



LUND UNIVERSITY

Value Relations

Rabinowicz, Wlodek

Published in:
Theoria: a Swedish Journal of Philosophy

DOI:
[10.1111/j.1755-2567.2008.00008.x](https://doi.org/10.1111/j.1755-2567.2008.00008.x)

2008

[Link to publication](#)

Citation for published version (APA):
Rabinowicz, W. (2008). Value Relations. *Theoria: a Swedish Journal of Philosophy*, 74, 18-49.
<https://doi.org/10.1111/j.1755-2567.2008.00008.x>

Total number of authors:
1

General rights

Unless other specific re-use rights are stated the following general rights apply:
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Value Relations

by

WLODEK RABINOWICZ

Department of Philosophy, Lund University

Abstract: The paper provides a general account of value relations. It takes its departure in a special type of value relation, parity, which according to Ruth Chang is a form of evaluative comparability that differs from the three standard forms of comparability: betterness, worseness and equal goodness. Recently, Joshua Gert has suggested that the notion of parity can be accounted for if value comparisons are interpreted as normative assessments of preference. While Gert's basic idea is attractive, the way he develops it is flawed: His modeling of values by intervals of permissible preference strengths is inadequate. Instead, I provide an alternative modeling in terms of intersections of rationally permissible preference orderings. This yields a general taxonomy of all binary value relations. The paper concludes with some implications of this approach for rational choice.

Keywords: value, parity, incomparability, buck-passing account of value, Chang, Gert.

THE OBJECTIVE OF this paper is to provide a general account of value relations. But where it takes its departure is with regard to a special type of value relation, parity. The notion of parity is due to Ruth Chang. According to her, two items may be evaluatively comparable even when neither is better than, worse than, or equally good as the other. There is a fourth kind of comparability: being on a *par*. Recently, Joshua Gert has suggested that this somewhat elusive notion of evaluative parity can be easily accounted for if one interprets value comparisons as normative assessments of preference and allows for two levels of normativity: requirement and permission. The distinction between equality in value and parity is easily made on this approach. As I show below, if appropriately extended, the approach in question also allows for a straightforward distinction between parity and incomparability.

However, while Gert's basic idea is attractive, the way he develops it is flawed. He takes it that rationally permissible preferences one might have for an item can vary in strength and then models value comparisons by comparing intervals of permissible preference strengths for different items. As will be seen, however, such an interval modeling has features that make it unfit for the representation of the structure of value relationships. Instead, I provide an alternative modeling in terms of intersections of rationally permissible preference orderings, use it to delineate a general taxonomy of binary value relations, and conclude

with some tentative remarks about the implications of this approach for rational choice.

1. Introducing Parity

In “The Possibility of Parity”, Ruth Chang (2002b) argues that two items may be evaluatively comparable even though neither is better than, worse than, or equally good as the other. Instead of being related in one of these standard ways, they may be on a *par* (see also Chang, 1997, 2002a). As an example, consider two great artists; say, Mozart and Michelangelo. They are comparable in their excellence, but we might want to deny that one of them is better or worse than the other, or that they are equally good. It seems appropriate, however, to treat them as being on a par.

How is the existence of this fourth kind of comparability established? Chang considers cases in which we confront two different items, x and y , neither of which in our view is better than the other. In some respects, one is better; in other respects, it is the other way round; but neither is better *tout court*. That x and y are not equal in value can, in such cases, be established by what she calls “the Small Improvement Argument”, in which we envisage a third item, x^+ , very similar to x , such that x^+ is slightly better than x without being better than y . Obviously, this would have been impossible if x and y had been equally good (cf. Chang, 2002b, section 1¹). That x and y nevertheless may well be comparable in value, rather than *incomparable*, is shown by Chang as follows. In cases like this, we can often think of some item, z , worse than both x and y , but of the same kind as y . In addition, we can envisage a finite sequence of items starting with z and then going all the way to y , in which every successive item in some respect slightly improves on its immediate predecessor, while being equal to it in all the other relevant respects. We might call such an improvement “unidimensional”. Clearly, if z is worse than y in several respects, improvements in the sequence need to be made in each of these respects as one moves from z to y . But in every step in the sequence there is a (slight) change in one respect only. Now, it would seem that a small unidimensional improvement should not affect comparability: it should not take us from an item that is comparable with x to one that is not comparable. Consequently, since the first element in the sequence is supposed to be comparable with x (by hypothesis, z is worse than x), the same should apply to every element that follows, up to and including the last

1 In that section, Chang also presents an argument for the claim that very diverse items normally will never be equally good (cf. pp. 671f.).

element, y . Chang calls this “the Unidimensional Chaining Argument” (see Chang, 2002b, section 2).

Chang admits, though, that the principle which underlies her chaining argument is only meant to apply to a certain class of cases. That a small unidimensional improvement cannot effect a switch from comparability to incomparability is a principle for cases in which value comparisons are not made in accordance with some algorithmic rule, but instead are “a matter of balancing or trading off” different relevant respects of comparison against each other (Chang, 2002b, p. 676). Algorithmic rules may well allow for sharp breaks in comparability occasioned by small unidimensional changes. On the other hand, in the case of informal balancing procedures, the chaining argument can be accused of not taking into consideration possibilities of vagueness in judgements of comparability: this argument then becomes dangerously similar to a *sorites*.² Chang admits that the argument for the possibility of parity as a fourth type of evaluative comparability remains incomplete until it is shown, as she endeavours to do in her article, that parity phenomena cannot be explained away as cases of vagueness in evaluative comparisons, or as mere gaps in our evaluative knowledge.

2. Value and Rational Preferences

Chang takes the possibility of parity to show “the basic assumptions of standard decision and rational choice theory to be mistaken: preferring X to Y, preferring Y to X, and being indifferent between them do not span the conceptual space of choice attitudes one can have toward alternatives” (Chang, 2002b, p. 666). Joshua Gert (2004b) questions this claim and suggests that there is no need to revise the traditional trichotomy of preference relations in order to account for parity.

On the face of it, Gert also wants to make another claim, which applies to value rather than to preference. He denies that cases of parity necessitate giving up the traditional trichotomy of *value* relationships between comparable items: better, worse and equally good:

2 As is easily seen, Chang’s assumption that unidimensional improvements do not make comparability disappear might well be questioned if we allow for vague comparability. Her chaining argument can then be dismissed as just a version of a *sorites*. The starting point of the sequence (z) might be clearly comparable with x , the end point (y) might be clearly not comparable with x , and we might have cases of vague comparability in between. As it is not my purpose to defend Chang’s arguments for the existence of parities, but rather to show that parity is conceptually possible, I will not discuss this objection any further.

The trichotomy thesis holds that if two items are comparable, it is because one of the items is better than the other, or because they are equally good . . . This article 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 (Gert, 2004b, p. 493).

However, as one continues to read Gert's article, it becomes clear that this promissory note must be based on a misunderstanding. As it turns out, Gert does accept that the three traditional relations do not exhaust all the possible ways in which two comparable items can be related to each other. Like Chang, he takes parity to be a positive value relation that can only obtain when none of the three traditional value relations are present. Nor does his approach allow for defining all the positive value relations, including parity, in terms of the trichotomy of 'better', 'worse' and 'equally good'. What can be shown, however, even though Gert puts it in a misleading way, is that the traditional trichotomy of *preference* relations suffices to account for all value relationships, parity included (see also Chang, 2005).³ More precisely, as will be shown below, the taxonomy of all value relations can be constructed in a framework that, along with the traditional triad of preference relationships, also allows for preferential gaps.

Gert's positive solution is based on an analysis of the notion of betterness that has attractive and influential antecedentia in the philosophy of value. While he never mentions it, he follows a long tradition. According to the view that goes back at least to Franz Brentano and counts among its proponents such philosophers as A. C. Ewing, John McDowell, David Wiggins, Allan Gibbard and Thomas Scanlon, to be valuable is to be a fitting object of a pro-attitude. More precisely, an object is valuable insofar as it has features that make it fitting or appropriate to favour that object in some way. 'Fitting', 'appropriate', 'ought', etc. stand for the normative component in this type of analysis; the features of the object that make favouring appropriate are what we call its value-making properties; and different kinds of favouring – desire, admiration, liking, cherishing, etc. – correspond to different kinds of value: desirable, admirable, likeable, precious and so on. For the relation of betterness, the relevant kind of favouring is preference: an item is better than another if and only if it ought to be preferred. Or, as Brentano put it: "When we call one good 'better' than another, we mean that the one good is preferable to the other. In other words, it is *correct to prefer* the one good, for its own sake, to the other" (1969 [1889], p. 26; cf. Ewing, 1947; McDowell, 1985; Wiggins, 1987; Gibbard

3 Apart from questioning Gert's apparent adherence to the traditional trichotomy of value relations, Chang (2005) also criticizes Gert's interval modeling, as I do, using partly similar arguments. She does not provide any alternative modeling, though. Indeed, Chang is not prepared to accept Gert's basic idea that value comparisons can be analysed in terms of normative assessments of preference. In this important respect, her view differs from mine.

1990, 1998; Scanlon, 1998).⁴ On this format of analysis, then, a claim of betterness consists in a normative assessment of preference.⁵

Gert specifies the normative component of the analysis in terms of the notion of rational requirement or, equivalently, in terms of its dual: the notion of rational permissibility:

An item x is *better* than another item y if and only if it is rationally required to prefer x to y .

Or, what amounts to the same, x is better than y if and only if it is not rationally permissible not to prefer x to y .

