Australasian Journal of Philosophy
Vol. 79, No. 3, pp. 313–332; September 2001

AXIOLOGICAL ATOMISM

Graham Oddie

Virtue is good. Furthermore, it is better for the virtuous to be happy than to be miserable. Since the addition of happiness to virtue makes a positive difference to value it seems to follow that happiness itself is valuable. But even though happiness is a valuable addition to virtue it is not clear that happiness always and everywhere adds value. Many (famously Kant) have denied that it is better for the vicious to be happy than for them to be miserable. Rewarding viciousness with happiness makes a bad situation worse. So even if happiness is valuable it does not invariably add value.

This little collection of claims about virtue and happiness, if correct, appears to be an instance of the organic unity of value. In general we have organic unity where a complex whole is not simply the sum of its parts. Value exhibits organic unity if the value of a complex, whether a complex state or complex quality, is greater or less than the sum of the values of its components or parts. Even if virtue and happiness both possess value, then it seems that the value of a complex whole involving combinations of virtue or vice, happiness or misery, is not invariably the sum of the values of its components. Rather, these components interact in such a way that components of positive value do not always add value.

Whether or not value is additive might be thought to be of purely metaphysical interest. It certainly is of deep metaphysical interest, but it is also connected with important aspects of evaluative reasoning. Additivity is closely connected with principles of bare difference and separability which are often tacitly assumed in value theory. I begin by spelling out these principles and trace their connections with additivity and organic value. I then develop an unpleasant paradox of separability and additivity. Additivity and separability apparently entail nihilism: that nothing is more valuable than anything else. Additivity involves a kind of axiological atomism—that complexes decompose into components or *factors*; that these factors possess value independently of their role in valuable complexes; and that the factors do not interact in their production of overall value. In order to avoid the paradox it seems as though the factors have to be akin to the metaphysically privileged states of logical atomism—a doctrine that has not enjoyed widespread support since the early part of this century.

The paradox poses a problem not only for the notions of organic unity and additivity, but also for the closely related bare-difference principles which lie at the heart of value theory and of its application. I propose a way of eliminating the paradox, and thereby saving additivity and separability, without presupposing an unpalatable variant of logical atomism. I close with a proposal about how to regard and work with the principles of additivity and of separability.

I. Separability

At the beginning of the *Metaphysical Foundations of Morals*, Kant tries to show that the only thing which is 'good without qualification', of 'intrinsic, unconditional worth' or 'good for itself' is a good will. The general form of Kant's argument is this. Take something which we might be inclined to think is intrinsically good, like happiness or prosperity. We can imagine pairs of situations which differ solely in the presence and absence of happiness or prosperity, and in at least some of these the presence of happiness or prosperity makes things worse.

> It need hardly be mentioned that the sight of a being adorned with no feature of a pure and good will, yet enjoying uninterrupted prosperity, can never give pleasure to a rational impartial observer. (Kant [9], pp. 9–10.)

The same kind of argument applies to particular features like courage or coolness:

> Moderation in emotions and passions, self-control, and calm deliberation not only are good in many respects but even seem to constitute a part of the inner worth of the person. But however unconditionally they were esteemed by the ancients, they are far from being good without qualification. For without the principle of a good will they can become extremely bad, and the coolness of a villain makes him not only far more dangerous but also more directly abominable in our eyes than he would have seemed without it. (Kant [9], p. 10.)

The assumed principle here seems to be this: if a feature is a good one, then if two wholes differ only over the presence of that feature, the complex whole which contains the feature must be better than the complex whole which lacks it. Let's call complexes which differ over the presence of just one feature *barely different*. Because it spells out a way in which the value of individual parts contributes to the relative value of barely different wholes, I will call the principle *bottom-up bare-difference*.[1]

To state the principle with a modicum of generality, let me introduce some useful concepts as well as some notation. We can think of wholes as involving a number of *factors*, like virtue and happiness. We can think of factors as determinables. So being *virtuous* and *vicious* (or degrees of virtue if there are such) can be thought of as determinate features of a determinable factor, a factor which I will simply dub the *virtue factor*. (I am not, of course, identifying *being virtuous* (a determinate) with the virtue *factor* (a determinable).) Similarly, being *happy* or *miserable* (or degrees of happiness) are the determinates of a factor which I will dub the *happiness factor*. Each factor is thus a collection of jointly exhaustive and mutually exclusive features. Features are *rivals* if they are incompatible determinates of such a determinable factor. Thus being happy and being miserable are rivals of the happiness factor, being virtuous and being vicious are rivals of the virtue factor.

[1] The *method* of bare differences was so-called in Rachels [20], pp. 111–14. Rachels, however, does not clearly distinguish the bottom-up and top-down principles.

The kinds of complex whole we are interested in are those which, given a range of relevant factors, specify a determinate feature for each factor. We can think of these complexes as conjunctions of determinates. If a complex W specifies a feature (*happy*, say) from some factor then that complex is said to be *over* that factor. Where W is a complex over the happiness factor, let W[*happy*] be the complex which differs from W in at most that *happy* replaces whatever is its determinate rival in W. So, if W is the complex which specifies being both miserable and virtuous, then W[*happy*] is the complex which specifies being happy and virtuous. Where W is the complex which specifies being happy and virtuous, then W[*happy*] is just W itself. In general, where W is a complex over the factor containing f, W[f] is the complex which differs from W in at most that it specifies f in the appropriate factor. Then W[f] and W[f′] are a pair of *barely different* wholes which differ only in that W[f] specifies f whereas W[f′] specifies f′. If W and V are distinct complexes over the same factors then they are *complex rivals*.

We can now state the principle of bottom-up bare-difference:

Bottom-up bare-difference.

> For any feature f of factor F, if f is better (in itself) than a rival f′ then (for every complex W over F), W[f] is better than W[f′].

Abbreviate *is better than* to > and the principle can be succinctly stated thus (where it is obvious, I will omit ‘over factor F’ and related terminology):

> If f > f′ then, for every W, W[f] > W[f′].

A similar principle takes us from evaluatively equivalent parts (≈) to evaluatively equivalent barely different wholes:

> If f ≈ f′ then, for every W, W[f] ≈ W[f′].

It will be useful to consider a range of objections to additivity which, if sound, would also impugn bottom-up bare-difference. Consider the following:

U Sam is sad **and** Pam is sad

We can easily relativise a factor (say the happiness factor) to individuals. We can entertain the Sam-happiness factor, each determinate of which will specify a level of happiness for Sam ($happy^S$, sad^S), and likewise the Pam-happiness factor. Now consider U[$happy^S$], *viz*.

> Sam is happy **and** Pam is sad.

