Brit. J. Phil. Sci. **54** (2003), 245–259, axg011

Quantitative Parsimony and Explanatory Power

Alan Baker

ABSTRACT

The desire to minimize the number of individual new entities postulated is often referred to as *quantitative parsimony*. Its influence on the default hypotheses formulated by scientists seems undeniable. I argue that there is a wide class of cases for which the preference for quantitatively parsimonious hypotheses is demonstrably rational. The justification, in a nutshell, is that such hypotheses have greater explanatory power than less parsimonious alternatives. My analysis is restricted to a class of cases I shall refer to as *additive*. Such cases involve the postulation of a collection of qualitatively identical individual objects which collectively explain some particular observed phenomenon. Especially clear examples of this sort occur in particle physics.

1 Introduction

Scientists tend to be frugal in their postulation of new entities. When a trace is observed in a cloud-chamber, physicists may seek to explain it in terms of the influence of a hitherto unobserved particle. But, if possible, they will postulate one such unobserved particle, not two, or twenty, or 207 of them. This desire to minimize the number of individual new entities postulated is often referred to as *quantitative parsimony*. Its influence on the default hypotheses formulated by scientists seems undeniable. Yet relatively little attention has been paid by philosophers to whether—and if so how—a preference for quantitative parsimony might be justified on rational grounds. Is the initial assumption that one particle is acting to cause the observed trace more rational than the assumption that 207 particles are so acting? Or is it merely the product of wishful thinking, aesthetic bias, or some other non-rational influence?

The central claim of this paper is that the preference for quantitatively parsimonious hypotheses is demonstrably rational, for a certain generally characterizable class of cases. I will first outline an historical case study from particle physics, then present an analysis of the role of quantitative parsimony in this and similar cases, and finally make some connections between my analysis and broader themes raised by other philosophers' discussions of parsimony and simplicity. My goal is to plug a striking gap in the philosophical literature concerning the role of quantitative parsimony in the methodology of science.

2 Particle physics: a case study

Physicists in the 1930s were puzzled by certain anomalies arising from experiments in which radioactive atoms emit electrons during so-called Beta decay. In these experiments, the total mass-energy of the system of particles before Beta decay is greater than the total mass-energy of the observed particles that are emitted following the decay, and the total spin of the particles in the system before decay exceeds by $\frac{1}{2}$ the total spin of the observed particles emitted following the decay. Being unwilling to give up the laws of conservation of mass-energy or conservation of spin, scientists concluded that there were particles being emitted following Beta decay which had not been detected by their instruments. Their response was to posit a 'new' fundamental particle, the neutrino, with variable mass-energy and with spin $\frac{1}{2}$, and to hypothesize that exactly one neutrino is emitted by each electron during Beta decay.

However, as Daniel Nolan points out in a recent paper, there is a wide range of very similar neutrino theories which can also account for 'the missing mass-energy, the variation of the mass-energies of the emitted electrons, and the missing spin' ([1997], p. 333). (The issue is also discussed by Bunge ([1963]) and by Schlesinger ([1963]).) If we focus for the moment on explaining the missing spin, then the following series of alternative neutrino hypotheses can be straightforwardly constructed:

H_1 1 neutrino with a spin of $\frac{1}{2}$ is emitted in each case of Beta decay

H_2 2 neutrinos, each with a spin of $\frac{1}{4}$ are emitted in each case of Beta decay

H_3 3 neutrinos, each with a spin of $\frac{1}{6}$ are emitted in each case of Beta decay

and, more generally, for any positive integer n,

H_n n neutrinos, each with a spin of 1/2n are emitted in each case of Beta decay.

Each of these hypotheses adequately explains the observation of a missing $\frac{1}{2}$-spin following Beta decay. Yet the obvious default hypothesis, both intuitively and from the point of view of actual scientific practice, is that exactly 1 neutrino is emitted in each case. Can anything substantive be said in defense of this intuition?

3 Three kinds of simplicity

A distinction that is often made is that between syntactic simplicity (roughly, the number and complexity of hypotheses) and ontological simplicity (roughly, the number and complexity of things postulated). Call these two facets of simplicity *elegance* and *parsimony* respectively. Philosophers who have discussed parsimony typically make a further distinction between *qualitative* parsimony (roughly, the number of types of thing postulated) and *quantitative* parsimony (roughly, the number of individual things postulated). It is this latter concept—of quantitative parsimony—which is operative in the neutrino example. Hypothesis H_1, which attributes the missing $\frac{1}{2}$-spin to a single unobserved particle, postulates fewer individual things than its rivals.

