

MODELING PARITY AND INCOMPARABILITY*

Wlodek Rabinowicz

Department of Philosophy, Lund University

Wlodek.Rabinowicz@fil.lu.se

ABSTRACT: According to Ruth Chang, 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: The two items may be 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: The distinction between equality in value and parity becomes unproblematic on this approach. As I show, the approach in question, if appropriately extended, also allows for a straightforward distinction between parity and incomparability.

However, while Gert's basic proposal is attractive, the way he develops it is flawed. He takes it that rationally permissible preferences for each 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. I provide an alternative modeling, in terms of intersections of rationally permissible preference orderings, delineate a general taxonomy of binary value relations, and conclude with some remarks about the connection between value comparisons and various concepts of choiceworthiness.

1. *Introducing parity*

In "The Possibility of Parity", Ruth Chang 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, the two items may be on a *par*.¹ An example might be the evaluative relationship between two great artists, say, between Mozart and Michelangelo. They are comparable in their excellence, but we might want to deny that

* 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, Jonas Olson, Johannes Persson, Björn Petersson, Martin Petersson, Toni Rønnow-Rasmussen, Walter Sinnott-Armstrong, Daniel Svensson, Folke Tersman and Michael Smith. In August 2004, this paper was presented at philosophy seminars of the Research School of Social Sciences at ANU in Canberra and at the University of Melbourne. In November the same year, I presented it at philosophy seminars in Luleå and in Lund. I want to thank the participants of all these events for rewarding discussions. I am grateful to the Research School for providing me with a stimulating work environment, to Geoffrey Brennan for being such an excellent host in Canberra, and to The Bank of Sweden's Tercentenary Foundation for a generous research grant.

¹ *Ethics* 112 (2002): 659-88. See also her "Introduction" to a collection she edited on *Incommensurability, Incomparability and Practical Reason*, Cambridge: Harvard UP, 1997, pp. 1-34, and her monograph *Making Comparisons Count*, Routledge 2002.

the former is either better or worse than the latter, or that they are equally good. Instead, it seems appropriate to treat them as being on a par.

How is the possibility of this fourth kind of comparability established? Chang considers cases in which we confront two very different items, x and y , none of which is better than the other. In some respects, one is better, in other respects, it is the other way round, but none 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^+ , that 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. section 1 of her article.²) That x and y nevertheless may well be evaluatively comparable, 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 some such cases, we can also 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 leaving all the other relevant respects unchanged. We may call such an improvement ‘unidimensional’. Clearly, if z is worse than y in more than one respect, improvements in the sequence need to be made in several different respects as one moves from z to y . But in each step in the sequence there is a change in one respect only. Now, it would seem that a small unidimensional improvement should not affect comparability: It should not change an item in the sequence from being comparable with x to being incomparable. 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 element, y . Chang calls this “the Unidimensional Chaining Argument” (see section 2 of her article).

Chang admits, though, that the principle which underlies her chaining argument is only meant to apply in 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 the way one relevant respect is borne [by the items that are compared] against the way another relevant respect is borne [by these items].” (Ibid., 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 considerations possibilities of vagueness in judgments of comparability: It then becomes dangerously similar to a *sorites*.³ Chang admits that the argument for the possibility of parity as a fourth type of evaluative comparability is not complete until it is shown, as she endeavors 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

² In that section, Chang also presents an argument for the claim that very diverse items normally will never be equally good. (Cf. *ibid.*, pp. 671f.)

³ As is easily seen, her assumption that unidimensional improvements don’t make comparability disappear might well be questioned if we allow for vague comparability. One might then dismiss Chang’s chaining argument as just a version of *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 in my paper.

indifferent between them do not span the conceptual space of choice attitudes one can have toward alternatives.” (ibid., p. 666) In his article, which is a follow-up on Chang’s, Joshua Gert questions this claim and suggests that there is no need to revise the traditional trichotomy of preference relations in order to account for parity.⁴

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.

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. (ibid., p. 493)

However, as one continues to read his article, it becomes clear that this promissory note must be based on a misunderstanding. As it turns out, Gert accepts that the three traditional relations do not exhaust all the possible ways in which two comparable items can be related to each other. He explicitly allows for parity in his taxonomy of value relationships, and, like Chang, he takes parity to be a positive value relation that can only obtain in the absence of the three traditional value relations. 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.⁵ 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 antecedents in the philosophy of value. While he doesn’t mention 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 favor 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 favoring appropriate are its value-making properties; and different kinds of favoring – 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 favoring is preference: An item is better 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.”⁶ On this format of analysis, then, a claim of betterness consists in a normative assessment of preference.⁷

⁴ See J. Gert, “Value and Parity”, *Ethics* 114 (2004): 492-520.

