

Infinite Utilitarianism: More Is Always Better*

Luc Lauwers

Center for Economic Studies, K.U.Leuven

Peter Vallentyne

Department of Philosophy, University of Missouri-Columbia

Economics and Philosophy 20 (2004): 307-330.

Abstract: We address the question of how finitely additive moral value theories (such as utilitarianism) should rank worlds when there are an infinite number of locations of value (people, times, etc.). In a recent contribution, Hamkins and Montero have argued that Weak Pareto is implausible in the infinite case and defended alternative principles. We here defend Weak Pareto against their criticisms and argue against an isomorphism principle that they defend. Where locations are the same in both worlds but have no natural order, our argument leads to an endorsement, and strengthening, of a principle defended by Vallentyne and Kagan, and to an endorsement of a weakened version of the catching-up criterion developed by Atsumi and by von Weizsäcker.

1. Introduction

Call a theory of moral value *finitely additive* just in case it morally ranks worlds with finitely many locations of value on the basis of the sum of the values at each location. Utilitarianism is a paradigm finitely additive theory of value (with people as the locations of value, and welfare as the conception of

value). There has been much debate about how finitely additive moral theories should rank worlds when there are *infinitely* many locations of value. The problem is that when there are infinitely many locations of value, the sums may be infinite or may not exist at all and thus may provide no guidance for ranking worlds with infinitely many locations. Many authors, however, believe that not all worlds with infinite sums are equally valuable. Vallentyne and Kagan (1997), for example, have argued that a world with an infinite number of people each having two units of value (e.g., happiness) is morally better than a world with the same people but each having only one unit of value. The sums in both cases, however, are infinite.¹

The core idea that Vallentyne and Kagan (and others) appeal to and defend is a Pareto principle:

Weak Pareto: *If two worlds U and V have the same locations, and each location has more goodness in U than it does in V , then U is*

¹ In the 1920s, Frank Plumpton Ramsey discovered the problem of aggregating utilities over an infinite time horizon. This problem was further analyzed in growth theory and in social choice theory; see for instance, Koopmans (1960), Atsumi (1965), Von Weizsäcker (1965), Lauwers (1998), and Fleurbaey and Michel (2003). In philosophy, the problem was independently discovered, and a solution proposed, by Segerberg (1976). More recently, the problem was rediscovered by Nelson (1991), and a solution was proposed by Vallentyne (1993).

*better than V.*²

This principle gives the desired judgement in the above example (that everyone with 2 units is better than the same people with only 1 unit).

Vallentyne and Kagan develop and defend various strengthenings of Weak Pareto.

In a recent contribution, Hamkins and Montero (2000) have argued that Weak Pareto is implausible in the infinite case (although quite plausible in the finite case), and defended alternative principles. We shall here defend Weak Pareto against their criticisms and argue against an isomorphism principle that they defend. Where locations are the same in both worlds, but have no natural order, our argument leads to an endorsement, and strengthening, of a principle defended by Vallentyne and Kagan, and to an endorsement of a weakened version of the catching-up criterion developed by Atsumi and by von Weizsäcker.

Throughout, it's important to keep in mind that our arguments are addressed to those who accept finite additivity of value (e.g., utilitarians). They are not meant to defend finite additivity against those who reject it (e.g., egalitarians). The issue concerns how finite additivity should be extended when there are infinitely-many locations of value.

2. Background

Throughout we focus on the ranking of worlds on the basis of the relation of

² This is a special case of their Basic Idea.

being morally at least as good. We assume, as do Hamkins and Montero (and virtually everyone else), that this ranking of worlds is transitive:

Transitivity: *If a world U is at least as good as V , and V is at least as good as W , then U is at least as good as W .*

As usual, a world is *better* than a second just in case it is at least as good but not vice versa, and a world is *equally good* as a second just in case it is at least as good and vice versa. We do not assume completeness (that, for any two worlds, one of them is at least as good as the other). When there are an infinite number of locations of value, completeness is difficult to achieve in a plausible explicitly stated ranking.³

Although we shall be concerned with cases where the total value in a world is infinite, we limit our attention throughout (as do Hamkins and Montero) to cases where the value *at each location* is finite. We shall, that is, only be considering cases where infinite total value arises from the sum of infinitely many finite values. Moreover, we assume throughout that the number of locations is countably infinite.

³ Fleurbaey and Michel (2003, p. 794) conjecture that there does not exist an explicit description of a complete rule that satisfies loose Pareto (see below, Section 3) and a weak infinite version of Suppes indifference (to wit, indifference to finite permutations). They conjecture that although such rules exist, proof of their existence involves some version of the axiom of choice (or the weaker ultrafilter axiom).

We further assume, again in agreement with Hamkins and Montero, the following principle:

Sum: *If, for each of two worlds, the sum of the values at their locations exists and is finite, then the first world is at least as good as the second world if and only if its sum is at least as great.*⁴

This principle asserts that, even in the infinite case, worlds should be ranked on the basis of their summed goodness—if these sums exist and are finite. The whole problem with ranking infinite worlds (i.e., worlds with infinitely many locations of value) is that often the sums are infinite or do not exist. Where finite sums exist, there is no problem. The ranking should be done on the basis of the sums (as in the finite case). Thus, for example, if a world has an infinite number of people with value (e.g., happiness) levels (on some arbitrary ordering of the individuals) of $1/2, 1/4, 1/8, 1/16$, etc., and a second world has an infinite number of people with value levels of $.9, 0, 0, 0$, etc., then the first world is better than the second world. It's a mathematical fact that the first sum is one and the second is $.9$.

In the infinite case, Sum must be applied with caution. If locations of value have a natural order⁵—as points in time do—then it can be applied

⁴ In Fishkind, Hamkins, and Montero (2002), Sum is considered a rock-bottom principle for additive theories.

⁵ See Vallentyne and Kagan (1997) for a discussion of the tricky notion of “natural” order.

straightforwardly using the natural order. If, however, locations have no natural order (e.g., as in the case of people), then the sum exists and is finite if and only if no matter how the values are ordered, the corresponding series converges to the same finite number (absolute convergence). For example, $1/2+1/4+1/8+1/16\dots$ converges absolutely to 1, but $1/2-1/3+1/4-1/5\dots$ converges only conditionally (to about .307). This latter convergence is conditional because rearranged as $(1/2+1/4)-1/3+(1/6+1/8+1/10)-1/5\dots$ the series diverges to positive infinity. Hence, there is no sum for this series. Furthermore, even if the order is kept fixed, the sum exists only if the grouping of terms (by parentheses) does not affect the result. For example, $1+(-1+1)+(-1+1)\dots$ converges conditionally to 1, but the convergence is not unconditional, since $(1-1)+(1-1)+(1-1)\dots$ converges conditionally to 0. Thus, no sum exists for this series.

Let us now turn to the criticism of Weak Pareto that Hamkins and Montero raise.

3. Isomorphism and Loose Pareto

Hamkins and Montero ask us to consider the assessment of soccer teams in the game of “infinite soccer”. Each team has an infinite number of players, all of whom play at any given time (there are no extras). Furthermore, each player is assumed—solely for reasons of simplicity of illustration—to be equally good at all positions. (The fact that there are specialized positions on a soccer team is irrelevant to the point of the example.) Consider a team for which the talent levels of the players are as follows:

Player:	...	a_{-2}	a_{-1}	a_0	a_1	a_2	...
Team A	...	-2	-1	0	1	2	...

Table 1

This table indicates that player a_k has a talent level equal to k units (on some scale of measurement).

The coach of the team wants to improve his team. There is a trainer able to raise the talent level of each player by one unit. Should the coach engage this trainer? According to Weak Pareto, the answer appears to be yes.

