

Value-Preference Symmetry and Fitting-Attitude Accounts of Value Relations^{*}

Johan E. Gustafsson[†]

ABSTRACT. Joshua Gert and Wlodek Rabinowicz have developed frameworks for value relations that are rich enough to allow for non-standard value relations such as parity. Yet their frameworks do not allow for any non-standard preference relations. In this paper, I shall defend a symmetry between values and preferences, namely, that for every value relation, there is a corresponding preference relation, and vice versa. I claim that if the arguments that there are non-standard value relations are cogent, these arguments, *mutatis mutandis*, also show that there are non-standard preference relations. Hence frameworks of Gert and Rabinowicz's type are either inadequate since there are cogent arguments for both non-standard value and preference relations and these frameworks deny this, or they lack support since the arguments for non-standard value relations are unconvincing. Instead, I propose a simpler framework that allows for both non-standard value and preference relations.

Some authors defend the possibility of non-standard value relations. Such relations might hold when none of the standard relations 'better', 'worse', 'equally good', and 'incomparable' holds. For example, Ruth Chang argues that there is a non-standard value relation, which she calls 'on a par'.¹ In order to analyse such relations, one needs a framework for value relations that is rich enough to allow for other value relations than the standard ones. Two such proposals are due to Joshua Gert and to Wlodek Rabinowicz.² Yet, while their frameworks allow for non-standard value relations, they do not allow for any non-standard preference relations, that is, preference relations that might hold when none of the standard pref-

p. 477

^{*} Published in *The Philosophical Quarterly* 63(252):476–491, 2013. The definitive version is available at <http://onlinelibrary.wiley.com>.

[†] I would be grateful for any thoughts or comments on this paper, which can be sent to me at johan.eric.gustafsson@gmail.com.

¹ R. Chang, 'The Possibility of Parity', *Ethics*, 112 (2002), pp. 659–88.

² J. Gert, 'Value and Parity', *Ethics*, 114 (2004), pp. 492–510; W. Rabinowicz, 'Value Relations', *Theoria*, 74 (2008), pp. 18–49; W. Rabinowicz, 'Incommensurability and Vagueness', *Aristotelian Society Supplementary Volume*, 83 (2009), pp. 71–94; W. Rabinowicz, 'Value Relations Revisited', *Economics and Philosophy*, 28 (2012), pp. 133–164.

erence relations preference, dispreference, indifference, and preferential gap holds.

In this paper, I shall argue for a symmetry between value and preference relations, namely, that for every value relation, there is a corresponding preference relation, and vice versa. The kinds of cases that have been taken as evidence of non-standard value relations seem to work equally well as evidence of non-standard preference relations. I claim that if the arguments that there are non-standard value relations are cogent, these arguments, *mutatis mutandis*, also show that there are non-standard preference relations. Hence frameworks of Gert and Rabinowicz's type are either inadequate since there are cogent arguments for both non-standard value and preference relations and these frameworks do not allow that, or they lack support since the arguments for non-standard value relations are unconvincing. As a substitute, I propose a simpler framework that allows for non-standard value and preference relations if the symmetry between value and preference relations holds. Finally, I shall try to dispel Gert and Rabinowicz's argument against this kind of framework.

1. Gert's and Rabinowicz's frameworks

According to the fitting-attitude account proposed by Franz Brentano and others, an object belongs to a certain axiological category if and only if it is fitting to have a certain attitude towards it or to act in certain ways in regard to it.³ On this approach, value relations between objects are determined by what preference relation is fitting to have towards the objects. Both Gert and Rabinowicz follow this approach and analyse value relations in terms of some kind of rationally permissible preferential attitudes.

Gert holds that for some items an agent is rationally permitted to have a preference with a range of different strengths for the item. These strengths of preference could be thought of as a measure of how much the agent wants the item. Gert then analyses value relations in terms of what ranges of strengths of preferences are rationally permissible. Betterness is analysed as follows:

³ See, e.g. F. Brentano, *The Origin of Our Knowledge of Right and Wrong* (Routledge and Kegan Paul, 1969), p. 18; A. C. Ewing, *The Definition of Good* (Macmillan, 1947), ch. 5; and T. M. Scanlon, *What We Owe to Each Other* (Harvard UP, 1998), pp. 95–100.