⁵ Chang makes a similar observation in her reply to Gert, “Parity, Interval Value, and Choice”, *Ethics* 115 (2005): 331-50. Apart from questioning Gert’s apparent adherence to the traditional trichotomy of value relations, Chang also criticizes Gert’s interval modeling, as I do, using partly similar arguments. She does not provide any alternative modeling, though. Indeed, unlike myself (see below), Chang is not prepared to accept Gert’s basic idea that value comparisons can be analyzed in terms of normative assessments of preference.

⁶ *The Origin of Our Knowledge of Right and Wrong*, translated by R. M. Chisholm and E. H. Schneewind, London, Routledge and Kegan Paul 1969 (1889), p. 26. Cf. A. C. Ewing, *The Definition of Good*, London: Macmillan, 1947 (this is the *locus classicus* for this format of analysis); J. McDowell, “Values and Secondary

Gert formulates the normative component of the analysis in terms of the notion of rational permissibility: An item x is better than another item y if and only if it is not rationally permissible not to prefer x to y . Or, what amounts to the same, if and only if it is rationally required to prefer x to y .

Gert writes:

The central piece of conceptual apparatus that will be exploited in this article is the distinction between the choices that a person is disposed to make, and the choices that a person is disposed to regard as rationally permissible. The former may plausibly be taken to reveal a person's preferences, but a person's valuations are more plausibly taken to be revealed by the latter. That is, preference and valuation are not the same thing, and preference is most closely related to choice, while valuation is more closely related to the assessment of choice. (Gert, p. 493)

In the first part of his paper, Gert applies the notion of rational permissibility and its cognates to choices rather than to preferences. Thus, he interprets "better" as meaning something like "to be chosen, on pain of having made a mistake" (ibid., 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 judgments of betterness.⁸ Note that preference in this context cannot itself be understood as a judgment 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.⁹

Qualities" In Ted Honderich, ed., *Morality and Objectivity*, London & Boston: Routledge & Kegan Paul, 1985, pp. 110-129; D. Wiggins, "A Sensible Subjectivism?", essay V in D. Wiggins, *Needs, Values, Truth: Essays in the philosophy of value*, Oxford: Blackwell, 1987; A. Gibbard, *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Oxford: Clarendon, 1990; A. Gibbard, "Preference and Preferability," in *Preferences*, Perspectives in Analytical Philosophy, ed. Christoph Fehige and Ulla Wessels, Berlin: de Gruyter, 1998; T. M. Scanlon, *What We Owe to Each Other*, Cambridge, Mass.: Harvard University Press, 1998. Scanlon calls this approach "the buck-passing account of value", because it transfers the reason for favoring the object from the object's value to its value-making properties (cf. ibid., p. 97) For a discussion of some difficulties facing this format of analysis, see Wlodek Rabinowicz and Toni Rønnow-Rasmussen, "The Strike of the Demon: On Fitting Pro-Attitudes and Value", *Ethics* 114 (2004): 391 - 423. One such difficulty is that we want to exclude cases when preferring an item may be required not because of the features that make it better, but because this very preference would be valuable or beneficial. 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 analyzed in terms of the concept of betterness. (For the interpretation of the notion of preference, see below.)

⁷ Note however that the multiplicity of different kinds of favoring 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 admired, and so on. These other kinds of relations will not be considered in what follows, but much of what I am going to say could be applied to them as well, *mutatis mutandis*.

⁸ This is especially clear when he provides his modeling, in which rationality assessments apply to strengths of preferences. As for the nature of such assessments of preference, Gert refers the reader to ch. 7 of his book, *Brute Rationality*, Cambridge University Press 2004, where he interprets them in a cognitivist way. As he points out, though, the analysis he uses could just as well be given a non-cognitivist twist if "rational" (or "rationally permissible") is taken to be an expression of approval rather than an attribution of a property.

⁹ That's one way of avoiding the circle in the analysis. Another alternative would be to interpret preferences as emotive attitudes rather than dispositions to choose. On such an emotivist approach, preferring x to y would consist in experiencing x as more appealing than y , or something along these lines. The circle in the analysis is thereby avoided, *if* emotions can be given an account that does not treat them as forms of evaluation. For a thorough discussion of the circularity issue in connection with emotions, see Daniel Svensson, "The Softhearted but Hardheaded Challenge: Sentimentalism, Emotive Cognitivism and the Circularity Problem", in W.

