Millian Superiorities

GUSTAF ARRHENIUS
Stockholm University

AND

WLODEK RABINOWICZ
Lund University

Suppose one sets up a sequence of less and less valuable objects such that each object in the sequence is only marginally worse than its immediate predecessor. Could one in this way arrive at something that is dramatically inferior to the point of departure? It has been claimed that if there is a radical value difference between the objects at each end of the sequence, then at some point there must be a corresponding radical difference between the adjacent elements. The underlying picture seems to be that a radical gap cannot be scaled by a series of steps, if none of the steps itself is radical. We show that this picture is incorrect on a stronger interpretation of value superiority, but correct on a weaker one. Thus, the conclusion we reach is that, in some sense at least, abrupt breaks in such decreasing sequences cannot be avoided, but that such unavoidable breaks are less drastic than has been suggested. In an appendix written by John Broome and Wlodek Rabinowicz, the distinction between two kinds of value superiority is extended from objects to their attributes.

I. INTRODUCTION

In this article, we distinguish between two forms of superiority relations that can obtain between valuable objects. If a single object of one type is better than any number of objects of another type, the former will be said to be 'superior' to any of the latter. An object is 'weakly superior' to another object if a *sufficient number* of objects of the former type are better than any number of objects of the other type. In an appendix, co-authored by John Broome and Wlodek Rabinowicz, this distinction between two kinds of superiority is extended from objects to their attributes.

Obviously, our discussion is to some extent inspired by John Stuart Mill's famous claim that some pleasures are 'superior in quality' to pleasures of other kinds, but we have chosen to approach the subject of value superiority in a more abstract and general way. The objects that are compared need not be pleasures or even mental states.

Now, suppose one sets up a sequence of less-and-less valuable objects in such a way as to make each element in the sequence only marginally worse than its immediate predecessor. Could one in this way reach an object that dramatically differs in value from the point of departure? Well, why not? Surely, one would have thought, something like this should be possible. However, it has been claimed in the literature that

© 2005 Cambridge University Press
doi:10.1017/S0953820805001494
Utilitas Vol. 17, No. 2, July 2005
Printed in the United Kingdom

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

if there is a radical value difference between the objects at each end of the sequence, then at some point there must be a corresponding radical difference between the adjacent elements: an abrupt break in the decreasing sequence. The underlying picture seems to be that a radical gap cannot be scaled by a series of non-radical steps. We show in what follows that this picture is incorrect if by the radical value difference we mean superiority, but correct if it is weak superiority that is in question. Thus, the conclusion we reach is that, in a sense, abrupt breaks in such decreasing sequences cannot be avoided but these unavoidable breaks are less drastic than it has been suggested.

II. PRELIMINARIES AND DEFINITIONS

Suppose a domain of objects is ordered by the relation of being at-least-as-good-as. This relation, let us assume, is both transitive and complete in the domain under consideration. (The completeness assumption is problematic, but we make it for the sake of simplicity.) Assume that the domain is closed under concatenation, by which we mean the operation of forming 'conjunctive' wholes out of any finite set of objects. Such wholes are themselves objects in the domain. We also take it that any object e in the domain is subject to what we might call 'self-concatenation': For any number m, the domain contains a whole composed of m mutually non-overlapping 'e-objects', by which we mean objects of the same type as e. We take object *types* to be understood in such a way that any two representatives of the same type are equally good and interchangeable in every whole without influencing the value of the whole in question. Intuitively, we might think of objects of the same type as being identical in all value-relevant respects.[1] In what follows, statements such as 'm e-objects are better than k e'-objects' should be read as claims about complex objects obtained by self-concatenation: 'A whole composed of m e-objects is better than a whole composed of k e'-objects.'

It will also simplify matters if we suppose that all the objects in the domain are positively valuable, by which we mean that for any object e and any $m > 1$, m e-objects are better than $m - 1$ e-objects. In other words, concatenating objects of the same type is value increasing. But we allow that the value of the objects in the domain may otherwise vary, and quite dramatically sometimes. Two kinds of such relatively radical value differences will be of special interest in this article: 'superiority' and 'weak superiority'.

[1] Object types should be distinguished from object *kinds*. Two wholes composed of objects of the same type can be said to be of the same kind. Thus, a whole composed of m e objects is of the same kind as a whole composed of n e objects, for any m, n, and e.

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

Definition 1: An object e is *superior* to an object e' if and only if e is better than any number of e'-objects.

Definition 2: An object e is *weakly superior* to an object e' if and only if for some number m, m e-objects are better than any number of e'-objects.[2]

In other words, e is superior to e' if it is better than any whole composed of e'-objects, however large.[3] It is weakly superior to e' if a sufficient number of e-objects are better than any whole composed of e'-objects, however large.[4] Consequently, if e is weakly superior to e', then a whole composed of a sufficient number of e-objects is superior to e'. Thus, the existence of weak superiorities entails the existence of superiorities in the domain, given closure under concatenation.[5]

Both superiority and weak superiority involve therefore violations of the so-called *Archimedean axiom* for betterness orderings. Roughly, that axiom implies that for any objects e and e', if e' is positively valuable, then there is some number k such that k e'-objects are better than a single e.[6]

Along with these two kinds of superiority relations between objects in the domain, we could define the corresponding relations between

[2] Note that, according to this definition, e may be weakly superior to e' even when e' is better than e. If this sounds unnatural, one might always add to the definition an additional requirement that e must be better than e' in order to be weakly superior to it. However, for simplicity's sake, we prefer to work with a more austere definition of weak superiority.