Range Rule: 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; otherwise the items are not traditionally comparable.⁴

p. 478

One counter-intuitive implication of Gert's Range Rule is that two items are 'traditionally comparable' if and only if one of them is better than the other. Yet two equally good items also seem traditionally comparable. Gert anticipates this worry and suggests that one modifies the Range Rule to allow for the following analysis of 'equally good':

[W]e can easily modify the rule to allow equality in the following circumstances: when two items each have the same unique rationally required strength of preference.⁵

This takes care of traditional comparability. In addition, Gert proposes that parity holds between items *A* and *B* in the following kind of case:

That *A* and *B* have exactly the same range. Thus, for any third item, *C*, the rational status of choosing *A* over *C*, or vice versa, will always be the same as the rational status of choosing *B* over *C*, or vice versa. This case might plausibly be called 'parity'.⁶

As pointed out by Rabinowicz, this account does not fit with the typical examples of parity.⁷ Suppose, for example, that a trip to California is on a par with a trip to Florida but the trip to California is worse than a California trip with an extra dollar.⁸ Still, it seems that the California trip with the extra dollar might not be better than the Florida trip. Yet, in Gert's framework, it must be better. Even more problematic is that Gert cannot account for the possibility that an object x^+ is better than an object x , an object y^+ is better than an object y , but x^+ is not better than y and y^+ is not better than x .⁹ Hence he cannot account for the possibility that both the California and the Florida trip would be better with an extra dollar

⁴ Gert, 'Value and Parity', p. 505.

⁵ Gert, 'Value and Parity', p. 506.

⁶ Gert, 'Value and Parity', p. 506.

⁷ Rabinowicz, 'Value Relations', p. 30.

⁸ The example is inspired by F. Restle, *Psychology of Judgment and Choice: A Theoretical Essay* (Wiley, 1961), pp. 62–63.

⁹ Rabinowicz, 'Value Relations', p. 33; S. Danielsson, 'Numerical Representations of Value-Orderings: Some Basic Problems', in C. Fehige and U. Wessels (eds.), *Preferences* (de Gruyter, 1998), p. 119.

but neither trip with the extra dollar is better than the other trip without the extra dollar. This is what motivates Rabinowicz to move to a slightly more complex intersection model instead.

Rather than in terms of permissible strengths of preference, Rabinowicz analyses value relations in terms of permissible preference orderings. Let K be the set of all permissible preference orderings. Rabinowicz assumes that in every ordering in K weak preference is a quasi-order, that is, it is reflexive and transitive.¹⁰ Given K , he defines ‘better’, ‘equally good’, ‘on a par’, and ‘incomparable’ as follows:

p. 479

x is *better* than y if and only if x is preferred to y in every ordering in K .¹¹

Two items are *equally good* if and only if they are equi-preferred in every ordering in K .¹²

x and y are *on a par* if and only if K contains two orderings such that x is preferred to y in one ordering and y is preferred to x in the other.¹³

x and y are *incomparable* if and only if every ordering in K contains a gap with regard to x and y .¹⁴

A preferential gap should be understood here as the absence of a preferential attitude.¹⁵ In addition to these fairly familiar value relations, Rabinowicz proposes some more exotic ones. In total, there are 15 atomic value relations in his framework.

2. Value-preference symmetry

Gert’s and Rabinowicz’s frameworks have room for more value relations than those mentioned above. But the mentioned relations suffice to see that in their frameworks there are value relations that lack a corresponding preference relation. This is because they both have the non-standard value relation parity but they have only the standard preference relations. Hence they violate the following principle, which I shall defend:

¹⁰ Rabinowicz, ‘Value Relations’, p. 37. For a comparison of types of orderings, see A. K. Sen, *Collective Choice and Social Welfare* (Holden-Day, 1970), p. 9.