Let A' be Team A after the training. The results would be as follows:

Player:	...	a_{-2}	a_{-1}	a_0	a_1	a_2	...
Team A	...	-2	-1	0	1	2	...
Team A'	...	-1	0	1	2	3	...

Table 2

Applying Weak Pareto to the talent levels of the players, Team A' (i.e. team A after the training) is better than team A. Each player has more skill and this makes the team better. Hamkins and Montero, however, believe this is the wrong answer. They believe that the training has no effect, that Team A' is equally good as Team A. Their claim is based upon the following principle:

Isomorphism: *Any world is equally as good as any isomorphic copy.*

Worlds are isomorphic just in case they have the same structure (or patterns) with respect to value at locations. If there is no natural structure to locations (as in the case of people), then one world is isomorphic to another just in case there is a one-to-one mapping from the locations of one world onto the locations of the other such that the value at a given location in one world is the same as the value at its corresponding location in the other world. (In this case, Isomorphism implies the well known Anonymity condition.) If there is a natural structure to locations (as in the case of points of time), the mapping must also respect that structure (e.g., if time t_1 is one year earlier than time t_2 , then the counterpart of t_1 must be one year earlier than the counterpart of t_2).

Hamkins and Montero argue that Isomorphism is plausible on the ground that it is the pattern of value and not the specific locations of value that matters. In particular, in the case of the infinite soccer team, the team has the same pattern of value after training as it did before. For any given level of talent displayed by a member of the team after training (e.g., that of a_k), there was someone with that level of talent before training (namely a_{k+1}). If the distribution of talent is the same, there is, they claim, no reason to think that the team is any better. Weak Pareto is thus, they claim, mistaken.

We shall argue against Isomorphism. First, however, it's important to note that Weak Pareto cannot be so straightforwardly applied in the above example. When dealing with the infinite case (unlike the finite case), it is extremely important to specify what the *basic locations of value* are (people, points in time, etc.). This is because in the infinite case it is possible, for example, for one world to be better for every single person but worse at every

single point in time.⁶ Hence, it's crucial to know whether persons or points in times are the basic bearers of value. (In the case of moral theory, the most natural assumption is, of course, that people are the basic locations.)

Throughout, the principles that we defend should be understood as applying to basic locations of value—and not to just any arbitrary locations.

In light of this, we need to reconsider the above assessments of the infinite soccer team. In applying Weak Pareto, it was implicitly assumed that individual players were the basic locations of values. This assumption, however, seems implausible. Suppose that each player improves by one unit and then takes on the position previously held by the player who previously had one unit more of talent. In this case, at each position, the talent of the player is unchanged. We fully agree that in this case the team is not improved. Weak Pareto applied to positions (and not individual players) is silent in this case. The values at each position are unchanged, and hence Weak Pareto (which requires that the value be increased at each basic location) does not apply. Consider, however, a version of the above example in which the talent of the player *at each position* is improved by one unit. In this case, assuming—as we shall—that positions are the basic locations of value for assessing the value (talent) of a soccer team, and that positions have at most a merely ordinal order (e.g., no natural center position), Weak Pareto does

⁶ For an ingenious example, see Cain (1995). For analyses emphasizing the importance of the prior determination of the basic locations of value, see Vallentyne (1995) and Mulgan (2002). We here appeal to Mulgan's analysis of the soccer game example.

indeed hold that the soccer team has improved. Isomorphism, on the other hand, holds that the team has not improved. Therefore, there is a genuine conflict between Weak Pareto and Isomorphism. In what follows, then, we shall focus on the “increase at each position” version of this example.

Hamkins and Montero argue that Weak Pareto should be replaced with a related principle (which they call “Fundamental Idea”, but which we re-label to highlight its Pareto character):

Loose Pareto (aka: Fundamental Idea): *If two worlds U and V have the same locations and each location has at least as much goodness in U than it does in V , then U is at least as good as V .*

Compared with Weak Pareto, this principle is weaker in some respects and stronger in others. It is weaker in that when each location has more in U than it does in V , it only requires that U be at least as good as V (as opposed to being better, as required by Weak Pareto). Loose Pareto is stronger than Weak Pareto in that it applies to all cases where each location has *at least as much* goodness in U as it does in V (and not merely to cases where each location has *more* goodness in U).

Loose Pareto is plausible, but so is, we shall argue, Weak Pareto. Isomorphism, on the other hand, is not plausible, and we shall argue against it.

4. Against Isomorphism

The starting point of our argument is the conjunction of Sum and Loose Pareto, which are two basic principles for finitely additive theories. Below we

introduce and defend a third principle that we label “Zero Independence”.

Those who accept Sum, Loose Pareto, and this third principle, must reject Isomorphism.

In our discussion, we restrict Isomorphism to cases where the locations in the two worlds are the same.⁷ Then, Isomorphism-like principles boil down to Suppes-indifference conditions: what type of rearrangements of the values over the set of locations keep a world equally good. Numerous authors have pointed out the conflicts between Suppes-indifference conditions and Pareto principles.⁸ Here, as already announced, we take another route against Isomorphism.

Consider a case involving a soccer team with infinitely many players. The players are all beginners and they each have a talent level equal to zero. A trainer shows up and improves the talent of one player (b_0) by one unit. *All*

⁷ If locations are not assumed to be the same, Isomorphism faces the following additional problem when (i) the (infinite) set of locations of one world constitute a proper subset of locations of a second world, and (ii) each location contains one unit of value. If the locations have no natural structure, the two worlds are isomorphic (because they can be put into a one-to-one correspondence) and thus Isomorphism holds that they are equally valuable. On the other hand, it seems quite implausible for finitely additive value theories to judge the two worlds equally good since one world has all the value of the first world plus much more.

⁸ See, for example, Vallentyne (1995), Van Liedekerke (1995), Ng (1995), and Lauwers (1998).

players stay in the same position. The talent levels of the initial team, Team B , and the team after training, Team B' , are thus as follows:

Position:	...	b_{-2}	b_{-1}	b_0	b_1	b_2	...
Team B	...	0	0	0	0	0	...
Team B'	...	0	0	1	0	0	...

Table 3

Does this training improve the quality of the team? The intuitive answer, without any doubt, is yes. The player at one of the positions has become a better player and no player at any position has become a worse player. Furthermore, Sum makes this judgement as well. The sum of the values of the players at each position in Team B' is one, and the sum for Team B is zero. Given that the sums exist and are finite, Sum says that Team B' is better than Team B .

Consider now some training applied to Team B' . The talent level of each position other than b_0 (the only position whose talent improved from the first training program) is increased by one unit. This produces Team B'' :

Position:	...	b_{-2}	b_{-1}	b_0	b_1	b_2	...
Team B	...	0	0	0	0	0	...
Team B'	...	0	0	1	0	0	...
Team B''	...	1	1	1	1	1	...

Table 4

Intuitively it is plausible that Team B'' is better than Team B' . After all, one position's talent did not change, and everyone else's talent improved. Here, however, we only insist on the weaker claim that Team B'' is at least as good as (as opposed to better than) B' . This judgement is endorsed by Loose Pareto (which is endorsed by Hamkins and Montero).

We have, then, the following intermediate result: B'' is at least as good as B' (by Loose Pareto), and B' is better than B (by Sum). It follows (by Transitivity) that B'' is better than B . Thus, in at least *some* cases, the fact that one world has more value at every location than a second world does guarantee—as claimed by Weak Pareto—that it is a better world. This does not, however, establish Weak Pareto, which makes the much stronger claim that this is always so. Nor does it challenge Isomorphism, since there is no isomorphism involved in this case. We shall argue, however, that if we add one plausible assumption, then Weak Pareto follows, and Isomorphism must be rejected.

