true when "failure to discriminate" means that the likelihoods are equal. However, I think that the razor of silence still has its point when likelihoods or priors are inscrutable. My assessments of the razors are summarized in Figure 2.

Figure 2. When do the razor of denial and the razor of silence make correct recommendations about what to believe? That depends on what it means for the observations to "fail to discriminate" between "X exists" and "X does not exist." Evidentialism is assumed.

	Likelihoods equal	Likelihoods inscrutable
Razor of denial	when Pr(X exists) << Pr(X does not exist)	never
Razor of silence	when Pr(X exists) ≈ Pr(X does not exist) or they are inscrutable	always

The razors go wrong in the same way that many epistemological principles go wrong—they offer advice about what to believe based just on the evidence at hand. The principle of induction does this, as does Reichenbach's principle of the common cause, and so do Fisherian significance tests and the Neyman-Pearson theory of hypothesis testing (Sober 2008b). No such principle is correct. This is a simple but important consequence of Bayesianism. The testimony of the evidence concerning various hypotheses is represented by their likelihoods. These likelihoods, by themselves, do not settle what the posterior probabilities are; you simply can't get a value for $Pr(H|O)$ from $Pr(O|H)$ alone. Frequentists decline to discuss prior probabilities because priors, they say, are insufficiently objective. This is often true, but not always. In any event, a conditional Bayesian point remains: *if you want posterior probabilities, you must have priors*. The two razors were doomed from the start.

Scientists have a saying—that absence of evidence isn't evidence of absence. The slogan is sometimes wrong, but there are many important cases in which it is exactly right or approximately so (Sober 2009). The razor of denial doesn't quite say that absence of evidence for the existence of an entity is *evidence* for its non-existence, but it comes close—close enough to run into trouble from the clear cases in which the scientist's slogan is correct. You have no evidence that there is a storm on the surface of Jupiter right now, but that hardly entitles you to deny that a storm is now occurring there. The razor of denial gives bad advice here.

Philosophers often think of tooth fairies when they think of Ockham's razor. We don't need to postulate the existence of these entities to explain what we observe, and that is a very good reason indeed for denying that they exist. Agnosticism in such cases seems like a pointless timidity. I have no problem with this inference, but the reason behind it does not generalize very far. Adults who place quarters under their children's pillows know that the coins didn't get there because a fairy was at work. This is why it is so clear to the adult mind that the tooth fairy is a myth. However, this is precisely the situation that does *not* obtain when scientists and philosophers reach for their razors. The relevant context for razoring is one in which the evidence does *not* settle which explanation is true and one wants the razor to provide an extra-evidential input.

The razor of denial makes sense if the putative entity under consideration can be assumed to have the following property:

> If *X* exists, then we'd have evidence that *X* exists.

Given this assumption, absence of evidence *is* evidence of absence. But I see no reason to think that all existence claims must have this property. Without this assumption, or something similar (Sober 2009), absence of evidence does not license *denial*. Notice, by the way, that silence, not denial, is what methodological naturalism$_s$ recommends to scientists with respect to the supernatural. This suggests that philosophical naturalists should view the razor of denial with suspicion.

Although I formulated the two razors so that they apply to the problem of evaluating "*X* exists" and its negation when the evidence fails to discriminate between them, it is important to see that both razors need to be generalized, and in two directions. Surely there is nothing special about existence claims and their denials; other pairs of hypotheses can differ in how simple or parsimonious they are, and so we should view the two razors as partial expressions of more general epistemological principles. The second limitation of the two razors is that they speak only to the case in which the evidence fails to discriminate. But if parsimony matters then, it presumably also matters when the evidence slightly (or substantially) favors one competitor over another. Here again, the razors need to be situated in a larger context.

Although philosophers who aren't philosophers of science often think of parsimony in connection with existence claims, philosophers of science often think about parsimony in connection with law-like statements (which they take to be universal generalizations). One example is Popper's (1959) argument that simpler theories are more falsifiable. Another is Harold Jeffreys' (1961) simplicity postulate. Jeffreys' subject was the class of differential equations; he sought to order these equations by their simplicity and then he recommended that simpler theories should be assigned higher prior probabilities. Jeffreys' proposal encountered some technical problems, but the main problem was that it was just a proposal. Jeffreys thought that his postulate gives a good summary of how scientists in fact reason, but it is entirely opaque why scientists *ought* to conform their reasoning to his postulate. In addition, there is Popper's point, noted earlier, that when models are nested, the simpler theory can't have the higher prior probability.

I think that parsimony arguments rarely can be justified by showing that the more parsimonious theory deserves the higher prior probability, but sometimes this does make sense. Suppose a physician is considering two possible diagnoses of a patient's observed symptoms. There is the hypothesis that the patient has disease D_1 and the hypothesis that he has D_2; each, if true, would explain the observations. Suppose further that the first disease is very common while the second is very rare. This fact may justify assigning D_1 a higher prior probability. Given all this, the physician may conclude that D_1 has the higher posterior probability. The physician may also find it natural to say that D_1 is the more parsimonious hypothesis. Why invoke the rare disease if the common disease can explain the data? Here we have a third parsimony paradigm; if you can justify assigning simpler theories higher prior probabilities, then you have shown that simplicity is epistemically relevant. I doubt that this thought will help in the philosophical arguments I'll now consider, but it deserves recognition, for the sake of completeness.

6. Theism versus Atheism

The problem of evil is often formulated as a deductive proof that there is no God, if God must be all-powerful, all-knowing, and all-good (all-PKG). However, the fact of the matter is that you can't deduce the nonexistence of such a God from a description of the kinds and quantities of evil that there are. There is no contradiction in the supposition that an all-PKG God exists but allows so much evil to exist for reasons that we cannot fathom. A fall-back position is the evidential argument from evil, which claims that the evils we observe are evidence against the existence of an all-PKG God (Rowe 1979). It is natural to formulate this argument in terms of likelihoods:

$$Pr(E \mid \text{there is no all-PKG God}) > Pr(E \mid \text{an all-PKG God exists}).$$

In this inequality, *E* is a fairly detailed description, not the bland statement that some evils exist. The question is not why there is some evil rather than none at all.[13]

Wykstra (1984) argues that the evidential argument makes the following assumption, which he thinks is untrue:

> If an all-PKG God had a reason for permitting horrendous evils to exist, human beings would know what those reasons are.

My response is that the likelihood formulation of the evidential argument from evil does not require so strong a premise. The argument is consistent with our having considerable uncertainty about what God's motives would be in allowing some horrendous evils to exist. The question is which type of universe makes *E* more probable—a universe in which there is no all-PKG God or a universe in which an all-PKG God exists.

The claim that God's motives are *completely* inscrutable to us is logically consistent with the hypothesis that he is perfectly good, but the conjunction of the two is an instance of a pragmatic paradox; the conjunction entails something of the form "*p* and we can't know that *p*." This is an example of what Sorensen (1988) calls a *blindspot*; the proposition, if true, is something we cannot know. Perhaps, then, we should consider a more modest variety of inscrutability, according to which *some* of God's motives are or may be inscrutable to us. The relevant likelihood comparison then becomes:

$$Pr(E \mid \text{there is no all-PKG God}) > Pr(E \mid \text{an all-PKG God exists, some of whose motives may be inscrutable to us}).$$

The likelihood version of the evidential argument from evil can be recast in terms of parsimony. A purely naturalistic explanation of *E* is more parsimonious than an explanation in which the naturalistic facts are supplemented by postulating a supernatural being who one might *prima facie* expect to have made *E* false, but who is sufficiently inscrutable that we need to allow that he may have reasons for permitting *E* to be true. Perhaps parsimony in this context is a surrogate for likelihoods.[14]

The argument from evil concerns the existence of an all-PKG deity. Suppose we drop the assumption that God, if he exists, must be all-PKG. What, then, does "God" mean? There are many options, but the one I now want to consider leaves the concept of God so unspecified that it says nothing about any observation that we might ever make. Not only does this God hypothesis not deductively entail any observation; it doesn't even confer a probability on any, not even when supplemented by independently plausible auxiliary assumptions. Atheists often yearn to apply Ockham's razor at this point, where the slicing away they have in mind leads to denial, not to silence. I have already commented on why I find this epistemology unconvincing.

