

Infinite Value and Finitely Additive Value Theory

Author(s): Peter Vallentyne and Shelly Kagan

Source: *The Journal of Philosophy*, Vol. 94, No. 1 (Jan., 1997), pp. 5-26

Published by: [Journal of Philosophy, Inc.](#)

Stable URL: <http://www.jstor.org/stable/2941011>

Accessed: 13-08-2014 16:45 UTC

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at

<http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Journal of Philosophy, Inc. is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Philosophy*.

<http://www.jstor.org>

THE JOURNAL OF PHILOSOPHY

VOLUME XCIV, NO. 1, JANUARY 1997

INFINITE VALUE AND FINITELY ADDITIVE VALUE THEORY*

Call a theory of the good—be it moral or prudential—*aggregative* just in case (1) it recognizes local (or location-relative) goodness, and (2) the goodness of states of affairs is based on some aggregation of local goodness. The locations for local goodness might be points or regions in time, space, or space-time; or they might be people or states of nature.¹ Any method of aggregation is allowed: totaling, averaging, measuring the equality of the distribution, measuring the minimum, and so on. Call a theory of the good *finitely additive* just in case it is aggregative, and for any finite set of locations it aggregates by adding together the goodness at those locations. Standard versions of total utilitarianism typically invoke finitely additive value theories (with people as locations).

A puzzle can arise when finitely additive value theories are applied to cases involving an infinite number of locations (people, times, and so on). Suppose, for example, that temporal locations are the locus of value, and that time is discrete and has no beginning or end.²

* Thanks to Dick Arneson, Tim Bays, John Broome, James Cain, Tyler Cowen, Chuck Cross, Michael Della Rocca, Jorge Garcia, Brad Hooker, Jerome Keisler, Mark Nelson, Alastair Norcross, Teddy Seidenfeld, Ted Sider, Howard Sobel, Roy Sorensen, and Paul Weirich for helpful discussion and comments.

¹ Here we follow John Broome's usage of the term 'location' as the generic term for the things with which local goodness is associated; see his *Weighing Goods* (Cambridge: Blackwell, 1991).

² Here and below, we shall assume for our examples that (1) there are only denumerably many locations in space, time, and space-time, and (2) they are discrete in the sense that for any location and any direction there is a well-defined "next" location. These assumptions are made for simplicity of presentation, and play no role in our argument.

How would a finitely additive theory (for example, a temporal version of total utilitarianism) judge the following two worlds?

Goodness at Locations (for example, times)

w1: ..., 2, 2, 2, 2, 2, 2, 2, 2, ...

w2: ..., 1, 1, 1, 1, 1, 1, 1, 1, ...

Example 1

At each time, w1 contains 2 units of goodness and w2 contains only 1. Intuitively, we claim, if the locations are the same in each world, finitely additive theorists will want to claim that w1 is better than w2. But it is not clear how they could coherently hold this view. For using standard mathematics, the sum of each is the same infinity, and so there seems to be no basis for claiming that one is better than the other.³ (Appealing to Cantorian infinities is of no help here, since for any Cantorian infinite N , $2 \times N = 1 \times N$.)

We shall argue that such theories can and should judge some worlds with an infinite number of locations as better than others. Total utilitarianism, for example, can and should judge w1 as better than w2 when they have the same locations. Moreover, we shall argue that there are some perfectly general metaprinciples governing how finitely additive theories should make judgments when an infinite number of locations are involved.⁴

³ This problem does not always arise when there are an infinite number of locations. First, if the locations are ordered (as in time or space) and bounded (entirely inside some finite region), then there may be no problem. For example, the "sum" (integral) of the finite values at all the infinitely many real-numbered locations in some bounded interval (for example, 1 to 2) is well defined, and finite. Second, if the locations are unbounded, but their values asymptotically approach zero "sufficiently quickly" (for example, as in $v(x) = 1/x^2$), then the "sum" (integral) may be well defined and finite. In these cases, the standard integration techniques deal "almost perfectly" with aggregation over an infinite number of locations. Strictly speaking, it is not perfect because standard integration holds that the aggregate value of 1 unit of happiness at each point in time inclusively between two points in time (with time being real valued) is the same as that of 1 unit at each time in this interval except with 2 units of happiness at the endpoint. Standard integration assigns the same aggregate value to these functions, but the second, we would argue, has a greater value. We agree, of course, that the results of standard integration are infinitesimally close to the correct answer, and for the purposes of this paper, we shall ignore the small errors. Thus, all references to sums should be understood as including integration when this is appropriate.

⁴ The problem of aggregating when time is infinite in length has been discovered and given roughly the same solution (on which we generalize in the present paper) at least three times. Most recently, it was discovered in Mark Nelson, "Utilitarian Eschatology," *American Philosophical Quarterly*, xxviii (1991): 339-47; and addressed in Vallentyne, "Utilitarianism and Infinite Utility," *Australasian Journal of Philosophy*, lxxi (1993): 212-17. Prior to that, it was discovered and addressed in (the important but unnoticed!) Krister Segerberg, "A Neglected Family of Aggre-

We shall assume throughout that the amount of goodness at each location is finite (so that the infinite goodness comes from the infinite number of locations and not from the values at the locations). We shall further assume that goodness is fully interlocationally comparable in the sense that it is possible to ensure that the same scale (both the zero point and the unit) for measuring goodness is used for all locations. This is a presupposition of finitely additive value theories. For if the numbers are not on the same scale, it makes no sense to add them together.⁵ All goodness numbers are thus assumed to be on the same scale.