For each of the three types of simplicity just mentioned, we can ask whether there are any substantive rational grounds for preferring hypotheses which are simpler in the corresponding sense. Philosophers who have pursued these questions have mostly focused on syntactic elegance and on qualitative parsimony. By contrast, quantitative parsimony has been largely ignored. One reason is that elegance and qualitative parsimony have generally been perceived as carrying more weight than quantitative parsimony when it comes to evaluating competing scientific hypotheses. Indeed a number of philosophers have expressed doubts about whether quantitative parsimony is a genuine theoretical virtue in science at all. A second, and related, reason is that many philosophers regard quantitative parsimony as irrelevant in the broader context of philosophical and metaphysical system-building. As traditionally construed, metaphysical theories are primarily concerned with the carving up of reality into various categories, as opposed to assessing how many things of each kind there are.[1] Thus David Lewis articulates the attitude of many metaphysicians when he writes,

[1] Contributing to this view of quantitative parsimony is the idea that establishing the number of individual things there are of a certain sort is a purely empirical matter, whereas establishing how many kinds of things there are is a project to which armchair philosophizing might in principle contribute.

> I subscribe to the general view that qualitative parsimony is good in a philosophical or empirical hypothesis; but I recognize no presumption whatever in favour of quantitative parsimony. ([1973], p. 87)

Partly for the above reasons, philosophers have tended to avoid analyzing the concept of quantitative parsimony further. Another reason is that the other two types of simplicity are easier to connect to independently recognizable theoretical virtues. Syntactic elegance in a theory tends to bring with it pragmatic advantages such as being more perspicuous, being easier to use and manipulate, and so on. Philosophers of science of a naturalistic bent are generally happy to take pragmatic considerations of this sort as rational grounds for discriminating between competing theories. Qualitative parsimony in a theory can be viewed as minimizing the number of 'new' kinds of entities and mechanisms which are postulated. This preference for old mechanisms may in turn be justified by the more general epistemological caution, or conservatism, which many philosophers take to be characteristic of normal scientific inquiry. Neither of these sorts of virtues automatically accompanies quantitative parsimony. The less quantitatively parsimonious 2-neutrino hypothesis is no less perspicuous, no more difficult to manipulate, than the rival 1-neutrino hypothesis. Nor does the 2-neutrino hypothesis postulate more new mechanisms. Its $\frac{1}{4}$-spin neutrino is a new kind of particle, as is the $\frac{1}{2}$-spin neutrino postulated by the more parsimonious 1-neutrino hypothesis. This lack of any obvious theoretical benefits accompanying quantitative parsimony contributes, I think, to philosophers' doubts concerning whether quantitative parsimony is a genuine theoretical virtue, either in science or in metaphysics.

Nolan ends his 1997 paper with a challenge to those interested in these issues concerning simplicity 'to work out why in general quantitative parsimony might be thought to be a good thing, and then see from there how wide its applicability is' ([1997], p. 342). This paper is intended as a partial answer to this challenge. My analysis in this paper is restricted to a class of cases I shall refer to as *additive*. Such cases involve the postulation of a collection of individual objects, qualitatively identical in the relevant respects, which collectively explain some particular observed phenomenon. The explanation is 'additive' in the sense that the overall phenomenon is explained by totaling the individual positive contributions of each object. I shall argue that in additive cases such as the neutrino case, it is rational to prefer quantitatively parsimonious hypotheses, not because quantitative parsimony is a primitive theoretical virtue, but because quantitative parsimony brings with it other independently recognized virtues. In particular, quantitative parsimony tends to increase the explanatory power of hypotheses compared to their less quantitatively parsimonious rivals.

4 Explanatory power

In the context of the neutrino case study, the general questions raised in the previous section boil down to the following: is there some rational justification for favoring H_1 over the competing hypotheses H_2, H_3, . . ., H_n other than just the brute fact that H_1 is more quantitatively parsimonious? There is nothing to choose between the various neutrino hypotheses on grounds of syntactic elegance. Nor do the hypotheses differ in their qualitative parsimony; each postulates exactly one new *kind* of particle.[2] Several other potentially relevant features are also on a par: for example, the predicates used in each hypothesis are equally 'natural',[3] and no one hypothesis is logically stronger than any other.