The crucial point about the above example is that for both B and B' the sums of the values at the locations are finite and could thus be used to rank them. Let us extend the example to see the significance of this. Suppose that the trainer returns to B'' and once again improves the talent of position b_0 by one unit (just as he did in improving B to B') to produce Team B^* . Suppose further that the trainer returns later to B^* and improves the talent of all

remaining positions by one unit (just as he did in improving B' to B'') to produce Team B^{**} . We have then:

Position:	...	b_{-2}	b_{-1}	b_0	b_1	b_2	...
Team B	...	0	0	0	0	0	...
Team B'	...	0	0	1	0	0	...
Team B''	...	1	1	1	1	1	...
Team B^*	...	1	1	2	1	1	...
Team B^{**}	...	2	2	2	2	2	...

Table 5

In comparing B'' , B^* and B^{**} , Loose Pareto yields the judgement that B^{**} is at least as good as B^* , which is at least as good as B'' —just as it did in comparing B , B' , and B'' . We cannot, however, here apply Sum to get the judgements that B^* and B^{**} are each better than B'' (as we did in the case of B , B' , and B''). This is because the sums are infinite for these three worlds. Surely, however, we should be able to make the comparable judgements here as well. After all, the improvements made in moving from B'' to B^* to B^{**} are the same as the respective improvements made in moving from B to B' to B'' . In both cases, the first improvement was to improve one position's talent by one unit, and the second improvement was to improve the talent of all remaining positions by one unit. So, if B' and B'' are better than B , then B^* and B^{**} are better than B'' .

The following principle captures the idea that the ranking of two

worlds is determined by the pattern of local *differences*. It appeals to the notion of *world value* “*addition*” understood as follows (where we also define world value “*subtraction*” for future reference). For two worlds, U and V , with the same locations, $U+V$ (respectively, $U-V$) is a world with the same locations, with the value at each location equal to value at U at that location plus (respectively, minus) the value at V at that location. For example, if U has 4 at each location, and V has 3, then $U+V$ has 7 at each location, and $U-V$ has 1 at each location. Here, then, is the principle:

Zero Independence: *If U , V , and W are worlds with the same locations, then U is at least as good as V if and only if the world $U+W$ is at least as good as $V+W$.*⁹

This principle has the effect of saying that the ranking of two worlds is determined by the pattern of differences in local value. To see this clearly, understand the *zero world* of a given world to be a world with the same locations but with zero units of value at each location. If we let W be $-V$ (the complement of V , i.e., $-k$ at each location at which V has value k), then Zero Independence says that U is at least as good as V if and only if $U-V$ is at least as good as its zero world. The ranking of U and V , that is, depends only on the pattern of *differences* in local value at the two worlds, and not on the particular values present at those locations. Zero Independence says, for example, that,

⁹ Hervé Moulin (1991, page 36) uses ‘Zero Independence’ to characterize finite utilitarianism.

above, the ranking of B^* relative to B'' is the same as that of B^*-B'' relative to its zero world. B^*-B'' just is B' and its zero world just is B . Hence, B^* is better than B'' because B' is, according to Sum, better than B .

Zero Independence on its own has no implications for how any two worlds are ranked. It is rather a consistency condition that requires that judgements by other principles be made in a certain way.¹⁰ The important point to note is that it allows judgements about a world ranking relative to its zero world to be applied to judgements about the ranking of two other worlds having the same locations and the same relative differences of local value. It thus allows the power of Sum to be exported to cases in which finite sums do not exist. In the above example, it permits B^* to be judged better than B'' even though neither world has a finite sum. It does this because B^*-B'' and its zero world do have finite sums.¹¹

Let us now apply Zero Independence to the original example given by

¹⁰ Of course, the imposition of Zero Independence rules out those rankings that do not satisfy the principle. Nevertheless it is clear that on the basis of Zero Independence alone, nothing can be said about how to rank two particular worlds. Finally, it's worth noting that Zero Independence entails the following Complement Principle, endorsed by Hamkins and Montero: If world U is at least as good as world V , then the complement of V is at least as good as the complement of U .

¹¹ Zero worlds always have a finite sum of zero, but the "value subtraction" of one world from another does not always have a finite sum. $B^{**}-B^*$, for example, does not have a finite sum.

Hamkins and Montero (where the talent level of each position is one unit greater in A' than in A). Team $A'-A$ has the following local values:

Position:	...	a_{-2}	a_{-1}	a_0	a_1	a_2	...
Team $A'-A$...	1	1	1	1	1	...

Table 6

Above, we established that Team B'' was better than Team B (using Sum and Loose Pareto). We can here use the exact same logic to establish that $A'-A$ is better than its zero world (indeed, if the locations are the same, $A'-A$ just is B'' and the zero world of $A'-A$ just is B). Applying Zero Independence we get the result that A' is better than A —as claimed by *Weak Pareto*, and as denied by *Isomorphism*.

More generally, the above reasoning shows that Sum conjoined with Loose Pareto and Zero Independence entails Weak Pareto. Indeed, that conjunction entails the following stronger version of Pareto:

Strong Pareto: *If two worlds U and V have the same locations, every location has at least as much goodness in U as it does in V , and at least one location has more goodness in U than in V , then U is better than V .*

Weak Pareto judges a world as better than a second world if it is better at all

locations, whereas strong Pareto makes this judgement as long as it is better at some locations and no worse anywhere else. Strong Pareto thus implies Weak Pareto but not vice-versa.

Generalizing the arguments of the previous examples, one obtains the following theorem (proof omitted):

Theorem 1: The conjunction of Transitivity, Sum, Loose Pareto, and Zero Independence entails Strong Pareto (and hence Weak Pareto).¹²

In the context of Sum and Zero Independence, then, Loose Pareto—which is endorsed by Hamkins and Montero—leads to Strong Pareto (as well as Weak Pareto). Given that Isomorphism is incompatible with Weak Pareto, the issue here boils down to a conflict between Isomorphism and Zero Independence. We claim that Zero Independence is more plausible, and hence that Isomorphism should be rejected. This claim is, of course, a controversial claim that can be consistently denied. We believe, however, that we can

¹² It might seem strange that we prove Strong Pareto on the basis of Sum and other conditions, given that in finite contexts Sum trivially entails Strong Pareto. In infinite contexts, however, the two principles are independent. Sum, but not Strong Pareto, is silent when one world has an infinite sum and the second world is a Pareto improvement. In any case, the crucial point here is that Hamkins and Montero accept Sum, and we are trying to appeal to conditions that they accept.

provide enough reasons for accepting it within the context of finitely additive value theory. As with most philosophical claims, our claims of plausibility (here and below) must be understood as claims put forward as part of an on-going debate and investigation.

Zero Independence holds that the ranking of two worlds is determined by the pattern of *differences* in local value. This, we claim, is highly plausible in the context of finitely additive value theories. In the finite case, finitely additive value theories always satisfy Zero Independence. Although they typically get expressed as judging a world as at least as good as another (having the same locations) if and only if its *total value* is at least as great, the reference to the total is not needed. An equivalent statement is that one world is at least as good as the second if and only if the *sum of the differences in value* is at least as great as zero. Only the pattern of differences matters. Even in the infinite case, Zero Independence is “partially” implied by Sum and Loose Pareto. Sum ranks U as at least as good as V if and only if Sum ranks $U-V$ as at least as good as its zero world. Moreover, if two worlds U and V satisfy the antecedent clause of Loose Pareto, then Loose Pareto ranks U as at least as good as V if and only if it ranks $U-V$ above its zero world. Zero Independence is thus, we claim, highly plausible for finitely additive theories.

