

Parity, Interval Value, and Choice*

Ruth Chang

In "The Possibility of Parity," I argued that the right thing to say about certain hard cases of comparison—cases in which one item is better in some relevant respects, while the other is better in other relevant respects, but there is no obvious truth about how the items compare in all relevant respects—is that the items are 'on a par'.¹ 'Parity', I said, is a fourth "positive" value relation of independent standing, not subsumable under the familiar trichotomy of relations 'better than', 'worse than', and 'equally good'. My argument took the form of an argument by elimination: the cases of interest are not cases of ignorance, in which one of the traditional three relations holds but we don't know which, nor cases of vagueness, in which the items occupy the borderline of one of the traditional three relations but are cases of determinate comparability; therefore, as cases of determinate comparability in which none of the traditional three relations holds, they must be cases in which a fourth relation of comparability holds—they are on a par.

In his interesting discussion of my article, "Value and Parity," Joshua Gert agrees that the cases I think are cases of parity are not cases of ignorance or vagueness and agrees that we cannot simply assume that because neither of two alternatives is better than the other and they are not equally good that they are thereby incomparable.² Nevertheless, he wants to challenge my claim that these cases are ones in which the items are related by a fourth positive relation "of the same sort" as the usual

* Many thanks are due to Kit Fine for invaluable comments on an earlier draft and to Erik Carlson, David Miller, Derek Parfit, and two anonymous referees at *Ethics* for useful criticisms that helped me to improve and clarify various points. I am grateful to Joshua Gert for his discussion of my earlier work and for lengthy e-mail correspondence on an earlier version of the present article.

1. Ruth Chang, "The Possibility of Parity," *Ethics* 112 (2002): 659–88.

2. Joshua Gert, "Value and Parity," *Ethics* 114 (2004): 492–510. All subsequent page references in parentheses are to this article.

Ethics 115 (January 2005): 331–350

© 2005 by The University of Chicago. All rights reserved. 0014-1704/2005/11502-0004\$10.00

three. He proposes what he says is an alternative interpretation of the hard cases that I say are cases of parity, an alternative, he says, that does not require us to give up what I call 'the Trichotomy Thesis', the claim that the conceptual space of comparability between two items is exhausted by the trichotomy of relations 'better than', 'worse than', and 'equally good'. He writes, "This paper will defend the trichotomy thesis, at least in one important sense: it will hold that any other positive value relations that we might wish to make use of can be defined in terms of the three traditional relations" (p. 493). As he says, he can "easily explain" putative cases of parity in terms of the usual three relations, and so there's no need to posit the existence of parity.

I believe, however, that Gert's argument does not succeed in "explaining away" parity and that it leaves the case for parity untouched. As we will see, although Gert presents his article as criticism, he is really a comrade in arms. Like me, he is interested in understanding cases of determinate comparability in which neither item is better than the other and nor are they equally good. I call these cases of 'parity', and although Gert shies away from the label, he gives no reason to think that such a relation does not hold in some hard cases.

Although Gert's argument does not undermine parity, it naturally suggests two interesting questions about it that I take up in this article. Gert proposes a model of comparability that employs an attractive representation of the value of an item by an interval range of real numbers. Unlike standard expected utility theory, which represents the value (or preferability) of an item by a single real number, interval representation attempts to capture the idea that the beauty of a vase, for example, cannot be adequately represented by a single point. Gert's interval model is supposed to account for the cases I think of as parity, but, he argues, his model shows that those cases are too diverse to fall under the rubric of a single unified value relation. This claim rests on the details of his particular model, but we will see that there are good reasons to reject the model: it does not satisfy the basic axioms of comparability that Gert himself should want the model to satisfy.

The failure of his model, however, raises the question, is there some model of parity consistent with an underlying interval representation of the value of items? It turns out that we can prove under certain reasonable assumptions that there is only one model of interval representation that could make room for the cases in which both Gert and I are interested, and according to this model, parity is given by a single, univocal kind of case. However, this model too is arguably problematic; there is a condition it implies that we might reasonably wish to reject. If this condition does not hold of items on a par, then we are left with a striking result: no reasonable interval model is adequate to capture

the possibility that items are on a par. As we will see, this result has interesting implications even for those who might be skeptical of parity.

The second interesting question that Gert's discussion raises has to do with the practical upshots of parity. In my original article, I said that parity is important because if there is parity, then it is a mistake to conclude from the fact that neither of two items is better than the other and that they are not equally good that they cannot be compared. And whether items can be compared is in turn important because it is plausible to think that if alternatives cannot be compared, there can be no justified choice between them. In this way, parity expands the range of cases in which justified choice is possible; choices between items about which practical reason might otherwise appear to be silent are in fact choices between comparable items and thus within the scope of practical reason. Parity, it might be said, is what gives practical reason a "voice" in hard cases. But what, exactly, does practical reason "say" in cases of parity? That is, what should one do when faced with items that are on a par? Gert usefully suggests that in the cases I think of as parity, it is rationally permissible to choose either alternative. I think this is roughly right. But there are three different senses of 'rationally permissible' that need to be distinguished, and their differences will help us to see why it is important to recognize parity as a relation apart from those of the traditional trichotomy.