Does the connection to preferences impose ontological restrictions on what entities can be related by 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 various properties of a 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.¹⁰ These restrictions on objects of preference would certainly be opposed by someone like Brentano.¹¹ But if some such restrictive view is correct, preferences regarding such entities as persons or material things would always at bottom 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 concrete things would be reducible to corresponding relations between states or properties. In what follows, however, I shall avoid taking stand on this issue of reduction.

An item *x* is *better* than another item *y* if preferring *x* to *y* is rationally required.¹² 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, required preferences need not track value relations. Indeed, under such conditions, it may 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. We cannot take it to mean complete familiarity with the properties of the items under consideration, at least for two reasons: (i) complete knowledge might be an ideal that is impossible to realize; (ii) on pain of circularity, we cannot assume that epistemic access extends to the evaluative features of the items.¹³ In what follows, these problems with the clarification of the assumption of epistemic access will be swept under the rug, however. Another problem that will be ignored is the question whether preferences and attitudes in general can be subject to rational requirements. I think they can, even if they are arguably outside our 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 preferring *y* to *x* is rationally required. Similarly, two items are *equally good* if what is rationally required in their case is

Rabinowicz and T. Rønnow-Rasmussen, *Patterns of Value*, vol. 2, Lund Philosophy Reports, Lund 2004: 261-290.

¹⁰ David Lewis, 'Attitudes De Dicto and De Se', *Philosophical Review* 88, (1979): 513 - 543.

¹¹ See R. M. Chisholm, *Brentano and Intrinsic Value*, CUP 1986, chs 2 and 3. Note, however, that for Brentano preferences are emotive attitudes. If one interprets preferences as dispositions to choose, it is much 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.

¹² 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 what is rationally required to prefer and what is merely rationally permissible to prefer. This is a crucial distinction in Gert's proposal. But since being preferable is an asymmetric relation, Gibbard's "rational" must be interpreted as 'rationally required' rather than as 'rationally permissible'. Clearly, the rational permissibility of a preference is logically compatible with the opposing preference being equally permissible.

¹³ Cf. C. D. Broad's cautious formulation of this format of analysis in *Five Types of Ethical Theory* (London: Routledge & Kegan Paul, 1930), p. 283: "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*." (my italics)

indifference, i.e. 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 . (Gert, p. 494)

Gert and Chang take it that comparisons between items always are made with respect to some covering value or consideration, which may differ depending on what we are interested in. Thus, we might ask of two persons who 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 the 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 satisficer, such as Michael Slote.¹⁴ 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 be incorrect. 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 on an example used by Slote. To an empty hotel comes an impoverished family 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 doesn't 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*. This requires, however, that the less luxurious room should not come up too low on the betterness scale for accommodations. We express this by saying that the ordinary room is good enough. It has to be good enough as an accommodation in order to be optimal as a shelter. The hotel manager's preferences can thus be accounted for if one keeps covering considerations apart. Many 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 may 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. I hope these sketchy remarks can suffice as a partial answer. The whole complex issue of satisficing cannot be adequately dealt with in this paper.

¹⁴ See his *Beyond Optimizing*, Cambridge: Harvard University Press, 1989. I am indebted to Jonas Olson for bringing this objection to my attention. Chang also raises this issue in her reply to Gert.

Let us go back to the analysis of evaluative relations. Before defining parity in a more precise way, I want to consider the notion of incomparability. Gert does not address this issue in his article, but the framework he works with allows for a straightforward extension that makes room for incomparabilities. As has been suggested above, preference is 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 them. Indifference is yet another type of choice disposition: One is indifferent insofar as one is equally prepared to make either choice. But then there might also exist pairs of items with respect to which a person *lacks* a choice disposition. Such a person would of course make some choice if necessary, but not because she is so disposed. Not all choices 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 after having inspected the options. (Buridan's ass is just a philosopher's fiction.) But in the absence of a choice disposition, we typically experience the situation as internally conflicted. We can see reasons on each side, but cannot (or simply will not) balance them off. Finally, we choose, since we have to, but the choice is made without the conflict of reasons being resolved.¹⁵

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.¹⁶ In principle, it seems, it is always possible to trace back my behavior to a configuration of some 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 I am confronted with.¹⁷ 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.¹⁸

Now, assuming that choice dispositions (in this qualified sense) can be absent, their absence can be subject to normative assessments. This allows us to accommodate 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, two items are *incomparable* if and only if it is not permissible to prefer one to the other or to be indifferent.¹⁹

¹⁵ An indirect evidence for the absence of choice dispositions may be in some cases provided by sequences of actions. 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 go on 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). However, such an erratic behavior could also be accounted for in other ways, for example by changes in preference or by preferential irrationality (cyclical preferences).