Zero Independence is equivalent to a condition in social choice theory known as Translation Scale Invariance when it is restricted to the case where locations are the same.¹³ This latter condition holds that interlocal comparisons of zero points are irrelevant to the ranking of worlds. The zero

¹³ See, for example, Bossert and Weymark (forthcoming 2003).

point for value at each location, that is, can be set independently of how it is set for other locations (although, of course, when comparing two worlds, the zero point used for a given location in one world must also be used for that location in the second world). For example, if a location has values of 10 in world U and 5 in world V , both measured on the basis of some particular zero point (the same for both worlds), those values could be changed to 7 and 2 (by making the zero point 3 units higher for that location), and this, according to Translation Scale Invariance, would not alter how the two worlds are ranked.

Zero Independence is equivalent to Translation Scale Invariance (restricted to the case where locations are the same), since any change in the zero points for the locations in worlds U and V can, for some W , be represented by $U+W$ and $V+W$. (For example, if there are just two people, and the first person's zero point is decreased by two units, and the second person's zero point is increased by one unit, then the resulting two representations of the value of U and V are simply $U+W$ and $V+W$, where W is $\langle 2, -1 \rangle$.) Zero Independence and Translation Scale Invariance thus each hold that $U \geq V$ if and only if $U+W \geq V+W$.

Translation Scale Invariance (and hence, Zero Independence) is highly plausible for finitely additive value theories. (Recall that our goal is to defend a particular extension of finite additivity, not to defend finite additivity against non-additive theories.) If there is no natural zero point that separates positive from negative value (if there is just more or less value with no natural separating point), then any particular zero point is arbitrary (not representing a real aspect of value). In this case, interlocational comparisons of zero-points are uncontroversially irrelevant. If, on the other hand, there is a natural zero

for value, it is still plausible for finitely additive value theories to hold that it is irrelevant for ranking worlds. What matters (e.g., from a utilitarian perspective), as argued above, are the *differences* in value at each location between two worlds—not the absolute level of values at locations. No interlocational comparison of zero points is needed for this purpose.

Isomorphism is not, we claim, as plausible as Zero Independence. In the finite case, Isomorphism is indeed plausible, but in the finite case it is not possible for there to be (1) a structure-respecting correspondence between the locations of two worlds such that the value at each location in one world is the *same* as the value at its counterpart location in the second world (i.e., an isomorphism), and also (2) a structure-respecting correspondence between the locations of two worlds such that the value at each location in one world is *greater than* the value at its counterpart location in the second world. In the infinite case, however, both kinds of correspondences may hold between two worlds (e.g., between the infinite soccer team before and after training). When both kinds of correspondences exist, it is implausible to hold that the two worlds are equally valuable.¹⁴

¹⁴ We do not deny that some isomorphism-like principles are valid. We believe, for example, that the following principle is valid (where an identity-preserving counterpart function is one that maps a location onto itself, if it exists in both worlds): if two worlds are such that *all* identity-preserving and structure-respecting counterpart functions are value-isomorphisms (i.e., ensure that counterparts have the same value), then they are equally valuable. We are skeptical that this principle remains valid if the “all” is weakened to “some”,

Zero Independence is, we conclude, more plausible for finitely additive value theories than Isomorphism. Given that they conflict in the context of Sum and Loose Pareto, Isomorphism should be rejected.

This concludes our defense of Weak (and Strong) Pareto and our criticism of Isomorphism. We shall now identify the ranking relation generated by the conjunction of Sum, Loose Pareto, and Zero Independence and show that it is slightly stronger than a principle defended by Vallentyne and Kagan and that it is equivalent to a weakened version of the catching-up criterion developed by Atsumi and by von Weizsäcker.

5. Differential Betterness and Indifference

The conjunction of Transitivity, Sum, Loose Pareto, and Zero Independence results in the following rule for ranking any two worlds U and V . In accordance with Zero Independence, this ranking can be done by comparing $U-V$ with its zero world. To do this, select those locations in $U-V$ at which there is a non-negative amount of goodness. If there are some such locations, set G equal to their sum. If there are no such locations, set G to zero. G is either a non-negative finite number or infinite. Next, select those locations with a negative amount. If there are some such locations, set L equal to their sum. If there are no such locations, set L equal to zero. L is either a non-positive finite number or minus infinity. Now, we distinguish the following cases:

1. If both G and L are finite numbers, then U is at least as good as V if and

but have not argued this point here.

only if $G+L \geq 0$.

2. If G is infinity and L is a finite number, then U is better than V .
3. If G is a finite number and L is minus infinity, then V is better than U .
4. If both G and L are infinite, then no judgement is made.

This principle can usefully be broken into the following two principles, where infinity is understood to be greater than any finite number and negative infinity understood to be less than any finite number:

Differential Betterness: For any two worlds U and V having the same locations, if the sum of the non-negative local values of $U-V$ is (finitely or infinitely) greater than the absolute value of the sum of the negative local values of $U-V$, then U is better than V .

Differential Indifference: For any two worlds U and V having the same locations, if the sum of the non-negative local values of $U-V$ is finite and equal to the absolute value of the sum of the negative local values of $U-V$, then U is equally as good as V .

We believe that these principles are plausible. The desirability of moving from world U to world V is determined by whether the gains exceed the losses. If the total gains G and the total losses L are finite, we follow the classical sum rule and judge V as better (respectively: equally good, worse,) than U just in case the gains exceed (respectively: equal, are less than) the losses. If the total gains (respectively: losses) are infinitely large and the total

losses (respectively: gains) are finite, then again, it seems plausible to conclude V is better (respectively: worse) than U . Finally, if both the gains and the losses are infinite, then silence is appropriate.

It's worth noting that Differential Indifference is considerably stronger than the principle (entailed by Sum) that worlds with equal finite sums are equally valuable. To see this, consider the following example:

Locations:	l_1	l_2	l_3	l_4	l_5	...
World X	-1	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{8}$	$\frac{1}{16}$...
World X'	1	1	1	1	1	...
World X''	0	$\frac{3}{2}$	$\frac{5}{4}$	$\frac{9}{8}$	$\frac{17}{16}$...

Table 7

Sum is silent about the ranking of X'' and X' , for each sum is infinite.

Differential Indifference, however, judges them equally good. It does this because $X'' - X'$ is simply X , and X is equally as valuable as its zero world (because for X the sum of its negative values is -1 and the sum of its positive values is 1). This is an example of the power of Zero Independence implicit in Differential Indifference.

In the appendix, we prove the following:

Theorem 2: The conjunction of Transitivity, Sum, Loose Pareto, and Zero Independence entails the conjunction of Differential Betterness and Differential Indifference, but not vice-versa. Nonetheless, they generate the

same ranking rule.

To intuitively explain the content of this theorem, we need to explain two things. One is why the conjunction of Differential Betterness and Differential Indifferences is strictly weaker than the conjunction of our four principles. The other is to explain how two non-equivalent sets of conditions can nonetheless generate the same ranking rule.

The conjunction of Differential Betterness and Differential Indifference is strictly weaker than the conjunction of our four principles. This is to say that each binary relation that satisfies Transitivity, Sum, Loose Pareto, and Zero Independence also satisfies Differential Betterness and Differential Indifference, but not vice versa. The vice-versa claim does not hold because the conjunction of Differential Betterness and Differential Indifference has no implications for the ranking of case 4 (above) pairs (i.e., where the sum of the non-negative differences in value is infinite and the sum of the negative differences in values is infinitely negative), whereas the conjunction of our four principles do have some (conditional) implications for such cases. . Let x and y , y and z , and z and x , each be a case 4 pair. The conjunction of the two differential principles is silent about each of these rankings, and is thus compatible with the supposition that x is better than y , y better than z , but z is better than x (a violation of transitivity). Our four principles—which include Transitivity—however, are incompatible with this supposition. Thus, the conjunction of Differential Betterness and Differential Indifference does not entail the conjunction of our four principles.