Finally, because for generality we make no specific assumption about what the relevant basic locations of goodness are (people, times, and so on), all references to locations should be understood as those which are specified by a given finitely additive theory under consideration.⁶

I. NONSTANDARD NUMBERS

Before proceeding, we need to mention a possible partial solution to the puzzle. In the discussion of the example above, we assumed that the sum of a denumerably infinite number of 2s is the same as the sum of an equal number of 1s. On standard mathematics, it would be more correct to say that neither sum is well defined, but the point would remain that neither sum would be defined and greater than the other. The important point to note here is that there is such a thing as nonstandard mathematics (for example, nonstandard analy-

gation Problems in Ethics," *Noûs*, x (1976): 221-44. And prior to that, it was discovered and addressed by Frank Ramsey in the 1920s and addressed by various economists, especially since the 1960s; see Carl Christian Von Weizsacker, "Existence of Optimal Programs of Accumulation for an Infinite Time Horizon," *Review of Economic Studies*, xxxii (1965): 85-104; and David Gale, "On Optimal Development in a Multi-Sector Economy," *Review of Economic Studies*, xxxiv (1967): 1-18. See also the discussions (and many other references) in D.E. Campbell, "Impossibility Theorems and Infinite Horizon Planning," *Social Choice and Welfare*, ii (1985): 283-93; L.G. Epstein, "Intergenerational Preference Orderings," *Social Choice and Welfare*, iii (1986): 151-61; and Luc Lauwers, *Social Choice with Infinite Populations*, Dissertation 101 (Catholic University of Leuven (Monitoraat E.T.E.W, Dekenstraat 2, B-3000 Leuven, Belgium), 1995).

⁵ If the locations (for example, people) are the same, then only unit comparability is needed, since the zeros will cancel out. But in the more general case, where the locations may be different in the worlds being evaluated, zero comparability is also needed.

⁶ For example, if, according to a given finitely additive theory of value, people are the only ultimate bearers of value, then only personal locations (and not temporal ones) are referenced. For a discussion of the problems that arise if locations are not restricted to ultimate locations, see James Cain, "Infinite Utility," *Australasian Journal of Philosophy*, lxxiii (1995): 401-04; and Vallentyne, "Utilitarianism and Infinite Utility," *Australasian Journal of Philosophy*, lxxi (1993): 212-17.

sis) dealing with infinitesimals and infinite numbers. Although such mathematics is not mainstream, it has been recognized as legitimate since the work of Abraham Robinson in the 1960s established that the existence of such numbers is perfectly consistent with standard mathematics.⁷ We shall now briefly explain how appealing to such mathematics may offer a solution to the puzzle.

Standard mathematics does not recognize infinitesimals and infinite numbers for the purposes of normal addition or multiplication. Nonstandard mathematics, however, does. Indeed, on nonstandard mathematics, all the basic arithmetic operations are well defined and operate in the "usual way" for both finite and infinite numbers. It recognizes, for example, that adding or multiplying two positive infinite numbers always yields a result that is greater than either of the two original numbers. It also says that, if N is some infinite integer and N 1s are added together, then the result is N . If one more 1 is added to the total, the result is $N+1$, which is 1 greater than N . If N 2s are added together, the result is $2 \times N$, which is twice as great as N . N , $N+1$, and $2 \times N$ are each nonstandard infinite numbers, and the first is smaller than the second, which is smaller than the third.

If the infinities involved in the puzzle cases are bounded nonstandard infinities, then nonstandard mathematics provides a straightforward resolution of the puzzle of the above example. For although the two sums each involve the same nonstandard infinite number of locations (call it N), the first nonstandard sum (of w_1) is greater than the second (of w_2), since $2 \times N > N$.

Is this the end of the story? No. For nonstandard mathematics, like standard mathematics, is silent in many cases where there are *unboundedly* many (that is, more than any nonstandard infinite number) positive numbers that are summed. Nonstandard mathematics says, for example, that 2 units of goodness at (nonstandard infinite) H locations (that is, $\langle 2, 2, 2, \dots, 2 \rangle$, where there are H 2s) has a greater total than 1 at H locations (that is, $\langle 1, 1, 1, \dots, 1 \rangle$, where there are H 1s). But it, like standard mathematics, is silent about the comparisons of 2, versus 1, at each of unboundedly many locations (that is, $\langle 2, 2, 2, \dots \rangle$ versus $\langle 1, 1, 1, \dots \rangle$).

So, nonstandard mathematics resolves the puzzle where the number of locations is a (bounded) nonstandard number, but not in general when it is unbounded. We shall therefore focus on the un-

⁷ See Abraham Robinson, *Non-Standard Analysis* (Amsterdam: North Holland, 1966). For an extremely intuitive and accessible introduction to nonstandard mathematics, see H.J. Keisler, *Elementary Calculus* (Boston: Prindle, Weber, and Schmidt, 1976), chapter 1.

bounded case, and for simplicity we shall assume that only standard numbers are involved. We shall develop and defend a metaprinciple for aggregating the goodness of infinitely many locations which makes no essential use of nonstandard mathematics.

II. THE BASIC IDEA: AGGREGATING WITHOUT APPEALING TO ANY ESSENTIAL NATURAL ORDER OF LOCATIONS

We start by formulating a metaprinciple that is applicable to all sorts of locations—whether or not they have any essential natural order (a notion that is explained below). In the following sections, we strengthen this metaprinciple for cases where there is some essential natural order (as arguably there is for spatial and temporal locations but not for people and states of nature).

Our most basic idea is very simple, and can be illustrated using example 1—in which w_1 has 2 units of goodness at each location and w_2 has 1 unit at each. Suppose that these two worlds contain exactly the same locations (for example, the same people). Finitely additive theories should, we claim, judge w_1 as better than w_2 for the following reason: no matter what finite set of locations one considers, relative to that set w_1 has a greater total than w_2 . That is, our most basic idea (BI) is:

BI (basic idea): if w_1 and w_2 have exactly the same locations, and if, relative to any finite set of locations, w_1 is better than w_2 , then w_1 is better than w_2 .⁸

This is the core of our approach, and we hope it is sufficiently plausible to need little discussion at this point. (We shall, of course, return to it below.) As it stands, it is incredibly weak: it has implications for finitely additive theories only where (1) the locations in the two worlds are exactly the same and (2) one world has more goodness at every single location (since just one countervailing singleton set renders the idea silent). It is hardly a robust metaprinciple, but it is fine, we claim, as far as it goes.

The rest of the paper is concerned with strengthening BI. There are two sorts of strengthening that we shall develop: one extends BI by weakening the requirement that one world be better, relative to all finite sets of locations; the other strengthens the metaprinciple to cover cases where the locations are not exactly the same.

⁸ For simplicity, we shall assume throughout that for the purposes of aggregation (1) all locations of a given type (for example, time, or spatial locations, or persons) have the same weight (both within and across worlds), and (2) the total weight of the locations in a given volume (or area, or interval, depending on the case) is—where defined—also constant within and across worlds. If these assumptions are dropped, then an appropriate condition needs to be added to our metaprinciples.