If bottom-up is correct, and $happy^S$ is better than sad^S, then U[$happy^S$] is better than U and that seems right. Now consider the following pairs:

A	*Sam is not sad **or** Pam is happy*	A*	*Sam is not happy **or** Pam is happy*
B	*Sam is sad **that** Pam is sad*	B*	*Sam is happy **that** Pam is sad*

A* is arguably worse than A. If A* is A[*happy*^{S}] then either bottom-up is false, or *happy*^{S} is not better than *sad*^{S}. Someone committed to the intrinsic value of happiness might regard this as an argument against bottom-up. Similarly, B* is arguably worse than B, and we face the same dilemma if B* is just B[*happy*^{S}]. However, A and A* don't pick out complex rivals since they are clearly not incompatible specifications of determinate features. As such our account of complexes does not tell us that A[*happy*^{S}] is A*. B*, on the other hand, is incompatible with B, and could be considered a complex rival of it. However, it specifies a determinate of a factor distinct from both Sam-happiness and Pam-happiness. For each proposition *P* we can consider the factor: *Sam's degree of happiness that P*. A and A* are rival determinates of such a factor where *P* is the proposition that *Pam is sad*. And the intuition about the ranking of those two determinates is quite compatible with bottom-up.

In Book X of the *Nichomachean Ethics* Aristotle sketches an argument, which he attributes to Plato, for the thesis that pleasure is not the (only) good.

> … a pleasant life … is the more desirable when combined with practical wisdom than without it; but if pleasure is better in combination with something else, it is not the good, since the good cannot become better by the addition of something to it. (Aristotle [1], pp 27–33.)

The implicit assumption here is that if the conjunction of wisdom and pleasure is better than the conjunction of stupidity and pleasure then wisdom possesses a value in itself, a value which is independent of, and separable from, whatever value pleasure possesses.[2] This principle is the converse of the one assumed by Kant. It tells us that given just one pair of wholes which differ only in respect of the presence or absence of a feature, if the former as a whole is better than the latter as a whole, then the presence of the feature is, in itself, better than its absence. Because the principle enables us to establish evaluative relations between components on the basis of the evaluative relations between wholes of which they are parts, I will call it *top-down* bare-difference.

Top-down bare-difference

> For any feature f and rival f′, if there is a W such that W[f] is better than W[f′], then f is better (in itself) than f′.

The principle can be succinctly stated thus:

> If, for some W, $W[f] > W[f']$, then, $f > f'$.

There is a corresponding principle governing equivalence in value:

> If, for some W, $W[f] \approx W[f']$, then, $f \approx f'$.

2 Lemos [10], p. 32, refers to this argument. Lemos, however, draws out a somewhat different principle.

If we put these two bare-difference principles together we get a powerful combined principle which I will call, simply, *separability.*[3]

We could approach the notion of separability from a slightly different direction. Kant suggests that goodness or betterness comes in both conditional and unconditional forms. Happiness is better than misery given (conditional upon) the presence of a good will. But misery is better than happiness, given the absence of a good will. So the ranking of happiness and misery varies according to the condition given. In general, let us indicate a ranking which is conditional upon holding fixed some feature g thus: $>_g$. Then we have two conditional betterness relations ($>_{virtue}$ and $>_{vice}$) the first conditional upon a good will (or virtue), the other conditional upon lack of good will (or vice). To capture an unconditional ranking we consider the relative merits of features f and f′ given various contexts. If every such conditional ranking gives the same result— say, $f >_g f'$, for every feature g— then f is unconditionally better than f′. Kant's thesis is that a good will is unconditionally good (unconditionally better than its rivals) and moreover it is the only such unconditional good. It is highly plausible that virtue is better *in itself* than vice just in case virtue is *unconditionally* better than vice. Thus Kant's apparent assumption, that the concepts *good in itself*, *unconditionally good* and *intrinsically good* are equivalent, can be vindicated.

The notion of a conditional ranking can be used to capture the concept of separability of an individual factor. A factor is separable if each feature of that factor makes the same contribution to value regardless of the context in which it is embedded. That is to say, if the features are ordered thus-and-so while holding other things fixed in a certain way, then they are so ordered no matter what way other things are held fixed. In other words, the factor does not interact with other factors in the production of value. Now suppose wholes can be decomposed into a range of factors. The collection of complexes as a whole satisfies separability just in case all the factors are individually separable.

Separability is a powerful and historically important principle. It reaches deep into our thinking about value. It underlies ordinary everyday talk about good and bad features as well as the ethicist's talk of *good-making* features. The fact that separability is tantamount to the conjunction of bottom-up and top-down bare-difference principles shows that it is presupposed in applications of the method of bare differences. For a recent and important example of the methodology of bare difference, consider James Rachels' seminal contribution to the debate over the relative value of killing and letting-die. Rachels shows that there are pairs of situations which differ only in that one of them is a killing and the other a letting-die, but which nevertheless seem evaluatively equivalent (Rachels [20]). An application of top-down bare difference yields the result that letting-die has the same value as killing. And then an application of bottom-up bare difference, yields the result that two members of any such barely different pair must have the same value.[4] Quite

[3] The term *separability* has been coined and rigorously defined by economists, although the results achieved concerning separability are not widely appreciated by philosophers of value theory. For an excellent introduction both to the mathematical theory and some of its important applications, see Broome [4].

[4] Alternatively, by bottom-up it follows (by a reductio) that killing is not, in itself, worse than letting-die and neither is letting-die worse than killing. Now assuming that killing and letting-die are connected by the better-than relation—either one is intrinsically better than the other or else they have the same intrinsic value—it follows that they are of the same value. Quite generally, connectedness together with bottom-up yields top-down. See Oddie [16].

generally, by top-down, a single pair of situations which differ over a pair of rival features yields a definitive evaluative ordering of those features. And by bottom-up, this ordering then extends to every pair of situations which differ only over that pair of rival features.[5]

Separability is obviously an enormously powerful tool for testing, systematising and extending evaluative judgements.

II. Additivity

What is it for an evaluative ordering to be additive? The notion of additivity is straightforward if we have, in addition to parts and wholes, a numerical measure of the value of both parts and wholes. But suppose that instead of such a numerical assignment all we have are non-numerical comparisons of value?[6] In the absence of numbers is there any way of cashing out talk about additivity or of organic unity?

Given a non-numerical evaluative ordering, the assumption of additivity clearly involves the possibility of representing the ordering numerically as additive. A *representation* of an evaluative ordering is an assignment of numbers to the elements (both wholes and parts) which preserves the ordering—the more valuable, the greater the numerical value. Such a representation is *additive* if the value of each whole is the sum of the values of its parts. So an ordering is additive just in case it is susceptible to an additive representation. Otherwise it is not additive, and exhibits organic unity. Note that even if the ordering only ranks wholes, without directly ranking their parts, we can still apply this account of additivity.

We can now sketch an argument to show that additivity entails both bottom-up and top-down bare difference. Suppose we have a collection of wholes each of which involves a factor F together with some other factors, which we will collectively label G. It will simplify things, without loss of generality, if we think of G as a single factor. So each complex whole is composed of a determinate realisation of the factor F and a determinate realisation of the (complex) factor G. Assume we have an ordering > both of individual features and of the complex wholes, and assume further that this ordering is additive. That is to say, there is an additive representation (call it Val) which, firstly, assigns numerical weights to wholes and their parts such that the value of each whole is the sum of the values of its parts, and which, secondly, replicates the qualitative ordering of both wholes and parts.