¹⁶ I am indebted to John Broome for pressing this point and emphasizing the need to clarify the notion of choice dispositions that is needed for my proposal.

¹⁷ A choice in this qualified sense is possible even in the case of indifference. That I am equally prepared to make either choice does not imply that the choice I proceed to make is unreasoned or non-deliberate.

¹⁸ See Wlodek Rabinowicz and Toni Rønnow-Rasmussen, "The Strike of the Demon", pp. 414-418, for a defence of the claim that the pro-attitudes to which one refers in the buck-passing account of value should be reason-based.

¹⁹ In conversation, Walter Sinnott-Armstrong has wondered why it is not enough for two items to be incomparable if the absence of a choice disposition with regard to them is rationally permissible, even if it is not required. Well, at least for linguistic reasons, such an account of incomparability would be rather awkward. To consider an analogy, note that 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.

Is it plausible to expect the existence of incomparabilities? To some extent, this depends on the item domain under consideration. If it 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.²⁰ In *Making Comparisons Count* (section 6.1), Chang distinguishes incomparability from what she calls ‘non-comparability’. To avoid terminological confusion, it is preferable, I think, to refer to the latter 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, we cannot compare Michelangelo to Mozart with regard to their excellence as sculptors. 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.²¹ However, when they do fall into different ontological categories, there may be no covering consideration at all with regard to which they could be compared.²² Preferring one to another then simply doesn’t make sense, however we might try to understand it.

In what follows, I do not distinguish essential incomparability as a separate type. 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. 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. However, in this paper, we do not need to take a definite stand on the actual existence of non-essential incomparabilities. For our purposes, it is enough to detail what’s logically possible.

What, then, about comparability? In a weak sense, two items are *comparable* if and only if they are not incomparable. But full comparability of items would mean that anyone who considers them 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 rationally required to either prefer x to y or prefer y to x or be indifferent between them. Which is not the same, of course, as to say that either preferring x to y is rationally required, or dispreferring, or indifference. What is required is the disjunction of these preferential attitudes.

²⁰ Unless we take a reductionist view and assume that a preference for a person at bottom 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.

²¹ 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.

²² I am indebted to Joshua Gert for this reminder.

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 stipulate that x and y are on a par if they are not equally good and none of them is better or worse than the other. For two such items, instead of being on a par, may not be comparable.

That the trichotomy of traditional 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 neither. 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 all those pairs regarding *required* preferential attitudes 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 pair 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.

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. Two items, x and y , shall thus be said to be *on a par*, in a broad sense, 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*.

As an aside, I should point out that Gert's own definition of parity is much narrower. 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 a too large category in his view.) They 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, p. 506). This would imply, in particular, that if it is required to prefer z to x then it must 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. As we have seen, in typical cases of parity, a small improvement x^+ of one item, x , need not be better than the other item, y .

3. Interval modeling

As an idealization, Gert assumes that the strength of possible preferences for different items is quantitatively measurable.²³ He then uses this idealization in his formal modeling of value relations. Since it may be rationally permissible to prefer a given item more or less strongly, we can assign to an item, 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 range. Now, Gert implicitly presupposes that any

²³ He does not specify the scale of measurement, but his discussion suggests that he must have in mind at least the so-called interval scale. Which 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.

combination of rationally permissible strengths of preference for different items is itself rationally permissible. For example, if items x and y are assigned partially overlapping ranges $[10, 40]$ and $[5, 30]$, respectively, then it is permissible to prefer x with, say, strength 20, and 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 is higher than the lower bound of the range for x .

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. *ibid.*, p. 505). Or, to put it shortly:

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 the better item.