[3] Mill's own label for this kind of superiority is 'superiority in quality'. According to his preference test of superiority, what decides the superiority issue among pleasures is that persons who are 'competently acquainted' with both kinds of pleasures would prefer a single pleasure of one kind to 'any quantity of the other pleasure which their nature is capable of' (J. S. Mill, *Utilitarianism*, in *On Liberty and Other Essays*, ed. J. Gray (Oxford, 1991 [1866]), p. 138).

[4] In J. Griffin, *Well-Being: Its Meaning, Measurement, and Moral Importance* (Oxford, 1986), p. 85, James Griffin seems to have something like weak superiority in mind when he considers a form of 'Discontinuity' in which 'enough of A outranks any amount of B'. On the other hand, superiority is similar to what Griffin (p. 83) calls 'Trumping', which 'takes the form: any amount of A, no matter how small, is more valuable than any amount of B, no matter how large'. For a different interpretation of Griffin's distinction, however, see appendix 3.

[5] We are indebted to Roger Crisp for pressing this point.

[6] More precisely, the Archimedean axiom has this implication provided that e' not only is positively valuable but also such that adding it to any whole does not worsen the value of that whole. In all generality, that axiom states the following:

The Archimedean axiom: For any objects e, e', e'' and e^*, if e'' is better than e^*, then there is some number k such that conjoining e with k e''-objects makes for a better whole than conjoining e' with k e^*-objects.

If we now let $e^* = e'$ and $e'' = 2e'$, then, if e′ is positively valuable, e'' is better than e^*. Therefore, it follows from the Archimedean axiom that for some k, $k+1$ e'-objects are better than a whole composed of e together with k e'-objects. Consequently, if adding e' to any whole never worsens the value of that whole, we reach the desired conclusion that, for some m, $(m = k+1)$, m e' objects are better than a single e.

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

object *types*, one being that any object of a certain type is better than any number of objects of another type, and the other being that a sufficient number of objects of one type is better than any number of objects of another type. In what follows, however, we shall restrict our attention to superiority relations between objects.

A different perspective on superiority relations would involve thinking of objects as exhibiting various value-relevant *attributes*, each of which can be present in an object in varying degrees. As an example, think of an object as a possible outcome that can be characterized in terms of such value-relevant attributes as, say, (the levels of) achievement, satisfaction, freedom, etc. We could then study superiority relations between attributes, rather than between objects themselves (or between object types). An attribute may be said to be superior to another attribute *relative* to an object e if and only if any improvement of e with respect to the former attribute is better than any change of e with respect to the latter attribute. Correspondingly, an attribute is weakly superior to another attribute relative to e if and only if some improvement of e with respect to the former attribute is better than any change of e with respect to the latter attribute. Apart from these superiority relations, which are relative to a specific reference-point (object e), one can also study global superiority relations between attributes, which hold for all reference-points. Such relative and global relations are considered in appendix 3 (written by John Broome and Wlodek Rabinowicz), where it is shown that the results we are about to prove for superiority relations between objects in a large measure extend to the corresponding relationships between attributes.

III. SUPERIORITY WITHOUT ABRUPT BREAKS

By a decreasing sequence $e_1, \ldots, e_n$, we shall in what follows mean a sequence of objects such that e_1 is better than e_2, e_2 is better than $e_3, \ldots,$ and e_{n-1} is better than e_n. As one of us has pointed out in a comment on a paper by Jesper Ryberg, there could exist decreasing sequences in which the first element is superior to the last one but no element is superior to the one that immediately follows.[7] Ryberg denies this possibility.[8] Following an influential tradition, he assumes that e_1 can be superior to e_n only if e_1 is *infinitely* better than e_n.[9] But if the latter is

[7] W. Rabinowicz, 'Ryberg's Doubts about Higher and Lower Pleasures – Put to Rest?', *Ethical Theory and Moral Practice* 6 (2003).

[8] J. Ryberg, 'Higher and Lower Pleasures – Doubts on Justification', *Ethical Theory and Moral Practice* 5 (2002).

[9] Cf. R. Crisp, 'Utilitarianism and the Life of Virtue', *The Philosophical Quarterly* 42 (1992); J. Riley, 'On Quantities and Qualities of Pleasure', *Utilitas* 5 (1993); J. Riley, 'Is Qualitative Hedonism Incoherent?', *Utilitas* 11 (1999); D. Parfit, 'Overpopulation and the

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

the case, then a decreasing sequence that starts with e_1 and ends with e_n must at some point involve an infinite drop in value, i.e. it must at some point reach an element e_i such that e_i is infinitely better than e_{i+1}.[10] This implies that e_i must be superior to its immediate successor e_{i+1}. As Ryberg puts it, '[i]f there is a discontinuity between the values . . . at each end of the continuum, then at some point discontinuity must set in'.[11]

This claim is incorrect. At least, it is incorrect in all domains in which *weak superiority does not collapse into superiority*. More precisely, the following can easily be proved:

> *Observation* 1: Consider any two objects e and e' such that e is better than e'. If e is weakly superior to e', without being superior to it, then the domain must contain a finite decreasing sequence of objects in which the first element is superior to the last one, but no element is superior to its immediate successor.

Proof: Suppose that e is better than and weakly superior to e', without being superior to it. By the definition of weak superiority, there is some

Quality of Life', *Applied Ethics*, ed. P. Singer (Oxford, 1986). It is arguable that infinite or lexical superiority also figures in the section on 'Discontinuity' in Griffin's *Well-Being*. There, he suggests that when 'enough of *A* outranks any amount of *B*' (as in the cases of weak superiority), 'we have positive value [*B*] that, no matter how often a certain amount [of that value] is added to itself, cannot become greater than another positive value [enough of *A*], and cannot, *not because with piling up we get diminishing value* or even disvalue (though there are such cases), but because *they are the sort of value that, even remaining constant, cannot add up to some other value*' (Griffin, *Well-Being*, p. 85, our italics). This can be read as a suggestion that, in those cases, one positive value is infinitely or lexically greater than another. However, as is clear from this quote, Griffin also recognizes cases in which the reason for 'Discontinuity' does not have anything to do with infinite (or lexical) value differences. As for Mill's own views on the matter, we leave that issue to the experts.