I. "EXPLAINING AWAY" PARITY

I suspect that many philosophers, upon being told that there is a fourth, independent positive value relation beyond 'better than', 'worse than', and 'equally good', will share Gert's reaction: there must be some way to account for the putative relation in more familiar terms.³ If, like Gert,

3. Another recent attempt to do away with parity is mooted by Nien-he Hsieh, "Equality, Clumpiness, and Incomparability," *Utilitas* (in press), who argues that a "clumpy" understanding of value allows us to treat cases of parity as cases of equality. But Hsieh's understanding of 'being equally good' as "belong[ing] to the same clump of value" (p. 12) is saddled with a dilemma. According to Hsieh, A and B can belong to the same clump of value even though something better than A, A+, also belongs to the same clump of value. This means that even though A+ is better than A, which is equally good as B, A+ is not better than B but is equally good as B. This is not equality as we know it. To avoid this result, Hsieh seems to suggest that we must relativize evaluative comparisons to a degree of precision. Thus, A and B are equally good to degree of precision p1, A+ is better than A to degree of precision p2, and A+ and B are equally good to degree of precision p1 (or perhaps p3). The trouble with this suggestion is that it commits us to denying the inferential links between evaluative comparisons that proceed with respect to the same covering consideration. If Mozart is better than Salieri with respect to creative genius, and Salieri is better than Talentlessi, a rotten sculptor, with respect to creative genius, it seems to follow that Mozart is better than Talentlessi with respect to creative genius. This is so even though the degree of precision according to which a comparison

they agree that cases in which parity supposedly holds are not cases of ignorance or vagueness, they must find some other way of "explaining away" parity. This is harder to do than it might at first seem, and an examination of Gert's discussion will help us to see why.

Gert's main claim is that all pairwise positive value relations can be defined in terms of "more basic" choice-theoretic relations involving the notion of a mistake in choice. Instead of talking about one thing being better than, worse than, as good as, or on a par with another, we can talk instead in terms of the more basic idea from which pairwise relations supposedly derive, namely, the rational permissibility of choosing one item over another. As he writes, "our valuations of two items are . . . given by our dispositions either to understand or to be puzzled by certain choices: to regard them as rational or irrational, mistaken or not mistaken" (p. 494). By understanding pairwise value relations in terms of mistakes in choice, we can, Gert urges, do away with parity. He summarizes his proposal like this, "The suggestion I would like to make, which need be only roughly true to indicate how we can understand Chang's hard cases without having to posit a fourth positive value relation, is that we use the word 'better' to mean something like the following: 'to be chosen, on pain of having made a mistake'. Correlatively, 'worse' can be taken to mean 'not to be chosen, on pain of having made a mistake'" (p. 499). And: "This article claims that all that is required [to explain putative cases of parity] are the same positive value relations that we have always been familiar with—although it is true that we need to understand these notions themselves in terms of the still more basic notion of a practical mistake in choice" (p. 501). Gert then appears to offer two definitions of parity in choice-theoretic terms: A is on a par with B if and only if "although in a choice between A and B, one could rationally choose either, A would have to be improved more than B, before it was no longer a mistake to choose it over C" (p. 508); and, in an alternative formulation, "It is not a mistake, or irrational, to choose A over B, or B over A, and . . . this may continue to be true even if one of the items is slightly improved" (p. 506).⁴

The thought here seems to be that because parity can, by hypothesis, be defined in the same "more basic" terms by which the usual three relations can be defined, it follows that parity is not "of the same sort" as the usual three. However, even if there are more fundamental terms in which all pairwise positive value relations can be defined, this does nothing to show that parity is not a relation on all fours with the usual

of the creative genius of musicians proceeds is different from the degree of precision according to which a comparison of a musician and a sculptor proceeds.

4 I say that these "appear" to be definitions of parity because although Gert says that he will offer definitions of parity that explain it away, he never explicitly says what these definitions are

three. Thinking otherwise would be like thinking that if we can define all colors in terms of the more fundamental notion of wavelength of light, it follows that red is not a color "of the same sort" as blue, green, orange, and so on. As Gert says in a concessive moment, he "will not exactly deny the possibility of parity, if parity is understood in a certain derivative way" (p. 493). But the way in which he makes parity derivative is exactly the same way in which he makes the usual three relations derivative—by defining each in terms of what he takes to be the more fundamental notion of a mistake in choice. Thus, the sense in which Gert succeeds in "explaining away" parity is the sense in which he succeeds in "explaining away" all pairwise positive value relations.

As a general matter, a definition of a value relation in other, nonvalue-theoretic terms does nothing to show that we do not have a genuine positive value relation. There is, of course, the more particular question of whether Gert's favored choice-theoretic definitions of parity give us any reason to think that it is not of the same sort as the usual three. But they do not. This is perhaps easiest to see if we imagine how Gert might define 'equally good' in choice-theoretic terms: A and B are equally good if and only if either can be chosen without making a mistake, and if one or other were improved, it would be a mistake not to choose it. This definition of 'equally good' is so similar to his definitions of 'parity' that it is hard to see how one can be a definition of a genuine positive value relation while the other is not. In short, why shouldn't there be more than one way in which one can fail to make a mistake in choosing either of two comparable alternatives?

