

Economics and Philosophy

<http://journals.cambridge.org/EAP>


Additional services for ***Economics and Philosophy***:

Email alerts: [Click here](#)

Subscriptions: [Click here](#)

Commercial reprints: [Click here](#)

Terms of use : [Click here](#)



VALUE RELATIONS REVISITED

Wlodek Rabinowicz

Economics and Philosophy / Volume 28 / Special Issue 02 / July 2012, pp 133 - 164
DOI: 10.1017/S0266267112000144, Published online: 29 August 2012

Link to this article: http://journals.cambridge.org/abstract_S0266267112000144

How to cite this article:

Wlodek Rabinowicz (2012). VALUE RELATIONS REVISITED. Economics and Philosophy, 28, pp 133-164 doi:10.1017/S0266267112000144

Request Permissions : [Click here](#)

VALUE RELATIONS REVISITED

WLODEK RABINOWICZ

Lund University, Sweden

wlodek.rabinowicz@fil.lu.se

In Rabinowicz (2008), I considered how value relations can best be analysed in terms of fitting pro-attitudes. In the formal model of that paper, fitting pro-attitudes are represented by the class of permissible preference orderings on a domain of items that are being compared. As it turns out, this approach opens up for a multiplicity of different types of value relationships, along with the standard relations of ‘better’, ‘worse’, ‘equally as good as’ and ‘incomparable in value’. Unfortunately, the approach is vulnerable to a number of objections. I believe these objections can be avoided if one re-interprets the underlying notion of preference: instead of treating preference as a ‘dyadic’ attitude directed towards a pair of items, we can think of it as a difference of degree between ‘monadic’ attitudes of favouring. Each such monadic attitude has just one item as its object. Given this re-interpretation, permissible preferences can be modelled by the class of permissible assignments of degrees of favouring to items in the domain. From this construction, we can then recover the old modelling in terms of the class of permissible preference orderings, but the previous objections to that model no longer apply.

My incentive to study value relations was provided by Joshua Gert’s ‘Value and Parity’ (2004). The subject of that paper was an interpretation of Ruth Chang’s well-known and controversial suggestion that ‘parity’ is

This paper is a much expanded and modified version of Rabinowicz (2011), which was my contribution to the electronic *Festschrift* for Kevin Mulligan. It has been long in the making and the list of people to whom I am indebted for helpful comments and suggestions is correspondingly long. Among them I especially want to mention Gustaf Arrhenius, Geoff Brennan, John Broome, Ruth Chang, Tom Hurka, Hannes Leitgeb, Christian List, Kevin Mulligan, Andrew Reisner, Toni Rønnow-Rasmussen, Andrew Sepielli, Sarah Stroud and the two no-longer-anonymous referees for this journal, Joshua Gert and Larry Temkin.

a fourth form of value comparability, along with the classical relations of better, worse, and equally-as-good-as (cf. Chang 1997, 2002a, 2002b).

Judgements of parity seem appropriate in some of the cases in which the assessment of items depends on a number of factors, coming from different 'dimensions', with one item being superior to the other on some of the dimensions but inferior on other dimensions. If the relative weights of dimensions might be set in different ways, with equal justification, several weight distributions could be seen as equally admissible. If such different admissible sets of weights would result in opposing relative assessments of the compared items, it is reasonable to deny that the items in question are equally good or that one is better than the other. But we might still view them as being comparable in value. In cases like this, it might be suggested that the items under consideration are on a par.

While this might help to clarify the scope of the notion of parity, it doesn't yet answer the analytical question. What does it mean to say of two items that they are on a par? Gert's suggestion was that the idea of parity could be made sense of if value judgements were interpreted as normative appraisals of preference. In this proposal, he followed the tradition of the so-called *fitting-attitudes analysis* (FA-analysis), sometimes also referred to as *the buck-passing account*.

On the FA-format of analysis, to say that an item is valuable is, roughly, to make a claim that it ought to be favoured, where 'favouring' is a place-holder for a pro-attitude (or, more generally, a pro-response). Disvalues are correspondingly interpreted in terms of fitting con-attitudes (con-responses). What's characteristic of this approach is its treatment of deontic notions as being prior to the axiological ones: value is explicated in terms of the favouring stance that *ought* to be taken towards the object. The features of the object in virtue of which it ought to be favoured are its value-making properties. That it is fitting to have a pro-attitude, that there are normative reasons for it, or that the attitude in question is appropriate, required or called for, can be seen as different ways of expressing the deontic component in FA-analysis. The buck-passing aspect of the analysis is the idea that normative reasons for a pro-attitude towards an object aren't to be found in the object's value, but in its value-making properties. Which of course is as it must be given that value itself consists in the existence of such reasons. (Cf. Rabinowicz, forthcoming.)

When the fitting-attitudes analysis is applied to value relations, the relevant attitude is preference. This suggestion goes back to Brentano (1969 [1889]). Thus, in particular, that an item *x* is better than another item *y* is analysed as the claim that *x* ought to be preferred to *y*. Analogously, *x* and *y* are equally good if and only if they ought to be equi-preferred. Gert's main innovation was to point out that the deontic component might appear in this kind of analysis either in a stronger or a weaker form: either

as an 'ought' or as a 'may', or – using another pair of deontic notions – either as a requirement or as a permission. Bringing in these two levels of normativity makes it possible to account for a broader range of possible value relations, parity included. Gert's own definition of parity was, as it happened, excessively complex and too demanding. But he was, in my view, on the right track there. Simplifying his definition, I suggested in Rabinowicz (2008) that items x and y are on a par if and only if it is permissible to prefer x to y but also permissible to have the opposite preference, for y over x . We can connect this FA-style analysis of parity with the idea of multidimensional value comparisons if we suppose that each of the opposing preferences with respect to the items in question could be justified by an appeal to a different admissible set of dimension weights. Thus, each of the opposing preferences is a preference all-things-considered, but it is based on a particular relative weighting of the relevant dimensions – a weighting that is justifiable, but not uniquely so.

Parity should be distinguished from *incomparability*, which I suggested obtains between two items when it neither is permissible to prefer one of them to the other nor to equi-prefer them both. We have an instance of incomparability when what is being required is a preferential 'gap' with respect to the items in question.

While Gert's suggestion that we should make use of two levels of normativity was fruitful, the formal model he proposed to represent the different value relations had serious disadvantages. In Gert's '*interval model*', each item in the domain is assigned a range – an interval – of the strengths of preferences one is permitted to have with respect to the item in question. An item is then taken to be better than another item if the weakest permissible preference for the former is stronger than the strongest permissible preference for the latter. It is easy to show that this interval approach is not sufficiently general: there are very natural structures of betterness relationships between items that cannot be represented in this way. In technical terms, this is just another way of saying that the relation of betterness needn't be an interval order, which means that it needn't be representable by an assignment of numerical intervals to items. Another problem with the model is that it lacks resources to represent incomparabilities. Given the interval representation, for all pairs of items, either one item will be better than the other, or they will be on a par, or they will be equally as good (this last possibility, however, will be a very rare phenomenon, which itself is a further problem with the interval approach).

In Rabinowicz (2008), I therefore proposed an alternative formal modelling, which instead of assigning intervals postulates the class of permissible preference orderings of the item domain. On this approach, an item x ought to be preferred to an item y , and thus is better than y , iff it is ranked above y in every permissible preference ordering. The

relation of betterness is in other words the intersection of all permissible preference orderings. In this *intersection model*, x and y are on a par, if some permissible preference orderings place x above y , while other permissible orderings place y above x . If permissible preference orderings are allowed to be incomplete ('gappy'), there is room in such modelling for value incomparabilities. As I show in that paper, the logical taxonomy of all possible binary value relations that the intersection model gives rise to is quite rich: there are 15 atomary types of such relations. Four of them are better-than, worse-than, equally-as-good-as and incomparable-with, while parity turns out to be a collection of four other atomary types. The remaining seven types are new additions. Thus, the logical space of value relations is even more multifaceted than Chang has envisaged.

Unfortunately, the account I have proposed leads to a number of problems. All of them have to do with the underlying notion of preference, which, following Gert, I interpreted more or less on the standard choice-dispositional lines, with some qualifications. As it turns out, this connection to choice spells trouble for value analysis. The problems I have encountered can be avoided if the relevant notion of preference is re-interpreted. Instead of treating preference as being essentially choice-related – as a 'dyadic' attitude directed towards two items at a time – I now want to think of it as a comparative relation between 'monadic' attitudes of favouring. Each such monadic attitude has just one item as its object and it exhibits a degree of strength. Preference for one item over another consists in it being favoured to a greater degree. Below, I will therefore present an interpretation of preference orderings in terms of assignments of *degrees of favouring* to items in the domain. On a given assignment, one item may be favoured to a greater degree than another, or to the same degree, or the degrees in question might be incommensurable. (It is not assumed that degrees are representable by numbers.) Each assignment of degrees determines a preference ordering on the domain of items. The class of permissible assignments of degrees therefore determines the corresponding class of permissible preference orderings, which means that we have what we need to analyse value relations on the lines of the intersection model. Formally speaking, that analysis is the same as before, but the underlying notion of preference is interpreted in a different way.

Here, then, is the plan of the paper. In section 1, I briefly present Gert's two-level approach to FA-analysis and suggest some emendations and extensions. Section 2 describes and criticizes his interval model. The intersection model and the taxonomy of value relations are presented in section 3, while section 4 lists different objections to this proposal. Section 5 presents an interpretation of preferences in terms of degrees of favouring, which allows a defence of the intersection model against the listed objections. Section 6 responds to some worries concerning this new approach and concludes.