[10] It does not really matter in the present context whether one interprets the superiority in question as infinite or as *lexical*. On the lexical view, a superior object carries a value that may be finite but is of a higher order than the value of an inferior object. These two constructions – the infinitistic and the lexical one – are closely related interpretations of the same idea. In particular, both imply that if the first element of the object sequence is superior to the last one, in the requisite sense, then at some point in a decreasing object sequence, a corresponding superiority relation must set in between the adjacent elements. (Parenthetically, however, one should add that the infinitistic and the lexical interpretations of superiority are not equivalent. In fact, the former lends itself to some counter-intuitive implications that the latter avoids. To see this, consider *lotteries*. On the standard view, the value of a lottery that gives a chance p to a prize e equals p times the value of e. In other words, the value of a lottery is its expected value. Consequently, lotteries that give different (positive) chances to an infinitely valuable prize have all the *same* value, independently of chance differences: they are all infinitely valuable. No such counter-intuitive conclusion is forthcoming on the lexical view. On that view, a lottery that assigns a positive chance to a higher-order prize may be lexically better than any lottery with a prize of a lower order, but at the same time still worse than lotteries that assign to the higher-order prize a higher chance.)

[11] Ryberg, 'Higher', p. 418.

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

$m > 1$ such that m e-objects are better than any number of e'-objects. This means (cf. section I above) that the whole composed of m e-objects is superior to e'. Now, consider the following sequence:

$$\begin{aligned} e_1 &= \text{the whole composed of } m \text{ } e\text{-objects},\\ e_2 &= \text{the whole composed of } m-1 \text{ } e\text{-objects},\\ &\dots\\ e_{m-1} &= \text{the whole composed of 2 } e\text{-objects}\\ e_m &= e,\\ e_{m+1} &= e'. \end{aligned}$$

The first object in this sequence is superior to the last one. Furthermore, since self-concatenation is value increasing, each element in the sequence is better than its immediate successor. Thus, the sequence is decreasing. At the same time, no element in the sequence is superior to its successor. In fact, as is easily seen, for all e_k such that $1 \leq k < m$, a whole composed of three e_{k+1} objects is better than e_k. (Such a whole consists of a larger number of e-objects than e_k. Consequently, by the assumption of value increasingness, it must be better than e_k.) The remaining case to consider is when $k = m$, but that e_m is not superior to e_{m+1} is true by hypothesis. Consequently, none of the objects in the sequence $e_1, \dots, e_{m+1}$ is superior to its immediate successor, despite the fact that e_1 is superior to e_{m+1}. This completes the proof.

Well, then, what about Ryberg's argument to the effect that decreasing sequences in which the first element is superior to the last one must contain elements that are superior to their immediate predecessors? As we have seen, his crucial assumption is that superiority involves infinite betterness. However, this assumption, which he shares with other writers on the subject, appears to rest on a presupposition that value is additive. A similar presupposition seems to lurk, for example, behind the following statement of Jonathan Riley:

> Given the hedonist claim that happiness in the sense of pleasure (including the absence of pain) is the sole ultimate end and test of human conduct, there are only two logical possibilities: either qualitative differences [between pleasures] may be reduced to finite amounts of pleasure (for example, one unit of higher pleasure might be deemed equivalent to ten units of lower pleasure), in which case the quality/quantity distinction is epiphenomenal because pleasure is at bottom homogeneous stuff; or qualitative differences are equivalent to infinite quantitative differences, in which case pleasure is a heterogeneous phenomenon consisting of irreducibly plural kinds or dimensions arranged in a hierarchy. The second alternative is embodied in my interpretation. (Riley, 'On Quantities', p. 292)

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

Here, Riley claims that, in the absence of infinite quantitative differences between higher and lower pleasures, a higher pleasure must be equal in value ('equivalent') to a finite number of lower pleasures. For if a higher pleasure only has a finite value, then that value sooner or later would be reached, if we started piling up lower pleasures. This reasoning is correct if value additivity is assumed, but without such an assumption it is a non-sequitur.

To make room for superiority between the extrema of a decreasing sequence without superiority setting in at any point in the sequence, we must give up the infinitistic interpretation of superiority, which requires giving up value additivity. We must allow that the aggregated value of several objects of the same type need not be the sum of the values each of them has on its own. That the value of a whole may differ from the sum of the values of its parts is of course an idea that should be familiar to the post-Moorean value theorists.[12] Among economists, an analogous phenomenon is referred to as *complementarity*. But the economists' standard examples of complementarity involve instrumental values rather than *intrinsic* ones. To illustrate, a knife may be more valuable, instrumentally, than a fork, but once you've got one knife getting another one is much less valuable than getting a fork instead. Moore′s important insight was that the phenomenon of non-additivity can also arise within the realm of intrinsic value.[13]

More precisely, giving up additivity is not enough; we must be even more radical. If we want to have a sequence in which the first element is superior to the last one, without it being the case that any element in the sequence is superior to its immediate successor, we must reject the idea that the value of the whole has to be a monotonically increasing function of the value of its parts. That is, we must give up the *independence axiom* for the betterness ordering:

> *Independence*: If an object e is at least as good as e', then replacing e' by e in any whole results in a whole that is at least as good.[14]

[12] Cf. G. E. Moore, *Principia Ethica* (Cambridge, 1993 [1903]), sects. 18–21, *et passim*. To be more precise, Moore does assume a form of additivity when he suggests that 'the value on the whole' is the sum of (i) the values of the parts *plus* (ii) 'the value of the whole, *as a whole*'. (The latter may be either positive or negative.) But it can be argued that this form of additivity is a purely arithmetical construct. His 'value of the whole, as a whole' could simply be interpreted as the arithmetical difference between 'the value on the whole' and the sum of the values of the parts, independently considered. Moore himself points out that the value of a whole, as a whole, may be 'expressed' as such a difference (ibid., sect. 129), but he seems to ascribe to it some independent significance.