Although the conjunction of Differential Betterness and Differential

Indifference is strictly weaker than the conjunction of Transitivity, Sum, Loose Pareto, and Zero Indifference, they generate the same ranking rule in the sense that their *binary* rankings are the same (one world is judged at least as good as another by one conjunction if and only if it is also so judged by the other conjunction). The extra strength that comes from our four principles (and Transitivity in particular) is “conditional”: they generate certain additional results when certain other assumptions are added. This extra strength, however, is not present when no extra assumptions are made. Thus, for example, on their own, our four principles—like Differential Betterness and Differential Indifference—are silent about the ranking of any case 4 pair. The difference between the two sets of conditions emerges only when additional assumptions are made. Thus, if we suppose that x is better than y , y is better than z , and that each is a case 4 pair, then our four principles (via Transitivity) entail that x is better than z , but the two differential principles are silent about this. This difference emerges only when such additional assumptions are made. On their own, they have the same implications for any pair of worlds. Thus, the two sets of conditions generate the same ranking relation.

Let us now compare the latter principles with some other principles in the literature.

6. Comparisons with Vallentyne and Kagan’s SBI1

Differential Betterness, rather interestingly, turns out to be equivalent to Vallentyne and Kagan’s SBI1:

SBI1 (Strengthened Basic Idea One): *If (1) U and V have exactly the*

same locations, and (2) for each finite set of locations there is a finite expansion (i.e., superset) and some positive number, k , such that relative to all further finite expansions U is k -better [defined below] than V , then U is better than V .

The basic idea here is that worlds having the same locations can be compared by comparing the values of *finite subsets* of their locations, which, given finite additivity, are evaluated by adding values together. Roughly, if, relative to certain finite sets of locations, one world is judged better than a second, and this judgement is also true no matter how one finitely expands the selected set of locations, then the first world is better than the second. More exactly, the principle states that if, no matter what finite set of locations you start with, it is possible to expand this set by adding finitely more locations so that at some point, no matter how one further finitely expands the set, relative to the finite number of locations selected, the first world is k -better (i.e., has a total that, on some specified scale, is at least k units higher, for some fixed real number k) than the second, then the first world is better than the second (considering all locations). (The notion of k -betterness is introduced to handle complications that arise when sums of infinitely many values have a finite value.)

Consider, for example, an earlier example:

Position:	...	b_{-2}	b_{-1}	b_0	b_1	b_2	...
Team B	...	0	0	0	0	0	...

Team B'	...	0	0	1	0	0	...
Team B''	...	1	1	1	1	1	...

Table 8

SBI1 agrees with Differential Betterness in judging B'' as better than B' , and B' as better than B . B'' is judged better than B' because, for any set of at least two locations (positions), the total goodness in B'' in the selected locations is at least one unit greater than that in B' . B' is judged better than B because, for any set that includes at least location b_0 , the total goodness in B' in the selected locations is at least one unit greater than that in B .

SBI1, it turns out, is equivalent to Differential Betterness. In the appendix, we prove:

Theorem 3: A ranking rule, for worlds with the same locations, satisfies SBI1 if and only if it satisfies Differential Betterness.

Thus, Transitivity, Sum, Loose Pareto, and Zero Independence entail Differential Indifference and Differential Betterness, the latter of which is equivalent to SBI1. Given that Differential Betterness is more intuitive and easier to understand than SBI1, all future discussion should, we suggest, focus on Differential Betterness. The complex and arcane language of SBI1 no longer needs to be considered.

7. Comparisons with the Catching-Up Criterion

In the economics literature, one of the best known finitely additive criteria for ranking worlds when there an infinite number of locations comes from the work of Atsumi (1965) and von Weizsäcker (1965):

Catching-Up: *One world, U , is at least as good as a second world (having the same locations), V , if and only if the lower limit, as T approaches infinity, of the sum of the values of $U-V$ at locations 1 to T is at least as great as zero (that is, if and only if*

$$\lim_{T \rightarrow \infty} \inf \sum_{t=1, \dots, T} (U_t - V_t) \geq 0).$$

This condition holds that (as with Zero Independence) the ranking of two worlds can be determined by considering the world that results from subtracting, at each location, the values of one world from those of the other and comparing this world with the value of its zero world. Roughly, the idea is that if, when enough locations are considered, the sum of these differences is at least as great as zero, and this remains so no matter how many additional locations are added in, then the first world is at least as good as the second world.

There are three ways in which Catching-Up may be too strong. First, it states necessary and sufficient conditions for being at least as good. Any two worlds that do not satisfy these conditions are deemed incomparable. A weaker view would state the conditions only as sufficient for betterness and leave open whether they may be other grounds as well. Although we shall endorse the stronger view below, we shall start by focusing on this weaker

view.

A second weakening of the criterion is needed to accommodate cases where locations have no natural structure. Atsumi and von Weizsäcker developed their criteria for the case where the locations were future times (assumed to be discrete for simplicity). They thus assume a natural structure for locations including a first location (the first future time relative to a time of evaluation). Once this assumption is dropped, their appeal to a specific numbering of locations (a specific way of ordering them) becomes arbitrary. Furthermore, one cannot simply reformulate this criterion so that it says that if there exists *at least one way* of numbering locations for which his original condition holds, then the first world is at least as good as the second. This approach is incoherent. Consider $\langle 1, 1, -1, 1, 1, -1, 1, 1, -1, \dots \rangle$ and compare it to its zero world $\langle 0, 0, 0, \dots \rangle$. Relative to the order in which the locations are listed, Catching-Up judges the first world as better (the partial sums of the differences are $\langle 1, 2, 1, 2, 3, 2, 3, 4, 3, \dots \rangle$, and thus are always greater than zero). If, however, the locations are listed in a different order, the opposite result can be obtained. For example, if locations have no natural structure, then the first world can also be specified by $\langle -1, -1, 1, -1, -1, 1, -1, -1, 1, \dots \rangle$, and, relative to this enumeration, Catching-Up says that the zero world is better. Hence, the mere existence of a particular enumeration that ensures that the original criterion is satisfied is not sufficient for being at least as good.

The other main way of responding to the absence of a natural order is to reformulate Catching-Up so that it applies only where the original condition is satisfied by *all possible* enumerations of locations. In what follows, we shall

invoke this weakening of the original criterion.¹⁵

A third way that Catching-Up may need to be weakened concerns the issue of whether the locations in the two worlds are the same. Atsumi and von Weizsäcker implicitly assume that they are (and this is plausible for the case of future times starting from a common point), but in general this need not be so. Ranking infinite worlds with different locations is a tricky business, and we cannot here discuss the many relevant issues.¹⁶ We shall therefore simply weaken Catching-Up so as to be conditional on the locations being the same in both worlds.

Taking into account the above three weakenings of Catching-Up leads to the following:

Weak Catching-Up: *If (1) U and V have the same locations, and (2) for all possible enumerations of locations, the lower limit, as T approaches infinity, of the sum of the values of $U-V$ at locations 1 to T is at least as great as zero (i.e., $\lim_{T \rightarrow \infty} \inf \sum_{t=1, \dots, T} (U_t - V_t) \geq 0$), then U is at least as good as V .*

In the appendix we prove:

¹⁵ We are indebted to an anonymous referee for suggesting that we pursue the connection between our conditions and the Weizsäcker condition weakened in the way just identified.

¹⁶ See, for example, the discussion in Vallentyne and Kagan (1997), where they limit assessment to worlds with a natural distance metric for locations.

Theorem 4: The conjunction of Transitivity, Sum, Loose Pareto, and Zero Independence entails Weak Catching-Up, but not vice-versa. Nonetheless, they generate the same ranking rule.