1. TWO LEVELS OF NORMATIVITY IN FA-ANALYSIS

On FA-analysis, value is analysed in terms of two components: a deontic component and an attitudinal one. Roughly, for an object to be valuable is for it to be such that we ought to favour it. Here, favouring is a stand-in for a pro-attitude, or – more generally – for a pro-response toward an object. To avoid the problem of the so-called *Wrong Kind of Reasons (WKR-problem)*, it has been argued that the supervenience base for the relevant ought needs to be restricted to the features of the object and exclude the features of the favouring itself (such as, say, that this attitude would have such-and-such consequences or that it would have such-and-such intrinsic features). Even if the features of favouring as such can sometimes make this attitude permissible or required, they do not thereby make its object valuable. Thus, to the extent they provide (normative) reasons for favouring, these reasons, however good they might be, are of the wrong kind from the point of view of the FA-analysis. But, and here comes the catch, the distinction between the features of the object and the features of favouring turns out to be difficult to draw in a principled way. This gives rise to the WKR-problem.¹ Here, however, we are going to ignore this difficulty.

When FA-analysis is applied to value relations, the attitudinal component is usually taken to consist in preference. Thus, Gert (2004) works with something like the following analyses of betterness and equal goodness:

Betterness: An item x is *better* than an item y iff it is required to prefer x to y .

Equal Goodness: x is *equally as good as* y iff it is required to equi-prefer x and y , i.e. to be indifferent between them.

Three remarks are in order at this point: (i) The notion of preference used by Gert is strongly connected to choice: He takes preferences to be dispositions to choose. Preference for x over y is a disposition to choose x rather than y , while equi-preference (indifference) consists in

¹ See Rabinowicz and Rønnow-Rasmussen (2004), for discussion. One might, however, try to argue that the WKR-problem is spurious, because the features of favouring an object, x , in contrast to the features of x itself, cannot ever be reasons for favouring x . They can only be reasons to favour the favouring of x . On this suggestion, then, we avoid the WKR-problem if we clearly distinguish between reasons for first-order pro-responses and reasons for pro-responses of the second order, which are directed to the first-order responses. For variants of this view see Skorupski (2010, esp. section 4.4), and Parfit (2011, vol. 1, Appendix A). Nevertheless, even if one were to accept this solution (which, if plausible at all, only applies to purely attitudinal forms of favouring, but not to behavioural pro-responses, such as promoting, protecting, emulating, etc.), one would still need to be able to draw the distinction between the features of the object and the features of favouring, as it is only the former that are supposed to provide reasons for favouring, if the view in question is correct. So, in this sense, the difficulty still remains.

being equally disposed to choose either. (ii) Either we take requiredness to be 'objective', i.e. independent of the agent's information, or we interpret it on subjective, information-dependent lines. In the latter case, preferential requirements must be understood as directed to agents who are epistemically familiar with the items that are being compared. (iii) Gert's notion of requiredness (oughtness) is qualified: he is interested in what is being *rationally* required. This qualification and the analysis of the concept of rationality need not concern us here; in what follows, I shall treat the notion of requiredness as primitive and abstain from discussing how it should be interpreted. This doesn't mean, of course, that the issue is unimportant. Thus, for example, interpreting the relevant requirement as a *moral* one would aggravate the WKR-problem: it might well be morally required to favour an object despite the fact that the object lacks value. Thus, it might be morally required to admire people that aren't admirable, or to love our neighbours even if they are not loveable. In fact, a moralized interpretation of requirements on attitudes is probably inappropriate in the context of FA-analysis.

Gert's main contribution to the FA-approach is his appeal to the distinction between two versions of the deontic component: the stronger and the weaker one. There are two levels of normativity: the stronger level of *requirement* and the weaker level of *permission*. In terms of the vocabulary of 'oughts', we express this difference by distinguishing between 'ought' and 'may'. The difference between normativity levels can, according to Gert, be put to use in the analysis of value relations. As he points out,

... 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 [...], could make different, but equally rational choices ... (Gert 2004: 494).

We might define Chang's notion of parity as being applicable precisely in those cases in which it is permissible to have each of the opposing preferences with respect to a pair of items. Which of course doesn't mean that it is permissible to have them jointly; in fact, entertaining them jointly would be outright impossible since what we here are talking about are preferences understood as dispositions to choose. One cannot be disposed to choose *y* rather than *x* and at the same time be disposed to choose *x* rather than *y*, in the same choice circumstances.²

² But even on a less behavioral interpretation of preference, I would argue that simultaneously entertaining opposing preferences is at least a form of incoherence, since what we here are talking about are preferences all-things-considered. It is at least incoherent, if not outright impossible, to prefer, all things considered, *x* to *y* and at the same

Parity: x and y are *on a par* iff it is permissible to prefer x to y and also permissible to prefer y to x .

Gert's own definition of parity is more demanding. In order to be on a par, x and y must according to him additionally satisfy the condition that for any third item z , 'the rational status' (i.e. the normative standing) of various possible preference attitudes towards x and z must be the same as that of the corresponding attitudes towards y and z (cf. Gert 2004: 506). This would in particular imply that if it is required to prefer z to x , then it must also be required to prefer z to y . Consequently, any item better than x would have to be better than y , and vice versa. This is obviously an excessively strong demand: in typical cases of parity obtaining between two items, x and y , a small improvement x^+ of x need not be better than y . Thus, to take an example, a trip to Australia (x) is intuitively on a par with a trip to South America (y), and the same still applies to a small improvement of the former alternative, say, to a trip to Australia with a discount of \$100 (x^+): That small improvement of one alternative is not better than the other alternative. In fact, for Chang, the possibility of such situations is quite central for the cases of parity. Typically, small improvements or small worsenings of one item in a pair do not make parity disappear.

I am therefore going to assume a relatively undemanding definition of parity: the one I presented above. In addition, I find it useful to extend Gert's framework in one important respect: along with different possible preferential attitudes concerning two items (preference for one item, preference for the other item, and indifference), I also want to allow situations in which a preferential attitude is absent, i.e. in which the agent neither prefers one item to the other nor is indifferent between them. Such situations seem to be important to take account of if we want to analyse the notion of incomparability in value. More precisely, I'd like to suggest that two items are incomparable if it is normatively inappropriate to prefer one of them to the other or to be indifferent between them. What's required in such case is a preferential 'gap' with respect to the items under consideration – an absence of a preferential attitude.

Incomparability: x and y are *incomparable* iff it is not permissible to prefer one of these items to the other or to be indifferent between them.³

time to prefer, all things considered, y to x . (I am indebted to Larry Temkin for pressing me on this point.)

³ Incomparability interpreted in this way should not be confused with the relation of *incommensurability* that obtains between two items when it is not required to prefer either item to the other nor required to equi-prefer them. The latter relation, which can be expressed as 'neither better than nor worse than, nor equally as good as', covers both parity and incomparability as its special cases.

This proposal invites a natural objection. In some cases in which the items intuitively are incomparable, the agent might still be required to choose between them. But then, if preferences are interpreted as choice dispositions, it would seem that the choice made reveals the preferential attitude. After all, doesn't it reveal what the agent is disposed to choose? To deal with this problem, but also because of other considerations, I suggested in Rabinowicz (2008) that preferences should be seen as choice dispositions in a narrower, qualified sense – as dispositions to make choices based on *balancing of reasons*. If the notion of preference used in the FA-style analysis of value relations is to be interpreted in choice-dispositional terms, then it is arguable that the relevant choice dispositions should be qualified in this way. It makes sense to maintain that pro-attitudes that figure in FA-analyses of value should be reason-based.⁴ This would mean that the agent who chooses without having resolved the conflict of reasons, either because she finds this conflict impossible to resolve or because she simply abstains from resolving it, does not thereby reveal a preference. In fact, under such circumstances she might have no preference at all, in this qualified sense. In other words, there is room for preferential gaps, if preferences are interpreted as dispositions to make reason-based choices.

2. INTERVAL MODEL

I now move to Gert's formal model of value relations. He assigns to each item x in the domain a closed interval of numerical values that represent all the permissible strengths of preference with respect to the item in question. We can let \min_x and \max_x be the lower and the upper bounds of that interval. The nature of the underlying scale for the measurement of strength is not clarified in his paper, but the use he makes of the scale is quite limited. In fact, for the analysis to follow, we only need to assume that the scale is invariant up to monotonic transformations.

The interval representation is then used to define the betterness relation:

x is better than y iff $\min_x > \max_y$.⁵

⁴ For the argument, see Rabinowicz and Rønnow-Rasmussen (2004: 414–418).

⁵ As Gert (2004: 505) puts it, 'One item is better than another in a certain respect if the lower bound of the range of the strengths of its relevant rationally permissible preferences is higher than the upper bound of the other's [...]. More roughly put, one item is better than another in a certain respect if and only if their relevant ranges are disjoint.' Here, I read the second occurrence of 'one item is better than another' as 'one of the items is better than another'. Gert's analysis concerns betterness 'in a certain respect', since he relativizes the notion of betterness to an underlying value that is of relevance in a given context of evaluation (such as artistic excellence, philosophical talent, moral goodness, or what have you).

The motivation behind this 'Range Rule', as Gert calls it, appears to be as follows: we already know that x is better than y iff it is required to prefer x to y . But then, if different preference strengths might be permissible with respect to each of the items, it would seem that the necessary condition of betterness is that the weakest permissible preference with respect to x must be stronger than the strongest permissible preference with respect to y . Or, what amounts to the same, $\min_x > \max_y$. So much for the (purported) necessity of this condition.⁶ Its sufficiency should be obvious. If $\min_x > \max_y$, then having a stronger preference for x than for y is required.