Consider the following example of two worlds having the same locations:

Goodness at Locations (for example, times)

w1: ..., 2, 2, 2, 2, 1, 2, 2, 2, ...

w2: ..., 1, 1, 1, 1, 2, 1, 1, 1, ...

Example 2

Here and throughout, when the locations are the same in the two worlds, we display them in the same column.

BI is silent here, since there is a finite set of locations, namely, the singleton set with the location that has 1 in w1 and 2 in w2, relative to which w2 is better than w1 (and, of course, there are sets relative to which w1 is better than w2). Are there no grounds for finitely additive theories to judge w1 as better than w2? Surely there are! For any finite set of locations containing at least three locations will be such that w1 is better than w2 relative to that set. And for cases such as the one above, that is surely sufficient to judge w1 as better.

The point here is that it is not necessary for judgments of betterness that one world be better than a second relative to all finite sets of locations. It is sufficient for this sort of case, it seems, that any finite set can be expanded sufficiently so that, relative to all further finite expansions, the first world is better. The rough idea is that, if, no matter what finite set of locations you start with, you can expand enough so that, relative to all further expansions, one world is better than another, then the former world is better tout court than the latter world. (In example 2, no matter what finite set one considers, if one expands it to include at least two locations, w1 is better than w2, relative to all further expansions.)

One strengthening of BI, which is very close to being plausible, is:

RSBI (rejected strengthened basic idea): if (1) w1 and w2 have exactly the same locations, and (2) for any finite set of locations there is a finite expansion (superset) such that, relative to all further finite expansions w1 is better than w2, then w1 is better than w2.

As indicated above, RSBI rightly judges w1 as better than w2 in example 2. Indeed, for almost all cases, RSBI gives clearly plausible directives. Unfortunately, there is one sort of technical case where RSBI's directives are incorrect. Consider the following example:

Goodness at Locations (for example, times)

w1: 1, 0, 0, ...

w2: 1/2, 1/4, 1/8, ...

Example 3

Here, assuming that the locations are the same in the two worlds, any finite set of locations can be expanded by adding the $1\text{-over-}1/2$ location (if it is not already included), and, relative to any further finite expansion, w_1 is better than w_2 . Thus, RSBI directs that w_1 be judged as better. But according to both standard and nonstandard mathematics, the total of $1/2 + 1/4 + 1/8 \dots$ is 1, and so the two worlds contain the same amount of goodness.

The problem that arises for RSBI in example 3 is that, although no matter how one finitely expands a set containing the $1\text{-over-}1/2$ location, w_1 is better than w_2 relative to that set, the relative advantage of w_1 over w_2 decreases the more one expands the set. At the limit, when the set is expanded to include all (infinitely many) locations, the advantage disappears. To deal with this case, let us say that, relative to a finite set of locations, a world, w_1 , is *k-better* than a second world, w_2 , for some positive number k , just in case the overall goodness of these locations in w_1 is at least k units better than their overall goodness in w_2 . (Here, for simplicity, we leave implicit a reference to some fixed scale of goodness.)

We can now formulate a strengthening of BI that we endorse:

SBI1 (strengthened basic idea 1): if (1) w_1 and w_2 have exactly the same locations, and (2) for any finite set of locations there is a finite expansion and some positive number, k , such that, relative to all further finite expansions, w_1 is *k-better* than w_2 , then w_1 is better than w_2 .

SBI1 is silent in example 3. For no matter what number, k , is chosen (no matter how small), and no matter how an initial finite set is expanded (even if the $1\text{-over-}1/2$ location is added), one can further expand by adding enough of the remaining locations with the largest values in w_2 so that the total in the w_1 locations is *not* at least k units of goodness greater than that in w_2 (and vice versa). And in example 2, SBI1 correctly directs finitely additive theories to judge w_1 as better than w_2 (since for any finite set containing at least three locations w_1 will have a total that is at least 1-better than w_2).

In what follows, we shall use the '*k-better*' terminology in the official statements of our metaprinciples, but for stylistic reasons we shall leave it to be understood implicitly in our discussions. Readers unconcerned with cases where betterness disappears at the limit may safely ignore *k-betterness* and think solely in terms of betterness.

For locations (like people and states of nature, but perhaps unlike spatial and temporal locations) that have no essential natural order (a notion that we shall explain below), SBI1 says all that we have to

say. We shall now strengthen SBI1 for cases where locations have certain sorts of essential natural order.

III. AGGREGATING WHEN LOCATIONS HAVE AN ESSENTIAL
NATURAL ORDER AND ARE THE SAME

SBI1 has bite only when every finite set of locations can be expanded enough so that, relative to all further finite expansions, one world is better than another. We think that SBI1 can be plausibly strengthened by making it applicable to a wider range of cases. Consider the following example:

Goodness at Locations (for example, times)

w1: ..., 5, 1, 5, 1, 5, 1, 5, 1, 5, ...

w2: ..., 3, 2, 3, 2, 3, 2, 3, 2, 3, 2, 3, ...

Example 4

In this example, no matter what finite set one selects, and no matter how one finitely expands initially, there are further finite expansions relative to which w1 is better (one just expands by adding enough 5-over-3 locations), and there are finite expansions relative to which w2 is better (one just expands by adding enough 1-over-2 locations). Thus, SBI1 is silent.

Nonetheless, it may be appropriate for finitely additive theories to judge w1 as better than w2. For if the (for example, temporal) locations have an essential natural order (as characterized below), then not all expansions of a given set are normatively relevant. For example, if one started with a singleton 5-over-3 set, then no expansion that includes another 5-over-3 location without including the intermediate 1-over-2 locations is normatively relevant. For such expansions violate the essential natural order of the locations (they skip over locations). Thus, we claim, if the locations have an essential natural order, as arguably spatial and temporal locations do, then we can strengthen our metaprinciple by considering only certain kinds of order-respecting expansions.

In order to develop this idea, we need to explain the idea of locations having an order and the idea of this order being both natural and essential. We postpone discussing the notions of naturalness and essentialness until after it is clear how they will be used, and start by explaining the idea of locational order.