In the first part of his paper, Gert applies the notions of rational requirement and permissibility to choices rather than to preferences. Thus, he interprets “better” as meaning something like “to be chosen, on pain of having made a mistake” (2004b, p. 499). But as one reads on, it becomes clear that it is preferences, understood as choice dispositions, that on his view are the primary object of the rationality assessments expressed in judgements of betterness.⁶ It should be noted that preference in this context cannot itself be understood as a judgement of betterness, as this would make the analysis of betterness in terms of required preferences circular. This circle is avoided if we instead take preferences to be dispositions to choose.⁷

4 Ewing (1947) is the *locus classicus* for this format of analysis. Scanlon calls this approach “the buck-passing account of value”, because it transfers the reason for favouring the object from the object’s value to its value-making properties (1998, p. 97). For a discussion of some difficulties facing this format of analysis, see Rabinowicz and Rønnow-Rasmussen (2004). One such difficulty (the so-called “Wrong Kind of Reasons” problem) is that preferring an object may be required not because of the features that make that object better, but rather because the preference would itself be valuable, as such or in virtue of its effects, or because this preference would be appropriate for deontological reasons not having to do with the value of its object. Cases like this must somehow be excluded if the analysis is to be acceptable. Another difficulty is that there is a danger of circularity in this approach if either the normative component (requirement) or the attitudinal component (preference) themselves need to be analysed in terms of the concept of betterness. (For some remarks on how the notion of preference can be analysed in order to avoid circularity, see below.)

5 Note, however, that the multiplicity of different kinds of favouring that might be fitting with respect to different items implies that this general format of analysis could be used not just for betterness but also for other kinds of asymmetric value relations. (I am indebted to David Alm and Daniel Svensson for this useful reminder.) Thus, an item is *more admirable* than another item if it ought to be more admired, it is *more desirable* if it ought to be more desired, and so on. These other kinds of relations will not be considered in this article, but much of what follows could be applied to them as well, *mutatis mutandis*.

6 This is especially clear when he provides his modeling, in which rationality assessments apply to preferences of variable strength. As for the nature of such assessments, Gert refers the reader to ch. 7 of his book, *Brute Rationality* (2004a), in which he interprets them in a cognitivist way. As he points out, though, his value analysis could just as well be given a non-cognitivist twist if “rational” (or “rationally permissible”) were interpreted as an expression of approval rather than as an attribution of a property.

7 That is one way of avoiding the circle in the analysis. Alternatively, one could try to avoid circularity by interpreting preferences as emotive attitudes rather than choice dispositions. On this emotivist approach, preferring x to y would involve experiencing x as more appealing than y , or more pleasing, or something

Does the connection to preferences impose any ontological restrictions on the possible relata of the betterness relation? According to an influential view, the objects of preference can only be states of affairs (particular or generic). On another view, preferences instead are directed towards properties of the preferrer: I prefer to eat rather than to drink, to listen to Mozart rather than to Frank Sinatra, to live in a world in which Hitler was defeated rather than in a world in which he was victorious, etc. (Lewis, 1979). Such restrictions on the objects of preference would certainly be opposed by someone like Brentano (Chisholm, 1986, chs. 2 and 3).⁸ But if some such restrictive view were correct, preferences regarding such entities as persons or material things would always at their base consist in attitudes towards some states of affairs, or towards some properties of the preferrer. On the analysis of betterness in terms of fitting preferences, this would mean that betterness relations between persons or between concrete things would ultimately be reducible to corresponding relations between states or properties. Whether this reducibility claim is correct or not is a difficult matter, which I will not try to deal with in this paper.

An item x is better than another item y if preferring x to y is rationally required.⁹ There is here a tacit assumption that a potential preferrer who is the subject of this requirement is familiar with the items under consideration. Needless to say, in the absence of epistemic access to the compared items, one's preferences need not be required to track value relations. Indeed, under such conditions, it might be rationally required not to have any preference at all with regard to the items in question. Epistemic access will be presupposed throughout in what follows.

What this assumption of epistemic access exactly amounts to is not easy to specify. For at least two different reasons, we cannot take it to mean complete familiarity with the properties of the items under consideration: (i) complete knowledge might be an ideal that is impossible to realise; (ii) on pain of circularity,

along these lines. The circle in the analysis is thereby avoided, *if* emotions at issue can be given an account that does not make use of evaluative concepts. For a thorough discussion of the latter avenue, see Svensson (2004).

⁸ Note, however, that for Brentano preferences are emotive attitudes. If one interprets preferences as dispositions to choose, it is more difficult to resist the conclusion that objects of preferences have to be state-like or property-like in nature. I am indebted to Björn Petersson for pressing this point.

⁹ Cf. Gibbard (1998, p. 241): "To be desirable, we might say, is to be desired fittingly, or justifiably, or rationally. Or since a desirable thing might not be desired at all, we should speak hypothetically: something is desirable if it *would* be reasonable to desire it. It is desirable if desiring it would be *warranted*, if it would *make sense* to desire it, if a desire for it would be *fitting* or *rational*. Likewise, the preferable thing is the one it would be rational to prefer." Gibbard does not clearly distinguish in this context between required and merely permissible preference. This is a crucial distinction in Gert's proposal. But since preferability is an asymmetric relation, Gibbard's expressions – "warranted", "fitting" and "rational" – must be interpreted as cognates of "rationally required" rather than of "rationally permissible". For the permissibility of a preference is logically compatible with the opposing preference also being permissible.

we cannot assume that epistemic access extends to the evaluative features of the items.¹⁰ In what follows, however, these problems with the clarification of the assumption of epistemic access will be swept under the rug. Another problem that will be ignored is the question of whether preferences and attitudes in general can at all be subject to rational requirements. I think they can, despite the fact that they are arguably outside our direct voluntary control, but I cannot pursue this discussion here.

Let us continue with the analysis of evaluative relations. Being worse is simply the converse of being better. Thus,

x is *worse* than *y* if and only if it is rationally required to prefer *y* to *x*.

Similarly,

x and *y* are *equally good* if and only if it is rationally required to be indifferent between *x* and *y*.

In other words, two items are equally good if they ought to be equi-preferred.

On this view, it is easy to see where parity comes in: if *x* and *y* are on a par, it is rationally permissible to prefer *x* to *y*, but it is also rationally permissible to have the opposite preference. Gert describes situations like this as follows:

... only very rarely do we think of our particular personal preferences as the uniquely rational ones. This view of preference and value allows that two people in the same epistemic situation, who have the same perfectly precise standards for assessing the value of items with respect to V, and who take the same interest in whether or not something has value V, could make different, but equally rational choices between two items, when the relevant value is value V (2004b, p. 494).

Gert and Chang take it that comparisons between items are always made with respect to some covering value or consideration, which may differ depending on the context of comparison. Thus, when comparing two persons, we might ask which of them is the better artist, the better swordsman, or the better lover. The covering consideration is important when we inquire what preferences are rationally permissible. When it is a question of, say, Michelangelo's and Mozart's relative merits as artists, we want to know whether it is permissible to prefer one to the other *as an artist*, and not, say, as a conversationalist. Thus, the preferences that are at issue always are relative to some such more or less specific covering consideration.

In what follows, this reference to the covering consideration will be suppressed to make the exposition simpler, but before we leave this matter, let me take up an objection to the analysis of betterness as preferability that might be raised by a

10 Cf. Broad's (1930) cautious formulation of this format of analysis: "I am not sure that 'X is good' could not be defined as meaning that X is such that it would be a fitting object of desire to any mind which had an adequate idea of its *non-ethical* characteristics" (p. 283, my italics).

satisficer, such as Michael Slote (1989).¹¹ On one interpretation of satisficing, it is sometimes permissible to prefer a worse item to a better one, or at least to be indifferent between the two. Such preferential attitudes are permissible if the worse item is “good enough”. Obviously, if this view is correct, the analysis of betterness in terms of required preference would not be viable. I am inclined to think, however, that the *prima facie* appeal of satisficing rests on the conflation of different covering considerations that might be involved in comparisons between items. Let me illustrate this with an example used by Slote. An impoverished family comes to a hotel looking for shelter. The hotel manager gives them a room to stay in, not one of the best but one he thinks is good enough. While he considers the presidential suite to be a better accommodation, he does not prefer it to an ordinary room in the problem at hand. There is no mystery here, I think, and no deep problem for the analysis. The presidential suite is better, i.e., more preferable, *as an accommodation*, but an ordinary room is at least as good, or on a par, *as a shelter*. The latter requires, however, that the less luxurious room does not come too low on the betterness scale for accommodations. We express this by saying that the ordinary room is good enough: it is good enough as an accommodation to be perfectly good as a shelter. The hotel manager’s preferences can thus be accounted for if one keeps covering considerations apart. Many of the examples that satisficers come up with could, I think, be dealt with along similar lines.

In yet other cases, the idea of satisficing reduces to the distinction between the goodness of *options* and the goodness of *outcomes*. Thus, according to the satisficing version of consequentialism, an option might be optimal even if it leads to a sub-optimal outcome, provided that the outcome is good enough (i.e., sufficiently good to make the option optimal). Such a non-standard form of consequentialism might be coherent, but even if it is, it does not pose any threat to the analysis of betterness in terms of required preferences, as long as one is clear as to whether it is the options or outcomes that are being compared. I hope these sketchy remarks can suffice for now. The whole complex issue of satisficing cannot be adequately dealt with in this paper.

Let us go back then to the analysis of evaluative relations. Before defining parity in a more precise way, I want to consider the notion of incomparability. How is value incomparability to be analysed on the present approach? Gert (2004b) does not address this issue, but the framework he works with allows for a straightforward extension that makes room for incomparabilities. As has been suggested above, preference can be seen as a disposition to choose. To prefer *x* to *y* is to be disposed to choose *x* rather than *y* when one has to make a choice between the two items in

11 I am indebted to Jonas Olson for bringing this objection to my attention. Chang (2005) also raises this issue.

question. Indifference is yet another type of choice disposition: one is indifferent insofar as one is equally prepared to make either choice. But then, it seems, there could also exist pairs of items with respect to which a person *lacks* a choice disposition. If necessary, a person like this will of course make some choice, but not because she is so disposed. Not all choices we make are manifestations of choice dispositions.