Parity, on the definition of this concept that I have proposed, gets interpreted in the interval model as follows:

x is on a par with y iff $x \neq y$, $\max_x > \min_y$ and $\max_y > \min_x$.

Thus, it is permitted to prefer x more than y and it is permitted to prefer y more than x .⁷

The problems with the interval model are easy to discern.⁸ To begin with, equal goodness between distinct items x and y reduces in the interval model to the case in which the intervals for x and y both shrink to a single point – the same point for x and for y . For it is only in such a degenerate case that the permissible preference for each of them must be the same, i.e. it is only then that it is required that items be equi-preferred. This implication of the interval model is clearly unwelcome, partly because it makes equal goodness a very rare phenomenon and partly because it has weird implications for the interactions between equal goodness and other value relations.⁹

⁶ If open intervals were allowed as assignments to items, Gert would probably be willing to modify his range rule, in the following way: x is better than y iff every numerical value within the interval for x is greater than any numerical value within the interval for y . Or, to put it differently, iff every permissible preference with respect to x is stronger than any permissible preference with respect to y . For open or partially open intervals, this could be the case even if $\min_x = \max_y$, as long as at least one of these bounds lies outside the interval it delimits.

⁷ Gert, who assumes a much more demanding account of parity, takes x and y to be on a par iff x and y are distinct items that are assigned exactly the same intervals. Such an interpretation is obviously untenable: it would make the relation of parity transitive (*ex alio*, i.e. it would imply that x and z must be on a par if $x \neq z$ and there is some y such that x is on a par with y and y is on a par with z), contrary to our intuitions concerning this notion.

⁸ Here, I follow Rabinowicz (2008), but see also Chang (2005) (which partly draws on an early draft of Rabinowicz 2008).

⁹ I am indebted to Christian List for this point. For example, the implications for interactions between equal goodness and parity are weird: The interval model cannot represent situations in which two items, x and y , are on a par, but each of them has a copy which is just as good; say, x is equally as good as its copy x' , while y is equally as good as its copy y' (with $x \neq x'$ and $y \neq y'$). Proof: If x and x' are equally good, the interval for x is just a

Another problem is that the model leaves no room for incomparability understood as the required preferential gap. For any two items x and y , either their ranges are disjoint, in which case it is required to prefer one of these items to the other, or they overlap, in which case it is permitted to be indifferent between x and y .¹⁰

Finally and most importantly, certain structures of betterness relationships aren't representable in the interval model. To see this, consider a case with four items, x , x^+ , y and y^+ , in which the only betterness relations that obtain are those between x^+ and x on the one hand and y^+ and y on the other. Intuitively, we can think of x^+ and y^+ as small improvements of x and y , respectively, with x and y being on a par. While x^+ is better than x , it doesn't improve on x so much as to become better than y , and similarly for y^+ : the latter is better than y , but it doesn't improve on y so much as to become better than x . We can illustrate this situation with our example of two trips, to Australia (x) and to South Africa (y), with x^+ being the trip to Australia with a small discount and y^+ the similarly slightly discounted trip to South Africa. It is easy to show that one cannot assign intervals to the four items in question in such a way that this intuitive structure of betterness relationships is preserved. If we set up the intervals so as to make x^+ better than x and y^+ better than y , then either x^+ will come out as better than y or y^+ will come out better than x . Formally speaking, for any interval assignment on which $\min_{x^+} > \max_x$ and $\min_{y^+} > \max_y$, we have either (i) $\max_x > \max_{y^+}$, in which case $\min_{x^+} > \max_y$, or (ii) $\max_y \geq \max_{x^+}$, in which case $\min_{y^+} > \max_x$.

Why do we get this untoward result? Consider again the comparison between a trip to Australia and the same trip with a small discount. The latter is better, but is it reasonable to suppose – as Range Rule would have it – that every permissible preference for this alternative is stronger than any rationally permissible preference for the former alternative? Surely, this cannot be right. For every permissible preference for the better alternative, a slightly weaker preference for the worse alternative should be permissible. It is thus to be expected that there will be a significant

single point: $\max_x = \min_x$. Likewise, if y and y' are equally good, the same must apply to the interval for y : $\max_y = \min_y$. But then, if we compare the intervals for x and y , either $\max_x = \min_x > \max_y = \min_y$, in which case x is better than y , or $\max_y = \min_y > \max_x = \min_x$, in which case y is better than x , or the single point intervals for x and y are identical, i.e. $\max_y = \min_y = \max_x = \min_x$, in which case x and y are equally good. Which means that parity between x and y is out of the question. Obviously, this implication is undesirable, to say the least. More generally, in an interval model in which every item has an equally good copy, all intervals are degenerate: they all reduce to single points.

¹⁰ To make room for incomparabilities in this model, we would need to allow items for which there don't exist (non-empty) ranges of permissible preference strengths. Such items, however, would then be incomparable with all other items in the domain, which is not what we want. If two items are incomparable, it should normally still be possible for each of them to be comparable with other items.

overlap between the intervals for both alternatives, which means that Range Rule cannot be correct: the condition it specifies is sufficient, but *not* necessary for betterness. What *is* necessary is instead some sort of required *connection* between preferences for the two alternatives. We would need to require that whatever preference one has for one alternative one should prefer the other alternative – the one which is worse – somewhat less. But the interval modelling lacks resources for prescribing or forbidding *combinations* of preference strengths for various items. This means that we need a model that is able to represent such ‘holistic’ conditions on permissible preferences with respect to different items.¹¹

A word of caution: when I speak of holistic conditions on permissible preferences, this has obviously nothing to do with the so-called value holism, i.e. the idea that the value of a whole need not be the sum, or any other function for that matter, of the values of its parts. My usage of the term is instead related to its use in epistemology: just as, according to the so-called ‘confirmation holism’, the proper object of empirical confirmation is a theory as a whole rather than a single hypothesis (this is of course the famous Duhem–Quine thesis), so one needs to consider a preference ordering as a whole in order to judge its permissibility, rather than focus on preferences for each item separately.

3. INTERSECTION MODEL AND TAXONOMY OF VALUE RELATIONS

In this model, which implements the idea of holistic conditions on permissible preferences, we work with the class **K** of permissible preference orderings on the domain of items (cf. Rabinowicz 2008). It

¹¹ Gert himself recognizes that his model should be enriched by conditions on combinations of preferences (cf. Gert 2004: 504; I am indebted to him for drawing my attention to this passage), but not for the reasons mentioned above. His worry instead is that an agent’s preferences for a given item might vary in strength within the permissible interval when the item in question is being compared with different items, thereby giving rise to irrational preference structures. Gert doesn’t notice, however, that imposing conditions on combinations of preferences necessitates giving up his Range Rule. Thus, here’s what he writes about an example he uses in his discussion of preference combinations: ‘the fact that the ranges ... for [avoiding] the dental pain, the day of itching, and the month of itching are, respectively, [50, 140], [20, 60], and [130, 250] should not be taken to mean that the very same agent might rationally have preferences of strengths 50 and 60 when considering the first two of these evils but have preferences of strengths 140 and 130 when considering the first and third. [...] [S]omeone who prefers the trip to the dentist to the day of itching ... would not be rational to prefer the month of itching to the trip to the dentist. These kinds of restrictions [on combinations of permissible preferences] do not argue in any way against the proposal being made here, however, since the strength of the agent’s preference for an item is a very different thing from his valuation of that item.’ (Gert 2004). This is true, of course. Preference and valuation are not the same thing, but – as I have argued – a valuation of items imposes restrictions on permissible preference patterns.

is convenient to think of a preference ordering as a relation of weak preference (i.e. preference-or-indifference). In terms of this relation, both (strict) preference and equi-preference (indifference) are definable in the standard way: x is preferred to y iff x is weakly preferred to y , but not vice versa; x and y are equi-preferred iff each is weakly preferred to the other. Every *permissible* preference ordering, i.e. every ordering in class \mathbf{K} , is assumed to be a quasi-order, i.e. transitive and reflexive. It is not assumed that the orderings in \mathbf{K} must be complete. They might contain preferential gaps: for some items x and y , it might be that, in a given ordering, none of them is weakly preferred to the other.

In terms of \mathbf{K} , it is then straightforward to define betterness, equal goodness, parity and incomparability. In particular, betterness is the intersection of preferences obtaining in the permissible orderings (and thus corresponds to required preference), equal goodness is the intersection of indifferences obtaining in all such orderings (required indifference), while incomparability is the intersection of gaps obtaining in all of them (required gap).

x is *better* than y iff x is preferred to y in every ordering in \mathbf{K} .

x is *equally as good* as y iff x is equi-preferred with y in every ordering in \mathbf{K} .

x and y are *incomparable* iff every ordering in \mathbf{K} contains a gap with respect to x and y .

Parity, on the other hand, consists in the presence of opposing permissible preferences:

x and y are *on a par* if x is preferred to y in some orderings in \mathbf{K} and dispreferred to y in some other orderings in \mathbf{K} .

The example with trips to Australia (x) and to South America (y) and their small improvements (x^+ and y^+ , respectively) is now easily representable. We just need to stipulate (i) that in all orderings in \mathbf{K} , x^+ comes above x and y^+ above y , and (ii) that \mathbf{K} contains an ordering in which x comes above y^+ , and another ordering in which y comes above x^+ . These two assumptions are easily jointly satisfiable¹² and they give us the betterness structure we have been after: x^+ is better than x and y^+ is better than y , but all other pairwise relations between the four items in questions are instances of parity.