A sufficiently unspecified God will have no effect on the probability of E, once a naturalistic explanation of evil (N) is taken into account:

$$\Pr(E \mid N \text{ and God exists}) = \Pr(E \mid N \text{ and God does not exist}).$$

This likelihood equality allows the razor of denial and the razor of silence to draw their separate conclusions about whether God probably exists; which is right depends on what the prior probabilities are. For those who are disinclined to assign priors in this case, there is only silence.

7. The mind/body identity theory and dualism

Mind/body identity theorists argued for their theory, and against dualism, by invoking a principle of parsimony (Smart 1959; Brandt and Kim 1967). Suppose, to use an example much in vogue in the 1950s and 1960s, that a perfect correlation is discovered between having the c-fibers in one's brain fire and being in pain. How does this observation bear on the two philosophical theories? The identity theory advances the claim that

> (AE) For every mental property M, there exists a physical property P, such that $M=P$.

This proposition (so named because of the order of its quantifiers) is what dualism denies. Neither AE nor its negation mentions c-fiber firings. If pain turns out to be less than perfectly correlated with c-fiber firings, this does not refute the identity theory. Identity theorists can merely search for some other physical property with which to identify the mental property of being in pain.

Is AE more parsimonious than its negation? Consider what each says about how many properties there are. Dualism counts mental and physical properties separately. The identity theorist regards this as double counting and counts physical properties alone. If the number of physical properties is finite, the identity theorist will have the shorter list. And even if there are infinitely many physical properties, the point remains that the items on the identity theorist's list comprise a proper subpart of the items on the dualist's. But notice how disconnected this numerology is from the two parsimony paradigms described earlier. When parsimony is a surrogate for likelihood, the more parsimonious hypothesis confers on the observations a higher probability than the less parsimonious hypothesis does. And when parsimony is a consideration in model selection, we fit models to old data and attempt to estimate how accurately the models will predict new data. Property counting is another story entirely.

Perhaps, then, we should evaluate the two mind/body theories by focusing, not on AE and its negation, but on an instantiation of each. In the example at hand, we need to consider the hypothesis that pain and c-fiber firing are identical and the hypothesis that they are not. Why is the identity hypothesis more parsimonious? I set to one side the fact that 1 is a smaller number than 2. Identity theorists often point out that their theory can explain why pain and c-fiber firings are correlated: they are correlated because they are one and the same property. Dualists, on the other hand, must regard the correlation as a brute fact, which they can accept but cannot explain.[15] Where does parsimony enter this argument? Is the idea that the identity theory postulates fewer brute facts? If only we knew how to count brute facts and knew why the count is epistemically relevant.

I hope the suggestion that the identity theory *entails* the observed correlation, while dualism can only *accommodate* that correlation, reminds the reader of the example about corn plants and the two models NULL and DIFF. Can the identity theory and dualism be represented as models in the sense required by model selection theory? To this end, let's consider pain and c-fiber firing as quantitative, not dichotomous, characteristics. Pains have their intensities and c-fiber firings have their rates. What does the claim that pain is identical with c-fiber firing entail about their mathematical relationship? I suggest that the identity claim is committed to monotonicity, not linearity; see Figure 3. Pain and c-fiber firing are determinables and their specific values are determinates. If the determinables are identical, then each determinate of the one should be identical with precisely one determinate of the other. The symmetry of the identity relation rules out the relationship shown in (*iii*) of Figure 3.

If the identity thesis were committed to a linear relationship, then it would be easy to represent that thesis as a model. Linear relationships have the form $y = mx + b + e$ (where e is an error term). In contrast, the thesis of monotonicity is not a model in the required sense; the claim that the slope of a curve is everywhere positive (or everywhere negative) leaves open which of many models one wishes to consider. And what are we to say about Dualism? What is the model to which it corresponds? The statement that x and y are related *non*linearly is not a model, nor is the statement that their relation is *non*monotonic. These statements aren't models in the sense relevant to model selection theory. The negation of a model need not be a model.

Figure 3. If intensity of pain and rate of c-fiber firing are identical quantitative characteristics, what is the mathematical form of their relationship? In (i) the relationship is linear; in (i) and in (ii), it is monotonic.

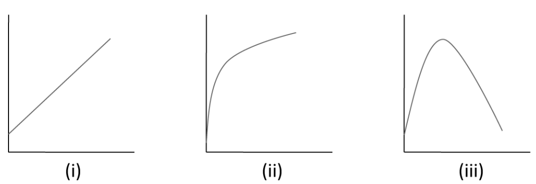

A way forward may be found in the fact that the identity theory says that an individual who is in the same c-fiber firing state at two different times will be in the same pain states at those two times (and vice versa), while dualism leaves open whether this is so.[16] Consider an experiment in which each of several experimental subjects is subjected to two kinds of pain.[17] One of them involves being hit over the head with a hammer while the other involves being subjected to loud music. We measure rate of c-fiber firing by a neural detector and we measure pain by asking subjects to locate the intensity of their pain on a numerical scale. For each subject, we vary the hammer blows and the music until we find a setting of each that leads him or her to report that their pains both have a value of 1. We then see what difference, if any, there is in the rates of c-fiber firing that occur. The identity theory takes the form of a null hypothesis:

> (IT) If the hammer blows and the loud music both are rated as 1 on the pain scale, then the difference in the rates of c-fiber firing in the two treatments = 0.

Dualism, on the other hand, contains an adjustable parameter, since it is consistent with any amount of difference that might arise in the two c-fiber firings:

> (DUAL) If the hammer blows and the loud music both are rated as 1 on the pain scale, then the difference in the rates of c-fiber firing in the two treatments = d.

Notice that IT can't fit the data better than DUAL does, which means that IT can't have a higher likelihood than L(DUAL). But that does not settle the question of which model can be expected to make more accurate predictions about new data. Depending on the data at hand, IT may have the better AIC score; DUAL may not fit the data *sufficiently* better than IT does to overcome the former's greater complexity.

The difference in parsimony that separates IT and DUAL is modest, but a more elaborate set of experiments can widen the gap. Suppose there are ten experiments. The first is as described above—the hammer blows and the loud music both are arranged so that each elicits a pain rating of 1 from each subject. In the second experiment, the two treatments are configured so that both produce pain scores of 2. And so on. The identity theory still takes the form of a null model:

> (IT-10) If the hammer blows and the loud music receive the same pain scores, then the difference in the rates of c-fiber firing in the two treatments = 0.

Dualism, on the other hand, now corresponds to a model with ten parameters:

> (DUAL-10) If the hammer blows and the loud music both receive a pain rating of i, then the difference in the rates of c-fiber firings in the two treatments = d_i (where $i = 1, 2, \ldots, 10$).

This move from one experiment to ten is just the tip of the iceberg. As we increase the number of experiments that differ in their levels of pain, the identity theory remains a null model, while dualism grows ever more complex. It becomes easier and easier for the identity theory to obtain a better AIC score than dualism.