It is important to distinguish the absence of a choice disposition from indifference. In the latter case, the subject smoothly proceeds to choice. Buridan's ass, after all, is just a philosopher's fiction. But in the absence of a choice disposition, we typically experience the situation as involving an internal conflict. We can see reasons on each side, but cannot (or simply will not) balance them off. We do make a choice if we have to, but this choice is made without the conflict of reasons being resolved.¹²

At this point, one might object to this whole idea of absent choice dispositions and point out that, whatever I do, I must be doing it because I am in some sense so disposed.¹³ In principle, it seems, it is always possible to trace back my behaviour to external stimuli plus something dispositional: a configuration of internal factors that make me react to external stimuli in a certain way. So in *this* sense I always have a disposition to choose, when I choose. But what I am trying to get at is a stronger sense of a choice disposition – the sense in which this disposition is present only if I am disposed to make a deliberate and reasoned choice among the items with which I am confronted.¹⁴ In this stronger sense, of course, not everything one does is due to a choice disposition, since not everything one does is based on a reasoned choice. At the same time, it is arguable that the notion of preference used in the analysis of comparative value relations should be understood as a choice disposition in this stronger sense.¹⁵

Now, assuming that choice dispositions (in this qualified sense) can be absent, their absence can be subject to normative assessments. This allows us to accom-

12 An indirect evidence for the absence of choice dispositions is in some cases provided by sequences of choices. Thus, for example, someone who prefers x^+ to x but who does not prefer x^+ to y nor y to x , might first exchange x^+ for y , and then decide to exchange y for x , thereby being left with an item (x) she disprefers to the one she has started with (x^+). We can explain this sequence of actions if we assume that the subject lacks choice dispositions with regard to pairs (x^+ , y) and (x , y) and, in addition, is myopic, i.e., makes choices without predicting her own future behaviour. However, such an erratic choice sequence could also be accounted for in other ways; for example, by the combination of myopia with changes in preference or with preferential irrationality (cyclical preferences).

13 I am indebted to John Broome for pressing this point and emphasising the need to clarify the notion of choice dispositions that is needed for my proposal.

14 A choice in this qualified sense is possible even in the case of indifference. When two options come out as equal in my balancing of reasons, my choice of one of them is reasoned and deliberate even though I could just as well have chosen the other option.

15 See Rabinowicz and Rønnow-Rasmussen (2004, pp. 414–418) for a defence of the claim that the pro-attitudes to which one refers in the “fitting attitudes”-analysis of value should be reason-based.

moderate incomparabilities in our analysis. More precisely, if the absence of a choice disposition with regard to a pair of items is not just rationally permitted but is rationally required, then the items can be said to be incomparable. In other words, x and y are *incomparable* if and only if it is not rationally permissible to prefer one to the other or to be indifferent.

One might wonder whether this definition might be too demanding. Should it not be enough for incomparability that the absence of a choice disposition with regard to the items under consideration is rationally *permissible*, even though it is not required?¹⁶ Well, at least for linguistic reasons, such a lenient criterion would seem rather awkward. To give an analogy, we do not say that something is undesirable if it is merely permissible not to desire it. It is undesirable only if desiring it is *impermissible* in some sense.

However, perhaps we might still stipulate that:

x and y are *weakly incomparable* if it is rationally permissible neither to prefer one to the other nor to be indifferent.

Is it plausible to expect the existence of incomparabilities? To some extent, this depends on the item domain under consideration. If that domain contains items from different ontological categories, incomparabilities will not be hard to find. When we consider, say, a person and a state of affairs, it does seem just as irrational to prefer any of them to the other as to be indifferent between them. Preferring one to the other or being indifferent simply does not make sense.¹⁷ In fact, even within one and the same ontological category, incomparability might sometimes be expected. In *Making Comparisons Count*, Chang (2002a, section 6.1) introduces the notion of “non-comparability”. To avoid terminological confusion, it is preferable, I think, to refer to that relation as *essential incomparability*. Two items are essentially incomparable with respect to a given covering consideration if at least one of them does not fall into the domain in which that consideration is applicable. In this sense, for example, since Mozart was not a sculptor, he is essentially incomparable to other persons with regard to his excellence as a sculptor. Thus, two objects do not have to belong to different ontological categories in order to be essentially incomparable with respect to a given covering consideration.¹⁸

16 This question was posed to me by Walter Sinnott-Armstrong.

17 Unless we take a reductionist view and assume that a preference for a person at its base consists in a preference for some state of affairs involving that person; but the apparent absurdity of preferring a person to a state itself argues against such a reduction.

18 This assumes that comparing Mozart with other persons in regard of his excellence as a sculptor would require Mozart to *be* a sculptor in the first place. As Dan Egonsson has pointed out to me, this assumption might well be questioned. Similar criticism might be raised against other examples of purported essential incomparabilities that do not cross ontological borderlines.

In what follows, I do not distinguish essential incomparability as a separate type of value relation. Still, if one needed a criterion that picks out essentially incomparable items, here is one: for any such pair of items, at least one of them is incomparable to any other item, with regard to a given covering consideration. (The reason is that the item in question falls outside the domain in which that consideration is applicable. The criterion above is necessary for essential incomparability; it is also sufficient provided that every item to which the consideration applies is comparable to at least some other items, with regard to that consideration.)

What about items that are *not* essentially incomparable? Could they be incomparable anyway? One might doubt this for the following reason. It may well be permissible, one might suppose, to have no preferential attitude regarding two items that both fall into the domain of the covering consideration. In other words, weak incomparabilities might easily obtain with such a domain. But can the absence of a preferential attitude be positively *required* in such cases? Well, logically it is possible, of course, but it is not clear whether this logical possibility has actual instantiations. Probably, the most promising examples would be some cases of tragic dilemmas, such as Sophie's Choice. It is arguable that when you must choose which of your children should be saved, preferring one of the options is as impermissible as being indifferent. But is it a rational impermissibility or rather a moral one?

Still, for the purposes of this paper, it is not necessary to take a definite stand on the actual existence of instances of non-essential incomparability. It is enough to draw a map of conceptual possibilities.

What, then, about comparability? In one sense,

x and y are *comparable* if and only if they are not incomparable.

But in this sense, comparability and weak incomparability are not mutually exclusive. Full comparability of items would mean more than that: it would mean that the items are not even weakly incomparable; i.e., that anyone who considers the items should prefer one to the other or be indifferent. Absence of a choice disposition is disallowed in this case. Thus,

x and y are *fully comparable* if and only if it is required to prefer one of these items to the other or to be indifferent.

Parity, as we have seen, is supposed to be a form of comparability. This makes it impossible to define this notion as simply the complement of the traditional triad of positive value relations: better, worse and equally good. We cannot take it that x and y are on a par if neither of them is better or worse than the other and they are not equally good, for two such items, instead of being on a par, may not be comparable.

But what about the possibility of defining parity in terms of the three traditional relations? This avenue is not available either. That the traditional trichotomy of value relations does not suffice to define parity can be seen from the following

argument: suppose that, for a given domain of items, we have already determined for every pair of items whether the first member of the pair is better than the second, worse, equally good or none of these; i.e., we have determined what preferential attitude, if any, is required regarding each such pair: preferring, dispreferring, indifference or none of these three. Consider a pair x and y for which it turns out that none of the preferential attitudes is rationally required. Clearly, from our information about *required* preferential attitudes regarding all pairs of items, nothing follows about whether it is rationally *permissible* to have any preferential attitude regarding x and y ; i.e., the relations of betterness, worseness and equal goodness in the domain do not determine whether pairs of items that fall outside the extensions of these relations are comparable or not. But this means that the notion of comparability is not definable from the traditional triad of value relationships. The same applies to parity, of course. Remember that, if x and y are on a par, both preferring the former to the latter and dispreferring are rationally permissible. From our information about the scope of *required* preferential attitudes, nothing follows about whether it is permissible to have opposing preferences regarding the pairs of items that are outside this scope.¹⁹

19 This argument, however, crucially depends on the assumption that value relations are analysable in terms of permitted and required preferential attitudes. Erik Carlson (2007), who does not make this assumption, thinks that parity can be defined in terms of betterness. Carlson's definition proceeds in two steps. First, he defines the notion of "almost better than" in terms of "better" and its converse – "worse" (x is worse than y if and only if y is better than x):

x is *almost better than* y if and only if x is not better than y but (i) every z that is better than x is also better than y , or (ii) every z that is worse than y is also worse than x .

Then the latter relation and its converse, which we might call "almost worse than", are used to define parity:

x is on a *par* with y if and only if neither of them is better than the other but there is some z that is (i) better than one of them and almost better than the other, or (ii) worse than one of them and almost worse than the other.

Discussion of Carlson's definition and its motivation would take us too far afield. It is obvious, however, that in the absence of special restrictions on the object domain his definition is not equivalent to ours. Thus, suppose that the domain of items consists of just three objects, x , y and z , such that z is better than x but otherwise none of the objects in the domain is better than the other. Carlson's definition then implies that x and y must be on a par: neither x nor y is better than the other, while z is better than x and (trivially) almost better than y . However, it may be the case that on our approach x and y are incomparable in this case: it may be that it is rationally impermissible to prefer one to the other or to be indifferent between them. Such a situation is at least logically possible.

What about the opposite direction? Can two items be on a par on our approach without being on a par according to Carlson's proposal? Again, the answer is yes, if we do not impose any special restrictions on the item domain. Suppose that the domain consists just of x and y , neither of which is better than the other. Then, clearly, Carlson's definition implies that x and y are not on a par: there is no z such that z is better (worse) than one of the items x and y and almost better (almost worse) than the other. However, x and y might be on a par on our proposal: it might be permissible to prefer one of them to the other and permissible to have the opposite preference. Again, such a situation is logically possible.