[13] Moore was not the first philosopher to make this point. Another standard reference is F. Brentano, *The Origin of our Knowledge of Right and Wrong* (London and Henley, 1969 [1989]).

[14] Qualification: If e'' is a whole in which e' is replaced by e, the restriction on the replacement is that e and e''-*minus*-e' are disjoint: No part of the former is a part of the latter.

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

As is easily seen, the independence axiom implies that each part makes a context-independent contribution to the value of the whole. In other words, the value contribution of a part does not depend on the other parts the whole in question is composed of. Clearly, value additivity presupposes independence, but the latter might hold even in the absence of additivity.

With Independence, we could not have had superiority between the first and the last element of the sequence without that superiority setting in at some point along the way.

> *Observation* 2: Suppose that the first element in a sequence $e_1, \ldots, e_n$ is superior to the last one. Then, provided that *at-least-as-good-as* is a complete and transitive relation on the domain under consideration, Independence implies that some element in the sequence is superior to its immediate successor.

For the proof, see appendix 1.

What if we give up Independence? Then the following becomes possible: Suppose that when we start adding more and more valuable objects of the same type, the marginal value contribution of each extra object sooner or later starts to decrease converging to zero at infinity. If this decrease is sufficiently steep, then adding extra objects of the same type will never get us above a finite value limit: For any object e of a finite value, there will exist some finite value level v_e such that the aggregated value of an arbitrarily large number of e-objects is always lower than v_e. But then nothing excludes that a single object e may be more valuable than any number of e'-objects: All it takes is that the value of e either equals or exceeds $v_{e'}$. The manoeuvre of letting the marginal value contribution of extra units of a given kind of good converge to zero is, of course, quite standard, even with respect to intrinsic values. To give just one well-known example, in *Reasons and Persons* Derek Parfit makes this suggestion in the area of population axiology, in order to avoid the (in)famous 'repugnant conclusion'. It would be repugnant to have to conclude that any possible world populated by people leading excellent lives must be worse than some world with a sufficiently large number of people all of whom have lives that are barely worth living. To avoid this conclusion, we do not have to give up the welfarist idea that the value of the world is an aggregate of the welfare values of the lives of its inhabitants. All we need is to assume that the aggregative operation is of an appropriate kind: Adding extra lives of a positive but low quality increases the value of a world but it will never increase that value beyond a certain finite limit.[15]

[15] Cf. D. Parfit, *Reasons and Persons* (Oxford, 1984), sect. 137. His application of this idea to the value of populations comes from T. M. Hurka, 'Value and Population Size',

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

Given this convergence of value to finite limits, it is easy to account for the possibility of a decreasing sequence $e_1, \ldots, e_n$, in which (i) the first element is superior to the last element, even though (ii) no element is superior to the one that comes next. As a simplest possible example, which for that reason is maximally artificial, assume that the sequence consists of three elements, e_1, e_2, e_3, with their values being, respectively, 5, 3, and 2, as measured on a ratio scale. Suppose now, unrealistically, that the value contribution of extra objects of the same type rapidly decreases, from the very beginning, with each new contribution being half as large as the preceding one. Thus, for example, while the value of one e_3-object equals 2, the value of two such objects equals $2 + 1$, the value of three e_3-objects equals $2 + 1 + \frac{1}{2}$, etc. It is easy to see that for each object type, there is a finite value limit that cannot be exceeded by a whole composed of the objects of that type. That limit can be defined as the sum of the infinite sequence in which the first term equals the value of a single object of the type under consideration and each successive term stands for the value contribution obtained from adding another object of the same type. In the example, these limits have been chosen in such a way as to guarantee that the sequence satisfies the required conditions (i) and (ii). The value of the first element (5) exceeds the value limit for the last element ($2 + 1 + \frac{1}{2} + \ldots = 4$). Consequently, the first element is superior to the last one. But for each element in the sequence, its value is lower than the value limit for the next object in the sequence. That is, no element is superior to the one that comes next. That such a construction is mathematically coherent is reassuring, since many cases of superiority are such that, intuitively, we take it to be possible to move from a superior e_1 to an inferior e_n by a *gradually* decreasing sequence in which at no point does there appear to occur a radical value loss.

IV. WEAK SUPERIORITY IS DIFFERENT

It is different with weak superiority. It can be shown, without assuming Independence, that any finite sequence whose first element is superior to its last element must contain some element that is *weakly* superior to the one that comes next. In other words, in a sequence in which no element is even weakly superior to its immediate successor, the first element cannot be superior to the last element.

This result can be strengthened. In a sequence in which no element is weakly superior to its immediate successor, the first element cannot

Ethics 93 (1983). For a discussion, see G. Arrhenius, *Future Generations: A Challenge for Moral Theory* (Uppsala 2000), ch. 4.

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

even be *weakly* superior to the last element. More exactly, we can prove the following:

> *Observation* 3: Suppose that *at-least-as-good-as* is a complete and transitive relation on the domain. Then, in any finite sequence of objects in which the first element is weakly superior to the last element, there exists at least one element that is weakly superior to its immediate successor.
>
> *Proof*: To establish this observation, it is enough to prove the following lemma:
>
> *Lemma* 1: Suppose that *at-least-as-good-as* is a weak order, i.e. a complete and transitive relation on the domain. For any objects e, e', and e'', if e is weakly superior to e'', e is weakly superior to e' or e' is weakly superior to e''.