In discussing Gert's argument, I went along with his definitions of parity and his assumption that value relations can be defined in terms of the "more basic" idea of mistake in choice. But both points are problematic. Gert's definitions are arguably overbroad; they pick out not only comparability when the usual three relations fail—that is, parity—but also incomparability—that is, the failure of any positive relation to hold. When items are incomparable, it is not a mistake to choose either—because reason fails to reach the question of which one should choose—and this may hold even if one of the items is improved.⁵ And for any two incomparable items, there may be some third item that is differentially related to the two—it might take more improvement in the one incomparable item to make it no longer a mistake to choose it over the third item than it might take in the other incomparable item. Perhaps by defining parity in a way that also holds of incomparability, Gert reveals himself to be a closet incomparabilist: cases in which I think

5. Moreover, Gert's first formulation of parity arbitrarily emphasizes one alternative over the other when in fact there is symmetry: not just one, but either of two items on a par might be improved without thereby becoming better than the other.

a fourth relation holds are really cases in which no relation holds. But in this case Gert would simply be begging the question at issue by assuming that if items are not related by the usual three positive relations, they are thereby not related by any positive relation.

There is also reason to balk at Gert's assertion that positive value relations can be defined in terms of the "more basic" notion of a mistake in choice. Whether, as a general matter, value concepts can be defined in terms of reason concepts is a large and controversial topic. My own view is that how items evaluatively compare is a conceptually distinct matter from how one rationally should respond to such items in the context of choice. While I believe that there are interesting connections between the two, these connections are, I believe, a matter of substantive argument and are not analytic truths given by the meaning of the terms 'value' or 'better than'. I cannot defend my "anti-buck-passing" view here. But perhaps it is enough to point out that Gert's particular form of "buck-passing" will not do. By defining 'better than' as "to be chosen on pain of making a mistake," Gert decrees by definitional fiat that satisficing, the view that it is sometimes not a mistake to forgo what is best, is conceptually incoherent. While I am no fan of satisficing, I doubt that it can simply be defined off the scene.

Perhaps Gert thinks that he can do away with parity, not by offering a choice-theoretic definition but, rather, by offering a value-theoretic definition, and in particular, a definition that involves only the usual three relations (even though those relations might in turn be defined in choice-theoretic terms).⁶ One possible definition of parity in terms of the usual three relations might be 'comparable, and yet neither is better than the other and one can be made better without thereby being better than the other'. A simpler one might be 'comparable, but neither better, nor worse, nor equally good'. And supposing that we can establish that there is no incomparability, we can even define parity simply in terms of the usual three relations: 'not better, not worse, and not equally good'. There is some question as to whether these definitions are legitimate. But it is easy to see that, even if they are, they provide no reason to think that parity is not a fourth positive relation of the same sort as the usual three. This is because the same sorts of "definitions" can with equal legitimacy also be given for the usual three relations: for example, 'better than' can be defined as 'not worse, not equally good, and not on a par.'

A definition of something is eliminative only if it captures what it is to be that thing. A "definition" that provides a coextensive description

6. As Gert writes, "This paper will defend the trichotomy thesis, at least in one important sense: it will hold that any other positive value relations that we might wish to make use of can be defined in terms of the three traditional relations" (p. 493).

in other terms—or even a necessarily coextensive description—need do nothing to eliminate what is “defined.” We can “define” parity in terms of the usual trichotomy of relations in the sense that we can give a coextensive description of parity in those terms. But it does not follow from the existence of a “definition” in this sense that parity lacks the status of a fourth relation “of the same sort” as the usual three. An analogy with color may help. Suppose we define red as “not blue, green, orange,” This “definition” constitutes a coextensive description of ‘red’ but does not give what it is to be red: being red is not simply a matter of not being blue, green, orange, and so on. It therefore does nothing to show that red is not “of the same sort” as blue, green, orange, and so on. In the same way, a “definition” of parity as “not better, worse, or equally good” does not give what it is to be on a par and, thus, does nothing to undermine its status as a positive value relation “of the same sort” as the usual three. Indeed, sometimes Gert talks not in terms of providing a definition of parity but in terms of “recasting,” “translating,” and “interpreting” it (pp. 494, 500, 508). But an equivalent description of one value relation in terms of others does nothing to eliminate the former in favor of the latter.⁷

Thus, Gert’s argument does nothing to undermine the thought that parity is a fourth positive value relation of the same sort as the usual three.

II. PARITY AND INTERVAL REPRESENTATION

Toward the end of his article, Gert presents an interesting model of comparability that is supposed to capture the cases that I think of as parity. On his own interpretation of the model, he admits that there are cases of comparability beyond the usual three relations that are plausibly called ‘parity’. But he hastens to add that he is not really admitting that there is parity because “if we are to be charitable to Chang, we should not hold that in arguing for the existence of a fourth positive value relation she means to be doing nothing more than putting a label to a phenomenon that can be easily explained in terms of the three traditional relations” (p. 507). As we have already seen, however, the way in which he “easily explains” parity in terms of the usual three relations relies on the mistaken thought that by finding a common denominator for all value relations, parity is thereby “explained away” in terms of the usual three. Thus, although Gert says he wants to defend

7. Sometimes I get the sense that what Gert really objects to is a reification of parity as something out there in the world. But that is no part of the claim that there is parity; there “is” parity in just the same sense in which there is betterness, worseness, and equality in value. What is at issue is parity’s status vis-à-vis the usual three relations, not its ontological status.

the Trichotomy Thesis, his real interest is in understanding how items might be related in ways beyond those allowed by the usual three relations. His own interpretation of his model shows that, like me, he thinks comparable items can be related by a fourth value relation.