For another application of our approach, consider how it can deal with the following objection to the concept of parity that has been put forward

¹² Thus, to give an example, \mathbf{K} might simply consist of two preference orderings, P_1 and P_2 , such that in P_1 the four items are ranked in the following order: x^+ , x , y^+ , y , while the order in P_2 is partly reversed: y^+ , y , x^+ , x .

by Martin Peterson. According to Peterson, parity would legitimize preferences that make the agent potentially vulnerable to money pumps:

Let x be on a par with y , and let the same be true of y and x^+ , where x^+ is a slightly improved, strictly better version of x . Parity is a non-transitive relation, so even though x is on a par with y and y is on a par with x^+ , it is not inconsistent to assume that x and x^+ are not on a par. Imagine that you are in possession of x^+ , and that you are offered to swap x^+ for y . Since the two objects are on a par, it seems reasonable to assume that you are permitted to swap. So you swap, and get y . You are then offered to swap y for x , which you do, since they are on a par. Finally, you are offered to pay a small amount of money for swapping x for x^+ . Since x^+ is strictly better than x , you accept the offer and end up where you started, the only difference being that you are now a little bit poorer. This is clearly irrational. (Peterson 2007: 507)¹³

Given our modelling, it is easy to see where this argument goes wrong. It does follow from x^+ and y being on a par that (i) it is permissible to prefer y to x^+ (and thus permissible to make the corresponding swap). Analogously, it does follow from y and x being on a par that (ii) it is permissible to prefer x to y (and thus permissible to make this swap). However, (i) and (ii) do *not* imply that (iii) it is permissible to prefer y to x^+ *and* to prefer x to y . Indeed, since in every permissible preference ordering x^+ is preferred to x (in view of the fact that the former item is better than the latter), no such ordering satisfies (iii). If x^+ is preferred to x and x is preferred to y , then x^+ is preferred to y , by the transitivity of permissible preference. Consequently, since permissible preference is asymmetric, it cannot be the case that y is preferred to x^+ . This means that no agent with permissible preferences will ever be pumped for money. While each swap in the sequence is permissible if considered in isolation, the swap sequence as a whole is impermissible.¹⁴ Note, by the way, that the same solution is unavailable on Gert's interval modelling, which again demonstrates the main weakness of this modelling: its inability to take into consideration holistic constraints on permissible preferences.

But is the intersection modelling we have presented really a novel one? Some readers might feel that this model is nothing but the standard supervaluationist approach, as applied to value comparisons, with the orderings in \mathbf{K} being different 'precisifications', or sharpenings, of relational value predicates, such as 'better than' or 'equally as good as'. This diagnosis, however, would be a mistake. The supervaluationist

¹³ In this money-pump, unlike as in the standard one, the agent is not required to make each of the swaps, but in Peterson's view it is enough for the money-pump argument that each swap in the money-pumping sequence is permissible.

¹⁴ Peterson (2007) also provides a non-sequential version of the money-pump argument against parity, but that version falls prey to the same objection as the one mentioned above.

approach is meant to account for indeterminacy, or vagueness. This is not the function of the intersection modelling. Thus, on the supervaluationist approach, a sentence is true iff it is true on all admissible sharpenings of its predicates, false iff it is false on all admissible sharpenings and indeterminate in truth-value iff it is true on some admissible sharpenings and false on others. If we apply this idea to value comparisons, it follows that such comparisons might be neither true nor false if relational value predicates allow of different sharpenings. Thus, if x is ranked above y in some admissible sharpenings of the value ordering but below y in other admissible sharpenings, then it is neither true nor false that x is better than y . In the intersection modelling, however, the orderings in \mathbf{K} are *not* potential sharpenings of the value ordering, but instead permissible *preference* orderings from which the value ordering is constructed. Consequently, if x is ranked above y in some orderings in \mathbf{K} but below y in other orderings in \mathbf{K} , then it is definitely *false* that x is better than y and definitely *true* that they are on a par.

Three comments may be in order at this point:

(i) Despite this difference in interpretation, there is of course a strong formal similarity between the intersection modelling and the supervaluationist approach to indeterminacy. These formal similarities were brought out in Rabinowicz (2008), where it was shown that the examples exhibiting the inadequacy of the interval modelling for the representation of betterness structures have exact analogues in formally similar examples that demonstrate the inadequacy of probability intervals as representations of *indeterminate probabilities*. We might have four propositions x , x^+ , y and y^+ such that x^+ is definitely more probable than x , y^+ is definitely more probable than y , but all other comparative probability judgements concerning the propositions in questions are indeterminate in truth-value. To account for such situations, the assignment of probability intervals to propositions needs to be replaced by a 'holistic' modelling of indeterminacy that postulates a class of admissible sharpenings – a class of sharp probability assignments to the propositions in the domain.

(ii) The intersection modelling which allows us to go beyond the traditional trichotomy of value relations ('better than', 'worse than', 'equally good as') and find room for such relations as parity is *not* meant to be a replacement for the approaches that account for indeterminacies in value comparisons. For a discussion of a more complex modelling that combines the intersection approach to value relations with a supervaluationist approach to indeterminacy in relational value judgements, the reader is referred to Rabinowicz (2009a, 2009c).

(iii) The intersection model should be sharply distinguished not only from supervaluationism but also from Amartya Sen's 'intersection approach', which consists in constructing the relation of betterness from

a class of value orderings that reflect different value commitments or different evaluative aspects of the items that are being compared. (See Sen 1973, ch. 3.) This is not what’s going on in the intersection modelling. Each of the preference orderings in **K** is supposed to take into considerations *all* the relevant aspects of the compared items. If it ignored some of them, it would hardly be permissible. In other words, it is not a matter of constructing betterness all-things-considered from different relations of being better in this or that respect.¹⁵

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

TABLE 1

The intersection model readily lends itself to a taxonomy of all possible dyadic value relations. See Table 1. The idea behind the taxonomy is simple. For each pair of items, the value relation that obtains between them is determined by the list of preferential positions that are permissible with respect to the two items in question. There are four such possible positions: preference for the first item (>), equi-preference (≈), preference for the second item (<), or preferential gap (/). Consequently, there are 2⁴ – 1 atomary types of value relations, where the number of such types is simply the number of ways a non-empty subset can be selected from the set consisting of four elements. We thus get the following table in which rows correspond to the four preferential positions and columns stand for the different atomary types of value relations. In each column, the permissible preferential positions are marked with a plus sign.

¹⁵ This means that the intersection model is not vulnerable to the standard objections to Sen’s intersection approach. The latter’s weakness is that it treats equally different respects in which one thing can be better than the other: it has no room for the gradation of their significance. Nor does it take into consideration *how much* each item scores on each respect or *how many* respects tell for one item as compared with the other. For a critical discussion of Sen’s intersection approach along these lines see, for example, Temkin (1993: 145f).

Better (**B**), Worse (**W**) and Equally-as-good-as (**E**) are all atomary types, and so is Incomparability (**I**). Parity (**P**), on the other hand, is a type in a broader sense: it is a collection of four atomary types, 6 to 9. It's interesting to note that types 4–5 and 10–14 lie outside the familiar classifications: they aren't instances of **B**, **E**, **W** or **I**, nor do they fall into the **P**-category. One should be aware, however, that the taxonomy lists logical possibilities. There is no guarantee that one can find real-life exemplifications of every atomary type of value relation. In fact, further conditions might exclude some atomary types. To illustrate, it might seem reasonable to expect that if two items are on a par, i.e. if preference for each is permissible, then it should also be permissible to be indifferent between the items in question. This 'convexity' condition would exclude types 6 and 8. If we were to stipulate, in addition, that in cases of parity preferential gaps should also be allowed, parity would reduce to type 9. Note, however, that such extra conditions importantly differ from, say, the condition of transitivity on preferences. The latter is a constraint on each ordering in **K**, while the extra conditions instead impose constraints on class **K** taken as a whole. They stipulate that **K** must contain orderings of certain kinds provided it contains orderings of certain other kinds. Whether such conditions are generally compelling is often difficult to decide.

One final remark. The taxonomy distinguishes between various ways in which two items can be permissibly preferentially situated vis-à-vis *each other*, but it disregards the preferential positions they are allowed to have with respect to *other* items in the domain. To put it differently, it is a classification of 'internal' dyadic value relations. External value relations, which obtain between items in virtue of their permissible preferential relationships to other items, are disregarded.¹⁶ This means, by the way, that the intersection model is less holistic than it might seem at first sight. If one only wants to study the 'internal' value relationships that obtain between items within some restricted set *X*, there is no need to take into consideration those parts of orderings in **K** that concern items not belonging to *X*. In other words, it is then enough to consider the restriction of **K** to *X*. If the domain of items is infinite, the class **K** of permissible preference orderings on that domain might also turn out to be infinite in size. But if the subset *X* of items under considerations is finite, the

¹⁶ Some such external value relations might be interesting in their own right. An example is congruence: *x* is *congruent* with *y* iff for every item *z* in the domain and for every preferential position (preference, dispreference, equi-preference and gap), that position is permissible with respect to *x* and *z* iff it is permissible with respect to *y* and *z*. Congruence doesn't necessarily imply equal goodness, but the opposite entailment holds: equal goodness entails congruence.

restriction of **K** to *X* will have to be finite as well. This radically simplifies the whole exercise.

4. PROBLEMS

The account I have proposed faces several problems, some of which were already mentioned in Rabinowicz (2008). Four problems, in particular, seem to be rather difficult to deal with.