¹¹ Rabinowicz, ‘Value Relations’, p. 38.

¹² Rabinowicz, ‘Value Relations’, p. 38.

¹³ Rabinowicz, ‘Value Relations’, p. 39.

¹⁴ Rabinowicz, ‘Value Relations’, p. 39.

¹⁵ See, Rabinowicz, ‘Value Relations’, p. 42.

Value-preference symmetry

For every value relation, there is a corresponding preference relation, and for every preference relation, there is a corresponding value relation.

The first part of my defence of this symmetry covers the standard dyadic value relations. Note that for each of these there is a preferential counterpart. They all have a corresponding dyadic preference relation that plays the same role preferentially as the value relation does axiologically:

p. 480

x better than y	preference for x over y
x worse than y	preference for y over x
x equally good as y	indifference between x and y
x and y are axiologically incomparable	preferential gap between x and y

Similarly, other relations that hold when certain combinations of the standard value relations hold, like *weakly better*, have preferential analogues that hold when the corresponding combinations of standard preference relations hold, in this case *weakly preferred*.¹⁶ Hence if the symmetry breaks down, it must be because of some non-standard value relation that lacks a corresponding preference relation, or vice versa.¹⁷ The second part of my defence of value-preference symmetry will be that for each remotely cogent argument for the existence of a non-standard value (pref-

¹⁶ I do not take a stand here whether ‘weakly better’ is a value relation or merely a disjunction of value relations. I claim merely that ‘weakly better’ is a value relation if and only if weak preference is preference relation (rather than merely a disjunction of preference and indifference).

¹⁷ One might object that preferential gaps are ruled out by some accounts of preference. If, for example, preferences are linked to choice behaviour in a certain sense, a case can be made for the claim that there are no preferential gaps. And likewise for some more mental accounts of preference. Suppose that ‘ P prefers x to y ’ means that P ’s hedonic tone would be raised by the news that x rather than y obtains, and ‘ P is indifferent between x and y ’ means that P ’s hedonic tone would be unaffected by the news. Again, this gives us an account of preference according to which there are no preferential gaps. If one of these accounts is correct and there is axiological incomparability, value-preference symmetry does not hold. But the standard arguments and examples given in support of axiological incomparability seem, *mutatis mutandis*, to support also preferential gaps. As I shall argue below, the stock argument for axiological incomparability—that is, the small-improvement argument—seems equally cogent when applied to preference relations as when applied to value relations. Hence it seems that if there is a cogent case for the possibility of axiological incomparability, there is a cogent case against accounts of preference that do not allow for preferential gaps. Furthermore, note that also Rabinowicz is in trouble if there is axiological incomparability but no preferential gaps.

erence) relation, there is an equally cogent analogous argument for the existence of a corresponding non-standard preference (value) relation.

To make a convincing case against value-preference symmetry, one needs a cogent argument for a value relation for which there does not exist a cogent parallel argument for a corresponding preference relation, or vice versa. The trouble for Gert and Rabinowicz is that the most prominent case for a value relation that holds when none of the standard value relations holds, the combination of the small-improvement argument and Chang's chaining argument, seems equally applicable for preference relations. At least Chang takes her case to be successful also for preference relations. She writes:

p. 481

Perhaps the most striking, the possibility of parity shows the basic assumption of standard decision and rational choice theory to be mistaken: preferring *X* to *Y*, preferring *Y* to *X*, and being indifferent between them do not span the conceptual space of choice attitudes one can have toward alternatives. Put another way, the "partial orderings" sometimes favoured by such theories will underdescribe the range of choice attitudes a rational agent can have toward alternatives.¹⁸

Unlike Chang, however, I shall defend neither the small-improvement argument nor the chaining argument. I shall just argue that these arguments are *equally* cogent in their axiological versions as in their preferential versions. The weaknesses of these arguments seem equally worrying when applied to preference relations as when applied to value relations.