To demonstrate bottom-up we have to show that for any feature f, if f is better in itself than a rival f′, then replacing f′ with f, leaving everything else the same, will improve things. So assume $f > f'$. By the fact that Val is a representation of the ordering, it follows that $Val(f) > Val(f')$. Now consider the barely different wholes W[f′] and W[f]. These do not differ over the G-factor, so let g be that determinate realisation of G which W[f] and W[f′] have in common.

[5] Throughout I will assume that the features may be complex as well as simple, thus yielding *strong* rather than *weak* separability. See Broome [4], p. 69.

[6] For discussion of this and related problems see Harman [7], Quinn [19], Oldfield [18], Carlson [5], Danielsson [6] and Oddie [17].

By additivity:	Val(W[f]) = Val(f) + Val(g) and Val(W[f′]) = Val(f′) + Val(g).
Since Val(f) > Val(f′):	Val(f) + Val(g) > Val(f′) + Val(g).
And so:	Val(W[f]) > Val(W[f′]).
By representation:	W[f] > W[f′], as desired.

Now for top-down: for any rival features f and f′, if substitution of f for f′ can improve things then f is better (in itself) than f′. So assume W[f] > W[f′], for some whole W, and again let g be the realisation of G that W[f] and W[f′] have in common.

By representation:	Val(W[f]) > Val(W[f′]).
By additivity:	Val(W[f]) = Val(f) + Val(g) and Val(W[f′]) = Val(f′) + Val(g).
Hence:	Val(f) + Val(g) > Val(f′) + Val(g) .
Thus:	Val(f) > Val(f′).
By representation:	f > f′.

We can illustrate the entailment by considering some rivals to Kant's ordering of complexes involving the virtue factor (**v** = virtuous, **v′** = vicious) and the happiness factor (**h** = happy, **h′** = miserable).

	virtue factor		*happiness factor*
$\mathbf{W}_1$	**v**	&	**h**
$\mathbf{W}_2$	**v**	&	**h′**
$\mathbf{W}_3$	**v′**	&	**h**
$\mathbf{W}_4$	**v′**	&	**h′**

The hedonist holds that happiness is the only unconditional good—virtue and vice are at best means to the end of happiness. The even-handed pluralist thinks that both virtue and happiness are good in themselves, and that losing one is as bad as losing the other. The virtuous pluralist thinks that both have value but that it is worse to lose your virtue than your happiness, and the happy pluralist thinks it worse to lose your happiness. Where ≈ is value equivalence, we can summarise the information thus:

Hedonist:	$\mathbf{W}_1 \approx \mathbf{W}_3 > \mathbf{W}_2 \approx \mathbf{W}_4$.
Pluralist_e:	$\mathbf{W}_1 > \mathbf{W}_2 \approx \mathbf{W}_3 > \mathbf{W}_4$.
Pluralist_v:	$\mathbf{W}_1 > \mathbf{W}_2 > \mathbf{W}_3 > \mathbf{W}_4$.
Pluralist_h:	$\mathbf{W}_1 > \mathbf{W}_3 > \mathbf{W}_2 > \mathbf{W}_4$.

All these orderings are additive, as the following numerical assignments demonstrate:

Parts	*Hedon*	$Plural_e$	$Plural_v$	$Plural_h$	**Wholes**	*Hedon*	$Plural_e$	$Plural_v$	$Plural_h$
h	1	1	1	2	$\mathbf{W}_1$	1	2	3	3
h′	0	0	0	0	$\mathbf{W}_2$	0	1	2	1
v	0	1	2	1	$\mathbf{W}_3$	1	1	1	2
v′	0	0	0	0	$\mathbf{W}_4$	0	0	0	0

By the entailment, all should be separable, as indeed they are. In the hedonist's case, happiness makes a difference to value whether we hold virtue or vice fixed, and virtue makes no difference to value whether we hold happiness or misery fixed. In the case of each of the pluralists, happiness makes a positive difference to value, whether we hold virtue or vice fixed. And virtue makes a positive difference to value, whether we hold happiness or misery fixed.

Is separability sufficient for additivity? The answer is *no*—at least for the kind of discrete case we have been considering here. What qualitative condition has to be added to separability to guarantee additivity is an interesting topic in itself, but here we need not dwell on it.[7] All we need to make use of is the fact that counterexamples to separability are counterexamples to additivity, and *ipso facto* examples of organic unity.

III. Unity

Precisely because of the power of separability it is perhaps not surprising that apparent counterexamples should arise. Because additivity entails separability, such counterexamples to separability will *ipso facto* be examples of organic unity.

The Kantian ordering is a case in point:

Kantian: $\mathbf{W}_1 > \mathbf{W}_2 > \mathbf{W}_4 > \mathbf{W}_3$.

(Kant explicitly endorses $\mathbf{W}_1 > \mathbf{W}_2$ and $\mathbf{W}_4 > \mathbf{W}_3$. But surely on Kantian precepts it is better to be virtuous and miserable than to be vicious and miserable: i.e. $\mathbf{W}_2 > \mathbf{W}_4$.) For the sake of a *reductio* assume Kant's ordering is separable. $\mathbf{W}_1$ is better than $\mathbf{W}_2$, and $\mathbf{W}_1$ differs from $\mathbf{W}_2$ only in that the former contains happiness and the latter doesn't. So, by top-down, happiness is more valuable than misery. But since $\mathbf{W}_3$ differs from $\mathbf{W}_4$ only in that the former contains happiness and the latter doesn't, by bottom-up, $\mathbf{W}_3$ should be better than $\mathbf{W}_4$—contradiction. Kant evidently accepts bottom-up—since it plays an essential role in his rejection of various purported goods—and so he must reject top-down.

The structure of this *reductio* of separability can be replicated in a *reductio* of additivity, reinforcing the close connection between the concepts. Assume that Kant's ordering is additive, let **Kant** be an additive representation of it:

By representation:	$\mathbf{Kant}(\mathbf{W}_1) > \mathbf{Kant}(\mathbf{W}_2)$.
By additivity:	$\mathbf{Kant}(\mathbf{W}_1) = \mathbf{Kant}(\mathbf{v}) + \mathbf{Kant}(\mathbf{h})$, and $\mathbf{Kant}(\mathbf{W}_2) = \mathbf{Kant}(\mathbf{v}) + \mathbf{Kant}(\mathbf{h}')$.
Hence:	$\mathbf{Kant}(\mathbf{v}) + \mathbf{Kant}(\mathbf{h}) > \mathbf{Kant}(\mathbf{v}) + \mathbf{Kant}(\mathbf{h}')$.
So:	$\mathbf{Kant}(\mathbf{h}) > \mathbf{Kant}(\mathbf{h}')$.
Hence:	$\mathbf{Kant}(\mathbf{v}') + \mathbf{Kant}(\mathbf{h}) > \mathbf{Kant}(\mathbf{v}') + \mathbf{Kant}(\mathbf{h}')$.
By additivity:	$\mathbf{Kant}(\mathbf{W}_3) = \mathbf{Kant}(\mathbf{v}') + \mathbf{Kant}(\mathbf{h})$ and $\mathbf{Kant}(\mathbf{W}_4) = \mathbf{Kant}(\mathbf{v}') + \mathbf{Kant}(\mathbf{h}')$.
Hence:	$\mathbf{Kant}(\mathbf{W}_3) > \mathbf{Kant}(\mathbf{W}_4)$.
By representation:	$\mathbf{W}_3 > \mathbf{W}_4$. contradiction.