The notion of locational order that we have in mind is that of a topological manifold. We shall not define it precisely, but the rough idea is that locations are connected to each other so that the notion of a (continuous, or unbroken) path is well defined and all locations are path connected. Loose marbles, for example, have no order, but they do when glued in a line on a rubber yardstick so that every mar-

ble touches its two adjacent neighbors. Movement from the marble at one end of the yardstick to the marble at the other end via adjacent marbles traces a path, but movement from one marble to a nonadjacent one is not a path (since there is a gap). This idea of locational order does not require there to be a well-defined, or fixed, distance between any two locations, since the paths (as on a rubber yardstick) may be contractible, expandable, or twistable.

If there is locational order (topological structure)—as with spatial and temporal locations, for example—the notion of a *bounded region* is well defined. Roughly, a bounded region is a set of locations that are all “inside a boundary.” For the one-dimensional case, a bounded region is a set of all the locations inclusively between some two locations (for example, the interval from 2 to 4 on the real line). For the two-dimensional case, a bounded region is a set of all the locations inclusively inside some simple closed path (for example, circle or rectangle). (A simple closed path (curve) is a path starting and ending at the same point without passing through any other locations twice.)

Locations with such topological order may either be *discretely connected* (that is, such that there are well-defined adjacent locations, and thus only finitely many locations between two locations on a given path), or *densely connected* (that is, infinitely many locations between any two locations on a given path).⁹ A bounded region contains finitely many locations in the discrete case but infinitely many in the dense case.

For expositional brevity, we are going to engage in some double talk. For from now we shall use the term ‘bounded region’ in a looser sense than that just defined. We shall understand a *bounded region* in this loose sense to be (1) a bounded set in the strict sense, if there is an essential natural locational order, and (2) any finite set, if there is no essential natural order. This will permit us to use the same language to cover both sorts of cases.

We claim that, where there is an essential natural order, not all logically possible finite expansions need to be considered. Only *bounded regional expansions* (that is, expansions that are bounded regions in the strict sense as defined by the topology) need be considered. SBI1 can, we claim, be strengthened as follows:

⁹ For the discrete case, we assume that each location has only a finite set of adjacent locations, and any simple closed path (that is, sequence of adjacent locations with no repetition except that the first location is also the last; for example, a circle or square) separates the set of all other locations into disconnected sets (that is, such that no location inside is adjacent to any location outside). This condition holds automatically in the case of dense topology.

SBI2 (strengthened basic idea 2): if (1) w_1 and w_2 have exactly the same locations, and (2) for any bounded region of locations there is a bounded regional expansion and some positive number, k , such that, relative to all further bounded regional expansions w_1 is k -better than w_2 , then w_1 is better than w_2 .¹⁰

Applied to example 4, this metaprinciple is silent, if the locations have no essential natural order (for example, if they are people), but it directs finitely additive theories to judge w_1 as better if the locations are the same, have an essential natural order (for example, if they are temporal), and are listed in their natural order. For no matter what bounded region one starts with, if one expands finitely leftward or rightward (or both) with no gaps (a bounded regional expansion) so as to include at least two 5-over-3 locations, then relative to any further such expansion w_1 is better than w_2 . The fact that it is logically possible to expand a given finite set of locations by adding *only* 1-over-2 locations (and thus produce a higher total for w_2) is, we claim, irrelevant, because this violates the essential natural order of temporal locations (by skipping over locations). SBI2 rightly directs that w_1 be judged better than w_2 .¹¹

Below, we shall further strengthen this metaprinciple. First, however, we need to say something about a location order being essential and natural.

Naturalness is the most difficult notion—indeed, we are embarrassed to say that we cannot give a crisp definition of what we mean by it! But the intuitive idea, we think, can be made reasonably clear. Locations can be ordered in all sorts of ways. People, for example, can be ordered by their dates of birth, their dates of death, their first names, their last names, and so on. At most one of these orderings is natural in the sense of reflecting an ontologically fixed order. In the case of people, we are inclined to think there is no natural order at all. Temporal and spatial locations, on the other hand, may well

¹⁰ SBI2 is a generalization of a principle (PMU*) formulated and defended in Vallentyne, "Utilitarianism and Infinite Utility." It is a generalization in three respects: first, it applies to any finitely additive theory of the good, and not just utilitarianism; second, it is applicable to any sort of location that has an essential natural order—not just time; third, it is applicable to worlds that extend infinitely in all directions, and not just to those which are infinite only "toward the future."

¹¹ When the locations are discretely connected, SBI2 can have implications without there being a distance metric. This is because a bounded region will have only finitely many locations, and their values can simply be added together. When the locations are densely connected, however, SBI2 has no implications without a distance metric. For a distance metric is required to aggregate the infinitely many locations in a dense bounded region (without a distance metric, the dense interval of 1 to 2 is indistinguishable in "size" from the dense interval from 1 to 20, for example).

have natural orderings in this sense. For at least on a common conception of time and space, it seems plausible to think that they are ontologically fixed in a certain order. You cannot “get” from year 1994 to year 2000 without “passing through” the intermediate years.

Strictly speaking, we want to allow the naturalness of an ordering to be world relative. But for simplicity we shall consider only natural orderings that are essential, that is, the same in all worlds. Although it seems somewhat strange to think of natural ordering as being nonessential, we do not see how to rule out this possibility. Instead, we shall simply restrict the scope of our metaprinciple below to essential natural orderings. At least where there is an essential natural order, it is, we shall argue, possible to strengthen the metaprinciple above.

We hope that these sketchy remarks on essential natural orderings are enough to indicate that there is something to this notion. We acknowledge that the notion remains somewhat mysterious, but shall not attempt here to make it more precise. Below, we shall typically assume that temporal and spatial locations do, and people do not, have an essential natural order; but this is for the sake of illustration only. Strictly speaking, our strengthenings of SBI1 apply to whatever locations have an essential natural order. If it turns out that no locations have such order, then the strengthenings have no applicability.