In conversations with scientists, I have encountered three arguments for preferring H_1 over its less parsimonious rival hypotheses in this sort of situation. However, I think that each of these arguments is flawed. One argument focuses on explanatory idleness. At first blush it seems that the extra neutrinos postulated by H_2, H_3, and the other less parsimonious hypotheses are explanatorily idle. For example, H_2 postulates that 2 neutrinos rather than 1 are emitted following each Beta decay. Doesn't this introduce an extra superfluous neutrino? If so, then this might offer grounds for preferring the 1-neutrino hypothesis, for example by introducing an injunction against postulating explanatorily idle entities. However, this objection is too quick, as Nolan ([1997], p. 339), Barnes ([2000], p. 355), and others have pointed out. Within the context of the explanation provided by H_2, neither of the neutrinos postulated is explanatorily idle; the $\frac{1}{4}$-spin of each neutrino is required to explain the overall missing $\frac{1}{2}$-spin.

A second argument relies on the historical background of the neutrino case. Prior to the postulation of neutrinos in Beta decay, other particles with spin $\frac{1}{2}$ (for example, electrons) were already known to exist, whereas particles with smaller fractional spins such as $\frac{1}{3}$, or $\frac{1}{4}$, or $\frac{1}{20}$ were not. Thus to postulate any spin value other than $\frac{1}{2}$ for the neutrino would have been to postulate a new *kind* of spin property, and this is what makes H_1 preferable to H_2. I have some sympathy with this line of reasoning, but I am doubtful that it is sufficient to explain away the alleged role of quantitative parsimony. Firstly, worries can be raised about the extent to which qualitative parsimony considerations actually favor $\frac{1}{2}$-spin neutrinos over $\frac{1}{4}$-spin neutrinos. In each case the neutrinos postulated are a novel kind of entity, thus qualitative parsimony at the *entity* level is neutral between the two hypotheses. It is only

[2] Whether we count the neutrinos postulated by each hypothesis as different kinds or not (because the spin is different in each case) does not affect this point. Either way, each individual hypothesis postulates exactly one new kind.

[3] For example, none of the alternative hypotheses use 'grue-like' or other gerrymandered or disjunctive predicates.

if qualitative parsimony also operates at the level of *properties* that a case can be made for favoring H_1 on this basis. Even if it does, there is a further question about whether $\frac{1}{2}$-spin is really a different *kind* of property than $\frac{1}{4}$-spin. Clearly spin is a different kind of property than mass-energy, or charge. But is each different fractional spin-value itself a separate property? (Imagine that we have measured many things which weigh 2 kg, but nothing that has weighed 2002 kg. When we come across an object of mass 2002 kg, are we attributing a new kind of property to it?) The issue of how to carve up kinds of properties seems to be even more problematic than how to carve up kinds of entities, and this threatens the clear-cut application of qualitative parsimony to cases of this sort.

Secondly, even if qualitative parsimony is operative in the neutrino example, it does not show that quantitative parsimony does not also play a role. While it was a feature of the historical postulation of neutrinos that other particles with spin $\frac{1}{2}$ were already known to exist, I suspect that scientists would still have preferred H_1 over less quantitatively parsimonious rival hypotheses even if no other $\frac{1}{2}$-spin particles had been discovered. Moreover, the prior discovery of other $\frac{1}{2}$-spin particles does nothing to explain other intuitions we have concerning the neutrino case. Compare hypothesis H_2, which postulates a pair of neutrinos each with spin $\frac{1}{4}$ in each case of Beta decay, and hypothesis H_{207}, which postulates 207 neutrinos each with spin $\frac{1}{514}$. H_2 seems clearly to be a more plausible hypothesis than H_{207}, yet each postulates a 'new' spin-value relative to the background theories at that time. Quantitative parsimony, but not qualitative parsimony, discriminates between these two hypotheses.