Functionalism is a third theory of the mind/body relation. It was much discussed in the 1960s and 1970s and was widely thought to be superior to both the identity theory and dualism. How does it fit into the model selection framework? In the experiments contemplated thus far, individuals were in the same pain state in the two treatments and data were gathered on what their c-fiber firing states were. Functionalism, because of its commitment to multiple realization, agrees with dualism about how these experiments should be modeled; subjects *may* be in the same c-fiber firing state, but they *need not be*. However, if we reverse the experimental design, functionalism is on the same page as the identity theory. If subjects are in the same c-fiber firing state in the two treatments, and the functionalist hypothesis is that this is a supervenience base for pain, then the prediction is that the individuals will be in the same pain state. In terms of models with adjustable parameters, functionalism is intermediate between the identity theory and dualism. Taking both experimental designs into account, the identity theory has zero adjustable parameters, functionalism has some, and dualism has more.[18]

Placing the mind/body identity theory and dualism within the context of model selection theory requires one to think of the contending theories in terms of their predictive accuracy, not their truth. Metaphysicians may balk at this, proclaiming that they don't care about predictive accuracy and want only to figure out what is true. I am not arguing against that preference. Rather, my point is that the parsimony argument for the identity theory finds a natural home in the model selection framework. If there is another treatment of the argument that establishes its connection with truth, I do not know what that treatment is.[19]

8. The causal efficacy of the mental

Assuming that human beings and other organisms have psychological properties, why think that the possession of those psychological properties causes behavior? If the behavior is caused by the organism's physical properties, isn't it unparsimonious to claim that there are psychological causes as well? Why postulate two causes when one will do? This line of questioning has been in the background of Kim's (1993, 1996) discussion of nonreductive physicalism and it traces back to Nozick's smart Martian problem, discussed in Dennett (1980) and in Sober (1999b). Should physicalists avoid the pomp of superfluous causes and embrace the hypothesis that our having the beliefs, desires, and sensations we do are epiphenomenal correlates of behavior, not their causes?

The idea that a one-cause model is more parsimonious than a two-cause model finds a natural representation within model selection theory (Forster and Sober 1994). For simplicity, let's consider two dichotomous properties A and B that each may be a cause of the dichotomous property E. The probability of E, conditional on different combinations of $\pm A$ and

$\pm B$, is shown in Figure 4. A model that says that both A and B are (or may be) causes of E will take the form.

(TWO)
$$Pr(E|+A \,\&\, +B) - Pr(E|-A \,\&\, +B) = a$$
$$Pr(E|+A \,\&\, -B) - Pr(E|-A \,\&\, -B) = a$$
$$Pr(E|+A \,\&\, +B) - Pr(E|-A \,\&\, -B) = b$$
$$Pr(E|-A \,\&\, +B) - Pr(E|-A \,\&\, -B) = b$$

This model allows that varying the state of A while holding fixed the state of B may make a difference in the probability of E, and that the same is true of varying the state of B while holding fixed the state of A.

The model that says that only $\pm A$ is (or might be) a cause of E takes the form:

(ONE)
$$Pr(E|+A \,\&\, +B) - Pr(E|-A \,\&\, +B) = a$$
$$Pr(E|+A \,\&\, -B) - Pr(E|-A \,\&\, -B) = a$$
$$Pr(E|+A \,\&\, +B) - Pr(E|+A \,\&\, -B) = 0$$
$$Pr(E|-A \,\&\, +B) - Pr(E|-A \,\&\, -B) = 0$$

This model says that varying the state of B while controlling the state of A makes no difference in the probability of E. Notice that TWO has two adjustable parameters while ONE has just one, and ONE is a special case of TWO. TWO will fit frequency data at least as well as ONE will, but ONE is more parsimonious. Depending on the data, a model selection criterion like AIC may award ONE the better score.

As an example, consider the question of how smoking and asbestos exposure are each related to lung cancer. Perhaps both are causes, or just one of them is. Indeed, there is the even simpler null hypothesis that says that neither of them makes a difference in the risk of lung cancer. Frequency data can be gathered that allows the models to be compared. Notice that the data you need here requires that some people smoke and others do not and that some people are exposed to asbestos and others are not. Without frequency data pertaining to the four cells of the two-by-two table depicted in Figure 4, there is no fitting of models to data, and no estimating predictive accuracy. What is tested here is not the hypothesis that people smoke or are exposed to asbestos, but whether these events are causes of lung cancer. This is important to bear in mind, for sometimes parsimony is taken to bear on the question of whether C exists, which differs from the question of whether C (assumed to exist) causes E.

Figure 4. A and B are possible causes of E. The probability of E, conditional on different combinations of $\pm A$ and $\pm B$, are as shown.

	$+A$	$-A$
$+B$	$x+a+b$	$x+b$
$-B$	$x+a$	x

Can this format be applied to the problem at hand in which a purely physicalistic explanation of a behavior is compared with an explanation that postulates both physical and mentalistic causes? If the physical explanation considered doesn't exhaust the physical causes that are at work, there is no difficulty. Perhaps thinking soothing thoughts helps reduce headache pain *and* so does taking aspirin. We could run an experiment to see whether there are two causes here or only one. However, this is not the case we need to contemplate. We need to consider a *physically complete* explanation and then ask whether it makes sense to supplement this physically complete story with the postulation of mentalistic causes. Now we run into a difficulty. The models ONE and TWO require that we consider what would happen if each putative factor were varied while holding fixed the other. But this can't be done in the case at hand if the mental *supervenes* on the physical—supervenience means that it is *impossible* to vary an individual's mental characteristics while holding fixed the individual's complete physical state.[21] Because of this, some of the conditional probabilities described in ONE and in TWO will fail to be well-defined (Shapiro and Sober 2007). And if they aren't well-defined, it won't be possible to estimate the values of the relevant parameters, and so it won't be possible to compute the AIC scores of models. The parsimony argument for epiphenomenalism, in which the one-cause model cites the physical property P and the two-cause model describes P and the mental property M where M supervenes on P, gains no purchase when placed in the context of model selection theory.

How, then, should one test the hypothesis that various mental characteristics are causes of behavior against the hypothesis that they are mere epiphenomenal correlates? A useful guide comes from a mundane and often-used example. Why think that barometer readings don't cause storms? The two hypotheses we need to consider are depicted in Figure 5. They agree that barometric pressure is a common cause of barometer readings and storms. And both predict that barometer readings and storms will be correlated. However, they disagree about whether barometric pressure *screens-off* barometer readings from storms. Epiphenomenalism (Ep) says it does:

Pr(storm|high barometric pressure & barometer reads high) =
Pr(storm|high barometric pressure & barometer reads low).
Pr(storm|low barometric pressure & barometer reads high) =
Pr(storm|low barometric pressure & barometer reads low).[22]

The causal hypothesis (C) denies this equality; it says that barometer readings affect the probability of storms even when one controls for the barometric pressure. This difference between the two hypotheses can be translated into the language of model selection theory. The Epiphenomenalist model is a null hypothesis; it says that there is no difference between various probabilities. In contrast, the Causal model will associate one or more adjustable parameters with various probability differences. Frequency data can then be used to estimate adjustable parameters and a model selection criterion like AIC can be applied.

Figure 5. (*C*) says that barometer readings cause storms; (*Ep*) says they do not. The two hypotheses agree that barometric pressure is a cause of barometer readings and of storms. And both predict that barometer readings and storms will be correlated.