Given Theorem 2, it immediately follows that the ranking rule generated by the conjunction of Differential Betterness and Differential Indifference is the same as that generated by Weak Catching-Up. Indeed, in the appendix, we prove the following slightly stronger result:

Theorem 5: A ranking rule, for worlds with the same locations, satisfies Weak Catching-Up if and only if it satisfies Differential Betterness and Differential Indifference.

8. Necessary Conditions for Being At Least as Good

The conjunction of Transitivity, Sum, Loose Pareto, and Zero Independence implies each of the following (equivalent) *sufficient* conditions for being at least as good: Weak Catching-Up, the conjunction of Differential Betterness and Differential Indifference, and the conjunction of SBI1 and Differential Indifference. Each of these is silent about the ranking of worlds that do not satisfy the imposed condition. More specifically, as discussed in Section 5, each is silent about the ranking of a world U and a world V if and only if (1) the sum of the non-negative values of $U-V$ is infinite, and (2) the sum of the negative values of $U-V$ is infinitely negative. Call two worlds that satisfy these two conditions a *double infinity pair*. Worlds $\langle \dots 1, 1, 1, 1, \dots \rangle$ and

$\langle \dots 2, 0, 2, 0, 2, 0, 2, \dots \rangle$ are, for example, a double infinity pair.

The above conditions are silent for double infinity pairs. This means that they are compatible with the two worlds being incomparable, but also compatible with one of them being at least as good as the other. We believe, however, that, *if locations have no natural structure*, then double infinity pairs should be judged incomparable. Hence, we believe, if locations have no natural structure, then the above conditions should be strengthened to necessary and sufficient conditions.

We shall not, however, attempt to defend this view here. Instead, we shall simply show that it follows from the above conditions if one final seeming plausible condition is added. The condition appeals to the notion of a *restricted transfer*, which is (1) a transfer of a positive amount of value from a location with positive value to a location with negative value such that (2) after the transfer, the donor location still has non-negative value and the recipient location still has non-positive value. For example, the move from $\langle -1, 3 \rangle$ to $\langle 0, 2 \rangle$ is a restricted transfer, but the move from $\langle -1, 3 \rangle$ to $\langle 1, 1 \rangle$ is not.

Consider then:

Restricted Transfers: *If locations have no natural structure, then, for any three worlds, U , U^* , and V , having the same locations, if (1) U is better than V , and (2) U^* is obtainable from U by some (possibly infinite) number of restricted transfers, then U^* is better than V .*

This condition is, we believe, plausible. Note first that the conjunction of Transitivity, Sum, Loose Pareto, and Differences entails that a *finite* number

of transfers (restricted or not) always preserve value rankings. The only issue, then, concerns infinite numbers of transfers. Clearly, many kinds of infinite transfers do *not* preserve value rankings (and are not covered by Restricted Transfers). For example, an infinite number of transfers from one positively valued location to another *positively* valued location does not always preserve value rankings. World $\langle \dots 1, 1; 1, 1, 1, \dots \rangle$ is (by Strong Pareto) better than $\langle \dots 1, 1; 0, 1, 1, \dots \rangle$, but the latter is not better than itself—even though it can be obtained from the former by transferring, for all locations to the right of the semi-colon (here used only to identify a specific location), one unit of value one location to the right. An unrestricted transfer principle is for these reasons implausible.

Restricted Transfers are uni-directional: they always have the effect of moving both the donor and recipient locations closer to zero. The effect of one restricted transfer on a given location cannot be reversed by some other restricted transfer. We shall not, however, here insist on the plausibility of Restricted Transfers. We wish merely to note that, in the presence of the other four conditions (Transitivity, Sum, Loose Pareto, and Zero Independence), it requires incomparability in all those cases where the four principles are jointly silent. For this purpose, consider the following condition:

Full Weak Catching-Up: *If U and V have the same locations, and locations have no natural structure, then U is at least as good as V if and only if, for all possible enumerations of locations, the lower limit, as T approaches infinity, of the sum of the values of $U-V$ at locations 1 to T is at least as great as zero (i.e., $\lim_{T \rightarrow \infty} \inf \sum_{i=1, \dots, T} (U_i - V_i) \geq 0$).*

This is exactly like Weak Catching-Up, except that (1) it asserts that satisfaction of the limit condition for all possible enumerations is necessary and sufficient (and not merely sufficient) for being at least as good, and (2) it is conditional on locations not having any natural structure. The second qualification is necessary, since where locations have a natural structure it is not necessary that the limit condition hold for all possible enumerations. It only needs to hold for certain kinds of (e.g., isometric) structure-preserving enumerations.

We can now characterize the Full Weak Catching-Up (proof in appendix):

Theorem 6: There is only one ranking rule, for worlds with the same locations, that satisfies Transitivity, Sum, Loose Pareto, Zero Independence, and Restricted Transfers. It is the Full Weak Catching-Up rule.

Where locations are the same and have no natural structure, satisfaction of the von Weizsäcker limit condition for all possible enumerations—which is equivalent to the conjunction of the conditions imposed by Differential Betterness and Differential Indifference—provides necessary and sufficient conditions for being at least as good. We believe that Restricted Transfers, and hence Full Weak Catching-Up, is plausible, but cannot here undertake a defense of Restricted Transfers.

9. Conclusion

We have argued that Hamkins and Montero are mistaken to reject Weak Pareto and endorse Isomorphism. We did this by defending Zero Independence, which ranks U and V on the basis of their differences in local value. The conjunction of Zero Independence with Sum and Loose Pareto (which Hamkins and Montero endorse) entails Weak Pareto (indeed Strong Pareto) and is incompatible with Isomorphism.

We also showed that the conjunction of Sum, Loose Pareto, and Zero Independence entails the conjunction of Differential Betterness and Differential Indifferences (each of which ranks two worlds U and V on the basis of $U-V$). We further showed that Differential Betterness is equivalent to Vallentyne and Kagan's SBI1, and that the conjunction of Differential Betterness and Differential Indifference is equivalent to a weakened version of Catching-Up developed by Atsumi and by von Weizsäcker. This general convergence of results from different sources thus provides some indirect support for each. In particular, it provides indirect support for Zero Independence, since, in the framework of finitely additive theories, Transitivity, Sum, and Loose Pareto are not very controversial.

Finally, we suggested that, where locations have no natural order, the above principles exhaust the plausible judgements. The conditions that they hold are sufficient for one world being at least as good as another are also necessary. We did not defend this suggestion, but we did show that it follows if one further accepts the Restricted Transfers principle that we formulated. Defense of this principle must await another occasion.

References:

- Atsumi, Hiroshi. 1965. Neoclassical Growth and the Efficient Program of Capital Accumulation. *The Review of Economic Studies*, 32: 127-136.
- Bossert Walter and John A. Weymark. Forthcoming 2003. Utility in social choice. *Handbook of Utility Theory*, Vol. II., *Applications and Extensions* (Dordrecht: Kluwer) edited by Barbera Salvador, Peter Hammond, and Christian Seidl.
- Cain, James. 1995. Infinite Utility. *Australasian Journal of Philosophy*, 73: 401-404.
- Fishkind, Donniell, Joel David Hamkins, and Barbara Montero. 2002. New inconsistencies in infinite utilitarianism: is every world good, bad, or neutral? *Australasian Journal of Philosophy*, 80: 178-90.
- Fleurbaey, Marc and Philippe Michel. 2003. Intertemporal equity and the extension of the Ramsey criterion. *Journal of Mathematical Economics*, 39: 777-802.
- Hamkins, Joel David and Barbara Montero. 2000. With infinite utility, more needn't be better. *Australasian Journal of Philosophy*, 78: 231-240.
- Koopmans, Tjalling Charles. 1960. Stationary ordinal utility and impatience. *Econometrica*, 28: 287-309.
- Lauwers, Luc. 1998. Social Choice with Infinite Populations. *The International Library of Critical Writings in Economics* 93: *Mathematical Economics I* (Cheltenham: Edward Elgar, 1998) edited by Graciela Chichilnisky, pp. 257-335.
- Moulin, Hervé. 1991. *Axioms of cooperative decision-making*. Econometrics Society monograph no. 15 (Cambridge: Cambridge University Press).