Let us return, then, to our explication of parity in terms of permissible preferences. We know that parity presupposes comparability. But does it require full comparability, or is comparability in the weak sense of this term sufficient? Either solution is possible. Thus, in the broad sense,

x and y are *on a par* if and only if it is rationally permissible to prefer x to y and also rationally permissible to prefer y to x .

If x and y in addition are fully comparable, they will be said to be *fully on a par*.

Note that this definition of parity is possible only because on the present approach we are supposed to distinguish between two levels of normativity: the *strong* level and the *weak* one. Gert's analysis of value relations in terms of rationally warranted preferences makes room for parity because warrant can be interpreted either strongly, as a requirement, or weakly, as a permission. This introduction of two levels of normativity is Gert's main contribution to the "fitting attitudes"-analysis of value. To my knowledge, none of the earlier theorists in that tradition made use of this distinction in their analysis. The standard approach was always to give the strong interpretation to the deontic component of the analysis.

As an aside, I should point out that Gert's own definition of parity is much narrower than the one suggested above. In his opinion, for x and y to be on a par it is not enough that preferring each is rationally permissible. This would make parity too large a category in his opinion. According to him, x and y must additionally satisfy the condition that for any third item z , "the rational status" of various possible preference attitudes towards x and z is the same as that of the corresponding attitudes towards y and z (cf. Gert, 2004b, p. 506). This would imply, in particular, that if it is required to prefer z to x then it must also be required to prefer z to y . In other words, any item that is better than x would have to be better than y , and vice versa. Surely this is an excessively strong demand: Gert must be mistaken on that point. In typical cases of parity obtaining between two items, a small improvement x^+ of one item, x , need not be better than the other item, y . As we have seen, Chang's Small Improvement Argument takes its departure from cases like this.

3. Interval Modeling

As an idealisation, Gert assumes that the strength of possible preferences for different items is quantitatively measurable.²⁰ He then uses this idealisation in his

20 He does not specify the scale of measurement, but his discussion suggests that he has in mind something like the interval scale. This means that what is arbitrary about the numbers representing preference strengths is at most the choice of the zero point and of the unit of measurement. Still, as far as I can see, his modeling, strictly speaking, requires much less than this: a purely ordinal scale would be fully sufficient. Thus, the important thing is only that a higher number stands for a higher strength of preference.

formal modeling of value relations. Since it may be rationally permissible to prefer a given item, x , more or less strongly, we can assign to x an interval of real numbers, $[x^{\min}, x^{\max}]$, which specifies the rationally permissible range of preference strengths with respect to x . x^{\min} is the lower bound of that range, while x^{\max} is its upper bound. Now, Gert implicitly assumes that any combination of rationally permissible strengths of preference for different items is itself rationally permissible. For example, suppose that items x and y are assigned partially overlapping ranges $[10, 40]$ and $[5, 30]$, respectively. This means that it is permissible to prefer x with, say, strength 20, and that it is also permissible to prefer y with strength 20. Therefore, Gert takes it, it is permissible to have these preferences simultaneously, i.e., to be indifferent between these two items. Preference strengths such as, say, 30 for x and 10 for y are permissible as well, which means that it is permissible to prefer x to y . However, it is just as permissible to prefer y to x , since the upper bound of the rationally permissible preference range for y (30) is higher than the lower bound of the range for x (10).

In terms of this interval representation of permissible preference strengths, Gert formulates his “Range Rule” that provides a definition of the notion of betterness. Betterness requires the absence of overlap between intervals. One item, x , is better than another item, y , if and only if the lower bound of the permissible preference range for x is higher than the upper bound of the corresponding range for y (cf. Gert, 2004b, p. 505); or, in brief:

The Range Rule: x is better than y if and only if $x^{\min} > y^{\max}$

In other words, even the weakest permissible preference for x is stronger than the strongest permissible preference for y . For instance, suppose that x is assigned range $[10, 40]$, as before, but the range for y is now given by $[5, 9]$. Since 10, the lower bound for x , exceeds 9, the upper bound for y , x is better than y .

On this modeling, both parity and equality in value between distinct items are to be found among those cases in which the ranges for the items that are being compared at least partially overlap. Gert himself notes that on his interval modeling, equality in value is a rare phenomenon. Items x and y are equally good if and only if it is rationally required to be indifferent between x and y . But, on the interval modeling, this is possible only if the ranges for x and y coincide and in addition have *zero length*, i.e., consist of a single point. Thus,

(i) the range for x must be the same as that for y ,

and

(ii) the lower bound of this range must equal its upper bound.

In other words, there is a unique rational strength of preference for x and for y , which is the same for both items. Condition (i) is obviously necessary for x and y

to be equally good. But if x and y are distinct items, then we need condition (ii) as well, because if (ii) did not hold, then it would be rationally permissible to prefer x with the strength in the vicinity of the upper bound of the common range and it would be rationally permissible to prefer y with the strength in the vicinity of the lower bound. Since nothing in the modeling hinders combining these preferences, it would be rationally permissible to prefer x to y . This is, however, excluded, if x and y are to be equally good. At the same time, as we have seen, Gert recognises that “only very rarely do we think of our particular personal preferences as the uniquely rational ones” (2004b, p. 494). In particular, only very rarely do we take the strength with which we prefer a given item to be uniquely rational. Therefore, condition (ii) can only very rarely be satisfied. This means that, on the interval modeling, equality in value between distinct items will obtain very seldom, if at all.²¹

This feature of the modeling should give us pause. Another problematic feature is that it is unclear how a modeling like this can account for incomparabilities. Two items x and y are incomparable, as we have seen, if it is rationally impermissible to prefer one of them to the other or to be indifferent. But on the interval modeling, this would require, as far as I can tell, that for at least one of the items the range of permissible preference strengths must be empty. For if there were some permissible preference strengths for each of the items, then it would either be permissible to prefer one to the other or it would be permissible to be indifferent between them. However, if the range for, say, x were empty, then x would be incomparable not just with y but with every other item as well! Surely, this cannot be right: an item that is incomparable with some items should normally be comparable with at least some other items in the domain.

Gert might reply at this point that the interval modeling is appropriate only in the absence of incomparabilities in the domain. He might also try to convince us that it is not as counterintuitive as it seems for equality in value between distinct items to be a very rare phenomenon. However, worse things are yet to come. Gert’s prime application of his modeling concerns cases in which an item x is worse than another item x^+ , but neither of them is either better or worse than some third item y . To use his own example, think of x and x^+ as suffering the itch of poison ivy for one week and for one day, respectively, and let y be the pain that is typically caused by getting a filling at the dentist’s. While x is worse than x^+ , Gert suggests that neither of these

21 Chang (2005, pp. 340f) goes as far as to suggest that Gert’s interval modeling fails to make equal goodness a reflexive relation for all those items with respect to which rational preference can vary in strength. This, however, seems to be a misunderstanding. Even if the interval of permitted preference strengths for an item is of a non-zero length, this does not mean that one is permitted to *simultaneously* have two preferences with different strengths with regard to that item. At any given time, a subject can only have one strength of preference for an item. Thus, it trivially follows that, for any x , one is required to be indifferent between x and x .

experiences is either better or worse than y . The two kinds of pain are too different from each other to make a straightforward comparison possible. Here is another possible example. Let x and y be trips to Australia and to South Africa, respectively, with x^+ being a trip to Australia with an added bonus of \$100. While the latter is better than the trip to Australia without a bonus, neither of these two alternatives might be better or worse than a trip to South Africa. As Gert shows, cases like this are easily representable in his interval modeling. When the lower bound of the range for x^+ exceeds the upper bound of the range for x , both these ranges might still overlap the range for y .

However, let us in addition envisage a *fourth* item, y^+ , which we can think of as a trip to South Africa with an added \$100 bonus (or, in Gert's example, a somewhat shorter dental treatment). y^+ is better than y , but, let us assume, it is not better than x . The relation of y and y^+ to x is thus the same as the relation of x and x^+ to y . (As for x^+ and y^+ , it follows from what we have assumed that neither of these two items is better than the other.) Now, it can be shown that this structure of value relations between the four items cannot be represented by the interval modeling.²² Here is the proof:

Since x^+ is better than x and y^+ is better than y , the Range Rule implies:

(i) $x^{+min} > x^{max}$ and (ii) $y^{+min} > y^{max}$.

Now, there are two possible cases: either (1) $x^{max} \geq y^{max}$, or (2) $y^{max} \geq x^{max}$.

But (i) and (1) together imply that $x^{+min} > y^{max}$, which contradicts our assumption that x^+ is not better than y , while (ii) and (2) imply that $y^{+min} > x^{max}$, which contradicts the assumption that y^+ is not better than x .

This is a general result. The interval modeling implies, for all items x^+ , x , y^+ and y , that

If x^+ and y^+ are better than x and y , respectively, then it must be the case that either x^+ is better than y or y^+ is better than x .²³

Since this general implication is unwelcome, as we just have seen, it follows that the interval modeling is unfit to represent value relations.

Gert motivates his use of the interval modeling by reference to similar approaches to imprecise subjective probabilities (2004b, p. 510). It is easy to see, however, that the objection we have presented applies just as well to probability comparisons.

22 For a similar example, see Danielsson (1998). In fact, I learned this lesson from Danielsson a long time ago, in the 1970s. He presented it in print as early as 1983, in "Hur man inte kan mäta välmåga" ["How one cannot measure well-being"].

23 If a betterness relation satisfies this condition, along with being transitive and asymmetric, then it is a so-called *interval order*. As is well known, interval orders are exactly those relations that are representable by interval modelings that use the Range Rule (cf. Fishburn, 1970, pp. 20–23; this result holds for all countable item domains).