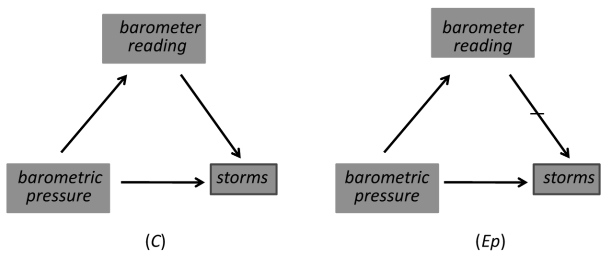

Parsimony is relevant here, and the epiphenomenalist model *is* more parsimonious.

Why do things run smoothly when it comes to barometer readings but we hit a rocky road in connection with mental causation? The reason is that we considered the *supervenience bases* of mental properties in the one case, but the *common causes* of barometer readings and storms in the other. The former formulation has the result that relevant conditional probabilities are not defined, but no such wreckage arises in the latter. This opens the door to a new and better version of epiphenomenalism. We should consider how mental properties and behaviors are related to their physical *common causes* (*not* to the mental properties' supervenience bases).[23] The epiphenomenalist and the causal models are both legitimate, and the fact that mental states are correlated with behavior is common ground, not evidence that favors one over the other. Epiphenomenalism *is* the more parsimonious model, but that does not suffice to show which model is better (Shapiro and Sober 2007).

9. Ethical realism

Let ethical realism be the thesis that some normative ethical propositions are true and that their truth is independent of anyone's thinking or saying that they are. So defined, ethical realism is incompatible with the divine command theory and with some forms of ethical relativism and existentialism. It isn't God's disapproving of an action, or society's, or the individual's that makes the action wrong. Ethical realism, as I understand it, does not say that ethical truths are independent of human psychology; hedonistic utilitarianism, for example, can be given a realist interpretation even though it says that ethical truths are made true by facts about pleasure and pain.

Harman (1977) argues against ethical realism[24] by way of a parsimony argument. He contends that there is no need to postulate the existence of an independent realm of normative ethical truths if we wish to explain human thought and behavior. A person's upbringing suffices to do the explaining. Harman concludes from this, not that we should remain silent on whether such independent ethical truths exist, but that there are no such things. Ruse and Wilson (1986) advance a variation on this line of argument, suggesting that evolutionary theory suffices to explain why we have the ethical beliefs and feelings we do and that, for this reason, we should be anti-realists about ethical truths. Although these anti-realist arguments focus on the task of explaining what we think and do, they, in fact, require a stronger thesis if they are to reach the desired conclusion. The thesis that needs to be established is that the postulates of ethical realism aren't needed to explain *any* observation.

The parsimony argument against ethical realism is sometimes given a *causal* formulation. We know that there are various descriptive facts (including facts about upbringing and evolution) that are causes of human thought and action. Why postulate a set of normative ethical facts, whose truth is independent of anyone's say-so, as a second set of causes? If there are normative ethical facts that *supervene* on descriptive facts (as, e.g., utilitarianism maintains), then this parsimony argument runs into the same trouble that derails the parsimony argument against mental causation discussed earlier. It is a mistake to expect supervening causes to have causal powers that go beyond those exhibited by their supervenience bases (Sturgeon 1984; Kim 1993).

What if we drop the causal version of this parsimony argument and stick to the concept of explanation? Harman says that upbringing provides the best or the only explanation of why we think and act as we do and that we don't need to postulate an independent realm of normative ethical facts to do the explaining. Sturgeon (1984) responds that normative ethical facts *are* sometimes explanatory; he argues, for example, that Hitler's depravity explains why he acted as he did. Although I disagree with Harman's argument, I don't want to endorse Sturgeon's response to it. The reason I hesitate to do so is that I think there are many correct explanations of a given behavior; which we should prefer depends on our interests (Sober 1999a).[25] The reply I prefer to the parsimony argument against ethical realism is to reject the requirement that ethical facts should explain human thought and behavior. Ethics is in a different line of work from psychology. Normative ethical propositions have the job of telling us how we *ought* to act, not of explaining why we *in fact* act as we do (Sober 1990; Shafer-Landau 2007). The distinction I have in mind is the same one that needs to be drawn between logic and psychology. Let logical realism be the view that some normative propositions about what we ought to conclude from given premises are true, and that these propositions are true independently of anyone's say-so. Logical realism may or may not be correct, but it is a mistake to evaluate it as if it were a descriptive psychological thesis.[26]

If normative ethical propositions should not be expected to explain our observations concerning what human beings think and do, are there other observations that these propositions can be expected to address? That depends on what we mean by observation. It has become standard in philosophy of science to say that observations are "theory laden." This means that the propositions that scientists properly treat as observation statements are knowable only by agents who have a relevant theory in their possession; a physicist who looks at the screen of a cloud chamber can be said to observe that an electron is moving from right to left, but a person with no knowledge of physics will not be able to see that this is so. The thesis that observations are theory-laden was once thought to lead to relativism, but, in fact, no such dire consequences are in the offing. Observations can be theory-laden and still provide a neutral basis for discriminating between competing hypotheses. If an observation statement *O* is to help us compare theories T_1 and T_2, it must be possible to know that *O* is true without already having to believe either T_1 or T_2. But that leaves it open that knowing *O* may require the use some other theory, T_3. What matters is that observations be *relatively* theory-neutral, not that they be *absolutely* theory-neutral (Sober 2008a).

Given this, I see no problem with regarding some normative ethical propositions as observation statements. Harman (1977) agrees and gives a good example. We see a gang of hoodlums set fire to a cat and the thought leaps to mind that what they are doing is wrong. If we regard the judgment that the act is wrong as an observation statement, we may ask different ethical theories to explain why it is true and then evaluate those theories by assessing the quality of the explanations they provide. If so, it is a mistake to maintain that normative ethical propositions aren't needed to explain *anything*; some of them may be needed to explain others. I do not suggest that this is an argument for ethical realism. Rather, the point is that if some normative propositions are true, others may be needed to explain them; a parallel claim applies to descriptive propositions (Sturgeon 1984).

In summary, the parsimony argument against ethical realism goes wrong on three fronts. First, the model-selection rationale for preferring models that postulate one cause over models that postulate two depends on the possibility of varying each cause while holding fixed the other; this cannot be done if one candidate cause supervenes on the other. Second, normative ethical propositions should not be evaluated by their ability to explain descriptive propositions about human thought and behavior. And third, even if normative ethical propositions aren't needed to explain what we think and do, it doesn't follow that they aren't needed to explain *anything*.

Just as evolution has been used to argue for ethical anti-realism, it also has been used in connection with an epistemological thesis—that we have no justification for the normative ethical propositions we happen to believe. Joyce's (2006) defense of this skeptical thesis is based on his contention that the evolutionary process causes us to have our ethical beliefs irrespective of whether those beliefs are true.[27] To understand this causal claim, let us begin with a thesis about screening-off:

(SO) $\Pr(\text{we believe } p \mid p \ \&\ \text{Ev}) = \Pr(\text{we believe } p \mid \textit{not-p} \ \&\ \text{Ev})$.

Here "Ev" represents a summary of the evolutionary processes that affected whether we'd believe that p is true. For the sake of specificity, let p be the proposition that we ought to take care of our children. It is important to see that (SO) does not suffice to establish

(J) $\Pr(\text{we believe } p \mid p) = \Pr(\text{we believe } p \mid \textit{not-p})$.