A third argument looks to inductive grounds for preferring H_1. According to this argument, what scientists are doing in the neutrino case, and in other similar cases, is choosing new hypotheses based partly on criteria that have been generated inductively from previous cases of theory choice. Choosing the most quantitatively parsimonious of the acceptable (and otherwise equally good) alternative hypotheses has tended to work in the past. Hence scientists continue to use this as a rule of thumb, and are justified in so doing on inductive grounds. One might try to bolster this point of view by considering a counterfactual world in which all the fundamental constituents of the universe exist in pairs. In such a 'pairwise' world, scientists might well prefer H_2 to H_1 in the neutrino case, and prefer pairwise hypotheses in general to their more parsimonious rivals. I find this third argument interesting but ultimately unsatisfactory, for two reasons. Firstly, one might legitimately wonder just how successful the choice of quantitatively parsimonious hypotheses has been; examples from chemistry spring to mind, such as oxygen molecules containing two atoms rather than one. Secondly, and more importantly, there remains the question of *why* the preference for quantitatively parsimonious hypotheses in science has been as successful as it has been.

The key, in my view, to explicating the role of quantitative parsimony in additive cases such as the neutrino example is to focus on the relative explanatory power of the alternative hypotheses, H_1, H_2, H_3, ... In one major respect these various hypotheses are explanatorily on a par: each entails (when conjoined with certain shared background assumptions) that there will be a missing $\frac{1}{2}$-spin following the Beta decay of an electron. This much seems clear. However, two authors who have discussed this issue in print, Daniel Nolan and Mario Bunge, both draw similar and unjustified conclusions from this fact. Nolan slides from the claim of explanatory equivalence in this specific—albeit important—respect, to the general claim that these alternative hypotheses 'explain the same phenomena equally comprehensively' ([1997], pp. 333–4). Bunge, in an earlier discussion of this same issue, also concludes that the alternative hypotheses are explanatorily equivalent. Indeed he uses this supposed equivalence as a basis for rejecting quantitative parsimony as a rationally defensible criterion of theory choice, claiming that

> there is no point in adopting any one of these more complicated hypotheses as long as we cannot distinguish experimentally among their consequences, and as long as they do not throw new light on the explanation of phenomena. ([1963], p. 108)

We have agreed that each of the competing neutrino hypotheses does explain the missing $\frac{1}{2}$-spin in Beta decay. Where both Nolan and Bunge go wrong, however, is in inferring from this that the explanatory power of the various hypotheses is therefore the same in all other respects. It is at this point that I part company with them; I shall argue that there are significant and relevant phenomena which are easier to explain when a more quantitatively parsimonious neutrino hypothesis is assumed.

My claim, in short, is that quantitatively parsimonious hypotheses allow the explanation of more things. This may seem implausible given our earlier remarks. After all, haven't we just agreed that each of the competing hypotheses, H_1, H_2, . . ., H_n, explains the observed missing $\frac{1}{2}$-spin in cases of Beta decay equally well? Given that spin is the only feature mentioned in any of the hypotheses, what other phenomena are there which they might potentially explain? To answer this question we need to broaden our focus. When neutrinos were first postulated in the 1930s, numerous experimental set-ups were being devised to explore the products of various kinds of particle decay. In none of these experiments did the total spin of the particles in the system before decay exceed the total spin of the observed particles emitted following the decay by any fraction smaller than $\frac{1}{2}$. Thus no cases of 'missing'

$\frac{1}{3}$-spin, or $\frac{1}{4}$-spin, or $\frac{1}{100}$-spin had been found. The absence of these smaller fractional spins was a phenomenon which competing neutrino hypotheses might potentially have helped to explain.

Imagine that we are trying to decide between the following two neutrino hypotheses:

H_1 1 neutrino with a spin of $\frac{1}{2}$ is emitted in each case of Beta decay

H_{10} 10 neutrinos, each with a spin of $\frac{1}{20}$, are emitted in each case of Beta decay

The principal background assumptions that I am taking to be operative in the neutrino case are that total spin is conserved for particles in a closed system, that spin quantity is a constant and fixed property of neutrinos (both in Beta decay and elsewhere), and that no otherwise undetectable particles with spin $\frac{1}{2}$ or less had been postulated prior to the Beta decay experiments.