Note: If at-least-as-good-as is transitive and complete, then, by Lemma 1, the complement of weak superiority, i.e. the relation of *not* being weakly superior, is transitive: If e is not weakly superior to e' and e' is not weakly superior to e'', then e is not weakly superior to e''. Since weak superiority by definition is asymmetric, its complement is a complete relation: e is not weakly superior to e' or e' is not weakly superior to e. Therefore, Lemma 1 implies that, given the transitivity and completeness of at-least-as-good-as, weak superiority is a so-called strict weak order, i.e. a relation the complement of which weakly orders the object domain.

It is easy to see that Observation 3 entails Lemma 1, since a triple e, e', e'' is an example of a finite object sequence. That Lemma 1 in its turn also implies Observation 3 is easy to show. Suppose that at-least-as-good-as is complete and transitive. Consider a sequence $e_1, \ldots, e_n$, in which e_1 is weakly superior to e_n. By Lemma 1, (i) e_1 is weakly superior to e_2, or (ii) e_2 is weakly superior to e_n. If (i) holds, Observation 3 is established. If (ii) holds, then we consider the reduced sequence $e_2, \ldots, e_n$, and repeat the argument above, i.e. either (iii) e_2 is weakly superior to e_3 or (iv) e_3 is weakly superior than e_n. If (iii), we are done, and if (iv), we consider the reduced sequence $e_3, \ldots, e_n$. Continuing in this way, we finally reach a two-membered sequence, e_{n-1}, e_n, and it is clear that Observation 3 trivially holds for such sequences.

It remains then to prove Lemma 1. Assume that (i) e is weakly superior to e'', but (ii) e is not weakly superior to e'. We need to show that, in such a case, e' is weakly superior to e''.

(i) means that there exists some number m such that

(1) m e-objects are better than any number of e''-objects.

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

(ii) implies that there is some number m' such that

(2) m e-objects are not better than m' e'-objects.

But then, given that at-least-as-good-as is a complete relation,

(3) m' e'-objects are at least as good as m e-objects.

By the transitivity of at-least-as-good-as, if one object is at least as good as another, which is better than some third object, then the first object is better than the third. Consequently, (3) and (1) imply that

(4) m' e'-objects are better than any number of e''-objects.

(4) implies that e' is weakly superior to e''. This completes our proof.[16]

Observation 3 shows that weak superiority cannot obtain between the extrema of a finite sequence without setting in at some point in that sequence. Now, if the elements in a decreasing finite sequence are chosen in such a way that each consecutive element is only *marginally* worse than the immediately preceding one, then it might seem that no element will be weakly superior to the element that comes next.[17] But then, as we just have shown, the first element will not even be weakly superior to the last element in the sequence, however long such a sequence may be. This is surprising, since one would intuitively expect that a sufficiently long series of small worsenings can sooner or later result in an element that is radically worse than the point of departure.

One interpretation of our result is that we should give up this intuition. We might want to deny that a sequence of small worsenings can ever yield an element that is radically worse than the original element. Since many examples of purported superior goods do admit of series of gradual worsenings by means of which we end up with something seemingly inferior to the point of departure, this option puts

[16] The proof above assumes that the relation 'at least as good as' is complete, which is a rather exacting requirement. What if completeness is not assumed? Well, even in the absence of completeness, we can prove a variant of Observation 3.

> *Definition* 3: An object e is *minimally superior* to an object e' iff for some number m, no whole composed of e'-objects, however large, is better than m e-objects.
>
> *Observation* 4: Suppose that *at-least-as-good-as* is a transitive relation. If the first element in a finite sequence of objects is minimally superior to the last element, then there must exist some element in that sequence that is minimally superior to its immediate successor.

For the proof, see appendix 2. For a similar result in the context of population axiology, see Arrhenius *Future Generations*, secs. 3.2 and 10.3–4. Note that weak superiority entails minimal superiority, but not vice versa. Still, just as it is the case with these stronger relations, minimal superiority is incompatible with the Archimedean axiom.

[17] Cf. the Quantity Condition discussed in Arrhenius, *Future Generations*, pp. 54–6, which captures this intuition in a population context. It is shown that this condition, together with some other weak conditions, implies the Repugnant Conclusion.

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

into question the existence of genuine superiority relationships in the domain. Alleged superiorities may dissolve upon reflection.

Another option is to hold on to the intuition of series of small worsenings leading to things that are radically worse than points of departure and instead to revise our pre-reflexive idea of a small worsening. On this second interpretation, there is an important difference between superiority and weak superiority. The former is a drastic form of a difference in value, but the latter is not. If an object is only slightly worse than another object, then the latter cannot be superior. But, contrary to appearances, it might still be weakly superior, even though it is better only by a small amount.[18] A sufficient number of objects of this type might, if conjoined, form a whole that is better than *any* whole composed of the objects of the other type.

As the authors of this article cannot agree on which interpretative option is to be preferred, we leave it to the reader to decide the issue.[19]

Gustaf.Arrhenius@philosophy.su.se
Wlodek.Rabinowicz@fil.lu.se

APPENDIX 1: THE IMPORTANCE OF INDEPENDENCE

Independence: If an object e is at least as good as e', then replacing e' by e in any whole results in a whole that is at least as good.

Observation 2: Suppose that the first element in a sequence $e_1, \ldots, e_n$ is superior to the last one. If *at-least-as-good-as* is a complete and transitive relation, Independence implies that some element in the sequence is superior to its immediate successor.