If screening-off had this consequence, then barometer readings wouldn't be able to give us evidence about storms, nor would we be able to gather evidence about a distal cause by looking at its effect when there is a screening-off proximate cause that lies in between. So even if (SO) were true, it still could be the case that our believing that *p* is evidence that *p* is true in the sense described by the law of likelihood. An additional premise is needed beyond (SO) to establish the truth of (J).

Before we get to that additional premise, it is important to note that Joyce's conclusion that we have no grounds for maintaining that *p* is true does not follow from the likelihood equality (J). In the first instance, what (J) asserts is something more modest: the fact that we believe that *p* is true does not discriminate between *p*'s being true and its being false. Is this psychological fact about our state of belief the *only* evidence we can bring to bear? If not, (J) is not enough. Here is a modest additional fact: not only do we believe that *p* is true; we also have discovered that *p* is logically consistent with the rest of what we believe. This may seem like a very small point in favor of *p*. It is, but it *is* a point in favor. Had we discovered that *p* is not consistent with the rest of what we believe, that would have counted against. The failure to obtain this bad news is itself good news for the truth of *p*, if only very weak good news.[28]

What has Joyce's argument to do with parsimony? If our evidence does not discriminate between *p* and *not-p*, the razor of silence recommends that we suspend judgment while the razor of denial recommends something more radical. As noted earlier, nothing follows from (J) about the posterior probability we ought to assign; everything depends on what one can say about the priors (Sober 1994). If we are unable to assign a prior probability (or range thereof, short of the all-inclusive unit interval), the razor of silence gives good advice—there is no posterior degree of belief (or range thereof…) that we ought to have concerning proposition *p*. What does not follow is what the razor of denial recommends—that we should assign *p* a low posterior probability. Joyce (2006, 181) is careful to note that his argument is for agnosticism, not for the claim that our moral beliefs are false.

What more is needed to establish (J) than the screening-off thesis (SO)? Let's begin with the fact that (J) is equivalent

$\Pr(\text{we believe } p \mid p \ \&\ Ev)\Pr(Ev \mid p) + \Pr(\text{we believe } p \mid p \ \&\ \textit{not-Ev})\Pr(\textit{not-Ev} \mid p) =$

$\Pr(\text{we believe } p \mid \textit{not-p} \ \&\ Ev)\Pr(Ev \mid \textit{not-p}) + \Pr(\text{we believe } p \mid \textit{not-p} \ \&\ \textit{not-Ev})\Pr(\textit{not-Ev} \mid \textit{not-p})$.

Given (SO), this simplifies to

$$\Pr(Ev|p) = \Pr(Ev|not\text{-}p),$$

which in turn is equivalent to

$$\Pr(p\ |Ev) = \Pr(p|not\text{-}Ev).$$

This last equality, which needs to be added to (SO) to complete the argument for (J), is part of Joyce's picture: "our moral beliefs are products of a process that is entirely independent of their truth" (Joyce 2006, 211). However, there is reason to doubt this equality (Brosnan unpublished). If the evolutionary process had been different, maybe we wouldn't have the same obligations. If evolution had not brought it about that human infants have prolonged periods of helplessness, perhaps human parents would not be obliged to care for their children. This thought does not require that evolution is guided by antecedent facts about what is right and what is wrong. Rather, the idea is just that facts about our biology affect what our obligations are.

10. Nominalism and Platonism about Mathematics

Discussion of Ockham's razor traces back to the problem of universals, so let us return to our roots for a moment and consider how the two parsimony paradigms discussed earlier bear on the question of nominalism and Platonism. This opposition can be considered in connection with the existence of properties and also with respect to the existence of mathematical objects (e.g., numbers). Platonism claims that properties and numbers exist and that they exist outside of space and time; nominalists deny the existence of properties and numbers (Platonistically conceived). I will focus on numbers.

Nominalists pursue a dual strategy in interpreting mathematics—translate what you can and deny the truth of the rest. The first line of attack is to show that mathematical propositions can be paraphrased into a nominalistically acceptable language. Consider, for example, the claim that

(Apples) There are two apples in the basket.

Nominalists suggest that (Apples) be understood as follows:

(N) There exist physical objects x and y such that x is an apple in the basket and y is an apple in the basket and $x \neq y$, and for all z, if z is an apple in the basket, then $z=x$ or $z=y$.

An alternative construal of Apples is Platonistic:

(P) There exists a number n such that n = the number of apples in the basket and $n=2$.

The first of these statements (N) quantifies only over physical objects while the second (P) quantifies over numbers. If (N) is a good enough paraphrase of (Apples), then we don't need to assert that (P) is true to say what we want. However, that provides no reason to think that (P) is false (Alston 1958). If we use the metric system to measure length, we don't need to talk about length in terms of inches and feet. However, that hardly shows that statements about inches and feet are false. Of course, there is an equivalence between length-in-meters and length-in-feet, and, arguably, no such equivalence connects (N) and (P). Maybe so, but the point remains: the fact that one doesn't need to say or write a sentence isn't evidence that what the sentence says is false.

Although (Apples) is often taken to be good news for nominalism, it is widely recognized that much of mathematics cannot be paraphrased in this way. Consider, for example:

(Prime) There are infinitely many prime numbers.

This is not equivalent to a claim about marks on paper or about what human beings can achieve by various constructions. Nor is (Prime) equivalent to

(Prime*) If numbers exist, then there are infinitely many primes.

Perhaps the best nominalist response to statements like (Prime) is fictionalism (Balaguer 2001), the thesis that this and other existence claims in mathematics are false. Fictionalists may seek to bolster their position by pointing out that what mathematics establishes are conditional results, like (Prime*). This *is* a defensible epistemic position, but it hardly shows that (Prime) is false. If we know that (Prime*) is true and nothing more, why not be agnostic about the logically stronger (Prime)? The agnosticism I have in mind is Carnapian. Mathematics assumes a framework of numbers just as physics assumes a framework of physical objects. We *assume* these frameworks; perhaps there is no compelling evidence that these frameworks are correct *or* that they are not. In the absence of such evidence, it makes sense to withhold belief, but also to withhold disbelief.

If the parsimony argument for nominalism is hard to justify, the same is true of what has become a standard Platonist reply to it—the indispensability argument of Quine (1953) and Putnam (1971). Their point is not just that mathematized natural science needs to quantify over numbers. That, after all, allows us to regard the existence of numbers as a useful (indeed, an indispensable) fiction. Rather, the indispensability argument claims that the empirical evidence that confirms a scientific theory also confirms the purely mathematical consequences that the theory has. This argument appeals to a form of *epistemological holism*: when a whole theory gets confirmed, so does each of its parts, even the parts that are propositions of pure mathematics.[29] For example, the observations that confirm relativity theory are said to confirm the existence of numbers, since relativity theory entails that numbers exist. I have criticized this indispensability argument on two grounds (Sober 1993; Sober forthcoming). First, it is guilty of *selective attention*; if the empirical success of relativity theory confirms the existence of numbers, why doesn't the empirical failure of many other mathematized theories disconfirm the existence of numbers? Second, I suspect that the indispensability argument falls into the trap of assuming that Hempel's (1965) *special consequence condition* is true. The special consequence condition says that if O confirms T, and T has the logical consequence C, then O confirms C. It has been known for a long time that the special consequence condition is false.[30] It remains to be seen whether the epistemology behind the indispensability argument can be retooled.