Given that H_{10} postulates particles with spin $\frac{1}{20}$, we might ask ourselves: why has no experimental set-up yielded a 'missing' spin-value of $\frac{1}{20}$? My claim is that H_1 allows a better answer to this question than H_{10} does, for H_1 is consistent with a simple and parsimonious hypothesis which explains why no cases of missing $\frac{1}{20}$-spin have been observed. This hypothesis is that there exist no particles with spin $\frac{1}{20}$ (or less). In the case of H_{10}, this potential explanation is ruled out because H_{10} explicitly postulates particles with spin $\frac{1}{20}$. Of course, H_{10} is consistent with *other* hypotheses which explain the non-occurrence of missing $\frac{1}{20}$-spin. For example, one might conjoin to H_{10} the law that neutrinos are always emitted in groups of ten. However, doing this makes the overall explanation less syntactically simple, and hence less virtuous in other respects. It also raises the further question of *why* this neutrino-law obtains, especially if there is no independent motivation for it.

I should stress that I am *not* arguing that H_1 fully explains the non-observation of missing $\frac{1}{20}$-spin. For H_1 alone does not entail that no particles with spin $\frac{1}{20}$ exist. My claim is rather that H_1 can serve as the basis for a better explanation of this non-observation than H_{10} can, because H_{10} entails the existence of particles with spin $\frac{1}{20}$, and hence a satisfactory theory incorporating H_{10} must say something about why such particles never seem to be emitted singly. Nor am I claiming that H_1 leaves nothing further to be explained. We might ask, for example, *why* there exist no particles with spin $\frac{1}{20}$ or less. Intuitively, however, it seems less problematic to take a non-existence claim of this sort as a brute fact than it does in the case of the claim that neutrinos are always emitted in groups of ten. Indeed I would claim that our natural attitude in the general case is to assume *prima facie* that a given kind of thing does not exist if its existence is not entailed by our best available theories and observations. (Why have there been no well-corroborated sightings of unicorns? Our favored explanation is that unicorns do not exist.

This explanation is consistent with our best available physical and biological theories, but it is not entailed by them.)

5 Explanation and non-observation

My argument linking quantitative parsimony to explanatory power gets off the ground only if explaining phenomena such as the non-observation of missing $\frac{1}{20}$-spin is of value in the context of scientific inquiry. Clearly we do not want to assume that arbitrary non-observations are all potentially in need of explanation. (It would be peculiar, for example, to seek an explanation for why physicists do not see tiny goblins during any of their Beta-decay experiments.) The point about the non-observation of missing $\frac{1}{20}$-spin is that the possibility of individual $\frac{1}{20}$-spin emissions is naturally suggested by the H_{10} neutrino hypothesis. This less quantitatively parsimonious hypothesis raises a question that is not raised by the most parsimonious hypothesis, H_1. If neutrinos with spin $\frac{1}{20}$ are emitted during Beta decay, then how come they are never emitted singly in other experimental set-ups?

H_{10} raises a question about why missing $\frac{1}{20}$ spins are not observed. But why assume that an explanation on theoretical grounds is required? Surely there may be a more mundane, practical explanation available for why smaller fractional missing spins were not observed, namely that our instruments are not (or were not in the 1930s) sensitive enough to detect such minute fractional spins. In this case the non-observation of missing $\frac{1}{20}$-spin, for example, can be explained by reference to the accuracy of our instruments and the sensitivity of our methods of detection. Since no 'theoretical' explanation is then required, this would undermine our stated reason for preferring H_1 over H_{10}.

My reply to this objection is twofold. First, in most actual cases the observed phenomenon does not occur right at the limits of experimental detection. In other words, this escape clause will not be available in all cases of comparative quantitative parsimony. Second, the non-observation of single-neutrino emissions (e.g. a missing spin of $\frac{1}{20}$ in the case of hypothesis H_{10}) is not the only phenomenon which the less quantitatively parsimonious neutrino hypotheses fail to explain. They also fail to explain the non-observation of missing spin corresponding to various multiple-neutrino emissions. Why, for example, did no experiments yield missing spins of $\frac{9}{20}$? H_1 allows a simple explanation for this fact; neutrinos are the only hitherto undetected subatomic particles that have spin, and each neutrino has spin $\frac{1}{2}$. Hence no combination of neutrinos can yield a missing spin of $\frac{9}{20}$. H_{10} rules out any such straightforward explanation; if each neutrino has a spin of $\frac{1}{20}$ then why (in principle) couldn't a cluster of 9 neutrinos yield a total missing spin of $\frac{9}{20}$? The same point applies to other fractional spins such as $\frac{11}{20}$, $\frac{12}{20}$, $\frac{13}{20}$, and so on. In general, the non-observation of fractional spins of the form $\frac{x}{20}$ which cannot be expressed as an equivalent fraction, $\frac{y}{2}$ (x and y integers), is explained by H_1 but not by H_{10}. The non-observation of fractional spins such as $\frac{1}{9}$ and $\frac{2}{7}$, whose denominators are not divisible by 2, is explained both by H_1 and by H_{10}. And the non-observation of fractional spins with a denominator of 2, such as $\frac{3}{2}$ and $\frac{19}{2}$, is explained by neither hypothesis. There are no non-observations of fractional spins which are explained by H_{10} but not by H_1. Thus the set of fractional spins explained by H_{10} is a proper subset of the set explained by H_1. This fleshes out my earlier claim that H_1, the most quantitatively parsimonious of the alternative hypotheses, has the greatest explanatory power.