7 For a highly readable exposition of sufficient conditions in the continuous case, see Broome [4]. For a discussion of necessary and sufficient conditions in the discrete case. see Oddie [17].

The notions of separability and organic unity enable us to throw considerable light on the following somewhat obscure passages from Bradley's *Appearance and Reality*:

> Where everything is connected in one whole, you may abstract and so may isolate any one factor. And you may prove at your ease that, without this, all the rest are imperfect and worthless; and you may show how, this one being added they all once more gain reality and worth. And hence of every one alike you may conclude that it is the end for the sake of which all the others exist. But from this to argue, absolutely and blindly, that some one single aspect of the work is the sole thing that is good, is most surely illogical. (Bradley [3], p. 405.)

> Deprived of any one aspect the Absolute may be called worthless. And thus, while you take your stand on some one valuable factor, the others appear to you to be means which subserve its existence. … Certainly your position in such an attitude is one-sided and unstable... They are factors not independent, since each of itself implies and calls in something else to complete its defects, and since all are over-ruled in that final whole which perfects them. But these factors, if not equal, are not subordinate one to the other, and in relation to the Absolute they are all alike essential and necessary. (Bradley [3], pp. 456–7.)

In the first passage, Bradley is arguing that if a complex whole is good, and altering a single feature destroys the good of the whole, then that feature can itself legitimately be regarded as good, or as a good. That is an endorsement of top-down separability. But he goes on to argue that one cannot conclude that this admittedly good feature is 'the sole thing that is good'. For there are by assumption a number of features such that without them the whole would not be good, and it would obviously be contradictory to conclude that each, *by itself*, is what makes the difference. This, in conjunction with the first claim, is tantamount to a rejection of bottom-up bare difference. That is to say, a good thing may not, just by itself, make a bad thing better. It may make a positive difference to value only by interacting with other goods. Bradley thus reverses Kant's apparent commitments, by endorsing top-down bare difference while rejecting bottom-up.

The second passage helps clarify Bradley's position. The Absolute (the unified whole) is perfect, but subtract any single component factor and you are left with something that is 'worthless'. If we were to model such an ordering (in a somewhat un-Bradleyan fashion) in our simple example space, the combination of virtue and happiness would be good, but separately each is good only conditional upon the presence of the other good. The ordering thus has following structure:

Absolutist: $\mathbf{W}_1 > \mathbf{W}_2 \approx \mathbf{W}_3 \approx \mathbf{W}_4$.

This ordering is clearly not separable. In some contexts, happiness alone can make the difference between perfection and worthlessness (compare $\mathbf{W}_1$ with $\mathbf{W}_2$) and in others it doesn't make any difference at all (compare $\mathbf{W}_3$ with $\mathbf{W}_4$). It is not hard to rework the little argument above to show that there is no additive representation of this ordering. (The first part of the argument is no different. In the second part, substitute $\approx$ for $>$ throughout.)

G.E. Moore, who explicitly rejects both bottom-up and top-down, appealed to several counterexamples, including the following:

> It seems to be true that to be conscious of a beautiful object is a thing of great intrinsic value; whereas the same object, if no one be conscious of it, has certainly comparatively little value and is commonly held to have none at all. But the consciousness of a beautiful object is certainly a whole of some sort in which we can distinguish as parts the object on the one and hand and the being conscious on the other. Now this latter factor occurs as part of a different whole, whenever we are conscious of anything; and it would seem that some of those wholes have at all events very little value, and may even be indifferent or positively bad. (Moore [13], p. 28.)

Moore takes this to be a straightforward refutation of additivity:

> Whatever the intrinsic value of consciousness may be, it does not give to the whole of which it forms a part a value proportioned to the sum of its value and that of its object. (Moore [13], p. 28–29.)

Moore's example, suitably massaged, does seem to constitute a reasonable counter-example to separability and thus to additivity. Let **b** be the fact that a certain object is beautiful, let **a** be the fact that one is aware of it, and let **b′** and **a′** be their negations. Then, by Moore's lights:

$$\mathbf{a}\&\mathbf{b} > \mathbf{a}'\&\mathbf{b}.$$

By this and top-down, $\mathbf{a} > \mathbf{a}'$. But then by bottom-up:

$$\mathbf{a}\&\mathbf{b}' > \mathbf{a}'\&\mathbf{b}'.$$

But this contradicts Moore's intuition that if an object is ugly then an awareness of the object either makes a bad situation worse ($\mathbf{a}'\&\mathbf{b}' > \mathbf{a}\&\mathbf{b}'$) or at best does not improve it ($\mathbf{a}'\&\mathbf{b}' \approx \mathbf{a}\&\mathbf{b}'$). W. D. Ross rejected Moore's examples of organic unity, but did endorse the Kantian example as a 'genuine illustration of the doctrine' (Ross [22], pp. 70–2, and [21], pp. 185–6). Others have detected in Ross's own treatment of duty what is effectively a thesis of inseparability (Slote [23], p. 33). Recent authors have endorsed Moore's attack on additivity and his endorsement of organic unity by means of similar examples (Lemos [10] and Kagan [8]).

The real problem with additivity, however, goes deeper than the critics of additivity have realised.

IV. Futility

Kant claims that it is better to be happy and virtuous than miserable and virtuous.

$$\mathbf{h}\&\mathbf{v} > \mathbf{h}'\&\mathbf{v}$$

By top-down, happiness in itself is better than misery.

$\mathbf{h} > \mathbf{h}'$.

By bottom-up this yields:

(*) $\mathbf{h}\&\mathbf{v}' > \mathbf{h}'\&\mathbf{v}'$.

This is, of course, incompatible with Kant's judgement, thus demonstrating Kant's repudiation of separability. Now let us say that a person *gets his just deserts* if and only if his happiness is proportionate to his virtue. That is, in our simple space, he gets his just deserts just in case either he is virtuous and happy, or he is vicious and miserable. It is not hard to see that getting one's just deserts (**j**) is tantamount to the following—*happy if and only if virtuous*:

$\mathbf{j} =_{\mathrm{df}} (\mathbf{h} \leftrightarrow \mathbf{v})$.

It is not hard to check that ($\mathbf{h}\&\mathbf{v}$) is tantamount to ($\mathbf{j}\&\mathbf{v}$), and ($\mathbf{h}'\&\mathbf{v}$) is tantamount to ($\mathbf{j}'\&\mathbf{v}$). Hence the initial judgement can be simply redescribed:

$\mathbf{j}\&\mathbf{v} > \mathbf{j}'\&\mathbf{v}$.