Why does the nominalist's parsimony argument go so badly wrong? Is it because its subject matter (numbers and universals) is a priori, whereas the two parsimony paradigms discussed earlier each involve empirical observations? That is not the central issue. Recall the relaxed understanding of what an observation is that I discussed in connection with ethical realism. Perhaps we can take this one step further. If a physicist can see that an electron is moving across a cloud chamber, can a mathematician see that the number 13 is prime? The word "see" has to do with vision in the first case, but not in the second; blind mathematicians are under no handicap in connection with their apprehension of this mathematical fact. If we nonetheless treat such singular judgments about particular numbers as observations, we can think of their relation to competing generalizations about numbers in something like the format that we use to think about the relation of observations and generalizations in the empirical sciences. Of course, mathematicians strive to find proofs or disproofs of these generalizations, but before they reach that point, they form judgments about the relative plausibility of different generalizations that fit the "observations" they have made of the properties of individual numbers. Such judgments might be guided by the relative simplicity of the competing generalizations; if so, it would be worth investigating how mathematicians' notions of simplicity connect with the ones used in empirical sciences.[31] This is a project for the epistemology of mathematics. However, the nominalist's parsimony argument is *philosophy* of mathematics, not mathematics proper, and it raises different questions.

11. Concluding Comments

Tolstoy begins *Anna Karenina* with the observation that "every happy family is alike, but every unhappy family is unhappy in its own way." This is not quite the pattern that pertains to parsimony arguments—the bad ones are bad for different reasons, but the same is true of the good ones. In science, parsimony is sometimes a surrogate for likelihood and sometimes it stands opposed to likelihood but is relevant to estimating a model's predictive accuracy. Ockham's razor offers advice about what to believe and therefore falls under neither of these headings. Here we must distinguish the razor of denial from the razor of silence; the first razor is right in what it says when the existence claim has a low prior probability. The second razor makes sense in a far more interesting set of circumstances; it gives the right advice for existence claims that are untestable or have prior probabilities that are equal to 0.5.

Philosophical parsimony arguments sometimes conform to the likelihood paradigm, and sometimes to the model selection paradigm, and sometimes to neither. The problem of evil may suggest that atheism is a more parsimonious explanation than theism; in this argument, parsimony is a surrogate for likelihood. In the mind/body problem, the identity theory is more parsimonious than dualism, and model selection theory provides a gloss of this judgment that makes the difference in parsimony epistemically relevant. Matters change, however, when it comes to parsimony arguments against mental causation, against ethical realism, and against mathematical Platonism. Model selection theory does describe a circumstance in which a one-cause model is better than a two-cause model, but this format applies only when each cause can be varied while holding the other fixed. This cannot be done if the mental and physical properties considered are such that the former supervenes on the latter. The same flaw may be found in parsimony arguments against ethical realism, if ethical facts supervene on nonethical facts. The parsimony argument against ethical realism is flawed for an additional reason; it says that postulating an independent realm of normative ethical truths is justified only if these truths are needed to explain descriptive propositions about human thought and action. But ethics isn't psychology, any more than logic is. As for the parsimony argument against Platonistic entities, the fact that "there are two apples in the basket" can be paraphrased without quantifying over numbers is no reason to deny that numbers exist. And the fact that mathematicians establish only conditional results (e.g., "if numbers exist, then there are infinitely many prime numbers") is likewise no reason to deny that numbers exist. Parsimony arguments against mental causation, against ethical realism, and against mathematical Platonism do not fit the two scientific formats I have described, but that does not prove that they are misguided. Perhaps we should reject naturalism$_p$. Or perhaps we need to delve more deeply into science to uncover new and convincing precedents. My expectation is that there is no rescuing these last three parsimony arguments, but that remains to be seen. Naturalism$_p$ is a fruitful source of problems concerning the relationship of science and philosophy, whether or not we decide, in the end, that it is true.

Appendix: A Reichenbachian Proof concerning the Likelihoods of Common and Separate Ancestry

Figure 6. A common ancestry and a separate ancestry explanation for why Humans (H) and monkeys (M) are in character state 1. It is assumed that the ancesters Z, Z_1, Z_2 are in state 0; the states of the ancestors A, A_1, A_2 are unspecified.

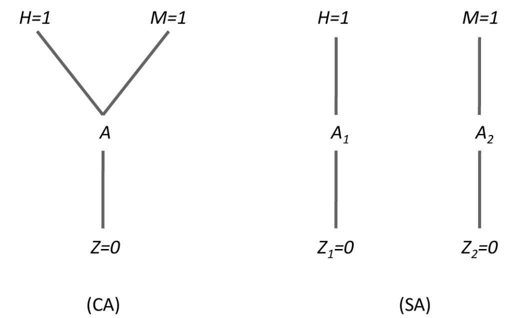

Consider the following likelihood inequality:

$$Pr(H=1 \,\&\, M=1 \mid CA) > Pr(H=1 \,\&\, M=1 \mid SA).$$

If we assume that the hypotheses incorporate the assumptions about the states of the ancestors shown in Figure 6, this inequality takes the form

(*) $$Pr(H=1 \,\&\, M=1 \mid CA \,\&\, Z=0) > Pr(H=1 \,\&\, M=1 \mid SA \,\&\, Z_1=0 \,\&\, Z_2=0).$$

Let us parameterize the CA and the SA hypotheses as follows:

$$Pr(A=1 \mid Z=0) = Pr(A_1=1 \mid Z_1=0) = Pr(A_2=1 \mid Z_2=0) = e$$
$$Pr(H=1 \mid A=1) = Pr(H=1 \mid A_1=1) = w$$
$$Pr(H=1 \mid A=0) = Pr(H=1 \mid A_1=0) = x$$
$$Pr(M=1 \mid A=1) = Pr(M=1 \mid A_2=1) = y$$
$$Pr(M=1 \mid A=0) = Pr(M=1 \mid A_2=0) = z$$

Given the Reichenbachian assumptions that all these probabilities are strictly between 0 and 1, that the states of the interior nodes (A, A_1, A_2) screen-off roots from tips, that common causes screen-off their joint effects from each other, and that separate causes do their work independently, (*) is true precisely when

$$ewy + (1-e)xz > [ew + (1-e)x][ey + (1-e)z].$$

This simplifies to

$$(w-x)(y-z) > o,$$

which is true if ancestors and descendants are always positively correlated, or always negatively correlated. The same assumptions also entail that the observation that humans and monkeys are both in state 1 favors *CA* over *SA* if the roots of the genealogies (Z, Z_1, and Z_2) are assumed to be in state 1.

In summary, the Reichenbachian framework entails that *matches favor CA over SA, regardless of whether the matches involve derived (apomorphic) or ancestral (plesiomorphic) character states*. Matters change if we stipulate a state for the ancestors A, A_1, and A_2. If these are stipulated to be in the same state, *CA* and *SA* are identical in likelihood.

Acknowledgments

I thank Mark Balaguer, Eric Barnes, Martin Barrett, Juan Comesaña, Jordan Ellenberg, Kenny Eswaran, Joshua Filler, Justin Fisher, Simon Fitzpatrick, Malcolm Forster, Suraya Gollob, Paula Gottlieb, Daniel Hausman, Casey Helgeson, Richard Joyce, Philip Kitcher, Steven Leeds, William Lycan, Alexander Paseau, Carolina Sartorio, Russell Shafer-Landau, Larry Shapiro, Nick Shea, Alan Sidelle, Walter Sinott-Armstrong, Dennis Stampe, Steve Sverdlik, Michael Titelbaum, Joel Velasco, Ralph Wedgwood, and Timothy Williamson for useful discussion.