Note that it is crucial for my analysis that spin be a fixed quantity for all neutrinos. As was mentioned earlier, the experiments on Beta decay also revealed 'missing' mass-energy. Since mass-energy was also assumed to be governed by a conservation law, another task was to explain this observation. However, unlike spin, the missing mass-energy varied seemingly continuously across a range of values. The explanation—as was eventually discovered—is that neutrinos themselves have variable mass-energy. The postulation of neutrinos in Beta decay explains this missing mass-energy, but H_1 does not explain the various observations and non-observations of missing mass-energy any better or worse than any of the other less parsimonious hypotheses. Each hypothesis allows for a continuous range of values of missing mass-energy, hence their explanatory power—in this respect—is on a par. Thus if missing spin were not also a feature which needed to be explained, there would perhaps be no rational grounds for preferring H_1 over less quantitatively parsimonious alternative hypotheses.

My analysis of the role of quantitative parsimony in additive cases is born out if we make certain changes to the neutrino example. Assume, for example, that (contrary to historical fact) undetectable particles—call them 'p-particles' and 'q-particles'—with spin $\frac{1}{2}$ and spin $\frac{1}{20}$ respectively had already been postulated prior to the experiments involving Beta decay. Thus, according to the best available theories at the time, both p-particles with spin $\frac{1}{2}$ and q-particles with spin $\frac{1}{20}$ exist. In this situation, the following two hypotheses would naturally suggest themselves:

H_1^* 1 p-particle with a spin of $\frac{1}{2}$ is emitted in each case of Beta decay

H_{10}^* 10 q-particles, each with a spin of $\frac{1}{20}$, are emitted in each case of Beta decay

Neither H_1^* nor H_{10}^* postulates a new kind of particle, but H_1^* is more quantitatively parsimonious than H_{10}^*. Nonetheless, I would argue that our intuitive preference for H_1^* over H_{10}^* is much weaker than for H_1 over H_{10} in the standard neutrino example, if indeed we would have any significant preference at all. My proposed analysis of quantitative parsimony helps to explain its failure to motivate a preference for H_1^* in this example. The point is that regardless of which of the above hypotheses is adopted to explain Beta decay, the overall best theory will still be ontologically committed both to p-particles and to q-particles. Hence our failure to observe missing spins of some given fractional value will not be easier to explain using H_1^* rather than H_{10}^*. In general, when it comes to a choice between 'old' mechanisms, it is unclear what—if anything—might provide rational grounds for favoring the more quantitatively parsimonious hypothesis.

6 Parsimony and scientific methodology

As I mentioned in my discussion of simplicity in Section 3, little attention has been paid in the philosophical literature to quantitative parsimony *per se*. Discussions have typically either focused specifically on qualitative parsimony, or have run together qualitative and quantitative versions under some more general parsimony principle. In a useful recent summary of the philosophical literature, E. C. Barnes ([2000], pp. 355–7) identifies five broad approaches to the rational justification of parsimony. I shall argue that three of these five approaches simply fail to apply in the case of additive quantitative parsimony. Of the remaining two approaches, my own analysis is closest in spirit to that of Elliott Sober ([1994]), although there are also important distinctions to be drawn between the two.