Statements such as “a proposition A is more probable than a proposition B ” cannot be interpreted by assignments of probability intervals to propositions together with the analogue of the Range Rule for the representation of the relation “more probable than”. Interval modeling is inadequate for this purpose as for the representation of value relations. To see this, we can use the same kind of structure as the one above. Thus, let A and B be two propositions about different issues, for which we do not have definite probability assignments. In particular, we do not consider them to be equiprobable, nor do we take one to be more probable than the other. Now, let C be some highly probable proposition that is logically independent of both A and B ; say, the proposition that the next throw of a die will not result in a six. A is slightly more probable than $A&C$, while B is slightly more probable than $B&C$. At the same time, it may well be that A is not more probable than $B&C$ nor that B is more probable than $A&C$. By the same argument as above, it then follows that no assignment of probability intervals to the four propositions A , B , $A&C$ and $B&C$ can account for their mutual probability relations.

What has gone wrong in such cases? Let us go back to betterness comparisons. Consider again the comparison between a trip to Australia and the same trip with a bonus of \$100. The latter is better, but is it reasonable to suppose that even the weakest rationally permissible preference for this alternative is stronger than the strongest rationally permissible preference for the former alternative? Surely this cannot be right. If we suppose that the range for the worse alternative is [10, 30], then the range of the better alternative should be, say, [11, 31], or something like that. It is thus to be expected that there will be a significant *overlap* between the two ranges. But the weakest permissible preference for the better alternative will be stronger than the weakest permissible preference for the worse alternative and the strongest permissible preference for the better alternative will likewise be stronger than the strongest permissible preference for the worse alternative.

Exactly the same observation applies to probability comparisons in our example above: the lowest permissible probability assignment to the more probable alternative A should be higher than the lowest such assignment to the less probable $A&C$, and similarly for the highest permissible probability assignments to these propositions. But their probability ranges should be expected to overlap.


Does this mean, then, that what is needed is just an appropriate weakening of the Range Rule? Should we say that for an item to be better than another it is sufficient if the range for the former item has upper and lower bounds that exceed the upper and lower bounds, respectively, for the latter item? This would mean accepting the following criterion:

The Weakened Range Rule: x is better than y if and only if (i) $x^{max} > y^{max}$ and (ii) $y^{min} > y^{min}$.

Probability comparisons could be dealt with in the same way.

Unlike Gert's Range Rule, this weakened criterion does not require the lower bound for the better item to be higher than the upper bound for the worse item. Unfortunately, such a weakening of the criterion of betterness would not preserve the intuition that it is rationally required to prefer the better item. For if the ranges for the better item x and the worse item y are allowed to overlap, then a relatively weak permissible preference for x might be weaker than a relatively strong permissible preference for y . To avoid the undesired conclusion that it is permissible to prefer the worse item to the better one, we would need to forbid *combining* a strong preference for the former with a weak preference for the latter. But the interval modeling lacks resources for forbidding or prescribing particular combinations of preference strengths for various items. There is nothing in the model to ensure that whatever preference one might have for one alternative, one is rationally required to prefer the other alternative *even more*.

As a matter of fact, there is also another, more direct objection to this weakened version of the interval modeling. There are possible betterness structures that cannot be represented by the interval modeling, even in its weakened version. Here is an example, with six items, x, x^+, y, y^+, z and u . The first four are related to each other as in the previous example, while z is better than both x and y , and u is worse than both x^+ and y^+ . In addition, z is neither better nor worse than x^+ and y^+ , while u is neither better nor worse than x and y . Diagrammatically, we can represent this structure as follows:



Downward paths in the diagram represent betterness relations. Now, it can be shown that, even with the Weakened Range Rule, there is no possible assignment of ranges to items that could represent this structure of value of value relations.²⁴

In his book on interval orders, Peter Fishburn (1985, p. 78) provides this example and several other instances of betterness structures that cannot be given an interval representation, due to their high “dimensionality”.²⁵ The notion of dimensionality is defined as follows. A betterness structure that contains some gaps (i.e., pairs of

24 An even weaker criterion would only require, for x to be better than y , that (i) $x^{max} \geq y^{max}$, (ii) $x^{min} \geq y^{min}$ and (iii) *at least one* of the bounds for x (the upper or the lower one) positively exceeds the corresponding bound for y . This criterion, however, is just as unfit to represent the betterness structure specified above as the Weakened Range Rule. Chang (2005) notes this, with a reference to my paper.

25 I owe this example and the Fishburn reference to Erik Carlson; cf. Carlson (2006).

items neither of which is better than the other) can be extended in various ways to linear betterness orderings, in which the gaps are filled in one way or another and all the items in the structure are linearly ordered by the betterness relation. Now, let a *base* for a betterness structure S be any set of its linear extensions such that the intersection of that set coincides with S . Different bases for S may contain different numbers of extensions. The *dimensionality* of S is defined as the number of extensions in a smallest base for S . It can be proved that interval representations using the Weakened Range Rule are possible for all structures with dimensionality up to 2, but not higher.²⁶ The dimensionality of the structure in our example equals 3 and thus exceeds the limit of the interval representation.

For a more specific illustration of that betterness structure, think of each item as being exhaustively characterised by different amounts of three good-making attributes, A , B and C . In item z , the three attributes are present in amounts a , b and c , respectively. Thus, z can be represented as a triple (a, b, c) . Let a^+ be a slightly increased amount of A and a^- a slightly decreased amount of that good-making attribute. Similarly for the other two attributes, B and C . Now, suppose the six items are characterised as follows:

$$x = (a^-, b, c), x^+ = (a^-, b, c^+), y = (a, b^-, c), y^+ = (a, b^-, c^+), \\ z = (a, b, c), \text{ and } u = (a^-, b^-, c^+)$$

Value comparisons between items that only differ with respect to one attribute are easy: the larger amount of a good-making attribute is always the better; or so, at least, let us assume. But comparisons that involve changes in several attributes are more difficult. For example, if there is no way to determine whether a small gain in attribute C compensates a small loss in attribute A , we might want to deny that x^+ and z are equally good or that either is better than the other. This should explain why the value relations that obtain between the six items in our example have the structure described above. As an aside, I should point out that essentially the same kind of example can be used to construct a *probability* structure that evades an interval representation.²⁷

26 For this result, see Fishburn (1985), ch. 5, theorem 9 (pp. 85f). The theorem itself was originally proved by Dushnik and Miller (1941). It should be noted that the dimensionality restriction does not apply to the original Range Rule. That rule is adequate for the representation of all “interval orders”, of any dimensionality (see above, fn. 23). For an example of an interval order with dimensionality higher than 2, see Carlson (2006, figure 4). But, on the other hand, as illustrated by the example of a trip to Australia or to South Africa, with or without a bonus, there are betterness structures of dimensionality as low as 2 that do not satisfy the characteristic condition on interval orders.

27 For the probability case, think of a , b and c as propositions about three unrelated subject matters, A , B and C , respectively. Let a^+ and a^- be propositions about A that are, respectively, slightly more and slightly less probable than proposition a . Make similar assumptions for b and c . Finally, let items be conjunctions

4. Intersection Modeling

If not intervals, then what? As we have seen, the interval modeling lacks resources to determine permissible *combinations* of preference strengths for different items. The remedy, therefore, is to think of permissible preferences in a holistic way. Instead of determining the range of permissible preference strengths separately for each item, the right solution is to consider the whole domain of items that are to be compared and to delimit the class of permissible preference *orderings* of that domain. In what follows, I will refer to this class as K . K can be assumed to be non-empty; i.e., there should be at least one permissible preference ordering of the items in the domain. I will allow, however, that the orderings in K need not be so well-behaved as to be representable by quantitative preference measures: it might not be meaningful to specify the relative strengths with which different items are preferred in a given ordering. In fact, it may not even be meaningful to assign numerical values to items that represent their positions in the ordering. A minimum condition for representing a preference ordering by an assignment of numbers to items is that the ordering is complete, i.e., that it contains no gaps. In a complete preference ordering, for every pair of items in the domain, either one of them is preferred to the other or both are equi-preferred. Since we need to make room for incomparabilities and thus have to allow for gaps in permissible preference orderings, completeness cannot be assumed. What we can assume, however, is that all the orderings in the ‘permissible’ class K are at least *partial* in the following sense: in every such permissible ordering, (i) preference is a strict partial order, i.e., an asymmetric and transitive relation, (ii) equi-preference (= indifference) is an equivalence relation, i.e., it is transitive, symmetric and reflexive, and (iii) for all items x and y , if x and y are equi-preferred, then any item preferred/dispreferred to one of them is respectively preferred/dispreferred to the other.²⁸

In terms of K , we can now immediately define what it means for one item to be better than another. Betterness is simply the *intersection* of all permissible preferences:

of propositions, with one conjunct for each subject matter. Thus, for example, $z = a \& b \& c$. While probability comparisons between items that only differ with respect to a proposition dealing with one subject matter are easy, comparisons across subject matters are problematic. This is why the structure of probability comparisons between items might only be partial, as in our example. The example shows that not even the weakened version of the interval modeling is fit to represent probability relations.

28 This is a rather cumbersome characterisation. We could simplify it if we instead used *weak preference* (i.e., preference-or-indifference) as our primitive notion, in terms of which both preference and indifference could then be defined in the standard way: preference as weak preference obtaining in just one direction, and indifference as weak preference in both directions. Then our three conditions on permissible preference orderings would be equivalent to the assumption that permissible weak preference is what is usually called a *preorder* (or a *quasi-order*), i.e., a transitive and reflexive relation.

(B) x is *better* than y if and only if x is preferred to y in every ordering in K .²⁹

To exemplify how this works, consider again the example with six items, x , x^+ , y , y^+ , z and u , from the preceding section. Suppose, for simplicity, that only the following three preference orderings with respect to these items are permissible. In each column, which represents one such ordering, the items are ordered from the most preferred at the top to the least preferred at the bottom. Equi-preferred items are placed on the same level. In this toy example, all permissible preference orderings are complete. Obviously, this need not be the case in general.