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 them. But on the interval modeling, this is possible only if the ranges for x and y coincide and have *zero length*, i.e. reduce to 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 furthermore 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. For if (ii) didn't 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 recognizes that “[o]nly very rarely do we think of our particular personal preferences as the uniquely rational ones.” (*ibid.* 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.²⁴

²⁴ Chang in her reply to Gert (“Parity, Interval Value, and Choice”, pp. 340f) goes as far as to suggest that his 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 me 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

This feature of the modeling should give us a 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 preferring either or being indifferent between them is rationally impermissible. 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 can't 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 suggest that it is not as counterintuitive as it seems to claim that equality in value between distinct items is 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 none 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 z be the pain that is typically caused by getting a filling at the dentist's. While x is worse than x^+ , none of these experiences is either better or worse than y . The pain in both cases is too different to make a straightforward comparison possible. Another example could be like this: 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, none of these two alternatives might be better or worse than a trip to South Africa. As Gert shows, such cases 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 hundred-dollar bonus (or, in Gert's case, a somewhat less painful 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 none 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.²⁵ Here is the proof:

Since x^+ is better than x , while y^+ is better than y ,

(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) 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 .

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 .

²⁵ For a similar example, see Sven Danielsson, "Numerical Representations of Value-Orderings: Some Basic Problems", in Ch. Fehige and U. Wessels (eds.), *Preferences*, Berlin and New York: W. de Gruyter, pp. 114-122. I learned this example from Danielsson long time ago, in the 70-ies. He presented it in print already in 1982, in an article "Hur man inte kan mäta välmåga" (How one cannot measure well-being), in *Filosofisk tidskrift*.

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

Gert motivates his use of the interval modeling by reference to similar approaches to imprecise subjective probabilities (ibid., p. 510). It is easy to see, however, that the objection we have just presented just as well applies to probability comparisons. Statements such as ‘a proposition A is more probable than a proposition B ’ cannot be interpreted by assigning probability intervals to propositions and then using the analogue of the Range Rule for the representation of superiority in probability. Interval modeling is as 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 any 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 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 hundred dollar bonus. The latter is better, but is it reasonable to suppose that even the strongest rationally permissible preference for the former is weaker than the weakest rationally permissible preference for the latter? Surely, this can’t 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, $[11, 41]$, or something like it. It is thus to be expected that there will be a significant *overlap* between the two ranges. But the weakest permissible preference for the worse alternative will be weaker than the weakest permissible preference for the better alternative and there will be a similar relationship between the strongest permissible preferences for these two alternatives.

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 respective bounds 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) $x^{min} > y^{min}$.

Probability comparisons could be dealt with 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 lower item. Unfortunately, such a weakening of the criterion of betterness would not preserve the intuition that a better item is rationally required to be preferred. 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 that 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 that modeling, not 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 additions, 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 there is no possible assignment of ranges to the six items in this structure that makes an item better than another whenever the upper and the lower bounds for its range exceed, respectively, the upper and the lower bounds for the other item.²⁶

In his book on interval orders,²⁷ Peter Fishburn provides this example and several other instances of betterness structures that cannot be given an interval representation, due to their excessively high ‘dimensionality’. The notion of dimensionality is defined as follows. A betterness structure that contains some gaps (i.e. pairs of items none of which is better than the other) can be extended in various ways to linear betterness orderings, in which the gaps are filled in 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.²⁸ The dimensionality of the structure in our example equals 3 and thus exceeds the limit for the interval representation.

²⁶ 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. Ruth Chang notes this, with a kind reference to my paper, in her reply to Gert.

²⁷ See Peter C. Fishburn, *Interval Orders and Interval Graphs - A Study of Partially Ordered Sets*, New York: John Wiley & Sons, 1985, p. 78. I owe this example and the reference to Fishburn to Erik Carlson. See his excellent paper “Incomparability and the Measurement of Value”, forthcoming.

²⁸ For this result, see Fishburn, Ch. 5, Theorem 9 (pp. 85f). The theorem was originally proved by B. Dushnik and E. W. Miller in “Partially ordered sets”, *American Journal of Mathematics* 63 (1941): 600-610. 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 so-called ‘interval orders’, of any dimensionality (cf. P. C. Fishburn, *Utility Theory for*

For a more specific instantiation of that betterness structure, think of each item as being exhaustively characterized by varying amounts of three good-making attributes, A , B and C . In item z , these amounts are 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. Now, suppose the six items are characterized 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, we can assume. But comparisons that involve changes in several attributes are more difficult. For example, if we cannot say whether a small increase in attribute C compensates a small decrease in attribute A , we might want to deny x^+ is either better or worse than z or that they are equally good. This should explain why the value relations that obtain between the six items in our example could fit 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 cannot be given an interval representation.²⁹