Three of the approaches which Barnes mentions are what he terms the 'pragmatic justification', attributed to Quine ([1966]) and to Walsh ([1979]); the 'unification justification', attributed to Friedman ([1983]), and the 'anti-free parameters justification', attributed to Lange ([1995]), and to Forster and Sober ([1994]). I have already argued, in Section 3, that the alternative hypotheses in the neutrino example do not vary from a pragmatic point of view, and on similar grounds I doubt that any difference in their relative unificatory power could be established. As for the third approach, although H_{10} postulates more individual things than H_1, it does not have more free, or adjustable, parameters. Each of the alternative hypotheses postulates a different number of neutrinos, but in each case all the neutrinos postulated have the same, fixed spin-value. Hence there is only one free parameter in each hypothesis, and the 'anti-free parameters justification' does not apply.

The remaining two approaches mentioned by Barnes are, I shall argue, each applicable (at least in principle) to cases involving quantitative parsimony. One is the 'general background knowledge' justification, which traces back at least as far as Newton. He famously included a parsimony principle in the introduction to the *Principia mathematica*, justifying it with the claim that 'Nature is pleased with simplicity and affects not the pomp of superfluous causes.' Newton's reasons for making this claim were as much theological as scientific, but the assumption that the world is naturally parsimonious is certainly one way to motivate the preference for theories which themselves exhibit parsimony. However modern philosophers have been understandably reluctant to follow Newton down this path, preferring to see parsimony as grounded in general methodological principles rather than in substantive empirical assumptions about the world.

The final approach, the 'local background knowledge justification', is also due to Elliott Sober ([1988], [1994]). As the label suggests, Sober rejects the Newtonian demand for a global assumption of parsimony in nature. However, he argues that appeals to parsimony always depend on specific factual assumptions for their rational justification. Thus he writes:

> The legitimacy of parsimony stands or falls, in a particular research context, on subject matter specific (and *a posteriori*) considerations. [. . .] What makes parsimony reasonable in one context may have nothing in common with why it matters in another. ([1994], p. 141)

This remark seems potentially applicable to quantitative parsimony also. However, while I support the claim that the specific justification for appeals to quantitative parsimony may vary from context to context, I conjecture that 'subject matter specific considerations' may be neither necessary nor sufficient to legitimize such appeals.

Let us first review the background assumptions operative in the neutrino case. Principal among them are the assumption that total spin is conserved for particles in a closed system, the assumption that spin quantity is a constant and fixed property of neutrinos, and the assumption that no otherwise undetectable particles with spin $\frac{1}{2}$ or less had been postulated prior to the Beta decay experiments. I want to argue two points concerning these assumptions. The first is that structurally similar assumptions will tend to operate in almost any example involving additive quantitative parsimony, and hence there is an important sense in which such assumptions are not 'subject matter specific'. Take the assumption that total spin is conserved. Without this assumption, there is nothing for the alternative neutrino hypotheses to explain; we could simply take the missing $\frac{1}{2}$-spin as something that is lost during Beta decay. As for the assumption that neutrinos have fixed spin, this simply narrows the array of alternative hypotheses that are being considered as potential explanations of the 'missing' $\frac{1}{2}$-spin, and ensures that these hypotheses are additive in the sense defined in Section 1. Finally, the assumption that no undetectable particles with fractional spin had previously been postulated implies that all the alternative neutrino hypotheses involve new mechanisms. Taken together, these assumptions limit the choice in the Beta decay example to competing explanations which are additive and which postulate new mechanisms. Of course, these assumptions are couched in subject matter specific language. However my claim is that their role is not subject matter specific. *Any* situation in which we face a range of alternative additive hypotheses differing only in their quantitative parsimony will have functionally analogous background assumptions in place.

This brings me to my second point against Sober's view. My claim is that the background assumptions in the neutrino example do not in themselves *justify* appeal to quantitative parsimony. In other words, these assumptions do not distinguish between the alternative neutrino hypotheses, H_1, H_2, H_3, ... The use of quantitative parsimony in additive cases such as this one is not dependent on the specifics of the subject matter. For example, no assumptions are made about the likelihoods of different sorts of subatomic particle emissions in Beta decay, or about the range of values for neutrino spin. In summary, the implication of my analysis is that *any* situation involving competing additive hypotheses is one in which appeal to quantitative parsimony considerations is rationally justified, other things being equal.