Gert's model is interesting not just because it seems to allow for the possibility of parity but because it employs an attractive representation of the value of items. This mode of representation, familiar from measurement theory, holds that the value of an item is to be represented not by a single real number but by an interval range of reals. So, for instance, the beauty of a vase might be represented by an interval of real numbers [10, 60], and the deliciousness of a meal by an interval [3, 5]. These intervals might be glossed as giving the permissible value-points an item might take.⁸

Interval representation appears to provide a promising tool for modeling parity. Standard expected utility theory represents the value of an item by a single real number. Its associated "rule of comparability" holds that one item is better than another just in case the number that represents its value is greater, that it is worse if the number is less, and that they are equally good if the numbers are the same. Since one real number can be related to another only by being greater, lesser, or equal—the analogues of 'better', 'worse', and 'equally good'—this model leaves no room for comparison by a fourth relation. One way to "loosen up" standard expected utility theory is to relax the representation of value from a single real to an interval of reals. Since intervals can be related in a variety of ways, perhaps room for comparison by a fourth value relation can then be found.

The question is whether there is any reasonable rule mapping possible relations between intervals onto comparisons of their underlying items that allows for the possibility of parity. Gert offers "the Range Rule" as a way of making room for the cases that I think of as parity:

Range Rule One item is better than another if the lower bound of its interval range of value is higher than the upper bound of the range of the other item, otherwise the items are not traditionally comparable. (P. 505)

In other words, if the intervals representing the value of two items do not overlap, the item with the interval that is higher is better than the item with the interval that is lower; if they do overlap, neither is better than the other, nor are they equally good. The idea that under-

8. Gert gives a choice-theoretic gloss of these intervals as representing the rationally permissible strengths of preference one might have for an item, but this is to import Humean assumptions about value that Gert recognizes we should ignore if we are to give his model its widest possible scope. My discussion of interval representation is neutral between different ways in which we might understand in what the value of an item consists.

writes this rule might be that what it takes for one item to be better than another is for every value-point it can permissibly take to be higher than every value-point that the other item can permissibly take. In this way, only nonoverlapping intervals can represent betterness. All other interval configurations, according to the rule, fall under the category of "not traditionally comparable." And, according to Gert, in these cases neither of the items is better than the other, and an improvement in one does not thereby make it better.

Drawing on his rule, Gert observes that, on his model, the cases I think of as parity are represented by all possible overlapping configurations of intervals—six all told, taking into account only the relative position of endpoints and discounting the corresponding "reverse" configurations that swap the intervals from left to right. He claims that these cases are too diverse to be "usefully" captured under the rubric of a single fourth unified value relation (p. 507). On the assumptions that both Gert and I share, there are only two possibilities that could hold of items in such cases. Either they are comparable, and by hypothesis related by some positive relation beyond the usual three, or they are incomparable, that is, not related by any positive relation. But Gert gives no argument for thinking that they are incomparable and recognizes that he cannot simply assume that they are, for that would be to beg the question against those who think that there is a fourth positive relation. Thus, he must think that they are cases of comparability. Once he admits that overlapping configurations are cases of nontrichotomous comparability, however, he must admit that there is comparability by a fourth relation beyond the usual three. His point then seems to be that since his model shows that there are many different configurations representing possible ways of being nontrichotomously comparable, it is a mistake to think that there are only four relations that span the conceptual space of comparability; by Gert's count there are (at least) nine!⁹

Gert's grounds for thinking that there are many positive relations beyond the putative four depend on the details of his model. But the model cannot be correct. There are two miniproofs showing that this is so. It will be useful to start by listing the six possible overlapping configurations that the Range Rule deems representations of "not traditionally comparable" (with their "reverse" cases—the intervals swapped from left

9. It is important to emphasize that when Gert suggests that there are more relations beyond the usual three plus parity, he is not claiming that there is any relation with logical properties different from what I call parity; his claim, rather, is that what I call parity can be "divided up" into many different relations. Whether he is right about this waits on an account of how value relations are properly individuated. In any case, Gert's only defense of the claim depends on the details of his model, and, as we will see, his model is problematic.




FIG. 1

to right—in parentheses; see fig. 1).¹⁰ These configurations exhaust the possible overlapping configurations of two intervals where the relative position of endpoints as higher or lower is what is relevant to their underlying value relation.

The first proof shows that, contra the Range Rule, case 6 cannot be a case in which the items are “not traditionally comparable.” Indeed, in what he takes to be a concession of sorts, Gert labels this case ‘parity’.¹¹ But there is a simple proof that this case must be a case of equality. Take any interval range. Let there be an item whose value is given by

10. If a configuration represents the first interval as better than the second, then its “reverse” will swap the intervals from left to right and represent the first interval as worse than the second. The reverse of betterness is worseness and vice versa. Since equality and parity are both symmetric, then reverses will represent equality and parity, respectively.