$P1$	$P2$	$P3$
x^+	y^+	x^+ y^+
z	z	u
x	y	z
y^+	x^+	x y
u	u	
y	x	

The intersection of $P1$, $P2$ and $P3$ gives us exactly the betterness structure of our example: x^+ and y^+ are better than x and y , respectively, and both are better than u , while z is better than both x and y . No other betterness relationships obtain between these items, just as we have stipulated.

Moving now to other value relations, it is easily seen how equality in value, full comparability, parity and incomparability are definable in this modeling. To begin with, equal goodness is defined as the intersection of all permissible equi-preferences:

(E) Two items are *equally good* if and only if they are equi-preferred in every ordering in K .

(FC) Two items are *fully comparable* if and only if in every ordering in K , one of them is preferred to the other or they are both equi-preferred.

²⁹ Essentially the same intersection approach as the one I delineate in this section can be used for other kinds of value to which the “fitting attitudes”-account is applicable. If, instead of betterness, we want to consider a kind of value that should be analysed not in terms of preference but rather in terms of some other pro-attitude such as, say, admiration, the obvious solution is to assume a class of all permissible “admiration orderings” of the items in the domain (allowing that some of these orderings might only be partial), and then define “ x is more admirable than y ” as the claim to the effect that x comes above y in every ordering in the class. In other words, an item is more admirable if and only if it ought to be more admired. Needless to say, other value kinds can be treated analogously.

- (P) x and y are *on a par* if and only if K contains two orderings such that x is preferred to y in one ordering and y is preferred to x in the other.
- (FP) Two items are *fully on a par* if and only if they are fully comparable, in addition to being on a par.

Finally,

- (I) x and y are *incomparable* if and only if every ordering in K contains a gap with regard to x and y , i.e., neither of these items is preferred to the other, nor are they equi-preferred.
- (WI) x and y are *weakly incomparable* if and only if some ordering in K contains a gap with regard to x and y , i.e., neither of these items is preferred to the other, nor are they equi-preferred.

This modeling is so straightforward that one might well wonder whether it adds anything to the original informal analysis of evaluative relations with which we started.³⁰ Does the claim that x is better than y if and only if x is preferred to y in every rationally permissible preference ordering add anything to the original analysis, according to which x is better than y if and only if it is rationally required to prefer x to y ? So far as I can see, it does not. This is just as well: it is always worrying if a formal modeling decides issues that have been left open by an informal analysis. Such decisions might introduce an element of arbitrariness and sometimes might lead us to pose spurious problems that are just constructs of an arbitrary formalisation. Having said this, though, I should point out that the intersection model is not quite innocuous, for two reasons. First, by assuming the class K of permissible orderings to be non-empty, we have excluded situations in which *nothing* is rationally permitted with respect to a given pair of items. Non-emptiness of K guarantees that it must be permissible to lack a preferential attitude as regards two items, if preferring one of them to the other as well as equi-preference are impermissible. Such an assumption might be questioned by philosophers who think it is possible to confront situations in which all options are forbidden. Second, we have imposed some formal restrictions on permissible preference orderings. This

30 The intersection modeling is based on an old idea, going back at least to Sen (1973, ch. 3; see also Atkinson, 1970). Sen has been arguing for this modeling since then, in various publications. But his “intersection approach”, as he calls it, does not provide an analysis of an evaluative relation such as betterness in terms of permissible preference orderings. Instead, it is a construction of the relation of *definite* betterness from a class of evaluative orderings that reflect different value commitments or different evaluative aspects of the items that are being compared. Also, on his approach, incompleteness only shows up in the resulting relation, but not in the underlying orderings. By contrast, our modeling allows preference orderings themselves to be gappy. This potential gappiness of the underlying preference orderings is essential if we want to distinguish parity from incomparability. (It should be added that Sen does discuss incomplete preferences in other places; for example, in Sen, 1997. But in those contexts he does not suggest applying intersection operation to sets of such incomplete preference orderings.)

has implications for value relations. Due to the features of the intersection operation, the modeling allows us to derive various formal requirements on evaluative relations from the corresponding requirements on preference orderings. Thus, it can now be shown (i) that betterness is transitive and asymmetric, (ii) that equal goodness is an equivalence relation, and (iii) that whatever is better than, worse than, on a par with or incomparable with one of equally good items must have exactly the same value relation to the other item.³¹ Thus, the modeling has its uses.³²

What is perhaps somewhat troublesome is that the model I propose makes formal features of evaluative relations less firmly established than one might wish them to be. To give an example, consider betterness. That this relation is transitive is, many would say, a conceptual truth. But in my modeling, this condition on equal goodness depends on the transitivity of preference. That the latter relation should be transitive in all permissible preference orderings may seem like a very reasonable requirement. But it may be doubted that it is a conceptual truth, as firmly established as the corresponding condition on betterness. Similar remarks apply to the comparison between the transitivity of equal goodness and the transitivity of equi-preference. I have to admit that this is a weakness in my proposal.

Let us move, however, to other matters. Chang (2002a, section 5.3.2) presents what she calls a “supervaluational interval model”, which exhibits some formal similarities to my intersection modeling. In that model, she postulates the existence of a class of legitimate *utility assignments* to the items in the domain. Each such utility assignment corresponds to a permissible evaluation of the items.³³ If one item is better than another, then all the assignments rank it higher. If all the assignments rank two items on the same level, the items are equally good. Parity obtains when two items are differently ranked in different assignments.³⁴ These

31 As an example, consider the proof that every z that is on a par with x must be on a par with every y that is equally good as x . If z is on a par with x , there exist some permissible preference orderings P and P' in K such that z is preferred to x in P and dispreferred to x in P' . If x and y are equally good, then in both P and P' these two items are equi-preferred. But then, since P and P' – as members of K – are partial orders, whatever is preferred (dispreferred) to x in these orderings must also be preferred (dispreferred) to y . Thus, z must be preferred to y in P and dispreferred to y in P' , which implies that z and y are on a par.

32 An *absolutely* innocuous modeling would instead simply specify, for each pair of items in the domain, which preferential attitudes are permissible with regard to these items, if any. Working with models of this kind would not allow us to draw any *a priori* implications concerning formal features of value relations.

33 “Each legitimate way of understanding the covering value might be represented by a standard utility function” (Chang, 2002a, p. 147).

34 While she calls her approach an “interval” model, the utility intervals she has in mind are not assigned to single items, but to pairs of items. For each pair x, y , we can define the interval $i(x, y)$ which has as its lower and upper bounds the minimum and the maximum of the set of differences $u(x) - u(y)$, for all legitimate utility assignments u . It is easy to see that x is better than y if the lower bound of $i(x, y)$ is positive, and it is worse than y if the upper bound of $i(x, y)$ is negative. x and y are equally good if $i(x, y)$ is a degenerate interval which has 0 as both its upper and lower bounds. As for parity, Chang defines it more broadly than I do (Chang, 2002a, p. 148). Instead of stipulating that x and y are on a par if and only if the

utility assignments are somewhat like the supervaluationist's sharpenings of a vague evaluative ordering, but Chang takes it that the existence of a plurality of legitimate assignments does not manifest any indeterminacy (vagueness) in the betterness ordering. Instead, it is a way to model such phenomena as parity. Thus, if one utility function ranks x higher than y , and another does not, then it is true on Chang's model that x is not better than y . Contrast this with the supervaluationist diagnosis, which would be that, in such a case, it is neither true nor false that x is better than y .

Technically, her approach is in some respects similar to my own. But there are important differences. (i) Chang takes utilities to be measured on an interval scale (where this scale is common to all the utility functions in the "legitimate" class). I do not make any such measurability assumptions. (ii) By working with utility assignments, she has no room for incomplete rankings. This leaves no scope for incomparabilities in her model. (iii) She interprets different utility functions as different legitimate *evaluative* orderings of the items, and not – as in my model – as different permissible preference orderings.³⁵ This makes her approach philosophically problematic. For if two items are on a par, Chang's modeling would allow that it is legitimate to evaluate one as better than another and also legitimate to have the opposite evaluation. However, how could such evaluations be legitimate if they are *incorrect*? After all, if the items are on a par, then neither of them is better than the other. Treating legitimate orderings as evaluative would have been appropriate if it were *indeterminacy* in evaluation that we wanted to model, along the supervaluationist lines, but it is not clear how it can be appropriate otherwise.³⁶ In my approach, I avoid this conceptual hurdle by replacing incompatible evaluations with opposing *preferences*.

lower bound of $i(x, y)$ is negative and the upper boundary is positive, she takes it that parity obtains as soon as the items are not equally good and neither of them is better than the other. In terms of intervals, this means that (i) $i(x, y)$ contains 0, just as in the case of equal goodness, but – in contradistinction to equal goodness – (ii) the upper and the lower boundaries of $i(x, y)$ do not coincide. Thus, in particular, she would say that x and y are on a par if x in every legitimate utility assignment is ranked at least as highly as y , and in some of them it is ranked higher. In my view, it is rather implausible to treat this asymmetric relationship as a case of parity. It would be much more natural to say that, in a case like this, x is at least as good as y , but not vice versa. For more on this issue, see the next section.

35 In this respect, as in the previous one, Chang's approach is strongly reminiscent of Sen's. See fn. 30 above.

36 This incoherence would be avoided if we instead interpreted these different legitimate orderings as reflections of different *aspects* of evaluation, with one ordering for each aspect. Thus, for example, when we compare different cars, one car might be ranked higher than another with respect to speed and lower with respect to comfort of travel. Clearly, such different aspect orderings can all be correct if the aspects they reflect are different. But then why assume – as Chang obviously does – that an item, x , can never be better than another item, y , if it is not the case that all legitimate evaluative orderings rank x higher than y ? This would disallow any possibility of *trade-offs* between various aspects of evaluation, which surely seems counterintuitive.