Although SBI2 is an improvement over SBI1, it can be strengthened further. To see this consider the following two-dimensional example (where there is one unit of goodness at each location except for two rows of -1):¹²

Goodness at Locations									
w1:									
...	1,	1,	1,	1,	1,	1,	1,	...	
...	1,	1,	1,	1,	1,	1,	1,	...	
...	-1,	-1,	-1,	-1,	-1,	-1,	-1,	...	
...	-1,	-1,	-1,	-1,	-1,	-1,	-1,	...	
...	1,	1,	1,	1,	1,	1,	1,	...	
...	1,	1,	1,	1,	1,	1,	1,	...	
.....									
Example 5									

Is w1 better than *its zero world* (that is, a world with the same locations but with 0 goodness at each location)? SBI2 is silent here, even assuming that the locations have an essential natural order (with each location being adjacent to just the locations immediately above, below, and to the side). For no matter what bounded region one starts with, and no matter

¹² The example is a modification of one provided by an anonymous editor of this JOURNAL, which helped us see the need for a strengthening.

how one finitely regionally expands initially, one can regionally expand along the two -1 rows enough to produce a negative total, and one can also regionally expand along two adjacent columns enough to produce a positive total. Consequently, SBI2 is silent. We think, however, that w_1 is indeed better than its zero world. After all, any two rows of 1s will cancel out the two rows of -1 and there will be infinitely many rows of 1s left.

SBI2 is silent because it considers all bounded regional expansions relevant—no matter how selective they are in the directions along which they expand. The expansion along the rows of -1 s, for example, is deemed relevant (and is enough to block SBI2 from directing that w_1 be judged better than w_2). We claim that SBI2 can be plausibly strengthened by restricting its appeals only to what we call uniform expansions, which expand in all directions “uniformly.”

If there is an essential natural distance metric for locations, then a uniform expansion is simply one that adds a band of constant width to the initial region. If there is no essential natural distance metric, but locations are discretely connected, then a uniform expansion can be defined as follows. For the one-dimensional discrete case a uniform expansion of a given region (interval) is simply an expansion that adds on the same number of locations to the right as to the left. For the two-dimensional discrete case a uniform expansion of a given region is a region obtainable from the initial one by successively adding on layers of locations, where a layer around a given region is a smallest (in terms of the number of locations) simple closed path that has the given region strictly on the inside. A *simple closed path* (for the discrete case) is a sequence of adjacent locations that starts and ends at the same location with no repetition of intermediate locations (for example, a circle or rectangle). There may be more than one smallest simple closed path with the initial region strictly on the inside, and in that case all such paths are layers.¹³

Where locations have no essential natural order, or where their essential natural order is dense but without an essential natural distance metric, there is no appropriate way of strengthening SBI2. For

¹³ For simplicity, we assume that the locations of a world are unbounded in all directions. Given that the problem of infinite aggregations can arise whenever the locations are unbounded in some directions (even if not all), this is a simplification. If the assumption is dropped, the definition of a layer around a given region would need to be modified to read “a smallest simple closed path that has the initial region strictly on the inside, except perhaps for those locations in the initial region that are on a world boundary” (to recognize that no expansion is possible in certain directions beyond a boundary). Also recall that for simplicity we are assuming throughout that all locations in all worlds have the same weight, and that the weight of the locations per unit of volume/area/distance is uniform. If this assumption is dropped, the definition of layers would need to appeal to expansions of uniform weight in all directions.

expositional brevity, in such cases we shall engage in some double talk, and understand a uniform expansion to be any finite superset.

We can now strengthen SBI2 to:

SBI3 (strengthened basic idea 3): if (1) w_1 and w_2 have exactly the same locations, and (2) for all bounded regions of locations there is a bounded uniform expansion and a positive number, k , such that, relative to all further bounded uniform expansions, w_1 is k -better than w_2 , then w_1 is better than w_2 .¹⁴

This is the same as SBI2 except that the appeal to regional expansions has been replaced by an appeal to uniform expansions.

In example 5 above, SBI3 rightly directs finitely additive theories to judge w_1 as better than its zero world (where all locations have zero goodness). For no matter what bounded regions one starts with, if one regionally expands uniformly until the region contains more 1s than -1s, then any further finite uniform expansion will contain more 1s than -1s (and thus be better than its zero world). Thus, SBI3 directs that finitely additive theories judge w_1 as better than its zero world. This judgment, we claim, is plausible.¹⁵

¹⁴ A slightly stronger principle is the following: if (1) two worlds have exactly the same locations, (2) there is a bounded region of locations, and a positive number, k , such that relative to all bounded uniform expansions, w_1 is k -better than w_2 , but (3) not vice versa (that is, (2) with w_1 and w_2 permuted), then w_1 is better than w_2 . This is like the principle in the text except that it requires only that there be *some* initial region relative to which all bounded uniform expansions favor w_1 , and in that it has a 'and not vice versa' clause (needed to avoid contradictory judgments). It directs finitely additive theories to judge $\langle \dots 1, -1, 1, -1, 2, -1, 1, -1, 1, \dots \rangle$ as better than its zero world (since, relative to all bounded uniform expansions of the $\langle -1, 2 \rangle$ region, w_1 is better than w_2 , and the not vice versa clause is satisfied). SBI3, on the other hand, is silent (since no bounded uniform expansion of a -1 location makes w_1 better for all further bounded uniform expansions). (Both principles are silent if the 2 is replaced by a 1, and both judge the world better than its zero world if the 2 is replaced by a 3.) Although this principle seems plausible to us, we have not been able to prove that it preserves transitivity (nor that it does not). We are indebted to Lauwers for suggesting the modification used in the text so as to ensure the preservation of transitivity.

¹⁵ We mention here a different sort of approach that agrees with many of SBI3's judgments. This approach was suggested to us by an anonymous editor of this JOURNAL, and was developed by Segerberg (his segmentation principle) in "A Neglected Family of Aggregation Problems in Ethics." Here is the principle: if (1) the locations of two worlds are the same, and (2) there is a partition of locations into regions such that w_1 is at least as good as w_2 in every region, but (3) not vice versa (that is, (2) with w_1 and w_2 permuted), then w_1 is better than w_2 . For finitely additive theories, this is a plausible metaprinciple as far as it goes. This metaprinciple agrees with our metaprinciples in the examples discussed. But it is silent (inappropriately, we think) about whether $\langle \dots -1, -1, -1, 2, 2, 2, \dots \rangle$ is better than its zero world. In contrast, SBI3 holds (appropriately) that it is better than its zero world. Note also that, unlike SBI3 (which evaluates larger and larger regions of worlds), the segmentation principle (which evaluates individual segments in isolation) would not be plausible for holistic theories of the good (such as strict egalitarianism).

Where the locations are the same in both worlds, SBI3 is the strongest metaprinciple that we shall defend. The following is an example where SBI3 rightfully remains silent:

Goodness at Locations (for example, times)

w1: ..., -1, -1, -1, 10, 1, 1, 1, ...

w2: ..., 0, 0, 0, 0, 0, 0, 0, ...