The small-improvement argument was first proposed by Ronald de Sousa under the title 'the case of the Fairly Virtuous Wife'. He writes:

I tempt her to come away with me and spend an adulterous weekend in Cayucos, California. Imagine for simplicity of argument that my charm leaves her cold. The only inducement that makes her hesitate is money. I offer \$1,000 and she hesitates. Indeed she is so thoroughly hesitant that the classical decision theorist must conclude that she is *indifferent* between keeping her virtue for nothing and losing it in Cayucos for \$1,000. [...] The obvious thing for me to do now is to get her to the point of clear preference. That should be easy: everyone prefers \$1,500 to \$1,000, and since she is indifferent between virtue and \$1,000, she must prefer \$1,500 to virtue by exactly the same margin as she prefers \$1,500 to \$1,000: or so

¹⁸ Chang, 'The Possibility of Parity', p. 666.

the axioms of preference dictate. Yet she does not. As it turns out she is again ‘indifferent’ between the two alternatives. The classical Utilitarian is forced to say that she is incoherent, because she violates his axioms of rationality. [...] I would prefer to say that the alternatives considered are *incomparable*.¹⁹

In de Sousa’s original rendition, the argument is purely about preferences, but Chang and others also use axiological versions of the argument. The common structure of all versions of the small-improvement argument is as follows: first we have a premise that states some comparisons, which are supported by an intuitive example, and then we have a transitivity premise from which it follows that none of the standard comparative relations holds—i.e. in the axiological case none of ‘better’, ‘worse’, and ‘equally good’ holds and in the preferential case none of preference, dispreference, and indifference holds. For an axiological version of the argument, we can replace de Sousa’s example with, for instance, the following from Chang:

p. 482

Suppose you must determine which of a cup of coffee and a cup of tea tastes better to you. The coffee has a full-bodied, sharp, pungent taste, and the tea has a warm, soothing, fragrant taste. It is surely possible that you rationally judge that the cup of Sumatra Gold tastes neither better nor worse than the cup of Pearl Jasmine and that although a slightly more fragrant cup of the Jasmine would taste better than the original, the more fragrant Jasmine would not taste better than the cup of coffee.²⁰

The axiological version can be stated schematically as follows:

The small-improvement argument (axiological version)

- (S1) There exist things x , y , and z such that x is not better than y , y is not better than x , z is better than x , and z is not better than y .
- (S2) For all things x , y , and z , if x and y are equally good and z is better than x , then z is better than y .
- (S3) So, there exist things x and y such that x and y are not equally good, x is not better than y , and y is not better than x .

¹⁹ R. de Sousa, ‘The Good and the True’, *Mind*, 83 (1974), pp. 534–551, pp. 544–545.

²⁰ Chang, ‘The Possibility of Parity’, p. 669.

A stock objection to the small-improvement argument is that the support for (S1) seems to trade on some indeterminacy of value and preference relations.²¹ Consider two borderline cases of baldness, Smith and Jones, with different baldness patterns. Smith has more frontal recession, while Jones's hair loss is mostly concentrated to the vertex. Because of the difference between their types of hair loss, you are unwilling to judge that Smith is balder than Jones nor that Jones is balder than Smith. Suppose, furthermore, that Jones loses some more hair but you are still unwilling to judge that Jones is balder than Smith. In this case, which looks analogous to those above, it seems unwarranted to conclude that neither were Smith and Jones equally bald nor was one of them more bald than the other. The unwillingness for making judgements is not due to *at least as bald* being incomplete; it is merely a product of indeterminacy.

The objection is that in order for the small-improvement argument to work we need support for the following version of (S1):

(S1') There exist things x , y , and z such that it is determinate that x is not better than y , determinate that y is not better than x , determinate that z is better than x , and determinate that z is not better than y .

p. 483

But our hesitation or perplexity in the cases that are offered to support (S1) seems to be equally well explained by the following weaker version:

(S1'') There exist things x , y , and z such that it is not determinate that x is better than y , not determinate that y is better than x , determinate that z is better than x , and it is not determinate that z is better than y .