[18] On this view, then, being weakly superior to an object is not sufficient for being much better than the object in question. In fact, it is not necessary for being much better either. It is easy to construct a case in which, in a descending sequence e, e', e'', (i) the value difference between the first element and the second one is larger than that between the second element and the third, (ii) the second element is weakly superior to the third, but (iii) the first element is not weakly superior to *any* of the elements that follow. While a sufficient number of e'-objects may be better than any number of e''-objects and e may be much better than e', there still may be no such number of e-objects that cannot be outweighed by a sufficiently large number of e''-objects.

[19] This article is a significantly revised and expanded version of G. Arrhenius and W. Rabinowicz, 'On Millian Discontinuities', in *Patterns of Value*, vol. 1, ed. Wlodek Rabinowicz and Toni Rønnow-Rasmussen (Lund, 2003). We are indebted to Roger Crisp for a helpful comment that led us to Observation 1 and to John Broome for useful ideas as to how to extend superiority relationships between objects to analogous relationships between attributes (cf. appendix 3, co-authored by Broome and Rabinowicz). Rabinowicz's work on this article was supported by a research grant from The Bank of Sweden's Tercentenary Foundation and was completed during his stay at the Swedish Collegium for Advanced Study in the Social Sciences in Uppsala.

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

Proof: Assume, for *reductio*, that none of the elements e_i in the sequence $(i < n)$ is superior to its immediate successor. By completeness, this means that for every such e_i there is some number m_i such that the whole composed of m_i e_{i+1}-objects is at least as good as e_i. Now, start with e_1 and replace it by a whole w_2 composed of m_1 e_2-objects. By assumption, w_2 is at least as good as e_1. If we replace any e_2 in w_2 by m_2 e_3-objects, Independence again implies that the resulting whole is at least as good as w_2. We can in this way replace every e_2-object in w_2, one after another, by m_2 e_3-objects, until we reach a whole, w_3, that is composed of $(m_1 \times m_2)$ e_3-objects. By Independence, w_3 is at least as good as w_2, and thus – by transitivity – it is at least as good as e_1. Continuing in this way, from w_2 to w_3, from w_3 to w_4, and so on, we finally reach a whole w_n that is composed of $(m_1 \times m_2 \times \cdots \times m_{n-1})$ e_n-objects. By Independence and transitivity, w_n is at least as good as e_1, which implies that e_1 is *not* superior to e_n.

APPENDIX 2: MINIMAL SUPERIORITY

Definition 3: e is *minimally superior* to e' if and only if for some number m, there is no such k that k e'-objects are better than m e-objects.

Observation 4: Suppose that *at-least-as-good-as* is a transitive relation. If the first element in a finite sequence of objects is minimally superior to the last one, there must exist some object in that sequence that is minimally superior to its immediate successor.

Proof: To establish Observation 4, we prove the following lemma:

Lemma 2: Suppose that *at-least-as-good-as* is a transitive relation. For any objects e, e' and e'', if e is minimally superior to e'', then e is minimally superior to e' or e' is minimally superior to e''.

Lemma 2 is implied by Observation 2 and it implies Observation 2 in exactly the same way as Lemma 1 implies Observation 1 (see above, section III). As for the proof of Lemma 2, it goes as follows:

Assume that (i) e is minimally superior to e'', but (ii) e is not minimally superior to e'. We want to show that, in such a case, e' is minimally superior to e''.

(i) means that there exists some number m such that

(1) for no number k, k e''-objects are better than m e-objects.

(ii) implies that there is some number k' such that

(2) k' e'-objects are better than m e-objects.

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

To prove that e' is minimally superior to e'', it is enough to establish that there is no such k that k e''-objects are better than k' e'-objects. Suppose, for *reductio*, that

(3) k e''-objects are better than k' e'-objects.

By the transitivity of at-least-as-good-as, better-than is transitive as well. Therefore, (3) and (2) imply that

(4) k e''-objects are better than m e-objects.

But (4) contradicts (1), which concludes the proof.

APPENDIX 3: SUPERIOR ATTRIBUTES

John Broome and Wlodek Rabinowicz

Consider a model that consists of (i) a set **X** of *attributes*: $\{A, B, \dots\}$ and (ii) a domain **D** of *objects*: $\{e, e', e'', \dots\}$. We take it that (iii) each attribute A in **X** determines an equivalence relation $=_A$ on the object domain. Intuitively, that relation obtains between any two objects that are exactly alike with respect to the attribute in question, which is the case, we could assume, whenever the attribute is present in both objects in the same degree. Relation $=_A$ is reflexive, transitive and symmetric, for every attribute A.

Finally, we have in the model (iv) a transitive and reflexive relation **R** on the object domain, with $e\mathbf{R}e'$ meaning that e is at least as good as e'. The relations of betterness and equal goodness are defined in terms of **R** in the standard way: e is *better* than e' if and only if $e\mathbf{R}e'$ but not $e'\mathbf{R}e$; e and e' are *equally good* if and only if $e\mathbf{R}e'$ and $e'\mathbf{R}e$.

Intuitively, we think of the attributes as 'natural' features of objects on which the value relation **R** between objects supervenes. Whether e is at least as good as e' depends on the degree in which various attributes are present in e and e', respectively. The assumption is that the attribute set **X** is exhaustive, in the sense that **R** is fully determined by the attributes in **X**. We express this assumption by the following formal requirement:

> For all objects e and e', if for all attributes A in **X**, $e =_A e'$, then e and e' are equally good.

An object e' will be said to be an *A-variant* of an object e if and only if e' differs from e at most with respect to attribute A. That is, if and only if for all attributes $B \neq A$, $e' =_B e$. Note that for any A and e, e is an A-variant of itself. As is easy to see, being an A-variant is an equivalence relation on the object domain.

We shall say that e' is an *A-improvement* of e if and only if (i) e' is an A-variant of e and (ii) e' is better than e. That is, if and only if e' is better than e, but differs from e only with respect to A.