11. Gert goes on to say that a case like case 6 in which the intervals, though the same, have a very narrow range is a case of “rough equality” (pp 506–7). But it is clear that which relation holds between two items should not turn on the size of the intervals representing them if all relational aspects of the intervals are held constant. Gert is misled by the fact that intervals with small ranges approach a degenerate interval, i.e., a point (with items represented by the same point being equally good). But approaching this degenerate interval and being identical to it are quite different matters.

that range. That item and itself are equally good. Therefore, any items whose values are given by the same interval range must be represented as being equally good.¹² Thus, the Range Rule must be rejected.

The second proof shows that, contra the Range Rule, the case of nonoverlapping intervals cannot be the only case of betterness. Take any case that, according to the Range Rule, is a case of being "not traditionally comparable," that is, a case from cases 1–6. According to Gert, in each such case we can improve one of the items without thereby making it better than the other. But we cannot do this if betterness is given only by nonoverlapping intervals. For in any of these cases take the interval whose top point is as high as or higher than the top point of the other interval. Call it A. In order to make A better, we must create an interval that does not overlap it, by making both its endpoints higher than the top point of A. But then it follows that the improved interval, A+, will not overlap with the other interval, B. So there is no way to improve A without thereby making the improvement, A+, better than B.¹³ Again, the Range Rule must be rejected.

These proofs rely on two uncontroversial axioms that Gert himself should accept: that 'equally good' is reflexive and that if items are comparable but not trichotomously so, either can be improved without thereby being better than the other, that is, 'the improvement condition'. If these two axioms show that the Range Rule must be rejected, what rule, if any, might succeed in its place?

In searching for an adequate interval model of comparability that makes room for parity, we might impose some further plausible conditions. First, we make some assumptions concerning the representations of value relations. We assume, following Gert, that nonoverlapping intervals are a case of betterness, that identical intervals are the only case of equality in value, and, finally, that "duality" holds. This last condition, although innocuous, needs some explaining. It states that betterness and worseness should not depend on the "direction" in which they are represented. If betterness is represented pointing upward, then worseness should be represented in the same way pointing downward, and vice versa. The 'dual' of a relation is the relation flipped upside down to indicate a change in direction. Duality holds that if a config-

12. In e-mail correspondence, Gert suggested that the proof can be resisted on the grounds that an item is not as good as itself. Without going into the plausibility of rejecting the reflexivity of equality, we can alter the proof to be one about an item and its twin: take any interval range and let the values of an item and its twin be given by that range. The item and its twin are equally good. Again, the Range Rule must be rejected.

13. It will not help to suggest, as Gert did in e-mail correspondence, that there is some way of making the improved item better than its unimproved ancestor beyond that sanctioned by the Range Rule since the Range Rule is supposed to provide a model of all relations of comparability.

uration represents betterness, then the same configuration flipped upside down must represent worseness and vice versa.¹⁴

We then add some assumptions about the intrinsic features of the value relations we are trying to model. In addition to the reflexivity of equality and the improvement condition on parity, we assume that 'better than' is transitive, that parity is symmetric, and that the four relations are mutually exclusive.

Given these assumptions, we can prove that there is only one rule of comparability that can make room for parity. This is what I will call 'the Pareto Interval Rule'.

Pareto Interval Rule: If the top point of interval A is at least as high as the top point of interval B, and the bottom point of A is at least as high as the bottom point of B, and the intervals are not identical, then A is better than B; if the respective top and bottom points of A and B are identical, A and B are equally good; and in a mixed case, A and B are on a par.

In other words, according to the Pareto Interval Rule, nonoverlapping intervals and cases 1-4 represent betterness; identical intervals, case 6, represent equality of value; and only case 5 represents parity. The intuitive idea behind this rule is that when the highest and lowest value-points of either item are not both at least as high as the corresponding highest and lowest value-points of the other item, the items are on a par. It is perhaps worth noting that on this interval model of comparability, parity is given not by a multitude of diverse configurations but by a single configuration. An abbreviated outline of an informal proof that only the Pareto Interval Rule is consistent with the above axioms is given in the appendix.

But there is a further problem. The Pareto Interval Rule implies a condition on parity that we might reasonably wish to reject. This condition is what we might call 'pair dictatorship'. It holds that for any three items on a par, two of them will be such that if they are worse than some other item, then so is the third. Put another way, it denies that there can be a trio of items on a par such that any two are worse than some other item without the third also being worse. Pair dictatorship maintains that there is always one pair of a trio of items on a par that dictates whether the third is better or worse than some other item.

14. If, e.g., in fig. 1, we flip configuration 3 upside down, we get the reverse of configuration 4. Duality implies that if the first interval of configuration 3 is better than the second, then the first interval of the reverse of configuration 4 must be worse than the second interval in that configuration. That is, if configuration 3 represents betterness, the reverse of configuration 4 must represent worseness. (And if the reverse of configuration 4 represents worseness, then configuration 4 must represent betterness.)




FIG. 2

It is easy to see that the Pareto Interval Rule is committed to pair dictatorship. Take three items on a par: A, B, and C. They will be represented as shown in figure 2. We can see that there will always be two intervals, namely, A and C, such that if there is an item better than both of them, that item will also be better than the third interval, B. This is because any item that is better than both A and C must have a bottom point at least as high as the bottom point of C and a top point at least as high as the top point of A. Since those top and bottom points will be respectively higher than the corresponding top and bottom points of B, that item will be better than B.