- Mulgan, Tim. 2002. Transcending the Infinite Utility Debate. *Australasian Journal of Philosophy*, 82: 164-177.
- Nelson, Mark. 1991. Utilitarian Eschatology. *American Philosophical Quarterly*, 28: 339-47.
- Ng, Yew-Kwang. 1995. Infinite utility and van Liedekerke's impossibility: a solution. *Australasian Journal of Philosophy*, 73: 408-412.
- Seegerberg, Krister. 1976. A Neglected Family of Aggregation Problems in Ethics. *Nous*, 10: 221-44.
- Vallentyne, Peter. 1993. Utilitarianism and Infinite Utility. *Australasian Journal of Philosophy*, 71: 212-7.
- Vallentyne, Peter. 1995. Infinite Utility: Anonymity and Person-Centeredness. *Australasian Journal of Philosophy*, 73: 413-420.
- Vallentyne, Peter and Shelly Kagan. 1997. Infinite value and finitely additive value theory. *The Journal of Philosophy*, 94: 5-26.
- Van Liedekerke, Luc. 1995. Should utilitarians be cautious about an infinite future? *Australasian Journal of Philosophy*, 73: 405-407.
- Von Weizsäcker, Carl Christian. 1965. Existence of Optimal Programs of Accumulation for an Infinite Time Horizon. *Review of Economic Studies*, 32: 85-104.

Appendix

Theorem 1: The conjunction of Sum, Loose Pareto, and Difference entails Strong Pareto (and hence Weak Pareto).

Proof: Omitted.

Theorem 2: The conjunction of Transitivity (T), Sum (S), Loose Pareto (LP), and Zero Independence (ZI) entails the conjunction of Differential Betterness (DB) and Differential Indifference (DI), but not vice-versa. Nonetheless, they generate the same ranking rule.

Proof: The proof is split up into two parts. Part A shows that the conjunction of T, S, LP, and ZI entails DB and DI. The ‘*not vice-versa*’ part is explained in Section 5 (the conjunction of DI and DB does not entail T). Part B shows that both sets of axioms generate the same ranking rule. Throughout, U and V denote two worlds with the same (countable infinitely many) locations, G is the sum of the non-negative local values (set at zero if there are none), and L the sum of the negative local values, of $U-V$ (set at zero if there are none).

A. The conjunction of T, S, LP, and ZI entails DB and DI.

A1. The conjunction of T, S, LP, and ZI entails DB.

Suppose that U and V satisfy the antecedent clause of DB. We have to show that (according to the conjunction of T, S, LP and ZI) U is better than V .

The hypothesis implies that L is finite and G is (finitely or infinitely) greater than $-L \geq 0$.

If G is finite, then $G+L > 0$ and S implies that $U-V$ is better than its zero world. Apply ZI and conclude that U is better than V . If G is infinite, let G^* be like $U-V$ except with 0 wherever $U-V$ has a negative value, and let L^* be like $U-V$ except that it has 0 wherever $U-V$ has positive value and has the absolute (hence positive) value where the values are negative. This ensures that $U-V = G^* - L^*$, and thus that U and V will be ranked as G^* and L^* are (by ZI). Given that L is finite, L^* is equally as valuable, by S, as any other world with the same finite (positive) total $|L|$. Let W be such a world, with (a) $|L|$ spread out over only locations which have positive value in G^* , and (b) so that the values in W are always non-negative and less than those in G^* . Such a construction is always possible, since G is infinite and L is finite. By S, W is equally good as L^* , and by Strong Pareto (which holds in virtue of Theorem 1), G^* is better than W . Hence, G^* is better than L^* . Consequently, U is better than V .

A.2. The conjunction of T, S, LP, and ZI entails DI.

Suppose that U and V satisfy the antecedent clause of DI. It follows that G and L are finite and $G+L=0$. According to S, we have that $U-V$ is equally good as its zero world. Apply ZI and conclude that U is equally as good as V .

B. Here we prove that the domain of the rule generated by T, S, LP, and ZI coincides with the domain of the conjunction of DB and DI. From Part A it

follows that the domain of the rule generated by T, S, LP, and ZI includes the domain of the conjunction of DB and DI. From their definitions it is clear that DB and DI are silent on a pair if and only if it is a double infinity pair (i.e., a pair, U and V , for which the sum of the non-negative values (G) and the sum of the negative values (L) are each infinite). Hence, it is sufficient to show that the rule generated by T, S, LP, and ZI is also silent on double infinity pairs.

We define a collection of rules that satisfy T, S, LP, ZI, and completeness. The definition is as follows: since we consider infinite worlds with the same countably many locations, the locations can be indexed by the set $N = \{0, 1, 2, \dots, n, \dots\}$ of natural numbers (of course, this numbering is not unique). Next, let F be a free ultrafilter on the set of natural numbers, i.e. F is a collection of subsets of N that satisfy:

- Finite sets do not belong to F ,
- If both A and B are in F , then the intersection $A \cap B$ is in F ,
- If A is in F and B is a superset of A (i.e. B includes A), then B is in F ,
- For each set of natural numbers A , either A or its complement $N \setminus A$ is in F .

The existence of a free ultrafilter follows from the axiom of choice (Zorn's lemma). Each free ultrafilter defines a limit operator.¹⁷ Let $a = (a_0, a_1, \dots, a_n, \dots)$ be a sequence of numbers, then $\lim_F a$ selects the limit point of

¹⁷ See, for instance, Robert E. Edwards, *Functional analysis, theory and applications* (New York: Holt, Reinhard and Winston, 1965) or James Dugundji, *Topology* (New Delhi: Universal Book Stall, 1990).

a subsequence of an indexed by a set that belongs to F . For example, let F contain the set $\{1,3,5,\dots\}$ of odd numbers and let $a = (1,0,2,0,3,0,4,\dots)$, then $\lim_{F} a = 0$. Indeed, the subsequence of the odd indexed locations reads $0,0,\dots,0,\dots$ and has a limit equal to 0. Similarly, in case F' is a free ultrafilter that contains the set of even numbers, then $\lim_{F'} a = +\infty$. Indeed, now we can restrict a to the subsequence $(1,2,3,4,\dots)$ of the even indexed locations. The limit operator \lim_F is well defined in the sense that each sequence obtains a unique value.

We define the following ranking rule:

Let U and V be two worlds with the same (countable infinite set of) locations. Then U is at least as good as V if and only if $\lim_F \Sigma(U_k - V_k) \geq 0$.

This rule satisfies reflexivity, T¹⁸, S¹⁹, LP, and ZI (and hence, by Part A,

¹⁸ Let F be a free ultrafilter and suppose that according to \lim_F world U is at least as good as V and V is at least as good as W . Hence, $\lim_F \Sigma(U_k - V_k) \geq 0$ and $\lim_F \Sigma(V_k - W_k) \geq 0$. Let A collect the locations n for which $\Sigma_{k=0,\dots,n} (U_k - V_k) \geq 0$. Let B collect the locations n for which $\Sigma_{k=0,\dots,n} (V_k - W_k) \geq 0$. Then A and B both belong to F . Therefore, the intersection $C = A \cap B$ also belongs to F . As a consequence, for each n in C , we have $\Sigma_{k=0,\dots,n} (U_k - W_k) \geq 0$. As C belongs to F , we obtain that $\lim_F \Sigma(U_k - W_k) \geq 0$. Therefore, U is at least as good as W .