Consequently, by top-down bare difference, getting one's deserts is better, in itself, than not getting one's just deserts.

$\mathbf{j} > \mathbf{j}'$.

And by bottom-up this yields:

(#) $\mathbf{j}\&\mathbf{v}' > \mathbf{j}'\&\mathbf{v}'$.

However:

$\mathbf{j}\&\mathbf{v}'$ is logically equivalent to $\mathbf{h}'\&\mathbf{v}'$
$\mathbf{j}'\&\mathbf{v}'$ is logically equivalent to $\mathbf{h}\&\mathbf{v}'$.

And so (#) is tantamount to:

(**) $\mathbf{h}'\&\mathbf{v}' > \mathbf{h}\&\mathbf{v}'$ But (*) and (**) are clearly in contradiction.

The problem here is not just that Kant's particular ordering is unsustainable or incompatible with separability. Rather separability, together with the assumption of just one barely different ordered pair, yields a contradiction. This is a problem with separability. By generalising the technique we can show that separability is not compatible with any reasonable evaluative ordering. The only complete ordering of complex wholes compatible with separability is thus the nihilist's: that nothing is more valuable than anything else.

To show that the above argument does not depend on any special features of the Kantian case, or of our simplified model of it, suppose that we have some set of complex wholes, and can divide each whole into two factors, the F-factor and the rest. The latter we can regard as a complex factor, the G-factor. Each whole specifies a determinate feature from each factor. We will assume that the factors are logically independent, so that the specification of the F-factor does not place any logical restrictions on the G-factor. (Later I will argue that such logical independence is part of the very notion of a factor.) Of course, I am leaving open the number of determinate features in each such factor. They can both be as large as you like.

Now suppose that some whole W[f] is superior to a barely different whole W[f′]. Let g be the G-feature which they share in common. Thus W[f] can be characterised as f&g, and W[f′] as f′&g:

So:	f&g > f′&g.		
By top down:	f > f′.		
By bottom-up:	f&g′ > f′&g′, where g′ is any G-feature distinct from g.		
For one such g′ define the following two features:	h	$=_{df}$	[~g′ → f] & [g′ → f′],
	h′	$=_{df}$	[~g′ → f′] & [g′ → f].

Consider the set of features H which consists of h, h′ together with all the F-features apart from f and f′. H is a determinable in the sense that every complex whole over F and G specifies one and only one element of H. Further H is independent of G. Lastly, every complex over F and G can be singled out (albeit indirectly) by the specification of a G-feature and an H-feature. (Symmetrically, every complex over G and H can be singled out by a specification of a G-feature and an F-feature.) Thus H and G are both factors. Now it is clear that:

f&g is equivalent to h&g
f′&g is equivalent to h′&g
f&g′ is equivalent to h′&g′
f′&g′ is equivalent to h&g′

By equivalence:	h&g > h′&g	
So by top-down:	h > h′.	
By bottom-up:	h&g′ > h′&g′.	
But by equivalence:	h′&g′ > h&g′,	Contradiction.

Thus no complex is better than a barely different complex. If two wholes differ over a complex feature and agree on some remainder, then we can always regard the complex feature as a realisation of a single (complex) factor, and the argument can be repeated.[8] Since no two complexes can differ on absolutely every feature, the argument can always be repeated. The only evaluative ordering compatible with separability is thus the one according to which all states have exactly the same value. Separability entails nihilism. (I am assuming that the wholes are comparable for value. If they are not then we have more or less incommensurability. Strictly speaking, the argument establishes that

8 Recall (footnote 5) that *separability* means *strong separability.*

separability entails either incommensurability or nihilism.) Thus additivity entails either incommensurability or nihilism. All non-nihilistic evaluative orderings must be organically unified.[9]

V. Relativity

We have a clash here. On the one hand we have some putative examples of interesting separable/additive orderings. On the other we have a general argument that separability entails a nihilistic levelling of value. We could, of course, simply reject those examples, and the various principles, like separability, which appear to capture them. But that would be a somewhat drastic dissolution of the problem, leaving us bereft of a number or attractive and useful argumentative strategies.

A tacit assumption of these derivations is the *equivalence principle*: that logical equivalents have the same value. Is this defensible? Let G be the proposition that *God is perfectly good*, and I the proposition that *1 is the smallest odd integer*. I is true in all possible worlds, and perhaps G is too. In that case they are logically equivalent. And yet G would arguably be a very good thing, while the trivial truth I is not.[10] Likewise, (*) may be compatible with (#) despite the relevant logical equivalences. If this is right then value is a hyperintensional concept, and there is simply no such thing as *the* value of states the identity conditions of which are given by logical equivalence. However, the G-I counterexample to equivalence is not compelling. If G is the proposition that a being would have to be morally perfect to count as God, then it may be true in all worlds, but on that construal G seems to have the same value as I—neither good nor bad in itself. On the other hand, construed as a proposition with existential import—in particular, entailing that a morally perfect being exists—G would seem to be genuinely valuable. But on that construal, G is true in only a tiny fraction of all possible worlds, and so is not equivalent to the valueless I.

The advantages of equivalence are, of course, overwhelming, especially considering the role of value in guiding action. Value makes a claim to being realised. That's its point. If S makes a claim to be realised, and it is logically impossible to realise one of S,S* without realising the other, then S* surely makes the same claim to being realised as S.

I think the problem with the paradox lies elsewhere. Bottom-up bare difference is intended to capture the idea that if a feature f is better than a rival f′ then, *other things being equal*, a complex containing f is better than one containing f′. And top-down bare-difference is meant to capture the idea that if, *other things being equal*, replacing f′ with f makes a positive difference to value, then f is better than f′. But what exactly does it take for *other things* to be *equal*? Call this the *ceteris paribus* problem.

Suppose two complexes differ over the factor F, say. One specifies f and the other ~f. Presumably things other than F are equal in that pair of complexes just in case the remainder of the complexes are the same. So, things other than F are equal in f&g and ~f&g because both specify g for the G-factor. It may seem obvious that if f and g are logically independent, and that a complex whole can be specified by their combination,

[9] This argument is related to Miller's paradox in the theory of truthlikeness. See, Miller [12] and Oddie [14], ch. 6.

[10] This argument against the equivalence principle was suggested by an anonymous referee.

then both f and g are parts of the complex, and that the pair f&g and ~f&g differ only over the F-factor. However, if f and g are logically independent so too are h and g, where h is defined as the biconditional state (f↔g). f&g specifies h while ~f&g specifies ~h. Is h (that is, (f↔g)) also a part of f&g, and ~h a part of ~f&g? We can call the collection consisting of h and ~h the H-factor. Does the fact that f&g and ~f&g differ over the H-factor count against all things other than F being equal?