It is doubtful, however, whether pair dictatorship need hold among items on a par. It seems possible for there to be a trio of items on a par such that any two of them is worse than something or other without the third also being worse than that other thing. Suppose, for example, that apple juice, orange juice, and pineapple juice are all on a par with respect to, say, tastiness. It is plausible to think that a cocktail of apple and orange juice is better than either apple juice on its own or orange juice on its own but not better than pineapple juice, which after all is a rather different taste. Similarly, it is plausible to think that pineapple-orange juice is better than either orange juice or pineapple juice on its own but not better than apple juice; and that apple-pineapple juice is better than either apple juice or pineapple juice but not better than orange juice. Here we have a case in which each pair of a trio of items on a par is worse than something else without the third also being worse than it. It might be worth noting that this case involves a generalization of the familiar improvement condition. That condition holds that it is possible to improve one of two items on a par without thereby making it better than the other. The generalized condition holds that it is possible to improve each of two items on a par without thereby making that improvement better than the third. Insofar as such a case is possible, the Pareto Interval Rule, which precludes it, must be rejected.¹⁵

15. In his "Modelling Parity," draft manuscript, Wlodek Rabinowicz cites a theorem employed by Peter Fishburn (and brought to his attention by Erik Carlson) that shows that interval modeling has insufficient "dimensionality" to cover a configuration that we

If this is right, we are left with a striking result. Parity requires giving up the interval mode of representation; there is no such form of representation that satisfies certain reasonable axioms. This conclusion is of interest not only to those who believe that items can be related by some fourth relation but also to those who think that items can be incomparable. We set aside incomparability by assuming that our model was not trying to capture it. But the conditions of "improvement" and "generalized improvement" are arguably conditions that hold also of incomparability. So it turns out that the interval model, like standard expected utility theory, makes no room for either parity or incomparability. Those who are skeptical of parity and yet think that representation of the value of items must be made by something more "relaxed" than a single real number must find a representative device other than intervals.

III. PARITY AND CHOICE

Given that Gert's attempt to "explain away" parity does not succeed and that the model of comparability he offers cannot be sustained, I suggest that we understand the thrust of his discussion as a demand for an explanation of what one should rationally do when faced with items on a par, along with a proposed answer. Indeed, Gert's definition of parity—that it's not a mistake to choose either and that this may continue to be true even if one is improved—roughly approximates what I take to be the practical consequences of being on a par. But his gloss does not distinguish between different ways in which it might be rationally permissible to choose either of two alternatives. As I now want to suggest, parity underwrites one distinctive way in which rational permissibility operates, one that requires, on certain reasonable assumptions, that there be a fourth positive value relation that may hold between items.

Gert usefully reminds us that we cannot read off our valuations of items from our choices. It might be that two bales of hay are equal in value, but since the ass is hungry, it goes for the one on the right. This does not imply that the ass or anyone else should think that it would

should expect sometimes to hold when there is parity. Dimensionality is a technical concept having to do with the least number of linear orderings whose intersection is the partial ordering. Using the same configuration that Fishburn relies on to reject interval modeling on the basis of dimensionality, I argue that interval modeling should be rejected on the basis of the somewhat more intuitive condition of par dictatorship. I am grateful to Rabinowicz for bringing Fishburn's result to my attention. See Rabinowicz, "Modelling Parity", Erik Carlsson, "Incomparability and the Measurement of Value" (draft manuscript, Department of Philosophy, Uppsala University); and Peter C. Fishburn, *Interval Orders and Interval Graphs—a Study of Partially Ordered Sets* (New York: Wiley, 1985), pp. 78 ff., as cited in Rabinowicz. Rabinowicz goes on to offer his own interesting suggestion as to how parity should be modeled. Rabinowicz's model, however, does not undermine the thought that parity is a genuinely distinct fourth value relation.

have been a mistake to go for the bale on the left. Or two courses of action might be incomparable, but one has to choose even if by doing nothing. The course of action one takes is not thereby the only rationally permissible action. Similarly, two careers might be on a par, and yet the fact that one chooses the one career does not show that it would have been a mistake to have chosen the other.

Thus, there are three different cases in which choice between either alternative is rationally permissible: when the alternatives are equally good, incomparable, and on a par. Now it might be thought that what is distinctive about parity, as opposed to equality or incomparability, is that when one is faced with choices on a par, it is not a mistake to choose either, *and* this may continue to be true even if one or other of the items is improved. As we have already noted, however, this condition also holds when items are incomparable. If there is no positive pairwise relation that holds between them with respect to what matters in the choice, then it is not a mistake to choose either. And if we improve one, it need not thereby become better, and so again there need be no mistake in not choosing the improved item over the unimproved one.

Of course, in order for parity to be a genuine fourth positive value relation, it need not be true that its practical upshots are different from those of any of the usual three relations or of incomparability. Indeed, those who have suggested related notions, such as Derek Parfit, who thinks there can be "imprecise equality," and James Griffin, who thinks there can be "rough equality," think that their notions have familiar practical upshots; in the case of Parfit, imprecise equality should be treated just as one should treat incomparable items, and in the case of Griffin, rough equality should be treated just as one should treat items that are equally good.¹⁶ If, however, there is practical work to be done that cannot be done either by the usual three relations or by incomparability, we will have a further argument—from practical reason—for the existence of parity.