For instance, in de Sousa's example, the reason for the Virtuous Wife's hesitation might be that her preference between \$1,000 and virtue is indeterminate. Similarly, in Chang's example, you might not want to judge that the cup of Sumatra Gold tastes determinately neither better nor worse than the cup of Pearl Jasmine. You might want to judge instead that the cup of Sumatra Gold tastes neither determinately better nor determinately worse than the cup of Pearl Jasmine. But this would support just (S1'') and not (S1').

My aim here, however, is not to defend this objection but merely to note that it seems equally worrying for the preferential and axiological

²¹ Rabinowicz, 'Incommensurability and Vagueness', p. 74.

versions of the small-improvement argument. Just as there might be indeterminate values, there might also be indeterminate preferences. Were the preferential versions of these cases just cases of indeterminacy, it would seem odd if the same were not true for the axiological versions. The same goes for other attempts to explain away the preferential examples. If we, for example, were to take the preferential examples as just cases where the agent has a preferential gap between the objects, it would seem strange if one did not also interpret the analogous axiological examples as just cases where the objects are axiologically incomparable.²²

If successful, the small-improvement argument shows that there are pairs of objects between which none of ‘better’, ‘worse’, and ‘equally good’ holds. But this does not allow us to conclude that some of these are related by a non-standard value relation. All these pairs might just be incomparable to each other. To rule out this possibility, Chang proposes the chaining argument. The point of the chaining argument is to show that for some pair of objects between which none of the standard comparative relations holds the objects are still comparable, according to some intuitive notion of comparability. The argument runs as follows:²³

p. 484

The Chaining Argument

- (C1) If x and y are comparable and the respects relevant to the comparison between them can be balanced against one another and z is like y except for a small, unidimensional change, then x and z are comparable.
- (C2) There exist things x , y , and z such that x and y are comparable and none of the standard comparative relations holds between x and z and there is a continuum of small, unidimensional changes connecting y with z .
- (C3) So, there exist things x and y such that x and y are comparable and none of the standard relations holds between x and y .

²² J. E. Gustafsson and N. Espinoza, ‘Conflicting Reasons in the Small-Improvement Argument’, *The Philosophical Quarterly*, 60 (2010), pp. 754–763 offer an objection that affects only the preferential version of the small-improvement argument. Nevertheless, a modified version of the small-improvement argument, proposed in E. Carlson, ‘The Small-Improvement Argument Rescued’, *The Philosophical Quarterly*, 61 (2011), pp. 171–174, avoids this objection. Hence the balance of cogency between the preferential and the axiological versions of the argument is intact.

²³ Chang, ‘The Possibility of Parity’, pp. 673–675.

Chang illustrates the chaining argument with an example with Mozart and Michelangelo, who might be thought to be neither more or less creative than each other nor equally creative:

According to this principle [(C₁)], if Mozart is comparable with Talentless_i, then he is also comparable with Talentless_i+, for the difference between Talentless_i and Talentless_i+ is a small unidimensional one, and by hypothesis, such a difference can't trigger incomparability between evaluatively very different items where before they were comparable. And if Mozart is comparable with Talentless_i+, then applying the principle anew, it follows that he is comparable with Talentless_i++, and so on. Comparability with Mozart is preserved through the continuum of small unidimensional differences, and thus we arrive at the conclusion that Mozart is comparable with Michelangelo. By hypothesis, Mozart is not more or less creative than Michelangelo, and nor are the two equally creative. And yet it seems that they are nevertheless comparable.²⁴

Like the small-improvement argument, however, the chaining argument does not seem more cogent in an axiological rendition than in a preferential one. Given that we grant the cogency of the preferential and the axiological versions of the small-improvement argument, (C₂) seems fairly plausible. Nonetheless, a common objection to the chaining argument is directed towards the small-unidimensional-difference principle, that is, (C₁). This principle conflicts with a version of the Pareto rule that says:

[U]nless we are equally well-off in each of two states of affairs, one state is better than another if at least one of us is better off than we would be in the other state and none of us is worse off, otherwise the states are incomparable.²⁵

For example, consider a state *a* of two individuals whose respective well-being is given by the ordered pair (2, 2) and a state *b* of the same individuals at (1, 2). The Pareto rule yields that *a* is better than *b* and, thus, comparable. A third state *c* with the individuals at (2, 1) is like *a* except for a small, unidimensional change.²⁶ But the Pareto rule yields that *b* and *c* are incomparable. Thus we have a counter-example to (C₁). To avoid

p. 485

²⁴ Chang, 'The Possibility of Parity', p. 674.

²⁵ Chang, 'The Possibility of Parity', p. 676.

²⁶ If you do not think that this change is small enough, note that the case would still work if we instead let *c* be the state of the individuals at (2, 2 - ϵ) where ϵ is an arbitrarily small positive number.

these kinds of counter-examples to the axiological version of the chaining argument, one might limit the scope of (C1). This is also Chang's move, which in turn raises worries about ad hocness. The same objection can also be mooted against a preferential version of the chaining argument, based on similarly troublesome trade-offs.²⁷

Another reason one might think that the combination of the small-improvement argument and the chaining argument works only in its axiological version is that one takes preferences to be value judgements.²⁸ Perhaps the hesitation or perplexity in the cases that are appealed to in the preferential version of the small-improvement argument is just due to uncertainty about what value relation holds between the objects. But even if we were to grant this conception of preference relations, it would not undermine value-preference symmetry. If preferences are just value judgements, it is trivial that every value relation has a corresponding preference relation, i.e. the judgement that the value relation holds. So, if the these arguments cogently show that there are non-standard value relations, each of these would have a corresponding non-standard preference relation.

This kind of value-preference symmetry is a problem for Gert and Rabinowicz since if the case for a non-standard value relation that might hold when none of the standard value relations holds is cogent if and only if a parallel case for a corresponding non-standard preference relation is cogent, then *either* their extended frameworks are unmotivated since there is no reason to think there are non-standard value relations *or* their frameworks are inadequate since there are non-standard preference relations and their frameworks do not allow that.

A related trouble is how to allow for axiologically conscientious, rational agents whose preferences are guided by their values. It does not seem irrational to let one's preferences be guided by one's values in the sense that one does not hold a certain preference between some objects unless the corresponding value relation holds between them. We should accept what we can call

p. 486

²⁷ Note that even if this problem were unique to the axiological version, this would not help Gert and Rabinowicz. This is because they need a cogent case for axiological parity that does not, *mutatis mutandis*, support also preferential parity—not the other way around.

²⁸ See, e.g. K. J. Arrow, *Social Choice and Individual Values* (Wiley, 1951), p. 23 and A. K. Sen, *Rationality and Freedom* (Harvard UP, 2002), p. 590.

Non-irrationality of axiological constrainedness

If none of a set of value relations holds between x and y , then one is not rationally required to have one of the corresponding preference relations.

The idea is that if, for example, Mozart is not better than Michelangelo, you should be rationally permitted not to prefer Mozart to Michelangelo. A reasonable practice seems to be to align one's preferences to one's values so one does not prefer something if it is not better. So, if Mozart is neither better nor worse than Michelangelo, you should be rationally permitted not to prefer either of them to the other.²⁹ Likewise, if no standard value relation holds between Mozart and Michelangelo, you should not be rationally required to have one of the standard preference relations.

The trouble is that Gert's and Rabinowicz's frameworks violate the above principle. Even if no standard value relation holds between two objects, you are, nonetheless, rationally required to hold one of the standard preference relations between them. On Rabinowicz's account, for example, the standard preference relations are exhaustive—for any pair of objects x and y , you cannot avoid holding one of these relations between x and y —and, no matter which value relation holds, some of the standard preference relations will be the only ones that are rationally permitted. If we have non-standard value relations and wish to allow for non-irrational axiological constrainedness, we need to also have some non-standard preference relations.