Endnotes

1. The words "inside" and "outside" are used metaphorically here. A more literal formulation is that metaphysical naturalism says that all existing things have spatio-temporal location; supernatural entities, if they exist, do not.
2. Mathematized evolutionary theory quantifies over numbers; if numbers are what Platonists say they are (entities that exist outside of space and time), then evolutionary theory violates methodological naturalism (Sober forthcoming). If intelligent design theory is formulated without specifying whether the postulated designer is a supernatural or a natural being, does it thereby obey methodological naturalism? I discuss this question in Sober 2007.
3. Here is a characteristic passage, from *Word and Object*: "What distinguishes between the ontological philosopher's concern and [the scientist's]...is only breadth of categories. Given physical objects in general, the natural scientist is the man to decide about wombats and unicorns. Given classes, or whatever other broad realm of objects the mathematician needs, it is for the mathematician to say...whether in particular…there are any cubic numbers that are sums of pairs of cubic numbers. …The philosopher's task differs from the others', then, in detail; but in no such drastic way as those suppose who imagine for the philosopher a vantage point outside the conceptual scheme that he takes in charge. There is no such cosmic exile" (Quine 1960, 275). It is curious that the alternative to naturalism$_p$ that Quine describes is one in which philosophers have a vantage point that is outside of the conceptual scheme that they in fact use. Surely rejecting naturalism$_p$ involves no such commitment.
4. Just as naturalism$_s$ does not entail metaphysical naturalism, so metaphysical naturalism does not entail naturalism$_p$. Metaphysical naturalists may want to appeal to parsimony, both in science and in philosophy, but it is a further question whether they should maintain that the same concept of parsimony applies in the two domains.

5. There may be good scientific arguments (and good ones using parsimony) that scientists have not thought of to date. And there may be philosophical arguments that appeal to parsimony that can be justified in terms of good principles of scientific reasoning, which scientists happen not to describe by using the term "parsimony" or its cognates. Naturalism_p is consistent with both possibilities.

6. See Sober 1988 (pp. 206-12) and Sober 2008 (pp. 278-83) for discussion of related arguments.

7. These examples concerning common and separate cause hypotheses are counterexamples to Lewis' (1973, p. 87) contention that parsimony is epistemically relevant only when it concerns *kinds* of causes, never when it pertains to the number of *token* causes.

8. The phylogenetic hypotheses that Edwards and Cavalli-Sforza (1964) considered all presuppose that the species we now observe trace back to common ancestors. Cavalli-Sforza and Edwards were not *testing* CA; rather, they were *presupposing* it. What they were testing were hypotheses concerning which species are more closely related to each other and which are related only more distantly.

9. Around the time that Edwards and Cavalli-Sforza published their paper, the work of Willi Hennig (1966) was translated into English and had a major impact on systematics. Cladists, as these followers of Hennig came to be called, regarded parsimony as the key to phylogenetic inference. They, like Edwards and Cavalli-Sforza, measured the parsimoniousness of a phylogenetic tree by counting the number of changes in character state that would have to take place to generate the data on tip species, but they did not regard parsimony as a surrogate for likelihood; more often, they sought to justify parsimony by linking it to Popperian ideas about falsifiability (see, for example, Eldredge and Cracraft 1980; Wiley 1981; Farris 1983; Sober 1988).

10. This is because $\Pr[(s_1 - s_2) = 4 \mid (m_1 - m_2) = 4] > \Pr[(s_1 - s_2) = 4 \mid (m_1 - m_2) = x]$, for any $x \neq 4$.

11. There are other criteria in statistics for evaluating models besides AIC and there is discussion about which criteria are best in which circumstances. Parsimony is part of all these approaches. One alternative is the Bayesian Information Criterion (BIC); its goal is estimating a model's average likelihood, not its predictive accuracy.

12. There is a third interpretation of "the evidence fails to discriminate between the two hypotheses." It might mean that "*X* exists" and "*X* does not exist" have the same *posterior* probability (namely, 0.5). The razor of silence is correct in this case, but now it is trivial.

13. Howard-Snyder (1996) is a useful anthology of recent work on the evidential problem of evil. Draper (1989) represents the argument in terms of likelihoods.

14. If the evil summary *E* is more probable under the hypothesis that an all-PKG God does not exist, are there other observations that tell in the opposite direction? It is here that various versions of the design argument need to be considered. See Sober (2008, chapter 2) for discussion of the *organismic* design argument. The fine-tuning argument is different; it claims that the

values of the physical constants that characterize our universe have a higher probability of falling in the narrow window that permits life to exist if an all-PKG God exists than would be the case if there was no such being. Perhaps *E* is evidence against the existence of God while the values of the physical constants are evidence for. I don't think so, since I think the fine-tuning argument is vitiated by an observation selection effect; see Sober (2009) for details.

15. Smart (1959, pp. 155-56) compares the identity theory and dualism with evolutionary theory, which postulates an ancient earth in which successive layers of fossils are gradually deposited, and the hypothesis "that the universe just began in 4004 BC, with sediment in the rivers, eroded cliffs, fossils in the rocks and so on."
16. The identity theory also entails that two individuals who are in the same c-fiber firing state are in the same pain state (and vice versa).
17. The experiments described here are of questionable morality. I describe them because the example of pain and c-fiber firing is standard in discussions of the identity theory.
18. How might a token identity theory and a token dualism be placed in the model selection framework? Again, there is an intuitive difference in parsimony. But how does the token identity theory constrain the data in a way that token dualism fails to do? And how can the idea of models' being fitted to old data and then predicting new data be applied to claims about token identity rather than type? Perhaps the parsimony argument needs to be understood in terms of type identity, not token.
19. What about parsimony as a surrogate for likelihoods? The problem is to figure out what value Pr(data |Dual) has. I see no way to do this that carries conviction.
20. For simplicity, I assume that this model says that *A* and *B* are related additively. To allow for a non-additive relationship, an additional parameter, an interaction term, would need to be introduced.
21. The supervenience thesis I'm considering is one of synchronic determination; it says that a system's complete physical state at time *t* determines its mental state at *t*. If cause must precede effect, supervenience, thus understood, is not a causal relation.
22. For the sake of simplicity, I treat screening-off as a relation among dichotomous propositions, rather than in terms of the values of continuous quantities. The latter would require shifting to equalities between expected values and would introduce irrelevant complications.
23. Notice that this version of epiphenomenalism is consistent with token and even type identities between the mental and the physical.
24. Harman holds that normative ethical statements are sometimes true, but when they are true, they are true because they reflect our inclinations to approve and disapprove of various actions; this makes Harman a *relativist*, not a *realist*.
25. In discussions of reductionism in the philosophy of mind, it is usually assumed that there is a true description of an organism's beliefs and desires

and another true description of its physical state; the question is which of these provides the better (or the only) explanation of behavior. We assume the descriptions are true, and then compare how explanatory they are. The situation is different when ethical realism is discussed. We don't assume that there is an independent realm of ethical facts and a set of descriptive facts about upbringing and then ask which provides the better explanation of behavior. Rather, the parsimony argument against realism asserts that we do not need to postulate such ethical facts to explain behavior and then concludes that no such facts exist.

26. This point about the "job" that normative propositions should be expected to perform applies to purely descriptive propositions as well. Consider the following spurious parsimony argument: "You don't need to postulate the existence of dinosaurs to explain why gold has the melting point it does. Therefore you have a reason to deny that dinosaurs exist."

27. Compare this with Ruse and Wilson's (1986, pp. 186-87) statement that "the evolutionary explanation makes the objective morality redundant, for even if external ethical premises did not exist, we would go on thinking about right and wrong in the way that we do. And surely, redundancy is the last predicate that an objective morality can possess."