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

The model we have assumed is deliberately very austere: It only contains what we need to define different kinds of superiority relations between attributes. Thus, for example, the framework we use does not allow us to make value comparisons between objects with respect to a given attribute, if these objects differ with respect to some other attributes as well. We could extend the model to make such value comparisons possible, but in the present context this is unnecessary.

Now we can define various superiority relations between attributes.

Superiority relative to a reference-point: A is superior to B relative to e if and only if all A-improvements of e are better than all B-variants of e.

Weak superiority relative to a reference-point: A is weakly superior to B relative to e if and only some A-improvement of e is better than all B-variants of e.[20]

In other words, one attribute is superior to another, relative to a reference point, if, at that point, any improvement with respect to the former attribute is preferable to any change with respect to the latter. On the other hand, one attribute is weakly superior to another, relative to a reference point, if, at that point, some improvement with respect to the former attribute is preferable to any change with respect to the latter.

Apart from these local, relativized relations of superiority, we can also define their *global*, non-relativized counterparts:

Superiority: A is superior to B if and only if A is superior to B relative to all e in **D**.

Weak superiority: A is weakly superior to B if and only if A is weakly superior to B relative to all e in **D** that admit of A-improvements.[21]

[20] It is arguable that James Griffin had something like these relations between attributes in mind when he distinguished between 'Trumping' and 'Discontinuity'. In the case of Trumping, *any* positive amount of one attribute, A, no matter how small, is preferable to *any* amount of the other attribute, B. In the case of Discontinuity, on the other hand, a *sufficient* amount of A is preferable to *any* amount of B. (Cf. Griffin *Well-Being*, pp. 83 and 85.) However, Griffin's definitions are framed in terms of *amounts* in which an attribute could be present rather than in terms of possible *improvements* with respect to a given attribute, as we have done. The latter approach is preferable if one wants to allow for the possibility of attributes that are negatively rather than positively value-relevant. For such attributes it is better if they are present in smaller rather than larger amounts. Furthermore, avoiding reference to amounts is advisable if one wants to leave room for attributes that cannot be measured on a ratio scale.

[21] Consider the following quote from Griffin: 'The mistake here seems to be to think that certain values – liberty, for instance – as *types* outrank other values – prosperity, for instance – as *types*. Since values, as types, can vary greatly in weight from token to token, it would be surprising to find this kind of discontinuity at the type – or at least fairly abstract type – level. It is much more promising to look at less abstract types, or at tokens' (ibid., pp. 85–6). It appears that what Griffin might have had in mind in this

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

The qualification 'relative to all e that admit of A-improvements' means that, in order to determine whether global weak superiority obtains between two attributes, we only compare them relative to such reference-points for which there exist A-improvements in the object domain. Note that, if e does not admit of an A-improvement, then relative to that reference point A cannot be weakly superior to any attribute. Consequently, without this qualification, we would get the counter-intuitive implication that no attribute can be globally weakly superior to anything else if some object in the domain is so perfect with respect to that attribute that it cannot be further improved.[22]

Local superiority obviously does not entail global superiority. In fact, nothing hinders that one attribute is superior to another relative to one reference point, while the opposite holds relative some other reference point. For example, think of the object domain as the set of possible life circumstances for a person and consider such circumstantial attributes as nourishment and entertainment. It may well be the case that the former attribute is superior to the latter relative to the circumstances in which nourishment is scant. But in the circumstances in which nourishment is abundant, the opposite might be the case: entertainment might be then be a superior consideration.

Superiority relative to a fixed reference point may well obtain between the extrema in a sequence of attributes, without obtaining between any element in the sequence and its immediate successor. Here is an example of a sequence of three attributes, A, B, C, which illustrates this possibility: Suppose, in this artificial example, that the domain of objects consists of all number triples in which at least two out of the three numbers are zeros and the remaining number, if distinct from zero, is either 1 or 2. As is easy to see, there are seven such triples, all in all. Suppose that the betterness relation orders them as follows (with the best triple on top and the worst triple at

passage is that it is more realistic to expect that superiority relations between relatively abstract attributes obtain relative to particular reference-points (i.e. on the level of 'tokens') rather than globally.

[22] However, the qualification comes at a cost. If an attribute A is such that no object in the domain admits of A-improvements (which will be the case if changes in A always require simultaneous changes in some other attribute), then, given the qualification, it follows that A is (vacuously) globally weakly superior to every other attribute. If no object admits of B-improvements either, then A is (vacuously) globally weakly superior to B, *and vice versa*. Thus, with the qualification, weak superiority stops being an asymmetric relation. Still, these problems do not arise in domains in which for every attribute there exist some objects that could be improved with respect to that attribute. Furthermore, the same problems exist for global superiority. If no objects can be A-improved or B-improved, because changes in each of these attributes always are correlated with changes in other attributes, then, vacuously, A is globally superior to B and B is globally superior to A.

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

bottom):

(2, 0, 0)
(0, 2, 0)
(1, 0, 0)
(0, 0, 2)
(0, 1, 0)
(0, 0, 1)
(0, 0, 0)

The intended interpretation is that the first term in a number triple e is meant to specify the degree in which attribute A is present in e, while the second and the third term specify the degrees of B and C, respectively. Ceteris paribus, the higher degree of an attribute makes the object better. It is assumed that the domain only contains items in which at most one of the attributes is present in a positive degree. It would be more reasonable, of course, to regard number triples not as objects in the domain but rather as the numerical representations of the objects in terms of the three attributes. But in this artificial example we can ignore such niceties.

If two triples are equal in their first term, they will be said to be exactly alike with respect to A. To be exactly alike with respect to B or to C, they have to be equal in their second or third term, respectively. Consequently, an A-variant of a triple e differs from e at most in its first term. Similarly for B- and C-variants.