Example 6

SBI3 does *not* direct that w1 be judged better than w2. For although starting with the 10 location, all uniform expansions are better under w1 than under w2, starting with any ten adjacent -1 locations all uniform expansions are better under w2 than under w1. Thus, SBI3 is silent, and this, we claim, is appropriate.

Before defending SBI3, we shall introduce one more strengthening.

IV. AGGREGATING WHEN LOCATIONS HAVE AN ESSENTIAL NATURAL ORDER BUT ARE NOT THE SAME

So far, then, we have strengthened the basic idea to SBI3. SBI3 is applicable only where the locations of the two worlds are the same. It is silent if there is no meaningful basis for the transworld identification of locations, or if there is such a meaningful basis but the worlds do not in fact have the same locations. We shall now generalize SBI3 so as to cover a certain limited range of cases where the locations are not the same.

As indicated earlier, the mere existence of an essential natural order (a topological condition) does not guarantee the existence of an essential natural distance metric. If, however, there is an essential natural distance metric for locations, then we can say that two worlds are *isometric* just in case there is a 1-1 mapping of the locations of one world onto the locations of the other, such that the distance between any two locations in one world is the same as the distance between their mapped counterparts in the other world. (Such mappings will automatically preserve the topological structure.¹⁶)

Where two worlds are isometric, we may appeal to *isometric counterpart functions*, which are simply distance-preserving 1-1 mappings of

¹⁶ For simplicity, we are assuming that locations have a simple and uniform sort of structure. More specifically, we are assuming that (1) within a given world all locations have the same weight and are distributed uniformly by distance, if there is an essential natural distance metric; (2) all worlds have the same structure with respect to the topology, distance metric, and weights of locations. We are restricting our attention to such worlds so as to be able to focus on the core ideas. Things would be a lot more complex without these assumptions.

the locations of one world onto the locations of the another. Typically, there will be many distinct isometric counterpart functions between isometric worlds. If, for example, the locations have the metric structure of the real line, then $f(x) = x$, $g(x) = x+1$, $h(x) = x+2$, and so on, are all isometric counterpart functions.

The idea, then, is that where two worlds are isometric but do not contain the same locations, some comparisons of the goodness of the worlds may be made by comparing isometric counterpart sets. For brevity, we shall formulate our metaprinciple in terms of *counterpart functions* where these are 1-1 mappings of the locations of one world onto the locations of the other (which, unlike isometric counterpart functions, need not preserve distances). Let us call a counterpart function *admissible* just in case (1) it is the identity function if the two worlds have the same locations, and (2) it is an isometric counterpart function if the two worlds are isometric but do not have the same locations.

Our generalized metaprinciple can finally be stated:

GM (general metaprinciple): if (1) w_1 and w_2 have exactly the same locations, or are isometric, and (2) for all admissible counterpart functions, and for all bounded regions of locations in w_1 , there is a bounded uniform expansion and a positive number, k , such that, relative to all further bounded uniform expansions, w_1 restricted to the expansion is k -better than w_2 restricted to the counterpart expansion, then w_1 is better than w_2 .¹⁷

To see the force of this generalization, consider once again example 4, repeated below:

Goodness at Locations (for example, times)

w_1 : ..., 5, 1, 5, 1, 5, 1, 5, 1, 5, ...

w_2 : ..., 3, 2, 3, 2, 3, 2, 3, 2, 3, ...

Example 4 (repeated)

If the locations are the same in each world, then GM, like SBI3, directs that finitely additive theories judge w_1 as better than w_2 . For in this case, the only admissible counterpart functions are identity func-

¹⁷ Throughout, we focus on metaprinciples governing when one world is *better* than a second. Each such metaprinciple has a natural companion metaprinciple governing when one world history is *equally good* as a second: one simply replaces all references to betterness with references to equal goodness. This will leave many worlds that are mutually incomparable (neither better than the other, nor equally good as the other). Alternatively, one could invoke a metaprinciple that judges two worlds equally good just in case neither is better than the other. This would have all worlds as comparable; but the relation of being equally good would not be transitive. We think the former approach more plausible, but shall not defend that view here.

tions, and the force of GM is no different than that of SBI3. If, however, the locations are not the same, and there is no essential natural distance metric, then GM (like SBI3) is silent. In this case, the worlds are not isometric and there are no isometric counterparts. Except in certain special cases,¹⁸ silence is appropriate, we claim, because without sameness of locations or isometry, there may be no way to rule out the possibility that w2 has the same locations as w1, plus a billion clones of each (or vice versa). And if this were the case, it would be a mistake to claim that finitely additive theories should judge w1 as better than w2.

If the locations are not the same, but there is an essential natural distance metric, then GM directs finitely additive theories to judge w1 as better than w2. For no matter what isometric counterpart function one considers (and there are infinitely many), relative to any uniform expansion of any bounded region, w1 restricted to that expansion is better than w2 restricted to the counterpart expansion.

We claim that GM represents a coherent and plausible way for finitely additive value theories to make judgments of goodness when aggregating over an infinite number of locations.

Before turning to a defense of GM, let us consider one final example that illustrates how GM has more power when the locations are the same than when they are merely isometric.

Goodness at Locations (for example, times)

w1: ..., -1, 0, 1, 2, 3, ...

w2: ..., -2, -1, 0, 1, 2, ...

Example 7

If the locations are the same, then GM (like BI) directs that finitely additive theories judge w1 as better than w2, since every single location is better under w1 than under w2. If the locations are merely isometric (not the same), then GM is silent. For although under the isometric counterpart function displayed every location in w1 is bet-

¹⁸ For simplicity, we shall not introduce two slight strengthenings of GM that are plausible. First, GM would remain plausible, we would argue, if it were revised to apply to worlds where all, except perhaps a finite number, of infinitely many locations are the same. For it is only where there are infinitely many nonshared locations that the problems of comparisons arise. Second, GM would remain plausible, if it were revised to apply also where the locations of one world are a subset of the locations in the other world. This is weaker in that it does not require that all locations be the same. It only requires that all the locations of one world (but not necessarily the other world) be in the other world. For finitely additive theories (but not for some types of nonadditive theories such as strict egalitarianism), this can be treated as if all the locations were the same in the two worlds by assigning "empty virtual locations" as counterparts for the locations are in one world but not in the other. For simplicity, we ignore these two strengthenings.

ter than its counterpart in w_2 , under the isometric counterpart function that pairs the 0 location in w_1 with the 1 location in w_2 (and the location with n in w_1 with the location with $n+1$ in w_2), every location in w_1 is worse than its counterpart in w_2 . Consequently, GM is silent if the locations are merely isometric. This silence, we claim, is appropriate.