Table 1

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
>	+	+				+	+	+	+	+	+				
≈		+	+	+			+		+		+	+	+		
<				+	+	+	+	+	+				+	+	
/								+	+	+	+	+	+	+	+
	B		E		W	FP	FP	P	P						I

5. Taxonomy of Binary Value Relations

We now have all we need for the general taxonomy of binary value relations. The taxonomy identifies different types of such relations by specifying the kinds of permissible preference relationships that can obtain *between* two items. We disregard, however, their potential permissible preference relationships to *other* items in the domain. This is an important restriction on our taxonomy, on which I shall comment below.

In table 1, each column specifies one type of value relation that can obtain between two items; i.e., each column specifies one possible combination of rationally permissible kinds of preference relations between the items. There are four kinds of such relations to consider: preferring (>), indifference (≈), dispreferring (<) and a gap (/), where the latter stands for the absence of a preferential attitude. There is a plus sign in each column for every preference relation between the items that is rationally permissible in that evaluative type. There must be at least one plus sign in each column, since for any two items at least one kind of preference relation between these items must be permissible.

This means that, to specify a type, we pick a non-empty subset out of the set of four possible preferential relations that can obtain between two items. As there are 15 such non-empty subsets, the table has 15 columns. Thus, for example, if an item x is evaluatively related to an item y as in type 7, then all preferential relations between these two items are permissible, except for the gap. Or, to take another example, if the items are related as in type 1, the only preferential relation that is permissible is preference, i.e., preferring one item to the other is required. In other words, type 1 stands for the betterness relation.

The columns in the table stand for *atomic* types. Unions of atomic types, such as, for example, full parity (types 6 and 7), full comparability (all types from 1 to 7) or weak incomparability (all types from 8 to 15), are types in a broader sense of the word.

Being better than (**B**), worse than (**W**), equally good as (**E**), and fully on a par (**FP**) are four mutually exclusive forms of full comparability. However, these four

types do not exhaust all the logically possible ways in which two items might be fully comparable. The remaining two forms of full comparability, type 2 and type 4, lack standard labels. Still, if x and y are related in such a way that it is rationally required to either prefer x to y or to be indifferent between them, then it seems appropriate to say that x is at least as good as y . This seems appropriate even when, as in type 2, x is neither better than nor equally good as y : both preferring x to y and indifference between the two are permitted in this type. Similarly, x can be said to be at most as good as y if what is required is that one either disprefers x to y or is indifferent. This holds even when, as in type 4, x is neither worse than nor equally good as y . Thus, the relations *at least as good* and *at most as good* are unions of atomic types: the former covers types 1, 2 and 3, while the latter covers types 3, 4 and 5. This implies, then, that the standard definition of the relation of being “at least as good as” as the union of “better than” and “equally as good as” must be given up. Nor can we continue to hope to define the latter two notions in terms of “at least as good as”. The modeling shows that the interrelations between these three notions are more complicated than that.

Apart from seven (atomic) types of full comparability and one type for incomparability (I, type 15), we have seven mixed types, 8 to 14. In these seven, the items are weakly incomparable. Parity in the broad sense of that term, in which it does not require full comparability, corresponds to types 6–9.

Fifteen atomic types is a lot, but it is lucky that we only consider binary evaluative relations. Suppose we were interested in ternary relations, such as, say, the relation that obtains between three items whenever preferring the first to the third rationally requires preferring the second to the third.³⁷ There are no less than 28 different ways in which three items can be related to each other in a partial preference ordering (as opposed to just four in the case of two items).³⁸ Consequently, the number of atomic types of ternary evaluative relations equals the number of non-empty subsets in a set with 28 elements. Thus, there are $2^{28} - 1$ such types. This is a staggering number.

37 An example might be a comparison between three artists, x , y and z , with x and y being quite similar to each other but very different from z . (Ex: x = Claude Monet, y = Paul Cézanne, z = Piet Mondrian.) All three artists might be on a par, but it might be still rationally required that whoever prefers x to z should also prefer y to z .

38 These 28 ways are complete specifications, for each pair of items in a given triple, of how the items in this pair are preferentially related to each other. An example of such a specification would be: the first two items in the triple are equi-preferred, and both are preferred to the third item. Or, the first item is preferred to the second and to the third, while there is a preferential gap between the second and the third item. Due to the formal constraints on preference orderings, certain specifications are excluded. For example, since preference is transitive, it is excluded that the first item is preferred to the second, which is preferred to the third, which is equi-preferred with the first. As can be shown, this leaves us with 28 possibilities to consider.

But even with binary evaluative relations, our taxonomy is kept relatively small because it only considers what kinds of permissible preference relationships can obtain *between* two items. We have disregarded potential similarities and dissimilarities between the compared items in their permissible preferential relationships to *other* items in the domain. If this had also been taken into account, we would have had many more types of binary evaluative relations to consider. In a sense, then, we have only considered “internal” binary relations between items and disregarded “external” ones, which depend on the permissible relative positions of the compared items vis-à-vis other items. To see this, consider as an example the “external” binary value relation of *indirect comparability* that obtains between x and y whenever, in each permissible preference ordering, there exists a chain from x to y such that every item z in the chain is either preferred, dispreferred or equi-preferred to its successor. Full comparability entails indirect comparability, but the opposite does not hold. Another example of an “external” relation is *congruence*: x and y are congruent if and only if in each permissible preference ordering (i) neither of these items is preferred to the other, and (ii) for any item z distinct from x and y , z has exactly the same preferential relation to x as to y . If x and y are related to each other in this way, then they are interchangeable in every permissible preference ordering. It is easy to see that if two items are equally good, then they are congruent. But the opposite does not hold. In principle, at least, there could even exist congruent items that are incomparable with each other. Whether this is a real possibility is another matter.³⁹

This leads us to the next issue. Consider again our taxonomy. The 15 atomic types we have listed are all *logically* possible. But it might be that some of these

39 If we were interested only in internal binary value relations, the modeling in terms of a class of permissible preference orderings would in fact contain more information than is necessary. Sometimes, two different such classes, K and K' , might induce exactly the same specification of the atomic types of internal binary value relations that obtain between the different items in the domain. As a very simple example, consider an item domain that consists of just three items, x , y and z , and suppose that the value relations in that domain are fully specified as follows: the value relation between every two items in the domain is of type 6. That is, for every pair of items, it is permissible to prefer one to the other and vice versa, but it is impermissible to be indifferent or to lack a preferential attitude with respect to the items in question. As is easily seen, this specification is equally induced by two different classes of permissible preference orderings, $K = \{P_1, P_2, P_3\}$ and $K' = \{P_4, P_5, P_6\}$:

	K			K'		
	P_1	P_2	P_3	P_4	P_5	P_6
x	x	x	z	y	y	z
y	z	z	y	x	z	x
z	y	y	x	z	x	y

However, in order to specify external binary value relations between the items, or their internal relations with an arbitrary number of arguments, we need the full power of our intersection modeling.

types do not represent “real” possibilities. For example, can there exist two items, x and y , that are related to each other in the way specified in columns 6 or 8? It might seem that whenever two items are on a par, i.e., whenever it is permitted to prefer one to the other and permitted to have the opposing preference, then it should also be permitted to be indifferent between them. In other words, we might require for all x and y that K should contain a preference ordering in which x and y are equi-preferred, if K contains a preference ordering in which x is preferred to y and another ordering in which y is preferred to x . This requirement, which – so to speak – imposes a constraint of *convexity* on the class of permissible preference orderings, would exclude types 6 and 8.⁴⁰ One might perhaps also require that the absence of preference with regard to items that are on a par should always be permissible. This requirement on K would exclude types 7 and 8. Given both requirements, only type 9 would be left for parity. To take another example, which brings in external value relations, one might question whether there could exist incomparable items that are indirectly comparable. Such extra constraints on permissible preferences might allow us to narrow the space of possibilities. Note that these extra requirements differ importantly from such conditions as, say, transitivity of preference or symmetry of indifference. The latter imposes separate constraints on each ordering in class K of permissible preference orderings. The extra requirements instead are holistic conditions: they impose constraints on class K taken as a whole.

6. Choiceworthiness

Suppose that the items in the domain under consideration are possible options, which are at least in principle available for choice. Let a finite subset A of the domain consist of the options that are actually feasible in a given situation. Assuming that the *choiceworthiness* of an option has to do with its relative value, as compared with the value of its alternatives, we might ask which of the options in A are choiceworthy. If all options in A are fully comparable, it is natural to identify choiceworthiness with “optimality”: an option x in A is *optimal* in that set if and only if x is at least as good as every other option in A . Remember that this might hold even when there are some y in A such that x is neither better than y nor equally

40 However, in private communication, David Braddon-Mitchell has offered a plausible and amusing example of a comparison in which opposing preferences might be permissible, but indifference is impermissible. Consider analytic and continental philosophy. One may prefer the former to the latter or have the opposing preference, but it does seem irrational to be indifferent between the two (assuming, as we always do, that the issue concerns preferential attitudes that are rationally permissible for someone who is familiar with both items that are being compared).

as good as y ; cf. type 2 in our taxonomy. Optimality requires that the relation between x and each y in A exemplifies types 1, 2 or 3.