Analyticity

The first problem has to do with some of the formal properties of value relations, such as the transitivity of betterness or of equal goodness. Intuitively, these properties are analytic, or at least this seems to be a common view. That betterness is transitive is thought to be a conceptual truth.¹⁷ However, in the intersection modelling such formal properties of value relations are derived from conditions on permissible preference orderings. Thus, betterness is transitive because permissible preferences are taken to be transitive. The case of equal goodness is analogous. On the interpretation of the intersection modelling I was assuming, according to which preferences are dispositions to make pairwise choices, such conditions on permissible preferences have to be treated as normative requirements of some sort – more precisely, as rationality constraints. If someone is disposed to choose *x* when confronted with *x* and *y*, and to choose *y* when confronted with *y* and *z*, then it is arguable that this person would be irrational if she at the same time were disposed to choose *z* when confronted with *x* and *z*. To have such a set of dispositions would not be impossible, but it would arguably be irrational, especially if we take the relevant choices to be reason-based. But the argument for the irrationality might not be absolutely compelling. It is conceivable that a non-transitive choice pattern could after all be rationalized, even in the case of reason-based choices. It is well known that changes in the alternatives available for choice might affect weightings of the dimensions that are considered to be choice-relevant and it is at least conceivable that such changes in weightings could be justified somehow.¹⁸ Be that as it may. What is important is that the formal properties of value relations are

¹⁷ There are exceptions to this consensus about the transitivity of betterness. See, for example, Rachels (1998), Temkin (1996), (2001) and (2012). For some criticisms of these attacks on transitivity, cf. Carlson (2003), Qizilbash (2005) and Arrhenius & Rabinowicz (2005). For a further discussion, see below, section 6.

¹⁸ For a suggestion that this kind of context-dependence might imply that a fitting-attitude analysis of betterness in terms of required preferences among alternatives should allow for violations of transitivity, see Arrhenius (2005), p. 210: 'Given this analysis, it should come as no surprise that 'better than' will be context sensitive and that transitivity fails ... This possible understanding of value-concepts might explain why some theorists, like Temkin and Rachels, have been willing to abandon the transitivity of 'better than'.'

in this approach derived from normative constraints on preferences. This seems incompatible with the claim that the transitivity of betterness is a conceptual truth.¹⁹

Elusiveness of incomparability

The second problem concerns value incomparability, which was equated with a required preferential gap. One would expect that incomparability shouldn't be a very rare phenomenon. However, if preferences are interpreted in partly behavioural terms, as dispositions to make reason-based choices, then it is difficult to provide examples of items (within one and the same ontological category) with respect to which such choice dispositions would be, normatively speaking, out of the question.²⁰ That it sometimes is *permissible* to lack a disposition to make a reason-based choice between two items is one thing. It doesn't seem to be especially problematic that sometimes, when a decision is called for, we make our choices without having resolved the conflict of reasons for and against the alternatives among which the choice is made. Making choices while remaining in such a conflicted state of mind might well be permissible in some cases. But that it sometimes should be *required* to proceed in this way does seem strange. Probably, the most promising examples would be instances of tragic dilemmas, such as Sophie's Choice. It is arguable; I suppose, that when Sophie had to choose which of her children was to be saved, it was impermissible for her to arrive at this choice by weighing reasons for saving one child against reasons for saving the other. It is arguable, in other words, that it was required of Sophie to remain in her conflicted state of mind while making a choice. But this requirement, if present, appears to be distinctively moral in nature. As I suggested above, it is probably inappropriate to give a moralized interpretation to the deontic component of the FA-analysis of value. It is therefore doubtful whether tragic dilemmas really are good examples of value incomparability. And I don't quite see what other examples would fit the bill, as long as preferences are interpreted as choice dispositions.

¹⁹ There is nowadays an influential view according to which rationality conditions are not genuinely normative (see Kolodny 2005; cf. also Broome 2005, 2008, who doesn't defend the view in question, but who still rejects extant arguments in favour of the normativity of rationality). But the philosophers who take this position have no intention of suggesting that rationality conditions on mental states instead are analytically or conceptually true. They are in their view at most analytically true of what rationality consists in.

²⁰ It's another matter if the items belong to different ontological categories. It doesn't make sense to choose between, say, a person and a state of affairs, or between an abstract object and a material thing.

Preferences vs. value judgements

The third problem has to do with the relation between preferences and judgements of value. Suppose the agent judges x to be on a par with y . Given the analysis I have suggested, this implies that, in her opinion, x is not better than y but preferring x to y is permissible. However, can she herself, given her judgement of value, have this preference for x over y , if preference is reason-based? It is not clear that she can, without inconsistency. It would mean that she can deny that x is better than y but still consistently view reasons in favour of x as being stronger than those in favour of y . This seems problematic.²¹

Domain of preference

The fourth problem concerns potential objects of preference. On the interpretation I have been assuming, preferences are essentially related to potential choices. However, for some item domains, the notion of choosing between items doesn't make much sense, if taken literally. We make choices between options. But what would it mean to choose, say, Mozart rather than Haydn? We might choose to listen to Mozart's music rather than to Haydn's. Or choose to rank Mozart higher than Haydn in a contest for the title of the greatest eighteenth-century composer. But this means that the two composers aren't themselves possible objects of choice, at least not literally. Consequently, if one analyses value relations in terms of permissible preferences among items, then connecting preferences too closely to choice spells trouble.

There may be ways to allay these worries, or at least some of them, while still retaining the interpretation of preference that I was using in my 2008 paper. Thus, in particular, the problem of Preferences vs. Value Judgements could be dealt with if we take seriously the idea of different admissible sets of weights for various respects or dimensions of comparison. If the weights are optional to some extent, the resolution of the conflict of reasons which an agent arrives at can go hand in hand with the recognition that this conflict might just as well be resolvable in a different way. Consequently, such an agent might take reasons in favour of x to be stronger than reasons in favour of y , but – to the extent that she is aware of the optional nature of this resolution – she would at the same time be willing to deny that x is better than y .

Possibly, one might also try to solve the problem of Analyticity by a suggestion that the transitivity of preference and of indifference, while being a normative requirement, is as such a conceptually necessary rationality constraint: One might claim that it is an analytic feature of permissible preferences (*pace* my worries that the rationality of the

²¹ I owe this objection to Andrew Reisner and Sarah Stroud (personal communication).

transitivity constraint might not be absolutely compelling). This would make the transitivity of betterness and of equal goodness come out as analytic, after all, despite their being grounded in normative requirements on preferences.²²

Still, other objections might be more difficult to deal with. A radical way to deal with them all is to make a fresh start by re-interpreting the notion of preference.

4. PREFERENCES RE-INTERPRETED

The main idea is simple. We need to give up the assumption that preference is an attitude that consists in comparing items to each other. We need to deny that it is an attitude directed to several items at a time. This treatment of preference as a polyadic, or – more specifically – a dyadic attitude is part and parcel of its interpretation in choice-dispositional terms. To be disposed to choose x rather than y is an intentional stance that takes both x and y as its objects. An alternative I want to consider is to interpret preference for x over y as *a relation that obtains between monadic attitudes*, where each of the latter takes a different item, x and y respectively, as its object. As long as we only consider items that are valuable, i.e. fit to be favoured, the relevant monadic attitudes should be thought of as the attitudes of favouring. (A more general account is sketched below.) I take it that for each such attitude we can talk about its *degree*, where these degrees of favouring can vary and are ordered by the transitive and asymmetric higher-than relation. An item x is preferred to another item y iff it is favoured to a higher degree than y . Or, to put it in a slightly different vocabulary, iff it is favoured more than y .²³ Two items are equi-preferred iff they are equally favoured, i.e. favoured to the same degree, and there is a preferential gap between the items iff the degrees to which they are favoured are incommensurable: they aren't the same but neither is higher than the other.

Incommensurability of degrees is possible if we suppose that, in some cases, the ways two items are favoured are very different. In one case,

²² I am indebted to Hannes Leitgeb for pressing this point. It should be noted that the conceptual character of rationality requirements is an important tenet of Donald Davidson's philosophy. In his view, it is conceptually impossible for an agent to be systematically irrational. As Davidson puts it: 'it is only by interpreting a creature as largely in accord with these principles [i.e. principles of rationality] that we can intelligibly attribute propositional attitudes to it, or that we can raise the question whether it is in some respect irrational.' (Davidson 1985: 352)

²³ To forestall possible objections, let me stress that this is not meant to be an analysis of 'preference', as this term is commonly used. Instead, my proposal should be seen as a stipulative definition. On this proposal, preference is to be understood as the relation of favouring more, whether or not this is the meaning of 'preference' in ordinary usage. (I don't think it is, by the way.) I am indebted to John Broome for pressing me on this point.

for example, it might be a matter of liking, while in the other of, say, admiring. Possibly, if one item is only mildly liked while the other is strongly admired, it is correct to say that the degree in which the latter is favoured is greater than the degree in which the former is favoured.²⁴ But if the differences in the intensity of favouring are not pronounced, then the difference in the *ways* of favouring might well lead to a situation in which the corresponding degrees of favouring should be judged to be incommensurable. Still, this kind of incommensurability is not really interesting for the problem we are discussing. When it is a matter of comparing items in value, we need, I think, to begin by fixing the kind of value we are focusing on. This is done by specifying the relevant kind of favouring. If items are compared with respect to such value as admirability, then the relevant kind of favouring is admiration and we focus on different degrees of admiration. If it instead is likeability that is at issue, then we focus on different degrees of liking. And so on. However, even if we restrict our attention to a specific kind of favouring, the ways in which items are favoured may still significantly vary within a given kind, provided that the kind of favouring we focus on is broad enough. To give an example, even though we greatly admire Mother Teresa and Isaac Newton, the nature of admiration is quite different in these cases. In situations like this, it would be hopeless to try to commensurate the degrees of admiration. Formally, there is no obstacle in allowing degrees of favouring to be incommensurable, since we don't need to assume that degrees are representable by numbers.