V. THE PLAUSIBILITY OF GM

In the previous section, we surveyed some examples and saw how GM applied to those cases. This discussion has served two purposes. One was simply to draw out the implications of GM. The other was to illustrate some cases in which GM gave answers that were plausible. We turn now to a general, and more abstract, defense of GM.

In assessing our defense of GM, two points should be kept in mind. The first is that there is nothing contradictory in combining GM with a finitely additive theory of value. By definition, finitely additive value theories are theories that evaluate worlds with a finite number of locations by adding up the total amount of local goodness. This is compatible with any sort of approach when there are an infinite number of locations.¹⁹ It is typically implicitly assumed that the most natural and plausible way of evaluating goodness in such cases is on the basis of the total amount of local goodness. The examples of the previous section were intended in part, however, to provide evidence that, in order to be plausible, a finitely additive theory must assess infinite cases in a way that appeals to more than standard mathematical totals. GM is being defended as a plausible basis of assessment for such cases.

A second point to keep in mind is that GM only states sufficient conditions for one infinite aggregation to be better than a second. It does not state necessary conditions. We would be very surprised if it cannot be strengthened.

The justification of the plausibility of GM can instructively be given in two steps. The first is to defend BI, of which GM is a strengthening. The second is to defend the strengthening of BI to GM.

BI, recall, is:

BI (basic idea): if w_1 and w_2 have exactly the same locations, and if, relative to any finite set of locations, w_1 is better than w_2 , then w_1 is better than w_2 .

¹⁹ This same issue arises, of course, for probability functions. By definition, probability is finitely additive when the events are exclusive. But it need not be additive for an infinite number of exclusive events. A common view is that the probability that any particular event of an infinite number of exhaustive and mutually exclusive events will occur is zero, but the probability that at least one of them will occur is one. This is incompatible with infinite additivity.

BI is a very weak metaprinciple. It has bite only if the locations are the same in the two worlds and every single location contains more goodness in one world than it does in the other.²⁰

The plausibility of BI rests on the view that no new goodness considerations emerge for infinite sets when all finite subsets agree on what is better. Of course, this is not true as a matter of logic. There is no contradiction in aggregating things one way for all finite cases and differently for the infinite case. We are simply claiming that it is plausible to hold no new goodness considerations emerge in the above manner at the infinite level. Example 1 supports this view.

Now, probably the most powerful objection to BI is that it requires (at least sometimes) that merely shifting goodness from some locations to others can change the assessment of how good a world is. But surely, it might be insisted, mere shifting does not change anything. To see this, consider example 1 once again:

Goodness at Locations (for example, times)

w1: ..., 2, 2, 2, 2, 2, 2, 2, 2, 2, ...

w2: ..., 1, 1, 1, 1, 1, 1, 1, 1, 1, ...

Example 1 (repeated)

Assuming that the locations are the same, but without any other special assumptions, BI directs finitely additive theories to judge w2 better than w1.

The point to note here is that the distribution of goodness in w1 can be obtained from the distribution in w2 by shifting around some goodness (but not adding any). More specifically, w1's distribution of goodness by location can be obtained from w2's by picking an arbitrary location in w2 and shifting 1 unit of goodness to it from its right neighbor.²¹ This yields 2 units of goodness in the selected location and 0 in its right neighbor. Then shift 2 (not just 1) units of

²⁰ For the purposes of expositional simplicity, we have ignored a metaprinciple that is weaker than BI. This weaker principle says that, if w1 and w2 have the same locations and every single location in w1 is better than every single location in w2, then w1 is better than w2. This metaprinciple would agree with BI in judging w1 as better than w2 in example 1 (where each location is 2-over-1), but it would not join BI in judging w1 better than w2 in a case where every location is either 2-over-1 or 4-over-3. For although each location is better in w1 (and thus BI judges w1 as better), some locations in w1 (namely, those containing 2 units of goodness) are worse than some locations in w2 (namely, those containing 3 units). Thus, the weaker principle is silent. In arguing that some sort of metaprinciple is plausible for distinguishing between some cases where two worlds have infinite value, the weaker metaprinciple is even more plausible (because weaker) than BI.

²¹ We are not presupposing any natural order here. The order of the listing of locations may be arbitrary. 'Rightward' is to be understood as rightward relative to this ordering.

goodness to the 0 location from its right neighbor. That yields 2 units in the formerly 0 location and -1 in its right neighbor. Then shift 3 (not just 2) units of goodness to the new -1 location, and so on. This leftward shifting process is repeated infinitely many times, and a similar, but rightward infinite shifting processes is done starting to the left of the original location. The net result is w_1 .

So w_1 just is w_2 with an infinite number of finite shifts of goodness. Surely, it might be claimed, BI is thus mistaken to judge w_1 as better than w_2 . Since no new goodness has been added, surely w_1 and w_2 are equally good.²²

BI is not mistaken. Note first that BI agrees that a finite number of shifts in goodness among locations does not change the goodness of a world (since a finite number of shifts will cancel out for any finite set that includes all the affected locations). It only holds that an infinite number of shifts can make a difference. When there are an unbounded infinite number of shifts, as illustrated above, one never has to settle one's accounts properly. One can keep borrowing unlimited amounts from the unbounded infinite pool of untapped creditors in order to give to others. Given this fact, w_1 is indeed better than w_2 .