At this point, one might wonder whether these worries are so problematic since Gert and Rabinowicz may just extend their frameworks to also accommodate some non-standard preference relations. But it is hard to see how one could extend their frameworks to allow for some non-standard preference relations without an even greater increase in possible value relations. Hence it seems like their frameworks would still violate value-preference symmetry after an extension. Furthermore, if value-preference symmetry holds, there is, as we shall see in the next section, a much simpler fitting-attitude analysis available.

p. 487

²⁹ This is at odds with Rabinowicz's account in Rabinowicz, 'Value Relations', p. 42 if his value relation *type 6* holds between the artists. According to him, *type 6* holds between x and y if and only if it is rationally permissible to prefer x to y , rationally permissible to prefer y to x , and rationally required to prefer one of them to the other.

3. A simpler approach

If we accept value-preference symmetry, as I think we should, and still wish to analyse some non-standard value relations, there is no need for complex frameworks like those of Gert and Rabinowicz. If we have value-preference symmetry, we can simply stick to a one-to-one pairing between value relations and their corresponding preference relations. In what follows, ‘fitting’ will be taken to be like ‘requirement’ in its normative strength rather than ‘permission’. That is, if it is fitting to prefer x to y , it cannot also be fitting not to prefer x to y . I propose the following general fitting-attitude analysis:³⁰

An atomic value relation V obtains if and only if the preference relation corresponding to V is fitting.³¹

Following this scheme, we can analyse the dyadic value relations we have considered as follows:

x is *better* than y if and only if it is fitting to prefer x to y .

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

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

x and y are *axiologically on a par* if and only if it is fitting to hold x and y preferentially on a par.

x and y are *axiologically incomparable* if and only if it is fitting to have a preferential gap between x and y .

³⁰ Note, however, that the thesis of value-preference symmetry does not need to be combined with some fitting-attitude account. As I argued earlier, value-preference symmetry also holds given some other accounts of the relation between preferences and values. It holds, for example, trivially on an account where preferences are just value judgements or on some subjectivist account where values are just preferences.

³¹ This account does not cover molecular value relations, e.g. disjunctions of value relations. It is not obvious, however, which relations should be counted as atomic. Take, for example, ‘at least as good’. If we grant it as atomic along with ‘better’ and ‘equally good’, we have to give up its standard interdefinability with these relations. Rabinowicz, ‘Value Relations’, p. 43 does not seem to have any qualms about giving up this interdefinability. Yet I wish to resist this move since it would considerably raise the price of adopting the approach. Better to leave open the possibility that ‘at least as good’ is just the molecular value relation ‘either better or equally good’.

This approach also provides an explanation of why a certain value relation holds when a certain preference relation is fitting. For example, on a standard fitting-attitude account x is good if and only if a pro attitude is fitting towards x . But why is goodness analysed in this manner in terms of a pro attitude rather than, for example, an anti attitude? A natural explanation is that this is because goodness is a pro value relation rather than an anti value relation. Goodness plays the same role axiologically as pro attitudes do preferentially. The above scheme generalizes this idea.

One might object that one feature of Gert's and Rabinowicz's frameworks is missing in mine. Their frameworks, unlike mine, provide an explanation of non-standard value relations in terms of the familiar standard preference relations. For example, if there is a relation such as axiological parity, my framework just analyses it in terms of preferential parity. We are left wondering what parity is in general. But neither Gert's nor Rabinowicz's account can explain parity in general since their accounts cannot be extended to explain preferential parity. Unlike Gert's and Rabinowicz's frameworks, my framework does not rule out that one can provide general explanations of such non-standard relations. Still, it does put some constraints on these explanations. It requires that the general explanation of a non-standard relation explains not only the axiological relation but also the corresponding preference relation. Explaining preference relations, however, falls outside of the scope of a fitting-attitude analysis of value.³²