¹⁹ Here we use the fact that each subsequence of a converging sequence

also DB and DI). In addition, this rule is complete, i.e. this rule is able to judge any pair U, V of worlds. The intersection of all these \lim_F principles includes the ranking rule generated by T, S, LP, and ZI. Furthermore, in defining a \lim_F principle there are two moments of choice (or arbitrariness): first, we number the set of locations, and secondly, we select an ultrafilter.

B1. Now consider a double infinity pair U and V (i.e. both G and L are infinite). In this case there are infinitely many locations at which $U-V$ is positive and infinitely many locations at which $U-V$ is negative. Select a finite set S of locations such that $\sum_S(U_k-V_k) > 1$ (this is possible because G is infinite). In numbering the locations, we start with this set S . Expanding S , we number the locations in such a way that at infinitely many locations the partial sum becomes greater than 1. Let F be a free ultrafilter that contains the set of these locations (where the partial sum is greater than 1). Then, according to the rule \lim_F , world U is better than V . If we repeat this construction with U and V interchanged, then we end up with a rule (that satisfies T, S, LP, and ZI) that considers V better than U (i.e., the opposite ranking). Hence, the conjunction of T, S, LP, and ZI is silent when G and L are each infinite.

Conclusion: The rule generated by S, LP, and ZI is silent if and only if G and L are both infinite, which is precisely when the conjunction of DB and

converges to the same limit. Note also that if the series $\sum(U_k-V_k)$ converges absolutely then the value of $\lim_F \sum(U_k-V_k)$ does not depend upon how the locations are numbered and upon which free ultrafilter is considered.

DI is silent.

Theorem 3: A ranking rule, for worlds with the same locations, satisfies SBI1 if and only if it satisfies Differential Betterness.

Proof: Let U and V be any two worlds, and let G be the sum of the non-negative local values of $U-V$ (set at zero if there are none) and let L be the sum of the negative local values (set at zero if there are none), of $U-V$.

A. DB entails SBI1.

Suppose that the antecedent condition of SBI1 (concerning finite expansions) is satisfied. By definition, L is negative, and there are two cases to consider:

(1) Suppose first that L is finite. In this case, it is sufficient to show that $G+L$ is positive (since DB will then judge U as better than V , as required). We do this by deriving a contradiction from the supposition that G is less than or equal to the absolute value of L . If so, then G and L are both finite and the sum of the values in $U-V$ absolutely converges to $G+L$, which is less than or equal to zero. If so, then for each positive k there exists a finite expansion such that (i) the restriction of $U-V$ to this expansion adds up to a number less than k and (ii) all further expansions generate a total less than k . This is in conflict with the assumption that the antecedent conditions of SBI1 are satisfied. Hence, $G+L$ is positive, and thus DB judges that U is better than V , as required by SBI1.

(2) Suppose next that L converges to minus infinity. If so, then there are infinitely many locations with a negative utility. Take any finite set S of

locations such that the restriction of $U-V$ to this set adds up to a positive number (given that antecedent conditions of SBI1 hold, such a set S must exist). Then add enough locations with a negative amount of utility and this superset will end up with a negative total. This, however, contradicts the assumption that the antecedent conditions of SBI1 hold. Hence, this case is impossible.

Conclusion: Only the first case is possible and in it DB ensures that U is judged better than V . Hence, the consequent clause of SBI1 holds.

B. SBI1 entails DB.

Suppose that the antecedent of DB is true, that is, $G+L$ is (finitely or infinitely) positive, say $G+L>\epsilon>0$. For this to be so, L must be finite, and G is either finite (and greater than $-L$) or infinite. In either case, let $k=\epsilon/2>0$. Then, no matter what finite set of locations one starts with, one can finitely expand it so that, relative to any further finite expansion, the sum of the values in U in this set is k -greater than the sum of the values in V in this set. Hence SBI1 ensures that U is better than V , as required by DB.

Theorem 4: The conjunction of Transitivity (T), Sum (S), Loose Pareto (LP), and Zero Independence (ZI) entails Weak Weizsäcker (WW), but not vice-versa. Nonetheless, they generate the same ranking rule.

Proof: First, observe that WW does not entail T (cf. proof of Theorem 2, Part A). Next, from the definition it follows that when applied to a pair of U and V with G , or L , or both, finite, WW agrees with the conjunction of T, S, LP, and ZI. Indeed, in this case $\lim_{T \rightarrow \infty} \inf \sum_{t=1, \dots, T} (U_t - V_t)$ does not depend upon the

particular enumeration of the locations and is equal to $\lim_{T \rightarrow \infty} \sum_{t=1, \dots, T} (U_t - V_t)$.

Hence, it is sufficient to show that WW is silent for double infinity pairs.

Now, if U and V is a double infinity pair, then one can define a \lim_F principle (cf. Part B1 in the proof of Theorem 2) such that $\lim_F \sum (U_t - V_t) < 0$. Therefore, with respect to the corresponding enumeration, one has that $\liminf \sum (U_t - V_t) \leq \lim_F \sum (U_t - V_t) < 0$. Repeat this construction with U and V interchanged, and obtain another enumeration for which $\liminf \sum (V_t - U_t) < 0$. Conclude that WW is silent for double infinity pairs and that WW coincides with the rule generated by T, S, LP, and ZI.

Theorem 5: A ranking rule, for worlds with the same locations, satisfies Weak Weizsäcker (WW) if and only if it satisfies Differential Betterness (DB) and Differential Indifference (DI).

Proof: Both WW and the conjunction of DB and DI are silent on a pair if and only if it is a double infinity pair. Therefore, the statement is implied by the previous theorem in combination with Theorem 2.

Theorem 6: There is only one ranking rule, for worlds with the same locations, that satisfies Transitivity (T), Sum (S), Loose Pareto (LP), Zero Independence (ZI), and Restricted Transfers (RT). It is the Full Weak Weizsäcker rule (FWW).

Proof: In view of Theorem 4, it is sufficient to show that (i) FWW satisfies RT, and (ii) in case of a double infinity pair, RT ~~turns~~ combined with the

conjunction of T, S, LP, and ZI ~~into~~ entails ~~silence~~ incomparability.

(i) Let (according to FWW) world U be better than V . Theorem 4 implies that L is finite and G is (finitely or infinitely) greater than $-L$. We show that by introducing axiom RT the ranking of U and V is not affected. The total goodness that RT allows to be transferred does not exceed $-L$. After a restricted transfer of K units, U is changed into U^* , G decreases to $G-K$, and L increases to $L+K$ with $-L > K \geq 0$. The inequality $G > -L$ implies $G-K > -L-K \geq 0$. Therefore, U^* is better than V . In this case, RT is entailed by the conjunction of T, S, LP, and ZI. Hence, FWW satisfies RT.

(ii) Consider a double infinity pair U and V , that is, for which G and L are both infinite. According to FWW such a pair is incomparable. We show that the conjunction of T, S, LP, ZI, and RT entails incomparability. This is done by contradiction: assume U is not worse than V . Fix a location l for which $n = U_l - V_l$ is negative. By RT, the local differences at all other locations can (with an infinite number of restricted transfers) be brought to zero without affecting the ranking of U and V . Hence, after these transfers, the local gains add up to 0 and the local losses to n . By S and ZI, world V should be better than U . This contradicts the initial assumption.

Conclusion: the conjunction of T, S, LP, ZI, and RT entails that double infinity pairs are incomparable and elsewhere it coincides with FWW.

* For valuable comments, we would like to thank Bart Capéau, Marc Fleurbaey, Joel Hamkins, Barbara Montero, Tim Mulgan, and two anonymous referees for this journal