Another way of making this point is to note that the claim that infinite shifts of goodness do not change the goodness of the world is incompatible with a fundamental metaprinciple of finitely additive theories, namely, that, if the locations are the same in two worlds and every location (for example, every person) has more goodness

²² Note that a special case of shifting utility is permuting utility (that is, simply switching the utility at two or more locations). BI agrees that a finite number of binary permutations make no difference, but it allows that an infinite number of binary permutations can. For example, where the locations are the same and have a complete essential natural order (as in time), BI directs finitely additive theories to judge $\langle 1,1,1,0,1,0,1,\dots \rangle$ as better than $\langle 1,0,1,0,1,\dots \rangle$ —even though the former is just an infinite permutation of the latter (namely, permute leftward all 1s in the latter that initially have a 0 to the left, then permute leftward again all 1s that still have a 0 to the left). Jorge Garcia and Mark Nelson object to this feature of the precursor of GM in their “The Problem of Endless Joy: Is Infinite Utility Too Much for Utilitarianism?” *Utilitas*, vi (1994): 183-92. Vallentyne replies in his “Infinite Utility and Temporal Neutrality,” *Utilitas*, vi (1994): 193-99. In Luc Van Liederkerke, “Should Utilitarians Bother about an Infinite Future?” *Australasian Journal of Philosophy*, LXXIII (1995): 405-07, it is shown that the principle that infinite permutations (and shifts) make no difference in how things are ranked is incompatible with a weak Pareto principle (there called *monotonicity*). (The example just given is taken from that article and shows the incompatibility.) In Vallentyne, “Infinite Utility: Anonymity and Person-Centeredness,” *Australasian Journal of Philosophy*, LXXIII (1995): 413-20, this result is used to support the view that infinite (but not finite) permutations can change how things are ranked (on the grounds that the Pareto principle is more compelling).

in one world than in a second, then the first world is better. (For finitely additive theories, this metaprinciple is equivalent to BI.) Example 1 shows very clearly this incompatibility. For in this example every location is better in w_1 —even though w_1 is “obtainable” from w_2 by means of infinite shifts of goodness. Faced with the choice of rejecting this fundamental metaprinciple or recognizing that infinite shifts can make a difference in goodness, extensions of finitely additive theories will be more plausible if they do the latter.

We conclude, then, that BI is plausible. We turn now to the question of whether our strengthening of BI to GM is plausible. Our most basic claim is that some sort of strengthening of BI is plausible. Again, all the examples of the previous sections other than example 1 illustrate this plausibility. Example 2 (in which all locations are 2-over-1 locations, except for one 1-over-2 location) is an especially compelling example. BI is silent here (since the singleton 1-over-2 set blocks the judgment by BI that w_1 is better, and any 2-over-1 set blocks the judgment that w_2 is better). But surely w_1 is better.

Of course, GM may be mistaken in the manner in which it strengthens BI, but we think it is not. For GM is a strengthening of BI in four ways, and each is plausible. The first (starting with SBI1) is that GM requires only that every set of locations be such that, if expanded enough, relative to all further expansions (and all relevant counterpart functions), the goodness of the expanded set in one world is greater than that of its counterpart set in the other. This permits GM (unlike BI) to judge one world as better than another even if there are a finite number of locations at which it is worse. Example 2 illustrates the plausibility of this increased scope. The second strengthening (starting with SBI2) is that GM judges only bounded regional expansions as relevant. For locations with no essential natural order this changes nothing, since all expansions of a given set are bounded regional expansions (by definitional stipulation) in that case. For locations with an essential natural order, however, this is a strengthening, since it judges some logically possible expansions as irrelevant. And surely this is right. For the fact that there is an essential natural order means that the locations are ontologically fixed in a certain order, and therefore that certain logically possible ways of expanding a given set of locations are incompatible with this natural order. Expansions that violate this natural order can thus be ignored. The third strengthening of GM over BI (starting with SBI3) is that only uniform regional expansions are deemed relevant. Roughly, these are regional expansions that expand in all directions by the same distance (if there is a distance metric) or by

the same number of locations (if there is no distance metric but locations are discrete). The restriction to uniform regional expansions is plausible because it rules out “nonrepresentative” expansions, which expand along some dimensions but not others (and thus fail to represent the structure of values of the world as a whole). The fourth strengthening is that GM is applicable where the locations are isometric, but not identical. Isometric worlds have exactly the same distance relations: the distance between any two locations in one world is exactly the same as the distance between their counterparts in the other. Consequently, it is plausible to make judgments whenever all isometric counterpart functions yield the same judgment. We thus conclude that both BI and the stronger GM are plausible.

VI. CONCLUSION

Finitely additive theories of value rank worlds with a finite number of locations (peoples, times, and so on) on the basis of the total goodness they contain. It is commonly supposed that the only way that such theories can apply to worlds with unbounded infinite numbers of locations is to rank them on the basis of the total goodness they contain. But this has the crazy result that where time (for example) is unbounded, a world with 2 units of goodness at each time is not better than a world with the same locations but only 1 unit at each location. Fortunately, this way of dealing with the infinite case is not necessary for finitely additive theories. And given the implausible rankings it generates, it is not plausible.

We have argued in favor of a different approach. At the most basic level we have defended BI, which says that, if two worlds have the same locations, and, relative to every finite set of locations (and a fortiori, every location), one world is better, then that world is better tout court. Somewhat more tentatively, we have defended the strengthening of BI to GM. GM is a strengthening in that, as long as the two worlds are location isometric, it applies even when the locations are not the same in the two worlds. Furthermore, it does not require that one world be better than the other, relative to all finite sets of locations, but only to all bounded uniform expansions of some bounded uniform expansion of any bounded region of locations. Although we can imagine that, due to a faulty understanding of the issues, we may be mistaken about the plausibility of these strengthenings, we are very confident that some strengthenings of BI are plausible. And it should be remembered that we are only claiming that GM is plausible as far as it goes. We certainly do not think that GM is the strongest plausible metaprinciple for dealing with unbounded locations.

Finally, in closing we note that, although we have only defended GM as a plausible metaprinciple for finitely additive value theories, it is also probably plausible for other kinds of aggregative value theory. For it makes no assumption about how bounded sets of locations are evaluated. It simply says that when a judgment of betterness holds for all bounded sets of a certain sort, then that judgment holds for the entire unbounded set as well. We think, for example, that GM is plausible for maximin theories of value (which evaluate the goodness of a world on the basis of the minimum goodness at any location), for average theories (which evaluate the goodness of a world on the basis of the average goodness at locations), and for strict egalitarian theories (which evaluate the goodness of a world on the basis of how equally goodness is distributed). But seeing how far GM is applicable to these and other theories of value is something that must await another occasion.

PETER VALLENTYNE

Virginia Commonwealth University

SHELLY KAGAN

Yale University