The second part of Sober's claim is that the justification of parsimony requires a piecemeal approach. My analysis indicates that this also holds true specifically for quantitative parsimony, since the analysis applies only to additive cases.[4] As we shall see, not all cases involving differences in quantitative parsimony are additive in this way. Imagine, for example, that perturbations are observed in the orbit of some planet which indicates that it is being acted upon gravitationally by other planets which have not previously been detected. It may turn out that the perturbations can be explained by postulating a single hitherto unobserved planet. If so, then they can also be explained by postulating 2 unobserved planets, or 3 planets, or 107 planets, whose gravitational forces act collectively to produce the observed anomalies. Should we prefer the most quantitatively parsimonious, 1-planet hypothesis in this case, and—if so—why? I think that the answer to this question is 'yes', and that an account can be developed which justifies favoring the most quantitatively parsimonious hypothesis in these sorts of cases. However, there are at least two important differences between the planets case and the neutrino case. First, planets vary greatly in size, mass, composition and so on, hence there is no reason to think that the various planets postulated will be qualitatively similar, as in the neutrino example. Second, the gravitational forces of several planets may partially or completely cancel one another out. It is not merely a question of summing up several

[4] Recall that the neutrino case study involves the postulation of a number of qualitatively identical individual particles which collectively explain some particular observed phenomenon. The explanation is additive in the sense that the overall phenomenon is explained by summing the individual positive contributions of each particle.

identical contributions. Hence the planets case is not additive. For these reasons, a rational justification of valuing quantitative parsimony in such cases will require a substantively different approach to that outlined in this paper.

7 Conclusions

To claim that quantitative parsimony is a theoretical virtue is to imply at minimum that quantitatively parsimonious hypotheses are to be preferred over their less quantitatively parsimonious rivals, other things being equal. My aim in this paper has been to go some way towards implementing the project which Nolan articulates at the end of his 1997 paper, 'to work out why in general quantitative parsimony might be thought to be a good thing, and then see from there how wide its applicability is' ([1997], p. 342). I have argued that rational grounds can be given for preferring more quantitatively parsimonious hypotheses, at least in additive contexts, but that this is because other things are *not* equal. Quantitative parsimony tends to bring with it greater explanatory power. Less quantitatively parsimonious hypotheses can match this power only by adding auxiliary claims which decrease their syntactic simplicity. Thus the preference for quantitatively parsimonious hypotheses emerges as one facet of a more general preference for hypotheses with greater explanatory power.

Department of Philosophy
Xavier University
Cincinnati, OH 45207
United States
bakerar@xavier.edu

References

Barnes, E. C. [2000]: 'Ockham's Razor and the Anti-Superfluity Principle', *Erkenntnis*, **53**, pp. 353–74.

Bunge, M. [1963]: *The Myth of Simplicity: Problems in Scientific Philosophy*, Englewood Cliffs: Prentice Hall.

Forster, M. and E. Sober [1994]: 'How To Tell When Simpler, More Unified, Or Less Ad Hoc Theories Will Provide More Accurate Predictions', *British Journal for the Philosophy of Science*, **45**, pp. 1–35.

Friedman, M. [1983]: *Foundations of Space-Time Theories*, Princeton: Princeton University Press.

Lange, M. [1995]: 'Spearman's Principle', *British Journal for the Philosophy of Science*, **46**, pp. 503–52.

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

Nolan, D. [1997]: 'Quantitative Parsimony', *British Journal for the Philosophy of Science*, **48**, pp. 329–43.

Nolan, D. [1999]: 'Is Fertility Virtuous In Its Own Right?', *British Journal for the Philosophy of Science*, **50**, pp. 265–82.

Quine, W. [1966]: 'On Simple Theories of a Complex World', in W. Quine (*ed.*), *The Ways of Paradox*, New York: Random House.

Schlesinger, G. [1963]: *Method in the Physical Sciences*, London: Routledge.

Sober, E. [1988]: *Reconstructing the Past: Parsimony, Evolution and Inference*, Cambridge, MA: MIT Press.

Sober, E. [1994]: 'Let's Razor Ockham's Razor', in *From A Biological Point of View*, Cambridge: Cambridge University Press, pp. 136–57.

Walsh, D. [1979]: 'Occam's Razor: A Principle Of Intellectual Elegance', *American Philosophical Quarterly*, **16**, pp. 241–4.