4. Intersection modeling

If not intervals, then what? As we have seen, interval modeling lacks resources to specify 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: there should be at least one permissible preference ordering of items. 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. 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 an asymmetric and transitive relation, (ii) equi-preference is an equivalence relation, i.e., it is

Decision Making, John Wiley & Sons, New York, 1970, pp. 20-3; this result holds for all countable item domains). To be an interval order, a betterness relation, in addition to being transitive and asymmetric, must satisfy the following condition: For all x, x', y and y' , if x' is better than x and y' is better than y , then x' is better than y or y' is better than x . For an example of an interval order with dimensionality higher than 2, see Carlson’s paper, figure 4. But, on the other hand, as the Australia-South Africa example illustrates, there are betterness structures of dimensionality as low as 2 that do not satisfy the characteristic condition on interval orders.

²⁹ Just think of a, b and c as probabilistically independent 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 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. In this way it can be shown that not even the weakened version of the interval modeling is fit to represent probability relations.

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.³⁰

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: x is *better* than y if and only if x is preferred to y in every ordering in K . In other words, whatever preferences one might have with respect to y , one is rationally required to prefer x even more.³¹

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.

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

The intersection of these orderings 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 than 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: Two items are *equally good* if they are equi-preferred in every ordering in K . They are *fully comparable* if in every ordering in K , one of them is preferred to the other or they are equi-preferred. x and y are *on a par* if K contains two orderings such that x is preferred to y in one ordering and y is preferred to x in the other. They are *fully on a par* if, in addition to being on a par, they are fully comparable. Finally, x and y are *incomparable* if every ordering in K contains a gap with regard to x and y , i.e., none of these items is preferred to the other, nor are they equi-preferred.

³⁰ This is a rather cumbersome characterization. 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 conditions on permissible preference orderings would reduce to the assumption that permissible weak preference is what is usually called a *preorder*, i.e. a transitive and reflexive relation.

³¹ Essentially the same intersection approach as the one I delineate in this section can be used for other kinds of value to which the buck-passing account is applicable. If, instead of betterness, we want to consider a kind of value that should be analyzed not in terms of preference but rather in terms of some other proattitude 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 is ranked above y in every ordering in the class. In other words, an item is more admirable if and only if it is ought to be more admired; i.e., if it is unfitting not to admire it more. Needless to say, other value kinds can be treated analogously.

This modeling is so straightforward that one might well wonder whether it adds anything to the original informal analysis of evaluative relations that we have started with.³² If we say that x is better than y if and only if x is preferred to y in every rationally permissible preference ordering, does this add anything to the original analysis according to which x is better than y if and only if preferring x to y is rationally required? So far as I can see, it doesn't. Which is just as well: It is always dangerous if a formal modeling decides issues that have been left open by an informal analysis. Having said this, though, I should point out that the model is not quite innocuous. By assuming that the class K of permissible orderings is 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 both preferring one of them to the other and equi-preference are impermissible. Such an assumption might be questioned by philosophers who think it might be possible to confront situations in which all options are forbidden. Also, 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.³³ Thus, the modeling does some work.³⁴

What perhaps is more 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 equal goodness. 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 equi-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 equal goodness. Similar remarks apply to the comparison between the transitivity of betterness and the transitivity of preference. I have to admit that this is a weakness in my proposal.

Let's move, however, to other matters. In *Making Comparisons Count*, section 5.3.2, Ruth Chang presents what she calls a "supervaluational interval model", which exhibits some formal similarities with 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

³² The intersection modeling is based on an old idea, going back at least to Amartya Sen's *On economic Inequality*, Clarendon Press, Oxford 1973, ch. 3. (See also A. B. Atkinson, "The Measurement of Inequality", *Journal of Economic Theory* 2 (1970).) Sen has been arguing for this modeling since then. But his "intersection approach", as he calls it, does not so much consist in an analysis of an evaluative relation such as betterness in terms of permissible preference orderings but rather is a construction of the relation of *definite* betterness from a class of evaluative orderings that reflect different value commitments. 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 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 his "Maximization and the Act of Choice", *Econometrica* 65 (1997): 745-779. But in those contexts he does not suggest applying intersection operation to sets of such incomplete preference orderings.)

³³ As an example, consider the proof that every z that is on a par with x must be on a par with every item y that is equally good as x . If z is on a par with x , there exist some permissible preference orderings P and P' such that z is ranked above x in P and below x in P' . If x and y are equally good, then both in P and in P' these two items are ranked on the same level. But then, since both P and P' are partial orders, z is ranked above y in P and below y in P' , which implies that z and y are on a par.