28. If the "moral faculty" (i.e., the part of our minds that generates normative ethical beliefs) operated *entirely* independently of all other mental faculties, that might be a reason to think that its deliverances are entirely unconnected with whatever normative ethical facts there may be. However, if this faculty takes advice from others, the situation is more complicated.

29. This is a distributive holism; holism can also be given a nondistributive formulation (Sober 2000).

30. Here is a simple example that shows why. You are playing poker and wonder whether the card you are about to be dealt will be the Jack of Hearts. The dealer is a bit careless and so you catch a glimpse of the card on top of the deck before it is dealt to you. You see that it is red. The fact that it is red confirms the hypothesis that the card is the Jack of Hearts, not in the sense of proving that the card will be the Jack of Hearts, but in the sense of raising the probability that it will be. The hypothesis that the card will be the Jack of Hearts entails that the card will be a Jack. However, the fact that the card is red does not confirm the hypothesis that the card will be a Jack.

31. This investigation would need to recognize that mathematicians are aware that inferences of the form "conjecture G holds for the first million integers, so it holds for all of them" have often failed. Maybe number theorists are more wary of simplicity as a principle of inference than physicists are.

References

Akaike, H. 1973. "Information Theory as an Extension of the Maximum Likelihood Principle." In *Second International Symposium on Information Theory*, edited by B. Petrov and F. Csaki. 267-81. Budapest: Akademiai Kiado.

Alston, W. 1958. "Ontological Commitments." *Philosophical Studies* 9: 8-17.

Balaguer, M. 2001. *Platonism and Anti-Platonism in Mathematics*. Oxford: Oxford University Press.

Brandt, and Kim, J. 1967. “The Logic of the Identity Theory.” *Journal of Philosophy* 64: 515-37.

Brosnan, K. (unpublished). “Dissolving a dilemma: why Darwinian considerations don’t confront moral realism with hard choices.”

Darwin, C. 1859. *On the Origin of Species*. London: Murray.

Dennett, D. 1980. “True Believers – the Intentional Stance and Why it Works.” *Herbert Spenser Lecture on Scientific Explanation*. Oxford University Press. Expanded version in *The Intentional Stance*. Cambridge: MIT Press, 1989, pp.13-42.

Draper, P. 1989. “Pain and Pleasure: An Evidential Problem for Theists.” *Noûs* 23: 331-50. Reprinted in D. Howard-Snyder (1996), pp. 12-29.

Edwards, A.W.F. and Cavalli-Sforza, L. 1964. “Reconstruction of Evolutionary Trees.” In *Phenetic and Phylogenetic Classification*, edited by V. Heywood and J. McNeill. New York Systematics Association Publication 6: 67-76.

Eldredge, N. and Cracraft, J. 1980. *Phylogenetic Patterns and the Evolutionary Process*. New York: Columbia University Press.

Farris, J. S. 1983. “The Logical Basis of Phylogenetic Analysis.” In *Advances in Cladistics – Proceedings of the 2nd Annual Meeting of the Willi Hennig Society*, edited by N. Platnick and V. Funk. 7-36. New York: Columbia University Press. Reprinted in E. Sober (ed.), *Conceptual Issues in Evolutionary Biology*, Cambridge: MIT Press, 1994.

Forster, M. and Sober, E. 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: 1-36.

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

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

Hempel, C. 1965. “Studies in the Logic of Confirmation.” *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press.

Hennig, W. 1966. *Phylogenetic Systematics*. Urbana: University of Illinois Press.

Howard-Snyder, D. (ed.) 1996. *The Evidential Argument from Evil*. Bloomington, IN: Indiana University Press.

Jeffreys, H. 1961. *The Theory of Probability*. Oxford: Oxford University Press, 2nd edition.

Joyce, R. 2006. *The Evolution of Morality*. Boston: MIT Press.

Kim, J. 1993. *Supervenience and Mind*. New York: Cambridge University Press.

Kim, J. 1996. *Philosophy of Mind*. Boulder: Westview Press.

Lewis, D. 1973. *Counterfactuals*. Oxford: Basil Blackwell.

Mayr, E. 1982. *The Growth of Biological Thought*. Cambridge: Harvard University Press.

Putnam, H. 1971. *Philosophy of Logic*. Boston: Allen and Unwin.

Quine, W. 1953. "Two Dogmas of Empiricism." In *From a Logical Point of View*. 20-46. Cambridge: Harvard University Press.

Quine, W. 1960. *Word and Object*. Cambridge: MIT Press.

Quine, W. 1963. "Carnap and Logical Truth." In *The Philosophy of Rudolf Carnap*, edited by P.A. Schilpp. LaSalle, Illinois: Open Court.

Reichenbach, H. 1956. *The Direction of Time*. Berkeley: University of California Press.

Rowe, W. 1979. "The Problem of Evil and Some Varieties of Atheism." *American Philosophical Quarterly* 16: 335-341.

Ruse, M. and Wilson, E. 1986. "Moral Philosophy as Applied Science." *Philosophy* 61: 173-92. Reprinted in E. Sober (ed.), *Conceptual Issues in Evolutionary Biology*, MIT Press, 2006.

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: 311-29.

Shapiro, L. and Sober, E. 2007. "Epiphenomenalism – The Do's and the Don'ts." In *Thinking about Causes*, edited by P. Machamer and G. Wolters. 235-64. Pittsburgh: University of Pittsburgh Press.

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

Sober, E. 1988. *Reconstructing the Past – Parsimony, Evidence, and Inference*. MIT Press.

Sober, E. 1990. *Core Questions in Philosophy*. New York: Prentice-Hall.

Sober, E. 1993. "Mathematics and Indispensability." *Philosophical Review* 102: 35-58.

Sober, E. 1994. "Prospects for Evolutionary Ethics." In *From a Biological Point of View*. 93-113. Cambridge: Cambridge University Press.

Sober, E. 1999a. "The Multiple Realizability Argument against Reductionism." *Philosophy of Science* 66: 542-64.

Sober, E. 1999b. "Physicalism from a Probabilistic Point of View." *Philosophical Studies* 95: 135-74.

Sober, E. 2000. "Quine's Two Dogmas." *Proceedings of the Aristotlean Society*, Supplementary Volume 74: 237-80.

Sober, E. 2007. "Intelligent Design Theory and the Supernatural – the 'God or Extra-Terrestrials' Reply." *Faith and Philosophy* 24: 72-82.

Sober, E. 2008a. "Empiricism." In *The Routledge Companion to Philosophy of Science*, edited by S. Psillos and M. Curd. pp. 129-38.

Sober, E. 2008b. *Evidence and Evolution*. Cambridge: Cambridge University Press.

Sober, E. 2009. Absence of Evidence and Evidence of Absence – Evidential Transitivity in Connection with Fossils, Fishing, Fine-Tuning, and Firing Squads." *Philosophical Studies* 144: 63-90.

Sober, E. (forthcoming) "Evolution without Naturalism." In *Oxford Studies in Philosophy of Religion*, volume 3, edited by J. Kvanvig. Oxford: Oxford University Press.

Sorensen, R. 1988. *Blindspots*. Oxford: Clarendon Press.

Sturgeon, N. 1984. "Moral Explanations." In *Morality, Reason, and Truth*, edited by D. Copp and D. Zimmerman. 49-78. Totowa, NJ: Rowman and Allanheld.

Wiley, E. 1981. *Phylogenetics – the Theory and Practice of Phylogenetic Systematics*. New York: Wiley.

Wykstra, S. 1984. "The Humean Obstacle to Evidential Arguments from Suffering: On Avoiding the Evils of 'Appearance'." *International Journal for Philosophy of Religion* 16: 73-93.