Some will doubtless be inclined to reject factors like H (or like *getting one's just deserts*) as logically complex. 'Biconditional' states, we will be told, along with 'disjunctive' and 'negative' states and properties, should be discounted. But what exactly is a 'disjunctive' and 'negative' state or property? If being picked out by means of a conjunction (disjunction, biconditional etc.) makes a state conjunctive (disjunctive, biconditional etc.) then *every* state is conjunctive (disjunctive, biconditional etc. etc.). f&g, f&(f↔g), ~(~f∨~g) all pick out the very same possible state of the world, although they are clearly three different ways of doing so, involving as they do different truth functions applied to various arguments. The fact that a state can be arrived at, or specified by, conjoining some of its logical consequences does not entail that the state itself is 'conjunctive'. (Similarly, the fact that a number, like 9, is the square of a negative number does not tell us that 9 somehow involves, as parts or constituents, negativity or the number 3.) If it did then the *ceteris paribus* problem would indeed be soluble, but in a way which would trivialise all *ceteris paribus* judgements. No two states could differ in some one state alone. (For suppose that W and V differ over f alone: W entails f and V entails ~f. Let g be some common non-tautologous consequence of both W and V. Then ~g ∨ f is a consequence of W but not of V. If there is no common non-tautologous consequence of both W and V then W is f and V is ~f. Let g be logically independent of W and V. Then W entails g ∨ f and V does not.)

Note that this observation stands even if we grant that f and g are metaphysically privileged atomic states. The state picked out by f&g can still be picked out in sundry ways using other truth functions. Still, if we *were* granted some privileged set of atomic states, whether metaphysically privileged or privileged in some other way, we could indeed solve the *ceteris paribus* problem. Two states which differ over some atomic state would be otherwise equal if they agreed in all other atomic states. This would appear to commit us to the metaphysics of absolutely atomic states. Maybe there are such. Certainly closely related ideas have become more popular in recent years. But I will pursue a different route, one without such a hefty presupposition.

Consider a space of possibilities, either a space of possible states or a space of possible attributes—a set of mutually exclusive and jointly exhaustive possible states for a world to be in, or possible conditions for an individual to exemplify . A *basis* is any way of carving up the possibilities into types of features. It is a list of factors, such that each member of the space can be completely characterised by a specification of one feature from each factor. Consider the Kantian space of four possibilities. There are a number of ways of carving this up. For example, each whole can be specified in terms of the virtue-factor and the happiness-factor. But another way is to specify each in terms of the virtue-factor and what we might call the *justice*-factor (consisting of **j** and **j′**).

$\mathbf{B}_1$ = <virtue, happiness>
$\mathbf{B}_2$ = <virtue, justice>

But there are others of course. For example:

$\mathbf{B}_3$	=	<justice, happiness>
$\mathbf{B}_4$	=	<virtue, justice, happiness>

and so on. $\mathbf{B}_4$ is interesting because it involves a certain redundancy. Specify a feature from any two of its factors and you thereby specify a feature from the third. There are two limiting kinds of basis. One is where the basis contains as many factors as there are ordered complexes. Each complex W is associated with the bivalent factor *being identical to W* ($\mathbf{B}_5$). The other is where the basis contains just one factor, the determinable of which each complex is a determinate realisation ($\mathbf{B}_6$).

With the concept of a basis in hand we can solve the *ceteris paribus* problem, at least in a basis-relative manner. Other things clearly can be equal or unequal relative to a basis. Select a basis B. Let W and U differ over some factor F in B: other things are equal (relative to B) just in case W and U agree on all the other factors in B. For example, relative to the first basis $\mathbf{B}_1$ the possibilities $\mathbf{W}_1$ and $\mathbf{W}_3$ differ over virtue, but are otherwise equal—they agree on happiness. However, relative to the second basis, $\mathbf{B}_2$, the possibilities $\mathbf{W}_1$ and $\mathbf{W}_3$ differ not only over virtue, but in the justice-factor as well.

We can now define the notions of separability, additivity and organic unity—all relative to some specified basis B. Relative to the basis <*virtue*, *justice*>, for example, Kant's ordering is separable, and can be captured by an additive value function. However relative to <*virtue*, *happiness*> and to <*justice*, *happiness*> it is not separable, and no assignment can represent the ordering additively. Some ways of decomposing complexes into components yield features which play the role of *axiological atoms* each one contributing a certain fixed amount of value to the wholes of which it is a part. But other ways of decomposing complexes do not yield such axiological atoms.

Armed with the basis-relative notions, the result that additivity entails separability holds—given a certain basis is held fixed. Further, the derivation of nihilism from separability can be blocked. The argument either changes bases in mid-stream, or else tacitly assumes that if an ordering is separable relative to one basis it is separable relative to all of them.

The important question now is this: Do we simply have to jettison absolute, basis-independent notions? Surely that would substantially undermine any interest these notions might have had. Look at Kant's ordering through the <*virtue*, *happiness*> spectacles and you see organic unity. Look at it through the <*virtue*, *justice*> spectacles, you see additivity. Neither set of spectacles is preferable. So there is just no truth of the matter as to whether orderings themselves are *really* additive or organically unified.

One proposal for restoring an absolute notion would be to adopt the Kantian line that there is a preferred set of spectacles for viewing the world—a unique preferred system of factors (or categories) with which we should operate. Absolute additivity is additivity relative to this preferred basis. After all, the idea that some properties or concepts are to be preferred to others in our characterisation of the world receives strong support in other domains. (Armstrong [1], Lewis [11] and Oddie [14].) Call this the *privileged* proposal.

The problem with the privileged proposal is that, independently of considerations of additivity and organic unity, there seems nothing in the two rival bases to force us to prefer one to the other. Getting one's just deserts or not seems no more nor less natural a factor than being happy or miserable, or being virtuous or vicious. It is not as though we were comparing properties like *green* and *grue* here. While *justice* can be defined biconditionally in terms of happiness and virtue, it doesn't seem any the less natural for all that. And even if it were slightly less natural, or less 'simple', would that in itself recommend banning it from being an axiological factor? Maybe not any old attribute can count as an axiological factor, but surely *this* one can.

A second proposal would be to allow all bases as legitimate, but to classify an ordering as additive only if it could be so represented under every such basis. Call this the *stringent* proposal. The problem with the stringent proposal is that it would leave us with a notion which is virtually useless. By recycling the argument of the last section we could show that no interesting ordering of more than two elements would be additive. Every ordering would be organically unified. The thesis of the organic unity of value would thereby lose all its interest.

A third proposal, like the second, allows all bases as legitimate, but classifies an ordering as additive just in case it can be so represented under at least one such basis. That is to say, if there is some way of carving the world up into features which can play the role of axiological atoms in the evaluative ordering, then the ordering is additive. If there is no way of doing so, then it exhibits genuine organic unity. Call this the *liberal* proposal.