Ideally, I would at this point provide an account of how comparative claims such as '*x* is favoured more than *y*, but equally as much as *z*' are to be understood, in psychological terms, together with a description of methods of assigning degrees to favourings, to represent comparative claims of this kind. This task, however, is beyond the purview of my paper. Psychologists have developed a whole host of methods of measuring attitudes, using not only self-reports of the subjects but also various indirect methods (delays in response time, Implicit Association Test, etc.) and measures of physiological responses or brain activity. Some of these measuring methods rely on one-dimensional numerical scales, but others use multidimensional approaches instead. A well-known example of a multidimensional scale is the classic semantic differential technique devised by Charles E. Osgood. (See Osgood *et al.* 1957, 1975, Heise 2010.) It consists in letting a subject assign numerical values to an item on a number

²⁴ Zimmerman (2001: 43) suggests that it might sometimes be possible to commensurate fitting pro-attitudes even when they belong to different kinds. What's needed for such commensuration is that 'the different attitudes have a common 'core''. I think this may be right, but it is important to note that such 'cross-kind' commensuration of attitudes has its limits: it does not allow of fine distinctions of degree.

of dimensions, each of which is determined by a different pair of bipolar adjectives. The vector of values on different dimensions is then seen as the measure of the subject's attitude: its degree. Multidimensional scales obviously leave ample room for incommensurabilities between attitudinal degrees.

It is arguable that specifying the degrees of different attitudes will at some stage require comparing these attitudes with attitudes directed towards other items. But this doesn't mean of course that the attitudes themselves are comparative in their intentional content. They are not, insofar as they are monadic.

Let's move on, then, to our formal model. x is weakly preferred to y iff x is preferred to y or equi-preferred with y . Thus, in order to determine the weak preference relation on the domain of items it is sufficient to specify to what degree each item is being favoured. As is easily seen, the so-defined relation of weak preference is transitive and reflexive (a quasi-order), but it need not be complete given the possibility of incommensurable degrees.

Thus, let us suppose that \mathbf{D} is a set of possible degrees of favouring (where favouring is assumed to be of some specific kind that we focus on: admiration, liking, desire, or some other, potentially more complex pro-attitude). The set \mathbf{D} is ordered by a transitive and asymmetric relation \succ (higher-than). Let f be a particular assignment of degrees of favouring from \mathbf{D} to items in the domain, i.e. f is a mapping from the item domain to \mathbf{D} . We shall refer to f as an *assignment of degrees of favouring* or – shorter – just as a *degree assignment*. A degree assignment straightforwardly determines a preference ordering on the domain of items:

x is weakly preferred to y in f iff

$f(x) \succ f(y)$ (in which case x is preferred to y in f)

or

$f(x) = f(y)$ (in which case x and y are equi-preferred in f).

Now, suppose that \mathbf{F} is the class of *permissible* assignments of degrees of favouring. In terms of \mathbf{F} , we can define the class \mathbf{K} of permissible preference orderings as the class of orderings that are constructible from assignments f in \mathbf{F} in accordance with the above definition of weak preference. In terms of \mathbf{K} , we then proceed to define various value relations, in the way we have done in the intersection model. Thus, to take just one example of such a relation, the definition of betterness now looks as follows:

x is better than y = it is required to prefer x to y = it is required to favour x more than y

$= x$ is preferred to y in every ordering in $\mathbf{K} =$ for every assignment f in \mathbf{F} , $f(x) > f(y)$.

Note that this approach is a bit more demanding than the old one: the resources it uses are stronger. By this I mean that the specification of class \mathbf{F} provides more information than the specification of \mathbf{K} . This is easily seen if one notes that two different degree assignments, f and f' , can sometimes give rise to exactly the same weak preference relation. This will be the case whenever f' is a monotonic transformation of f , i.e. whenever $f'(x) > f'(y)$ iff $f(x) > f(y)$, and $f'(x) = f'(y)$ iff $f(x) = f(y)$, for all x and y in the item domain. Consequently, if there is a mapping from \mathbf{F} to a class \mathbf{F}' such that every assignment in \mathbf{F} is mapped to its monotonic transformation in \mathbf{F}' , and a corresponding mapping from \mathbf{F}' to \mathbf{F} , then \mathbf{F}' and \mathbf{F} will give rise to exactly the same class \mathbf{K} of permissible preference orderings.

As long as we only are interested in defining those value relations that we have distinguished in the intersection model, the extra resources of the degree-assignment model do not increase the expressive power of our approach. From this point of view, therefore, the new model might be criticized for containing more than we really need. Still, this disadvantage is not especially serious. If we wish, we could simply decide to treat the choice of \mathbf{F} as arbitrary up to monotonic transformations. This would mean that we discount the extra information the degree-assignment model provides.

5. EXTENDING THE RE-INTERPRETATION OF PREFERENCE BEYOND FAVOURED ITEMS

I have been assuming that all the items in the domain are valuable, i.e. fitting to be favoured. This allows us to concentrate on cases in which both the preferred item and the dispreferred one are objects of favouring attitudes. Can we give up this assumption and still work with a model in which preference is accounted for in terms of a relation between monadic attitudes towards items? Yes, I think so. We can divide monadic attitudes towards items into three disjoint groups: favourings, disfavourings and attitudes of indifference. The attitudes in the first two groups may be assumed to exhibit varying degrees. Somewhat artificially, we may also assign a 'degree' to attitudes of indifference – the same degree to all of them – and refer to this degree as i . Now, let \mathbf{D}^+ and \mathbf{D}^- be the sets of possible degrees of favouring and disfavouring, respectively. Each of these sets is ordered by a transitive and asymmetric relation $>$. (It is possible that the two sets of degrees are essentially the same and only differ with respect to their signs, so to speak. But we don't need to make this assumption.)

On such a more comprehensive approach, a degree assignment f is a mapping from the domain of items into the union of possible attitudinal

degrees, $\mathbf{D}^+ \cup \mathbf{D}^- \cup \{i\}$. We now have all we need for the definition of preference:²⁵

x is preferred to y in f iff

either (i) $f(x), f(y) \in \mathbf{D}^+$ and $f(x) > f(y)$,

or (ii) $f(x), f(y) \in \mathbf{D}^-$ and $f(y) > f(x)$,

or (iii) $f(x) \in \mathbf{D}^+$, while $f(y) \in \mathbf{D}^- \cup \{i\}$,

or (iv) $f(x) \in \mathbf{D}^+ \cup \{i\}$, while $f(y) \in \mathbf{D}^-$.

In case (i), both x and y are favoured, but the former is favoured to a greater degree. In case (ii), both items are disfavoured, but y is disfavoured to a greater degree. In case (iii), x is favoured, while y is disfavoured or viewed with indifference, and finally in case (iv), y is disfavoured, while x is favoured or viewed with indifference.

Defining equi-preference is straightforward:

x and y are equi-preferred in f iff $f(x) = f(y)$.

Thus, two items are equi-preferred if they are both favoured or both disfavoured to the same degree or if they are both viewed with indifference.

We then define weak preference as the disjunction of preference and equi-preference, as previously. It is trivial to show that the so-defined weak preference is a quasi-order (a transitive and reflexive relation).

In what follows, however, I will for simplicity's sake hold on to the assumption that all the items in the domain under consideration are favoured, in varying degrees.

6. PROBLEMS SOLVED

Let us now consider how this new approach to preference can deal with the problems we have identified for the old approach. As for the problem of Analyticity, this difficulty is now avoided, since value relations are analysed in terms of permissible preferences and *all* preferences are now transitive by definition. Thus, transitivity is no longer just a property of permissible preferences, which means that it no longer is seen as a normative requirement of some kind. That equi-preference is transitive (which implies that equal goodness is transitive) immediately follows from the transitivity of identity: two items are equi-preferred iff they are favoured to the same degree. And that preference is transitive (which

²⁵ For a definition of preference along these lines, see Zimmerman (2001: note 94 on p. 128). Zimmerman (2001: 104) defends this account of preference against an objection by Brentano.

implies that betterness is transitive) follows from the transitivity of the higher-than relation on the set of attitudinal degrees. That the latter relation is transitive is, it seems to me, a conceptual truth. It is true solely in virtue of the meaning of 'higher', that whatever is higher than something that in its turn is higher than something else, must itself be higher than that something else. 'Higher' shares this analytical property of transitivity with other comparative predicates: 'longer', 'louder', 'lighter', etc.²⁶

Moving to the next problem on our list, the Elusiveness of Incomparability, we note that incomparability no longer has to be a very rare phenomenon on the new approach. It can obtain when two items call for significantly different attitudes of favouring, which is possible even if the items belong to the same ontological category and even if we focus on a specific category of favouring, provided the category is broad enough. I mentioned one such example above. The kind of admiration Newton and Mother Teresa respectively call for is very different, which makes it reasonable to expect that permissible degrees of admiration are mutually incommensurable in a case like this. These two individuals are both great human beings, deserving of admiration, but it is fair to say that they are incomparable in value.

As for the problem of Preferences vs. Value Judgements, and in particular the issue of apparent incompatibility between preferences and judgements of parity, we have already seen how this difficulty can be dealt with if one stays with the original interpretation of preference as involving some kind of a reason-based comparison between the items. If the balance of reasons is seen by the agent as being optional to some extent, then she can consistently deny that *x* is better than *y* while taking the reasons in favour of *x* to be stronger than those in favour of *y*. But on the new interpretation, the solution to this problem is even more immediate. If preference for *x* over *y* is not a dyadic attitude having both *x* and *y* as its intentional objects, but a relation between monadic attitudes directed separately towards these items, then the agent might have this preference without making any direct comparison between the items in

²⁶ For an argument that transitivity is an analytic feature of all comparative predicates, see Broome (2004: section 4.1). The argument is based on the assumption that each comparative 'Fer than' is a synonym of 'more F than' for some non-comparative predicate F. Broome's reasoning is compelling, but perhaps not fully convincing. Johan Brännmark has suggested that some comparative predicates are cyclical and therefore cannot be transitive. An example is 'later than' when applied to the times of day. Three in the morning is later than nine in the evening, which is later than three in the afternoon, which is later than nine in the morning, which is later than three in the morning. Broome tries to disarm this example (2004: 52ff), but I am not convinced that he succeeds. Still, even if he doesn't and transitivity cannot be assumed to be a feature of all comparatives, I find it intuitively very plausible that it is an analytic feature of 'higher' and, by implication, an analytic feature of 'better'. For a further discussion of the transitivity issue, see however the concluding section.

question. Consequently, on this new interpretation, her preference for x over y is not even *prima facie* inconsistent with her denial that x is better than y .