³⁴ An *absolutely* innocuous model would instead simply specify for each pair of items in the domain which preferential attitudes are permissible with regard to these items, if any. Such a model would not be of any use in drawing conclusions about formal features of evaluative orderings.

assignment corresponds to a permissible evaluation of the items.³⁵ 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.³⁶ These 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 don't 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. 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 can such evaluations be legitimate if they are *incorrect*? After all, if the items are on a par, then none 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. In my approach, I avoid this conceptual hurdle by replacing incompatible evaluations with opposing *preferences*.

5. Taxonomy of binary value relations

We now have all we need for the taxonomy of all 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, potential permissible preference relationships of the items that are being compared to the *other* items in the domain. This is an important restriction, to which I shall return below.

In the table below, each column specifies one type of an evaluative relation that can obtain between two items. That is, each column specifies the rationally permissible kinds of preference relations between the items. There are four kinds of such relations to consider:

³⁵ "Each legitimate way of understanding the covering value might be represented by a standard utility function.", *ibid.*, p. 147.

³⁶ 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 broader than I do (cf. *ibid.*, p. 148). Instead of stipulating that x and y are on a par if and only if the 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 none 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 next section.

preferring ($>$), indifference (\approx), dispreferring ($<$), and a gap ($/$), i.e. an absence of a preferential attitude. There is a plus sign in each column for every kind of 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 fifteen such non-empty subsets, the table has fifteen columns. Thus, for example, if an item x is evaluatively related to y as in type 7, then all kinds of preferential relations between these two items are permissible, except for the preferential gap.

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

The columns in the table stand for ‘atomic’ types. Unions of atomic types, such as full parity (**FP**), are types in a broader sense of the word. Being better than, worse than, equally good as, and fully on a par are four mutually exclusive forms of full comparability. However, these four types do not exhaust all the logically possible ways in which two items can be fully comparable. In the table, there are seven possible types of full comparability (types 1 to 7), two of which together correspond to full parity (**FP**, types 6 and 7) and three that correspond to better (**B**), equally good (**E**), and worse (**W**), respectively (types 1, 3 and 5). The remaining two types of full comparability (2 and 4), do not have 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 . Just like full parity, the relations *at least as good* and *at most as good* cover unions of types: The former covers types 1, 2, and 3, while the latter covers types 3, 4, and 5. Note that this implies that one item can be at least as good as another without being either better than or equally good as that item. Thus, the standard assumption that betterness and equally goodness are both definable in terms of ‘at least as good as’ must be given up. Nor can we hope to define the latter notion as the disjunction of the former two. The modeling shows that the relations 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 comparable in the weak sense, without being fully comparable. Parity in the broad sense of that term, in which it does not require full comparability, corresponds to types 6 - 9.

Fifteen types is a lot, but it is lucky 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 twenty eight 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 twenty eight elements. Thus, there are $2^{28} - 1$ such types. This is a staggering number!

Let us return, though, to binary evaluative relations. For these, our taxonomy turns on what kinds of permissible preference relationships that can obtain *between* two items. We have, however, disregarded potential similarities and dissimilarities in their permissible preference relationships to *other* items in the domain. If these 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’ binary value relations that depend on the compared items’ positions in permissible preference orderings with respect to other items. To see this, consider as an example the ‘external’ binary value relation of *congruence* that obtains between x and y whenever in all permissible preference orderings (i) none of them is preferred to the other, and (ii) for any item z that is distinct from x and y , z has exactly the same preference 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 exist congruent items that are incomparable or only weakly comparable with each other. Whether this is a real possibility though, is another matter.³⁹

Which leads us to the next issue. Consider again our taxonomy. The fifteen atomic types we have listed are all *logically* possible. But it might be that some of these types do not represent ‘real’ possibilities. For example, can there exist two items, x and y , that are related

³⁷ An example might be a comparison between three artists, x , y , z , with x and y being quite similar to each other and both of them being very different from z . All three artists may well be on a par, but it may be still required that whoever prefers x to z should also prefer y to z .

³⁸ These twenty eight ways are complete specifications, for each pair of items in a given triple, 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 of them are preferred to the third item. Or, the first item is preferred to the second and for each of them there is a preferential gap when we compare it with 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 following could hold: 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 twenty eight possibilities to consider.