Since we have already specified the betterness ordering on the object domain, the relation of an A-improvement is now well defined and the same applies to B- and C-improvements. Now, choose (0, 0, 0) as the reference-point. With respect to this reference point, A is superior to C: Every triple in which the first term is positive, i.e. every A-improvement of (0, 0, 0), comes higher up in the ordering than any triple that differs from (0, 0, 0) at most in the third term. At the same time, relative to (0, 0, 0), attribute A is not superior to B, as shown by the fact that (0, 2, 0) comes higher up than (1, 0, 0). Nor is B superior to C relative to that reference point, since (0, 0, 2) comes higher up than (0, 1, 0). Which completes the counter-example.

Note that, in this example, A is superior to C not only relative to (0, 0, 0), but also relative to all the other reference points. It is superior to c relative to (1, 0, 0), since (2, 0, 0), which is the only A-improvement of (1, 0, 0), is better than all the other triples. Relative to all triples that are distinct from (0, 0, 0) and (1, 0, 0), this superiority of A obtains trivially, as no such triple admits of A-improvements. Consequently,

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

the example shows that even global superiority can obtain between the extrema in a sequence of attributes, without obtaining between any element in the sequence and its immediate successor.

We now move to weak superiority. That relation behaves differently from superiority. The following can be proved:

Observation 5: If **R** is complete (in addition to being transitive), then, if weak superiority relative to a fixed reference point obtains between the extrema in a sequence of attributes, it must also obtain between some element in the sequence and its immediate successor.

To establish this observation, we need to prove the following lemma:

Lemma 3: Suppose that **R** is transitive and complete and consider any attributes A, B, and C. For every object e, if A is weakly superior to C relative to e, then A is weakly superior to B relative to e or B is weakly superior to C relative to e.

Note: If **R** is transitive and complete, then, by Lemma 3, the complement of weak superiority with respect to a fixed reference point is a transitive relation. Since weak superiority by definition is asymmetric, its complement is a complete relation. Lemma 3 therefore implies that, if **R** is transitive and complete, weak superiority relative to a fixed reference point is a strict weak order, i.e. a relation the complement of which weakly orders the set of attributes.

It is easy to see that Observation 5 entails Lemma 3, since A, B, C is an attribute sequence. That this lemma in its turn implies Observation 5 is also easy to prove. In an attribute sequence $A_1, A_2, \ldots, A_n$, in which A_1 is superior to A_n, Lemma 3 implies that (i) A_1 is weakly superior to A_2, or (ii) A_2 is weakly superior to A_n. If (i) holds, Observation 5 is established. If (ii) holds, then we consider the reduced sequence $A_2, \ldots, A_n$, and repeat the argument above. That is, either (iii) A_2 is weakly superior to A_3 or (iv) A_3 is weakly superior to A_n. If (iii), we are done, and if (iv), we consider the reduced sequence $A_3, \ldots, A_n$. Continuing in this way, we finally reach a two-membered sequence, A_{n-1}, A_n, and it is clear that Observation 5 trivially holds for such sequences.

Proof of Lemma 3: Let **R** be transitive and complete and suppose that (i) A is weakly superior to C relative to e, but (ii) A is not weakly superior to B relative to e. We need to prove that, under those circumstances, B must be weakly superior to C relative to e.

(i) means that

(iii) there is some A-improvement e' of e that is better than all the C-variants of e.

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

(ii) implies that

(iv) there is some B-variant e'' of e such that e' is not better than e''.

By the completeness of **R**, e'' is at least as good as e'. But then, given (iii) and the fact that **R** is transitive, (iv) implies that

(v) there is some B-improvement e'' of e that is better than all the C-variants of e.

This means that B is weakly superior to C relative to e. Which concludes the proof.

What about *global* weak superiority? Does the analogue of Observation 5 hold for that non-relativized relation? As is easily seen, the proof of Lemma 3, as given above, cannot be extended to global weak superiority. For if A is globally weakly superior to C but not to B, there is some reference point e relative to which A is not weakly superior to B. It can then be proved, if **R** is transitive and complete, that B is weakly superior to C relative to e. But that is a far cry from B being globally weakly superior to C, i.e. from its being weakly superior to C relative to *all* reference points.

In fact, it is quite easy to construct a counter-example in which global weak superiority obtains between A and C, but does not obtain between either A and B or B and C. Let us consider a domain that consists of eight number triples. These are, in the descending order of betterness:

$$(2, 3, 2)$$

$$(3, 2, 2)$$

$$(2, 2, 3)$$

$$(2, 2, 2)$$

$$(1, 0, 0)$$

$$(0, 0, 1)$$

$$(0, 1, 0)$$

$$(0, 0, 0)$$

The notions of an A-variant and an A-improvement are defined as in the previous counter-example.

In this domain, there are only two triples, (0, 0, 0) and (2, 2, 2) that admit of A-improvements. (Remember, in our technical sense, an A-improvement changes the object only with respect to A, leaving it exactly alike in all other respects.) Consequently, A is globally weakly superior to C if and only if it is weakly superior to C relative to each of

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press

these two reference-points. That this in fact is the case is easily seen. (1, 0, 0) is better than all the *C*-variants of (0, 0, 0), whereas (3, 2, 2) is better than all the *C*-variants of (2, 2, 2). At the same time, *A* is not weakly superior to *B* relative to (2, 2, 2): (2, 3, 2) is better than all the *A*-variants of (2, 2, 2). Neither is *B* weakly superior to *C* relative to (0, 0, 0): (0, 0, 1) is better than all the *B*-variants of (0, 0, 0). Consequently, neither is *A* globally weakly superior to *B* nor is *B* globally weakly superior to *C*. Which concludes the counter-example.

https://doi.org/10.1017/S0953820805001494 Published online by Cambridge University Press