Note, by the way, that this re-interpretation of the notion of preference need not imply that we give up the view that preferences have to be reason-based. While preference is now seen as a relation between monadic attitudes and thus does not involve any direct comparison of reasons in favour of one item with reasons in favour of another item, we might well require that the monadic attitudes themselves should be based on reasons (and on weighing of reasons) in order to be of interest from the point of view of FA-analysis.

As for the last problem, concerning the Domain of Preference, this difficulty now disappears. On the re-interpreted conception of preference, there is no essential connection between preferences and choice. The monadic attitudes of favouring, such as admiration, liking, etc., need not consist in choice dispositions. Some of them might be choice-related, but others need not be. Consequently, the items on which preference relation operates need not all be potential objects of choice. This domain restriction is now lifted altogether.

A nice feature of the new model is that it allows us to deal not just with generic value relations, better, equally good, etc., as we have done here, but also with specific relations, such as being more admirable or being more desirable, and so on. To account for them we need to specify the kind of favouring attitude that we are interested in (admiration, desire, or what have you) and then develop the degree assignment model for the attitudes of this kind. (For a proposal along this lines, see Rabinowicz 2009b.)

One attractive possible application of our model lies in the epistemic area: it is tempting to provide a fitting-attitude account of the notion of *probability*. Roughly, the idea would be that a proposition is probable to the extent that it is fitting to believe it. This opens up for a modelling of different types of probability relations between propositions, such as more/less probable than, equally probable as, being probabilistically on a par, etc., in terms of the set of permissible assignments of degrees of belief to propositions. Permissibility of an assignment would have to depend not just on its coherence but also on the support it receives from available evidence. The latter would normally allow for several alternative assignments rather than for exactly one. In this way, we get an interpretation of probability which combines elements of the subjective approach (connection to partial beliefs) with a normative perspective (embodied in the notion of fittingness of permissibility). On this interpretation, probability is seen as a value.

Needless to say, along with new applications and solutions to old problems, we might now encounter new problems instead. But, hopefully, the new problems will be possible to deal with.

7. CONCLUSION

In a way, the degree assignment model reminds of Gert's original proposal, which also assumed that for each item one could specify the degree to which that item is preferred (= favoured). But at the same time there are obvious differences between the two approaches. The three most important ones are:

- (i) In Gert's approach, preferences were thought to be dispositions to choose and the degree assignments were representations of these dispositions. This is not how the degrees of favouring are meant to be understood in our approach.
- (ii) Gert's model of permissible intervals of degrees for different items is replaced by a modelling that postulates a class of permissible degree assignments to all the items in the domain. Thereby, our definitions of standard value relations, such as betterness and equal goodness, can avoid problems that plague Gert's account.
- (iii) To allow for incommensurable degrees, we give up the idea that degrees must be representable by (single) numbers. This makes room for incomparabilities in value.

It seems that the degree assignment model gives us what we have been after: a versatile FA-framework for a study of value relations, which allows us to do everything we have been doing using the intersection model, while avoiding the problems the latter model had to face. These problems disappear when the notion of preference is re-interpreted – when the essential connection between preference and choice is severed and preference is treated as a difference in degree between monadic attitudes directed to separate items.

One of my referees, Larry Temkin, has questioned my reasons to give up the original intersection model, which was based on preferences understood as dyadic, comparative attitudes. My problems with that model aren't serious, according to him. Thus, to begin with, the Elusiveness of Incomparability is nothing we need to bother about. As Temkin puts it: 'for comparisons between alternatives of the SAME ontological sort – it seems to me that they should all be comparable, at least once one expands [the] domain of comparability beyond the standard trichotomy, to include the relation of parity.'²⁷ My other referee, Joshua Gert, concurs with Temkin on this point. He writes: 'Why should we expect incomparability not to be very rare? I would have thought that many theorists think that it is in fact so rare that it does not exist at all!'

²⁷ This quote and all otherwise unattributed quotes in what follows come from the referee reports.

Temkin is just as unimpressed with my Domain of Preference objection. He argues that we never make value comparisons between such items as, say, persons, or other concrete things, simpliciter. We only evaluate them in terms of some features we are interested in. Thus, we can ask: Who was the best composer in his age, Mozart or Haydn? Or: Whom would we rather listen to? The suggestion seems to be that asking questions in this way makes comparisons choice related: We are asking whom we should choose, Mozart or Haydn, if we were to pick someone to listen to or if we were to choose the best composer of their century.

Moving to the Preferences vs. Judgements objection, Temkin thinks that this objection is not serious either. He takes my own response to it to be fully satisfactory. As I suggested, an agent might deny that x is better than y and still, on balance, prefer x to y in a reason-based way, if she regards the weights she uses in her balancing of reasons to be optional to some extent. On this point, Gert does *not* concur with Temkin. He thinks my suggestion doesn't make sense. As he puts it: 'How can I take the reasons in favor of x to be stronger as those in favor of y , but at the same time regard this belief – since that is what it seems to be – to be optional in any strong sense?' Well, a way out might be to argue that assigning relative weights to reasons in the process of forming a preference should not be understood as forming a belief. A less cognitive interpretation of the weighing of reasons seems to be needed.

It is, however, the problem of Analyticity that Temkin especially targets in his comments. He is unhappy with my underlying assumption that the transitivity of betterness is a conceptual truth. In his writings, Temkin has repeatedly argued that the relation of betterness can be seen in two very different ways. On what he calls the Intrinsic (or Internal) Aspects View (IA), judgements of betterness are secondary to independent assessments of the levels of goodness for each of the items that are being compared. And how good an item is is determined by its internal aspects.

On the view that moral concepts are not comparative, how good or bad a situation is regarding some factor, f , will be an *intrinsic* feature of that situation – that is, it will not depend on the alternative that situation is compared with, but solely on features internal to the situation. On this view – henceforth, the *Intrinsic Aspect* view, or 'IA' for short – how a situation has come about, or who its members are, will be irrelevant to the abstract, impersonal judgement about how it fares regarding f . (Temkin 1986: 158f; cf. also Temkin 1997: 309, and 1999: 777)²⁸

²⁸ There is an ambiguity here that needs to be noted. When we evaluate an item independently of the alternatives it is being compared with and our goal is to determine its final value, i.e. its value for its own sake, it might be a mistake to concentrate on purely internal aspects of the item and disregard its other features. Its final value might well depend, to some extent, on its relational properties. Thus, to take an example used by Jonathan Dancy in

Whether one item, x , is better than another, y , is thus determined by simply comparing their independently given levels of goodness and checking which level is higher (if any).²⁹ Temkin contrasts this with a view according to which betterness is 'essentially comparative'.

On an *essentially comparative* view (EC), ... [t]here is no fact of the matter as to how bad a situation *really* is, considered just by itself. How bad it is depends on the alternative compared to it. (Temkin 1999: 777)

On EC, then, a judgement of betterness is not based on prior independent assessments of the compared items' goodness levels. Such independent assessments can't be made in the case of essentially comparative concepts.

As far as the transitivity of betterness is concerned, this feature trivially follows on IA (since higher-than relation between goodness levels is of necessity transitive), but on EC it is a matter of dispute whether better-than is a transitive relation. Now, in both my modellings, the earlier and the later one, the transitivity of better-than is grounded in the transitivity of preference, but in the earlier modelling, which relies on the interpretation of preference as a dyadic attitude, the transitivity condition on preference is only a requirement of rationality – a requirement that may well be disputed. The earlier modelling can be used even if we give up on that requirement. This makes the modelling in question attractive from Temkin's point of view. He has much sympathy for EC and he has been arguing against the transitivity of betterness for many years now. In his comments to my paper, he suggests (i) that my later modelling (unlike the earlier one) expresses the (highly controversial) IA view on betterness, and (ii) that I am wrong in my claim that transitivity is a conceptual feature of betterness. Transitivity trivially follows only if we accept the IA view of betterness, but not if the EC view is assumed instead. These competing

his discussion of Temkin (Dancy 2005: 6f), the goodness of a situation might partly depend on the historical context – on how that situation came about. For other examples and a further discussion of this issue, see Rabinowicz & Rønnow-Rasmussen (2000). Still, this potential supervenience of final value on the item's relational properties does not make such value dependent on the alternatives with which a given item is being compared.

²⁹ Here is how Temkin further characterizes IA: 'One natural way of representing IA is in terms of a numerical model. On such a model, each situation will merit a 'score' representing how good that situation is all things considered, where *that* score will be a function of other 'scores' for each factor relevant to preferability (for instance, inequality, utility, perfection, and so on). Naturally, scores will be based solely on the internal features of the situation, hence alternatives with the same internal features will be assigned the same scores, whatever their origins, or comparative alternatives. For complete concepts, perhaps a precise number could be assigned, at least in principle, for each situation, such that the better the situation was with respect to the concept, the larger the number it received. For incomplete concepts such precision would be impossible even in principle, but perhaps a rough range of numbers could be assigned along similar lines' (Temkin 1986: 159).

views do not concern the concept of betterness, but its nature, according to Temkin. Therefore, the issue of transitivity is substantive and cannot be decided on purely conceptual grounds.