I believe that there is a distinctive role for parity to play in understanding rational choice. But it is easy to overlook this role if we think of the rational permissibility of choice as given solely by internal features of a given choice situation. Sometimes whether it is rationally permissible to choose something depends on how that choice relates to one's other choices.¹⁷

16. See Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984), p. 431, and unpublished work; and James Griffin, *Well-Being* (Oxford: Clarendon, 1986), pp. 80-81, 96-98.

17. Compare Gert's claim that the rationality of choice is not simply a function of an agent's valuation of the alternatives. According to Gert, a rational agent might recognize that objects take a certain range of value-points, but to preserve "consistency," she must choose in a way that does not involve her taking those objects to have particular value-

Suppose we have a rational choice function that delivers rational choice over a set of alternatives in any such choice situation. For any given choice situation involving two alternatives, the function will yield one of two answers: either one alternative is what one should rationally choose or it is rationally permissible to choose either. Now if the choice function yields the answer that one of the alternatives and not the other is what one should rationally choose, then it seems plausible to think that the alternative favored by the function is better than the other. If, instead, it yields the answer that it is rationally permissible to choose either, it seems plausible to think that the alternatives are equally good. Finally, if the choice function is undefined, that is, if it does not yield an answer as to which of the alternatives one should rationally choose, then the alternatives are incomparable. Thus, it seems that there is no distinctive role in reason for parity to play.

But suppose we expand the choice function to include series of choices. A given pairwise choice can be understood statically, apart from any other choice, but also dynamically, as part of a series of choices. The key question is whether in a choice between two alternatives that is part of a series of choices there is a practical difference in choosing between items that are equally good and choosing between items that are on a par. If there is a difference, we will have identified a way in which it may be rationally permissible to choose either of two items, not because the alternatives are equally good and not because the items are incomparable.

That there is such a difference is most clearly seen in "value pump" cases, although there are others. Value pump cases are puzzling because through a series of putatively rational choices, one may end up with less value than one started with. But the cases arise only if the alternatives are on a par, not if they are equally good.¹⁸ If one faces a series of choices between items that are equally good, however one cycles the choices, one will always end up with something as good as what one started with. Not so in a series of choices between items that are on a par. Suppose A is on a par with B, B is on a par with A+, and A+ is better than A. This possibility is implied by the "improvement" condition

points that could lead to the charge of "inconsistency" in choice (p. 504). In short, to avoid the value pump problem, Geir adds a free-floating constraint that an agent not allow himself to be a value pump, while proponents of parity build such a constraint into the proper valuation of items as on a par.

18. Of course, as I have argued elsewhere, the value pump problem arises also when there is incomparability. But we can distinguish two sorts of value pump puzzles: those in which the value pump is created by choices delivered by a choice function, and those in which it is created by the failure of a choice function to deliver a correct choice. My focus here is on the puzzle arising when the choice function is not silent. See my "Introduction," in *Incommensurability, Incomparability, and Practical Reason*, ed. Ruth Chang (Cambridge, Mass.: Harvard University Press, 1997).

on parity. Now if one is faced with a choice between $A+$ and B , it is rationally permissible to choose either since they are on a par. Suppose one chooses B . Now suppose that one is offered a choice between B and A . Since they are on a par, again it is rationally permissible to choose either. Suppose one chooses A . But now one is left with A where before one might have had $A+$, which is better than A .

The rational permissibility of choosing either of two items on a par, then, must be constrained by one's other choices. If one chose B when offered a choice between $A+$ and B , one is thereby rationally prohibited from choosing A when offered a choice between B and A . This is true even though there is a sense in which because B and A are on a par, it is rationally permissible to choose either. This is the sense in which if one had not already chosen B over $A+$, it would have been rationally permissible to choose A over B . Sometimes, when items are on a par, it is both rationally permissible to choose either and also rationally impermissible to choose one of them. The air of paradox is dispelled once we see that the sense in which it may be rationally impermissible to choose one of two items on a par depends on understanding the rationality of choice against a background of other choices.

So we can distinguish three different senses in which it may be "rationally permissible" to choose either of two alternatives. Sometimes the rational choice function delivers the result that it is rationally permissible to choose either alternative regardless of one's other choices. This is the case in which the alternatives are equally good. Other times the rational choice function is undefined; it fails to give an answer as to whether one should rationally choose one alternative or whether it is rationally permissible to choose either. This is the case in which the alternatives are incomparable. In still other cases, the rational choice function delivers the answer that it is rationally permissible to choose either alternative, but only given certain assumptions about one's other choices.

If we think that there are choices in which whether it is rationally permissible to choose either alternative depends on our other choices, then there is a distinctive role for parity to play in practical reason. Indeed, I believe that without parity, it would be impossible for each of us to be the normatively distinct rational agents that we are, each of us rationally caring about different things in different ways. But that is an argument for another occasion. The point I want to emphasize here is that there is a perfectly ordinary sense in which it may be rationally permissible to choose either of two alternatives that cannot hold when items are equally good or incomparable. A fourth relation beyond the usual three is needed to do this practical work.

Appendix

I provide an informal proof that the Pareto Interval Rule is the only rule of comparison that can be given an interval representation and is consistent with certain plausible constraints.