Under the liberal proposal the notion of additivity will collapse in another direction: it classifies all orderings as additive. Consider the two degenerate bases, $\mathbf{B}_5$ and $\mathbf{B}_6$. It is not hard to see that in each case additivity is trivially satisfied, and so too is separability. Given the ordering of wholes we simply assign numbers which represent it. In the case of $\mathbf{B}_6$ each complex is identical to its sole part, so there is nothing further to do. In the case of $\mathbf{B}_5$ these numbers are assigned to the presence of the associated feature and 0 is assigned to its absence.

VI. Admissibility

Both the privileged and the liberal proposal contain a rational kernel which I will try to extract. The basic idea (as the liberal proposal suggests) is that absolute additivity is determined by additivity relative to some basis within a class of bases, but (as the privileged proposal suggests) not any old basis counts. Absolute additivity is determined by additivity relative to some *admissible* basis. What makes a basis admissible?

Consider again the six different ways we considered of carving up the complexes of the Kantian space. In each of $\mathbf{B}_1$, $\mathbf{B}_2$, and $\mathbf{B}_3$, the two factors are logically independent, in the sense that specifying one has no implications for the other. One might be tempted to think that, because *justice* is specified in terms of virtue and happiness, it is not logically independent of *virtue*. But that is not so. Specifying whether or not someone is virtuous carries no implications for whether or not they get their just deserts, and *vice versa*. The same goes for *happiness* and *justice*. In such bases the factors are *independent*. They do not logically interact. On the other hand, it is easy to see that the factors in $\mathbf{B}_4$ and $\mathbf{B}_5$ are not logically independent. It is trivial that any ordering is additive relative to some basis in which the factors are dependent. However, it is by no means trivial that it is additive relative to an independent basis. As it happens, we discovered one such independent basis for Kant's ordering (namely, $\mathbf{B}_2$), but we by no means established that there will always be one such. For a basis to be admissible it has to consist of independent features.

But independence is not enough. Consider the other degenerate basis: $\mathbf{B}_6$. Since there is just a single factor, trivially there is no logical interaction between factors. But just like $\mathbf{B}_5$, $\mathbf{B}_6$ trivialises additivity. We may not have interacting factors in $\mathbf{B}_6$ but only by virtue of the fact that $\mathbf{B}_6$ does not carve up the wholes at all. What is common to $\mathbf{B}_5$ and $\mathbf{B}_6$ is that they do not decompose the wholes into *genuine* factors.

Non-degeneracy (there being more than one factor) and independence (deciding one factor has no implications for other factors) are necessary conditions for a basis to be admissible. I leave open the possibility that these are jointly sufficient for admissibility, but it may be that a genuine factor, even if not simple, has to be 'natural' or 'unified' in some sense. An ordering is additive just in case there is some way of carving up the world up into non-degenerate, non-interacting, *natural* factors which can play the role of axiological atoms.

Call this the *admissibility* account. (Note that analogous accounts yield absolute notions of bottom-up and top-down bare difference.)

Does the admissibility account also collapse? Take the Absolutist's ordering—it has one good and three bad complexes. Any admissible basis will have to contain no fewer and no more two factors. (Fewer and it would be degenerate, more and we would have dependence.[11]) Each factor will have to contain two features apiece. The good complex has to share each of its features (say f and g) with one bad complex apiece. If both f and g add value, then both of those bad complexes which have one of those features would have to be better than the complex which has neither. If one (say f) adds value, and the other does not, then the bad complex with f will be better than the bad complex with g. If neither of the two features adds value then the good complex is no better than the complex that lacks both features. Thus Absolutist's ordering cannot be represented additively using an admissible basis. All the bases which additively represent the Absolutist's ordering involve either degeneracy or logical interaction amongst the factors. This result can be generalised to any space in which one element (the Absolute) is better than all the others, and all the others are on a par. Bradley was thus right: if the Absolute is good but any alteration would make it bad, then we have organic unity.

We can state an interesting fact about the invariance of additivity which goes a long way to vindicating the objectivity of the notions of absolute additivity and absolute organic unity:

> If two evaluative orderings have the same structure (*viz.* they are isomorphic) then either both are additive or both exhibit organic unity.[12]

Organic unity and additivity depend only on the overall *structure* of value. They are features of the structure of value, not of the idiosyncratic way in which we happen to decompose the wholes into factors for non-evaluative purposes.

[11] We could call this a 2×2 factorisation since it contains 2 factors each with two features. Note that the term *factor* is apposite. If a space ω of complexes contains n members then any factorisation of ω has p factors ($p \geq 2$), and there are numerical factors $m_1, m_2, \ldots, m_p$ of n such that the i^{th} factor has m_i features in it.

[12] Suppose that $>$ is isomorphic to $>^*$ under f, and that $\mathbf{B}$ is an additive factorisation of $>$. Then $f(\mathbf{B})$ is an additive factorisation of $>^*$.

VII. Ideality

This analysis of additivity and organic unity has interesting implications for both for the metaphysics of value and for the methodology of evaluative reasoning.

First, the metaphysical implications. Whether or not value is additive or organically unified depends on the structure of the ordering induced by value. It is by no means guaranteed that the structure of value will turn out to be additive. This much has to be conceded to the friends of organic unity. But the standard arguments against additivity and in favour of organic unity, like the one with which this paper began, do not establish their conclusion. At most they show that value is not additive relative to one particular way of slicing up complexes into factors. That is not surprising, since there will always be such non-additive bases, quite independently of the structure of value. In order to show that value exhibits organic unity in an absolute sense, one would have to show that the ordering cannot be captured by an additive factorisation. That can sometimes be done, at least with simple orderings (e.g. the Absolutist's ordering). But it is quite a tall order to expect it to be possible with complex orderings of large spaces.

Second, the methodological implications. If value is additive then there is a way of factoring out the axiological atoms; the bottom-up and top-down bare difference principles will be true; and the method of bare-differences *as applied to those factors* will work. However, the standard applications of the method of bare-difference (just like the standard objections to it) typically proceed without any awareness that factorisation is presupposed. Or if there is such an awareness, then perhaps it is tacitly assumed that, independently of questions of value, there is a unique, privileged, admissible factorisation of the set of possibilities, and that value is additive relative to that factorisation. As they stand the arguments are thus rather weaker than their authors might have realised.

The additivity of value is an hypothesis, one which it may be difficult to establish for any interestingly complex evaluative domain, but also one which may be equally difficult to disprove. But the additivity of value would be enormously useful in helping to organise and clarify our thinking and reasoning about value. In the absence of a proof, one way or the other, my proposal is that we should adopt additivity as a *heuristic tool*.

We have already seen how to use the assumption of additivity not only in direct applications (the bare-difference arguments) but also indirectly in the search for intrinsic goods. Kant tells us that a good will is the only thing of unconditional value, and that anything else is at best a conditional good. Kant tacitly assumes one factorisation. But by assuming additivity instead, we can cast about for a different way of factorising the space. We then identify *getting one's just deserts* as a plausible axiological factor alongside having a good will.