This is an interesting challenge. What do I have to say in response? Well, to begin with, I should point out that Temkin's suggestion (i) is incorrect. In my later modelling, in which I rely on preference interpreted as a relation between degrees of monadic attitudes, I do *not* understand judgements of betterness as comparisons between independently specified degrees of goodness. In fact, I make no attempt to define such degrees at all and my modelling lacks resources to do so. To this extent then, the later model – like the earlier one – treats betterness in an essentially comparative way. Both models take one item to be better than another iff the former is preferred to the latter in every permissible preference ordering.³⁰ The difference between the two models has to do with their treatment of preference. The *latter* is treated in an 'essentially comparative' way in the earlier model, but not in the later one, in which the relation of preference is grounded in independently specified degrees of favouring for each item. In both models, however, betterness (= preferability, i.e. required preference) is treated in a way that is incompatible with the IA view. As far as betterness is concerned, the IA approach to this relation is instead exemplified by Gert's modelling, in which one first determines the level of goodness for each item – the range of permissible preference strengths for the item in question – and then determines which item is better by comparing which level is higher, if any. (When levels are ranges, this means that the minimum of one range is greater than the maximum of the other range.) But – as we have seen – Gert's modelling cannot be upheld: It is unable to represent certain plausible betterness structures.

But what about Temkin's suggestion (ii)? Even if both my models may be seen as different versions of EC, Temkin might still object that these competing versions represent different views about the *nature*

³⁰ It might be questioned, however, whether my models are essentially comparative in the full-blown sense that Temkin has in mind. While it is clear that none of them instantiates the IA view of betterness, it might be that they lack some of the features Temkin would require from a genuine EC approach. Cf. the following quote: 'A principle (or moral ideal) *f* is *essentially comparative* if the *relevant and significant* factors for comparing two alternatives regarding *f* may vary depending on the alternatives being compared; and, more specifically, *f* is *essentially pairwise comparative* if one must directly compare two alternatives in order to determine their relative ranking regarding *f*' (Temkin 1997: 304). While my earlier model in principle does allow for variation in the 'relevant and significant' factors depending on what pairs of alternatives are being compared (but only as long as this doesn't lead to violations of the transitivity of preference), this kind of variation doesn't seem to be allowed for in my later model.

of betterness, but not about its *concept*. He might therefore insist that transitivity cannot be seen as a conceptual feature of betterness.

It is not at all clear to me how the line between the concept and the nature of betterness is to be drawn. The distinction between the nature and the concept of a phenomenon is standard in philosophy. Still, it is not easy to apply it when one moves beyond simple textbook examples, in which the concept is supposed to be a priori knowable while the nature is discoverable by empirical investigation. It is especially difficult to draw the distinction in those cases in which both the concept and the nature are supposed to be determinable by similar means: either in both instances by philosophical reflection, or in both instances by empirical investigation (as suggested by experimental philosophers). Anyway, *pace* Temkin, it doesn't seem to me that the difference between my two models has much to do with diverging views on the nature of betterness. It is not like, say, the debate between those philosophers who provide a hedonistic account of the nature of value and their opponents who present an objective list of values. The models I have considered have a much more conceptual flavour.

Let me, however, end this paper with a conciliatory suggestion: *There is room for several concepts of betterness*. In particular, if betterness is understood as preferability, as suggested by the fitting-attitudes analysis, it might be that different concepts of betterness correspond to different notions of preference. One concept of betterness is represented by my earlier model, which is based on preference understood as a dyadic attitude, while another concept of betterness, similar but different, is represented by a modelling that construes preference as a relation between monadic attitudes. Transitivity is a conceptual feature of the latter concept of betterness, but not of the former. Both concepts have their uses and we do well if we keep this in mind. I am not sure whether this conciliatory position will appeal to many antagonists in the debates concerning the formal properties of value relations, but – to me – it seems like a reasonable compromise.

REFERENCES

- Arrhenius, G. 2005. The paradoxes of future generations and normative theory. In *The Repugnant Conclusion*, ed. J. Ryberg and T. Tännsjö, 201–218. Dordrecht: Kluwer.
- Arrhenius, G. and W. Rabinowicz 2005. Value and unacceptable risk: Temkin's worries about continuity – reconsidered. *Economics and Philosophy* 21: 177–198.
- Brentano, F. 1969 [1889]. *The Origin of Our Knowledge of Right and Wrong*, transl. by R. M. Chisholm. London and Henley: Routledge & Kegan Paul.
- Broome, J. 2004. *Weighing Lives*. Oxford: Oxford University Press.
- Broome, J. 2005. Does rationality give us reasons? *Philosophical Issues* 15: 321–337.
- Broome, J. 2008. Is rationality normative? *Disputatio* 11: 153–171.
- Carlson, E. 2003. On some recent examples of intransitive betterness. In *Logic, Law, Morality*, ed. K. Segerberg and R. Sliwinski, 99–113. *Uppsala Philosophical Studies* 51. Uppsala: Department of Philosophy, Uppsala University.
- Chang, R. 1997. Introduction. In *Incommensurability, Incomparability and Practical Reason*, ed. R. Chang, 1–34. Cambridge, MA: Harvard University Press.

- Chang, R. 2002a. The possibility of parity. *Ethics* 112: 659–688.
- Chang, R. 2002b. *Making Comparisons Count*. New York & London: Routledge.
- Chang, R. 2005. Parity, interval value, and choice. *Ethics* 115: 331–350.
- Dancy, J. 2005. Essentially comparative concepts. *Journal of Ethics and Social Philosophy* 1: 1–15.
- Davidson, D. 1985. Incoherence and irrationality. *Dialectica* 39: 345–354.
- Gert, J. 2004. Value and parity. *Ethics* 114: 492–510.
- Gustafsson, J. E. and N. Espinoza. 2010. Conflicting reasons in the small improvement argument. *The Philosophical Quarterly* 60: 754–763.
- Heise, D. R. 2010. *Surveying Cultures: Discovering Shared Conceptions and Sentiments*. Hoboken, NJ: John Wiley & Sons.
- Kolodny, N. 2005. Why be rational? *Mind* 114: 509–563.
- Osgood, C. E., G. Suci and P. Tannenbaum 1957. *The Measurement of Meaning*. Urbana, IL: University of Illinois Press.
- Osgood, C. E., W. H. May and M. S. Miron 1975. *Cross-Cultural Universals of Affective Meaning*. Urbana, IL: University of Illinois Press.
- Parfit, D. 2011. *On What Matters*. Oxford: Oxford University Press.
- Peterson, M. 2007. Parity, clumpiness, and rational choice. *Utilitas* 19: 505–513.
- Qizilbash, M. 2005. Transitivity and vagueness. *Economics and Philosophy* 21: 109–131.
- Rabinowicz, W. 2008. Value relations. *Theoria* 74: 18–49.
- Rabinowicz, W. 2009a. Broome and the intuition of neutrality. *Philosophical Issues* 19 (supplement to *Noûs*): 389–411.
- Rabinowicz, W. 2009b. Values compared. *Polish Journal of Philosophy* 3: 73–96.
- Rabinowicz, W. 2009c. Incommensurability and vagueness. *Proceedings of the Aristotelian Society*, Supplementary Volume 83: 71–94.
- Rabinowicz, W. 2011. Value relations – old wine in new barrels. In *Philosophical Papers dedicated to Kevin Mulligan*, ed. Anne Reboul. www.philosophie.ch/kevin/festschrift/
- Rabinowicz, W. Forthcoming. Value: fitting-attitudes account of. In *International Encyclopedia of Ethics*, ed. Hugh LaFollette. New York, NY: Wiley-Blackwell.
- Rabinowicz, W. and T. Rønnow-Rasmussen 2000. A distinction in value – intrinsic and for its own sake. *Proceedings of the Aristotelian Society* 100: 33–51.
- Rabinowicz, W. and T. Rønnow-Rasmussen 2004. The strike of the demon: on fitting pro-attitudes and value. *Ethics* 114: 391–423.
- Rachels, S. 1998. Counterexamples to the transitivity of better than. *Australasian Journal of Philosophy* 76: 71–83.
- Sen, A. 1973. *On Economic Inequality*. Oxford: Clarendon Press.
- Skorupski, J. 2010. *The Domain of Reasons*. Oxford: Oxford University Press.
- Temkin, L. S. 1986. Intransitivity and the mere addition paradox. *Philosophy and Public Affairs* 16: 138–210.
- Temkin, L. S. 1993. *Inequality*. Oxford: Oxford University Press.
- Temkin, L. S. 1996. A continuum argument for intransitivity. *Philosophy and Public Affairs* 25: 175–187.
- Temkin, L. S. 1997. Rethinking the good. Moral ideals and the nature of practical reasoning. In *Reading Parfit*, ed. J. Dancy, 290–345. Oxford: Blackwell.
- Temkin, L. S. 1999. Intransitivity and the person-affecting principle: a response. *Philosophy and Phenomenological Research* 59: 777–784.
- Temkin, L. S. 2001. Worries about continuity, transitivity, expected utility theory, and practical reasoning. In *Exploring Practical Philosophy*, ed. D. Egonsson, J. Josefsson, B. Petersson and T. Rønnow-Rasmussen, 95–108. Aldershot: Ashgate.
- Temkin, L. S. 2012. *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford: Oxford University Press.
- Zimmerman, M. J. 2001. *The Nature of Intrinsic Value*. Lanham: Rowman & Littlefield.