Excluding their reverse cases, there are seven possible configurations of relations between two intervals, X and Y (fig. A1). We assume three axioms about representation: (1) that in case 0, X is better than Y; (2) that case 6 is the only case of equality; and (3) that "duality" holds, that is, if betterness is represented in one direction, then worse-ness is represented in the opposite direction.¹⁹

This leaves configurations 1–5, any of which might be a case of X being better than Y, X being worse than Y, or X being on a par with Y. Since the point of the representation is to model parity, at least one of configurations 1–5 should represent a case in which X is on a par with Y.

We also assume four axioms about the intrinsic features of the relations being modeled: (1) that 'better than' (and conversely, 'worse than') is transitive; (2) that parity is symmetric; (3) that if two items are on a par, then either can be improved without thereby being better than the other (the improvement condition); and (4) that the relations are mutually exclusive.

The proof can then be given in three steps.

Step 1: Employing the axioms of the transitivity of betterness (and worse-ness), exclusiveness, and duality, we prove that none of configurations 1–4 gives a case in which X is worse than Y. (Correspondingly, none of the reverses of configurations 1–4 gives a case in which the first interval is better than the second.)

Proof. Consider case 1 (fig. A2) as an example. Suppose that X is worse than Y. Then Y is worse than Z by the same token, and so, by the transitivity of worse than, X is worse than Z. But by case 0, we know that




FIG. A1

19. The assumption that configuration 6 is the only case of equality can be replaced by the weaker assumptions of the transitivity and substitutability of equality. The assumption that configuration 0 is a case of betterness can be dropped, in which case we should also allow the converse of the Pareto Interval Rule to hold. The duality assumption can be weakened so as not to involve the implication listed as fact 1. If, however, duality is dropped altogether, we no longer have interval modeling in the sense that both Gert and I have in mind, i.e., as modeling by a relaxation of the reals where direction has significance.




FIG. A2

X is better than Z. By the exclusivity axiom, however, if X is worse than Y, it cannot also be better than Y. Therefore, X is not worse than Y.

The same kind of reasoning can be employed in cases 2-4.

We are therefore left with the following constraint: cases 1-4 represent either betterness or parity, and case 5 represents any of betterness, worseness, or parity. There are forty-six possible rules of comparison that meet this constraint.

Step 2: We may prove the following four facts (the proofs are omitted for lack of space):

- Fact 1: Case 5 must be a case of parity (by duality).
- Fact 2: Cases 3 and 4 must represent the same relation (by duality).
- Fact 3: If case 2 is a case of betterness, then so is case 1 (by the transitivity of betterness).

Fact 4: If cases 3 and 4 are cases of betterness, then so is case 2. And by Fact 1, case 1 would also be a case of betterness (by the transitivity of betterness).

We can use these four facts to eliminate forty-two of the total forty-six rules of comparison. This leaves four possible rules of comparison, R1-R4:

	Case						
	0	1	2	3	4	5	6
R1	B	B	B	B	B	P	E
R2	B	B	B	P	P	P	E
R3	B	B	P	P	P	P	E
R4	B	P	P	P	P	P	E

Step 3: We now prove that R2, R3, and R4 must be rejected.

Proof. R2-R4 assign parity to case 3. If case 3 is a case of parity, we should be able, by the improvement condition, to improve X without thereby making it better than Y (fig. A3).

a) Suppose that the only case of betterness is case 0. Then any improved X must stand to X as X stands to Y in case 0, that is, any

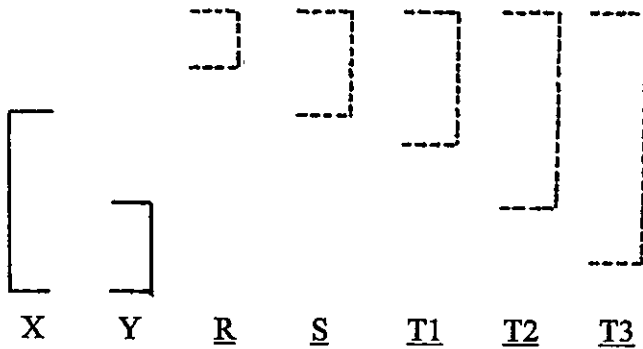


FIG. A3

improved X must be represented by an interval such as *R*. But *R* is better than *Y* by the betterness of case 0. Thus, there is no way to improve X without making the improved X better than *Y*; *R*4, which assigns betterness only to case 0, must therefore be rejected.

b) Suppose that there are two cases of betterness: cases 0 and 1. Given the reasoning of *a*, any improved X must stand to X as in case 1, that is, by an interval such as *S*. But *S* is better than *Y* by the betterness of case 0, and the improvement condition is again violated. Thus, *R*3 must be rejected.

c) Suppose that there are three cases of betterness: cases 0, 1, and 2. Given the reasoning of *a* and *b*, any improved X must stand to X as in case 2, that is, by any of the intervals such as those given by *T*1–*T*3. But *T*1 is better than *Y* by the betterness of case 0, *T*2 is better than *Y* by the betterness of case 1, and *T*3 is better than *Y* by the betterness of case 2; and the improvement condition is violated. Thus, *R*2 must be rejected.

This leaves only *R*1, the Pareto Interval Rule. It is a straightforward matter to show that this rule satisfies all the given axioms.