But what if some items in A are not fully comparable? Under such circumstances, no item in this set might be optimal. To deal with such cases, the standard solution has been to replace optimality with “maximality” as the criterion of choiceworthiness: an option x in A is *maximal* in that set if and only if no option y in A is better than x . This is a much weaker requirement than optimality. Type 5 is the only type of relation between x and other items in A that is excluded if x is maximal in A . It is easy to prove that every finite set must contain at least one maximal option.⁴¹

However, this is going too fast. We should note, first, that troubles with optimality as the candidate for a necessary and sufficient condition of choiceworthiness arise even in those cases in which the feasible set *does* contain some optimal alternatives. While all optimal options are arguably choiceworthy, the opposite need not hold. When the set contains some optimal option x , it might in addition contain a non-optimal option y such that it is rationally *permissible* to equi-prefer x and y . Since x is optimal but y is not, there is a permissible preference ordering in which y is dispreferred to some of the options, including x . However, if x is optimal and it is permissible to equi-prefer x and y , there is a permissible preference ordering in which y is among the highest ranked alternatives, together with x . But then, after all, why should y be unworthy of choice?⁴²

This observation might suggest the following explication of choiceworthiness:

An option is *choiceworthy* in an alternative set A if and only if it is ‘weakly optimal’, i.e., if there is a permissible preference ordering in which that option is preferred to or equi-preferred with any alternative in A .

A problem with this proposal is that it might again leave us without choiceworthy alternatives. As in the case of optimality, there is no guarantee that a given alternative set contains some weakly optimal options. It might turn out that every permissible preference ordering is incomplete regarding comparisons between its top alternatives. Furthermore, even if some permissible orderings happen to be complete, others might not be. This appears to make weak optimality an excessively strong criterion. For suppose that some option y is not weakly optimal, but it is still permissible to have a preference ordering in which no alternative is preferred to y . Would this not suffice to make y worthy of choice? Possibly, it would. It seems, then, that we might be well advised to replace weak optimality with a weaker criterion:

41 For a good discussion of the properties of maximality, see Sen (1997, section 5).

42 I owe this point to Joshua Gert (private communication).

An option is *choiceworthy* in an alternative set A if and only if it is “strongly maximal”, i.e., if there is a permissible preference ordering in which that option is not dispreferred to any alternative in A .

It is easy to see that every finite alternative set will contain at least one strongly maximal option. In this respect, strong maximality behaves like maximality. However, maximality is a logically weaker condition: while there may not exist any option that is preferred to x in every permissible preference ordering, in each such ordering there may be some options that are preferred to x . Under these conditions, x would be maximal, but not strongly maximal.⁴³ Whether this is more than just a logical possibility depends on what kinds of additional constraints we might want to impose on the class of permissible preference orderings. Some such additional restrictions were mentioned in the previous section.

A different conception of choiceworthiness is suggested by Chang (2002a, ch. 2). On her “comparativist” proposal, an option is choiceworthy if it is comparable with every option in the alternative set, without being worse than any of them. If comparability is interpreted weakly, as the contradictory of incomparability (rather than strongly, as full comparability), this would mean, in terms of our modeling, that an option x is choiceworthy in an alternative set A if for any option y in A , there is a permissible preference ordering in which x is preferred to or equi-preferred with y . As is easily seen, maximality is weaker than the comparativist notion of choiceworthiness, and weak optimality is stronger, while strong maximality is neither weaker nor stronger. It is evident that in the presence of incomparabilities, the alternative set might lack choiceworthy options in the comparativist sense. But this difficulty does not arise if we work with a domain of items that at most contains weakly incomparable options. A more serious problem with the comparativist conception is that an option might be choiceworthy on that view even though it is dispreferred to some alternatives in every permissible preference ordering.⁴⁴ This makes one wonder whether choosing such an option could be justified.

A thorough discussion of choiceworthiness would require a paper of its own. The main purpose of this paper was to show that the analysis of value comparisons in terms of normative assessments of preference makes it possible to provide a plausible modeling in which different types of evaluative relations, including parity and incomparability, can be clearly represented, classified, and distinguished from each other. If the reader finds the modeling attractive, the objective of this paper has been achieved.

43 As an example, consider the class of three preference orderings, P_1 , P_2 and P_3 , that we used to represent Fishburn’s example of a betterness structure with six items. In that example, item z is maximal but it is dispreferred to some items in each of the three preference orderings.

44 This possibility can again be illustrated with item z in Fishburn’s example, if P_1 , P_2 and P_3 are the only permissible preference orderings (see the preceding note).

Acknowledgements

I am indebted for helpful suggestions and criticisms to David Alm, David Braddon-Mitchell, Campbell Brown, John Broome, Johan Brännmark, Erik Carlson, Ruth Chang, Sven Danielsson, Dan Egonsson, Joshua Gert, Peter Godfrey-Smith, Lars Lindahl, Christian List, Jonas Olson, Johannes Persson, Björn Petersson, Martin Petersson, Christian Piller, Mozaffar Qizilbash, Joseph Raz, Andrew Reisner, Toni Rønnow-Rasmussen, Walter Sinnott-Armstrong, Michael Smith, Daniel Svensson, Folke Tersman, Ralph Wedgwood and the anonymous referee for *Theoria*. This paper is a substantially revised version of my “Modeling Parity and Incomparability” (2004), which appeared in W. Rabinowicz and T. Rønnow-Rasmussen (eds), *Patterns of Value – Essays on Formal Axiology and Value Analysis*, vol. 2, Lund Philosophy Reports, Lund, pp. 201–228. In August 2004, that paper was presented at philosophy seminars of the Research School of Social Sciences at ANU in Canberra and at the University of Melbourne. Later, it was discussed at philosophy seminars and conferences in Luleå, Lund, Uppsala, Montreal, Leiden, Lisbon, Oxford, London and York. I wish to thank the participants of all these events for very rewarding discussions. I am grateful to the Research School for providing me with a stimulating work environment, to Geoffrey Brennan for being an excellent host in Canberra, and to the Bank of Sweden’s Tercentenary Foundation and the Swedish Research Council for generous research grants.

References

- ATKINSON, A. B. (1970) “The Measurement of Inequality.” *Journal of Economic Theory*, 2: 244–263.
- BRENTANO, F. (1969 [1889]) *The Origin of Our Knowledge of Right and Wrong*, translated by R. M. Chisholm and E. H. Schneewind. London: Routledge & Kegan Paul.
- BROAD, C. D. (1930) *Five Types of Ethical Theory*. London: Routledge & Kegan Paul.
- CARLSON, E. (2006) “Incomparability and the Measurement of Value.” In K. McDaniel, J. R. Raibley, R. Feldman and M. J. Zimmerman (eds), *The Good, the Right, Life, and Death – Essays in Honor of Fred Feldman*, pp. 19–43. Aldershot: Ashgate.
- CARLSON, E. (2007) “Parity Defined in Terms of Betterness.” In T. Rønnow-Rasmussen, D. Egonsson and B. Petersson (eds), *Hommage à Wlodek – Philosophical Papers Dedicated to Wlodek Rabinowicz* [see <http://www.fil.lu.se/HommageaWlodek/site/abstra.htm>].
- CHANG, R. (1997) “Introduction.” In R. Chang (ed.), *Incommensurability, Incomparability and Practical Reason*, pp. 1–34. Cambridge, MA: Harvard University Press.
- CHANG, R. (2002a) *Making Comparisons Count*. London: Routledge.
- CHANG, R. (2002b) “The Possibility of Parity.” *Ethics*, 112: 659–688.
- CHANG, R. (2005) “Parity, Interval Value, and Choice.” *Ethics*, 115: 331–350.
- CHISHOLM, R. M. (1986) *Brentano and Intrinsic Value*. Cambridge: Cambridge University Press.

- DANIELSSON, S. (1998) "Numerical Representations of Value-Orderings: Some Basic Problems." In Ch. Fehige and U. Wessels (eds), *Preferences*, pp. 114–122. Berlin and New York: W. de Gruyter.
- DANIELSSON, S. (1983) "Hur man inte kan mäta välmåga" (How one cannot measure well-being). *Filosofisk tidskrift*.
- DUSHNIK, B. and MILLER, E. W. (1941) "Partially ordered sets." *American Journal of Mathematics*, 63: 600–610.
- EWING, A. C. (1947) *The Definition of Good*. London: Macmillan.
- FISHBURN, P. C. (1970) *Utility Theory for Decision Making*. New York: John Wiley & Sons.
- FISHBURN, P. C. (1985) *Interval Orders and Interval Graphs – A Study of Partially Ordered Sets*. New York: John Wiley & Sons.
- GERT, J. (2004a) *Brute Rationality*. Cambridge: Cambridge University Press.
- GERT, J. (2004b) "Value and Parity." *Ethics*, 114: 492–520.
- GIBBARD, A. (1990) *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Oxford: Clarendon.
- GIBBARD, A. (1998) "Preference and Preferability." In C. Fehige and U. Wessels (eds), *Preferences*, pp. 239–259. Berlin and New York: W. de Gruyter.
- LEWIS, D. (1979) "Attitudes De Dicto and De Se." *Philosophical Review*, 88: 513–543.
- MCDOWELL, J. (1985) "Values and Secondary Qualities." In T. Honderich (ed.), *Morality and Objectivity*, pp. 110–129. London and Boston: Routledge & Kegan Paul.
- RABINOWICZ, W. and RØNNOW-RASMUSSEN, T. (2004) "The Strike of the Demon: On Fitting Pro-Attitudes and Value." *Ethics*, 114: 391–423.
- SCANLON, T. M. (1998) *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- SEN, A. (1973) *On Economic Inequality*. Oxford: Clarendon.
- SEN, A. (1997) "Maximization and the Act of Choice." *Econometrica*, 65: 745–779.
- SLOTE, M. (1989) *Beyond Optimizing*. Cambridge, MA: Harvard University Press.
- SVENSSON, D. (2004) "The Softhearted but Hardheaded Challenge: Sentimentalism, Emotive Cognitivism and the Circularity Problem." In W. Rabinowicz and T. Rønnow-Rasmussen (eds), *Patterns of Value*, Vol. 2, pp. 261–290. Lund Philosophy Reports, Lund.
- WIGGINS, D. (1987) "A Sensible Subjectivism?" In D. Wiggins, *Needs, Values, Truth: Essays in the Philosophy of Value*, essay V. pp. 185–214. Oxford: Blackwell.