³⁹ If we were interested only in internal binary value relations, the modelling 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 exhaustive specification of the atomic types of internal binary value relations between all the items in the domain. As a very simple example, take an item domain consisting of just three items, x , y and z , and suppose that the value relations between them are fully specified as follows: For every pair of items in the domain, the value relation between them is of type 6. As easily seen, this specification is equally induced by two different classes of permissible preference orderings, $K = \{P_1, P_2, P_3\}$ and $K' = \{P_1, P_4, P_5\}$:

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

However, to determine 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 modelling.

to each other in the way specified in columns 6 or 8? It might seem that whenever two items are on a par, then it should be permitted that they are equi-preferred. If it is rationally permissible to prefer x to y and to have the opposite preference, then – it seems – it should also be rationally permissible to be indifferent between the two. This requirement would exclude types 6 and 8.⁴⁰ One might go further and require that the absence of preference with regard to x and y should be permissible whenever it is permissible to have opposing preferences with regard to these items. This would exclude type 7 in addition to types 6 and 8, which means that parity in the broad sense would reduce to type 9. To take another example which we have considered before, one might question whether there could exist incomparable items that are not essentially incomparable. Such extra requirements on evaluative relations between items, which are not based on logic but rather on considerations of rationality, might allow for some narrowings of the space of conceptual possibilities.

6. *Choiceworthiness*

Suppose that the items in the domain under consideration are possible options, which at least in principle are available for choice. Let a finite subset A of such items consist of the options that are feasible in a given situation. Assuming that the *choiceworthiness* of options has to do with their relative value, as compared with the value of their alternatives, we might ask which of the options in A that are choiceworthy. If all options in A are fully comparable, it might seem 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 option in A . Remember that this may hold even when there are some y in A such that x is neither better than y nor equally good; 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 by x being maximal. 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. For when the set contains some optimal option x , it might in addition contain a non-optimal option y such that equi-prefering x and y is rationally *permissible*. 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 equi-prefering x and y is permissible, 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

⁴⁰ However, in private communication, David Braddon-Mitchell has offered a plausible and amusing example of a case in which opposing preferences might be permissible, but indifference is impermissible. Think of a comparison between 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 we only consider what is rationally permissible given familiarity with the items that are being compared).

⁴¹ For a good discussion of the properties of maximality, see Amartya Sen, “Maximization and the Act of Choice”, section 5.

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

This observation might suggest the following explication of choiceworthiness:

An option is *choiceworthy* in an alternative set 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.

A problem with this proposal, however, is that it might leave us without choiceworthy alternatives. Just as it was the case with 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 may not be. This fact alone appears to make weak optimality an excessively strong criterion. For suppose that some option y is not weakly optimal, but it still is the case that it is permissible to have a preference ordering in which no alternative is preferred to y . Doesn't this suffice to make y worthy of choice? I would suggest that it does. It seems, then, that we should replace weak optimality with a weaker criterion:

An option is *choiceworthy* in an alternative set 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.

It is easy to see that every finite set will contain at least one strongly maximal option. In this respect, strong maximality behaves like maximality. However, maximality is a logically weaker condition than strong maximality: 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 .⁴³ Under these conditions, x would be maximal, but not strongly maximal. Whether this is more than just a logical possibility depends on what kinds of additional restrictions, beyond logics alone, we might want to impose on the class of permissible preference orderings. We mentioned some such additional restrictions in the previous section.

A different conception of choiceworthiness is suggested by Ruth Chang in ch. 2 of *Making Comparisons Count*. 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 as full comparability, this would mean, in terms of our modeling, that an option x in the alternative set A is choiceworthy if for any option y in A , there is a permissible preference ordering in which x is preferred to or equi-preferred with y . It is easy to see that in the presence of incomparabilities, the alternative set might lack choiceworthy options in this sense. Maximality is weaker than the comparativist notion of choiceworthiness, while strong maximality is neither weaker nor stronger. A difficulty with this 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 analysis of value comparisons in terms of normative assessments of preference makes it possible to distinguish parity from incomparability, and to provide a plausible modeling in which different types of evaluative relations can be clearly

⁴³ As an example, consider the class of three preference orderings P1- P3 that we used above 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 preference orderings P1 - P3.

⁴⁴ This possibility can again be illustrated with item z in Fishburn's example, if P1 - P3 are the only permissible preference orderings (see the preceding note).

represented and classified. If the reader finds these conclusions convincing, I consider my work being done.