Let me finish by applying the heuristic to a Moorean example like the one in section III. Moore claims that being admiringly aware of a beautiful object is better than failing to be admiringly aware of it, but being admiringly aware of an ugly object is worse than failing to be admiringly aware of it (Moore [13], p. 208). Let **b** = a certain object's being beautiful, let **a** = one is admiringly aware of the object, and let **b′** and **a′** be their respective negations.

Moore: $\mathbf{a}\&\mathbf{b} > \mathbf{a'}\&\mathbf{b}$ and $\mathbf{a'}\&\mathbf{b'} > \mathbf{a}\&\mathbf{b'}$.

Any ordering compatible with these core judgements is inseparable and hence organically unified (relative to <*admiring awareness*, *beauty*>). There are obviously a number of completions of the ordering. Moore argues that a world full of beauty of which nobody is aware at all is still better than a world full of ugliness of which nobody is aware, which is compatible with $\mathbf{a}'\&\mathbf{b} > \mathbf{a}'\&\mathbf{b}'$. Thus one extension of the Moorean ordering is this:

Moorean$_1$: $\mathbf{a}\&\mathbf{b} > \mathbf{a}'\&\mathbf{b} > \mathbf{a}'\&\mathbf{b}' > \mathbf{a}\&\mathbf{b}'$.

Like all completions of the core judgements, this ordering is organically unified relative to <*admiring awareness*, *beauty*>. But consider the concept of *aesthetic calibration*. Someone is aesthetically calibrated with respect to an object just in case they are admiringly aware of it if it is beautiful and not admiringly aware of it if it is ugly. Otherwise they are aesthetically *discordant*. Then the Moorean ordering can be respecified thus:

Moorean$_1$: $\mathbf{c}\&\mathbf{b} > \mathbf{c}'\&\mathbf{b} > \mathbf{c}\&\mathbf{b}' > \mathbf{c}'\&\mathbf{b}'$.

Relative to the basis <*calibrated*, *beauty*> this ordering is both separable and additive. Beauty is better than ugliness ($\mathbf{b} > \mathbf{b}'$), calibration is better than discordance ($\mathbf{c} > \mathbf{c}'$), and beauty trumps calibration (the value difference between beauty and ugliness is greater than that between calibration and discordance).

There are obviously other possible extensions of Moore's core evaluative judgements. One might argue that while the existence of beauty is a good thing, failing to be admiringly aware of such beauty as a whole is no better (and no worse) than the existence of an ugly thing of which nobody is admiringly aware (Moorean$_2$). Or that the failure in the first case makes the situation worse overall (Moorean$_3$).

Moorean$_2$: $\mathbf{a}\&\mathbf{b} > \mathbf{a}'\&\mathbf{b} \approx \mathbf{a}'\&\mathbf{b}' > \mathbf{a}\&\mathbf{b}'$.

Moorean$_3$: $\mathbf{a}\&\mathbf{b} > \mathbf{a}'\&\mathbf{b}' > \mathbf{a}'\&\mathbf{b} > \mathbf{a}\&\mathbf{b}'$.

Again, while both are inseparable and organic, we can switch factorisations to obtain separable and additive orderings relative to the new basis:

Moorean$_2$: $\mathbf{c}\&\mathbf{b} > \mathbf{c}'\&\mathbf{b} \approx \mathbf{c}\&\mathbf{b}' > \mathbf{c}'\&\mathbf{b}'$.

Moorean$_3$: $\mathbf{c}\&\mathbf{b} > \mathbf{c}\&\mathbf{b}' > \mathbf{c}'\&\mathbf{b} > \mathbf{c}'\&\mathbf{b}'$.

In the both cases, beauty and calibration are both good. In the second ordering beauty and calibration are evaluatively equivalent (the difference between $\mathbf{b}$ and $\mathbf{b}'$ is the same as the difference between $\mathbf{c}$ and $\mathbf{c}'$), and in the third ordering calibration trumps beauty.

My proposal, then, is to treat the hypothesis of additivity as a *regulative ideal* in the theory of value. We have not proved it false, we can hope that it is true, and we know how to reason evaluatively if it is true. So we can and should employ it as a powerful heuristic tool in our search for additive factors and by so doing bring simplicity, order, and clarity to our inquiries into value.[13]

University of Colorado

Received: February 2000
Revised: July 2000

REFERENCES

1. Aristotle, *Nichomachean Ethics*, trans. Martin Ostwald (Indianapolis: Bobbs-Merrill Co., 1962).
2. Armstrong, D.M., *A Theory of Universals* (Cambridge: Cambridge University Press, 1978).
3. Bradley, F.H., *Appearance and Reality* (New York: Macmillan, 1893).
4. Broome, John, *Weighing Goods* (Oxford: Blackwell, 1990).
5. Carlson, Erik, 'The intrinsic value of non-basic states of affairs', *Philosophical Studies* 85 (1997), 95–107.
6. Danielsson, S., 'Harman's equation and the additivity of intrinsic value' in Lars Lindahl *et al.* (eds.), *For Good Measure* (Uppsala, 1997).
7. Harman, G., 'Towards a theory of intrinsic value', *Journal of Philosophy* 64 (1967), 792–804.
8. Kagan, Shelley, 'The additive fallacy', *Ethics* 99 (1988), 5–31.
9. Kant, I., *Foundations of the Metaphysics of Morals,* trans. Lewis White Beck (Indianapolis: Bobbs-Merrill, 1959).
10. Lemos, N.M., *Intrinsic Value: Concept and Warrant* (New York: Cambridge University Press, 1994).
11. Lewis, D., 'New work for a theory of universals', *Australasian Journal of Philosophy* 61 (1983).
12. Miller, D., 'The accuracy of predictions', *Synthese* 30 (1975), 159–91.
13. Moore, G.E., *Principia Ethica* (London: Cambridge University Press, 1960).
14. Oddie, G., *Likeness to Truth* (Dordrecht: Reidel, 1986).
15. Oddie, G., 'Addiction and the value of freedom', *Bioethics* 7 (1993), 373–401.
16. Oddie, G., 'Killing and letting-die: from bare differences to clear differences', *Philosophical Studies* 88 (1997), 267–87.
17. Oddie, G., 'Recombinant Values' (forthcoming).
18. Oldfield, E., 'An approach to a theory of intrinsic value', *Philosophical Studies* 32 (1977), 233–49.
19. Quinn, W.S., 'Theories of intrinsic value', *American Philosophical Quarterly* 11 (1974), 123–32.
20. Rachels, J., *The End of Life* (Oxford: Oxford University Press, 1986).
21. Ross, W.D., *The Foundations of Ethics* (London: Oxford University Press, 1939).
22. Ross, W.D., *The Right and the Good* (London: Oxford University Press, 1963).
23. Slote, M., *From Morality to Virtue* (New York: Oxford University Press, 1992).

[13] I have given this paper in several places and it has been read by others. I would like to thank all of the following for helpful critical comments and suggestions: John Broome, Erik Carlson, Jack Copeland, Alan Goldman, Risto Hilpinen, Roy Sorenson, Roy Perrett, Michael Tooley and two very helpful referees for the journal.