4. Strong and weak levels of normative strength

Perhaps Gert would object that accounts like mine, where fitting is taken to be like requirement in its normative strength, link values and preferences too closely. I have yet to take into account his argument for the need for an analysis in terms of fittingness with two levels of normative strength—one strong, interpreted as what is rationally required, and one weak, interpreted as what is rationally permissible. He writes:

[O]nly 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

³² For an account of some non-standard preference relations, see J. E. Gustafsson, 'An Extended Framework for Preference Relations', *Economics and Philosophy*, 27 (2011), pp. 360–367.

have the same perfectly precise standards for assessing the value of items with respect to V, and who take the same interest in whether or not something has value V, could make different, but equally rational choices between two items, when the relevant value is value V.³³

p. 489

Similarly, Rabinowicz writes concerning his definition of parity, which holds between two objects if and only if it is permissible to prefer one of them and also permissible to prefer the other,

Note that this definition of parity is possible only because on the present approach we are supposed to distinguish between two levels of normativity: the strong level and the weak one. Gert's analysis of value relations in terms of rationally warranted preferences makes room for parity because warrant can be interpreted either strongly, as a requirement, or weakly, as a permission.³⁴

They seem to argue that there are cases with two objects x and y such that it is rationally permissible to prefer x to y and also rationally permissible to prefer y to x . And an approach like mine that analyses value relations with just a strong level of fitting cannot account for such cases.

This argument, however, does not seem entirely successful. One way out would be to interpret the strong level of fitting in some other way than as 'rationally required'.³⁵ But let us, for the sake of the argument, grant

³³ Gert, 'Value and Parity', pp. 494–495.

³⁴ Rabinowicz, 'Value Relations', p. 30.

³⁵ There might also be independent reasons to do so. Any theory that yields that x is better than y if and only if it is rationally required to prefer x to y , is vulnerable to a version of the well-known wrong-kind-of-reason problem. R. Crisp, 'Review of *Value ... And What Follows* By Joel Kupperman', *Philosophy*, 75 (2000), pp. 458–462, p. 459 presents the following example:

Imagine that an evil demon will inflict severe pain on me unless I prefer this saucer of mud; that makes the saucer well worth preferring. But it would not be plausible to claim that the saucer of mud's existence is, in itself, valuable; rather, my pain will be 'disvaluable'.

One can easily construct a variation where the demon demands instead that one prefers the saucer to something that should be more valuable. Yet, since the demon will inflict one with severe pain unless one prefers the saucer, it seems that one should be rationally required to prefer the saucer. So, if we hold that something is better if and only if it is rationally required to prefer it, we have to conclude that the saucer is better. A standard response to the wrong-kind-of-reason problem is to drop the analysis of what is fitting as what is rationally required. This is also more in line with Brentano's original fitting-attitude analysis of betterness in Brentano, *The Origin of Our Knowledge of Right and Wrong*, p. 26, i.e. that x is better than y if and only if it is *correct* to prefer x to y . For a re-

that ‘fitting’ in my scheme above should be read as ‘rationally required’.³⁶ An implicit assumption in Gert and Rabinowicz’s argument seems to be that if one preference relation is rationally required, no other preference relation is rationally permissible. But there seems to be no reason to accept this. There might be preference relations—for example, *preferentially on a par*—that do not rule out that other preference relations hold. Even though Chang’s technical definition of parity rules out that any of the standard relations holds, this does not seem to be the case with our ordinary, pre-theoretic notion of parity. If, for example, Mozart and Michelangelo are equally good, it seems counter-intuitive that they would not be on a par. Similarly, it seems natural to allow that Mozart and Michelangelo are on a par even if one of them really is slightly better than the other. The same should hold for preferential parity. To avoid confusion, we may call this pre-theoretic notion *weak parity* and call Chang’s notion *strong parity*. On my approach, we then have the following value relation:

p. 490

x and y are *axiologically weakly on a par* if and only if it is fitting to hold x and y preferentially weakly on a par.

Note, furthermore, that in order for Gert’s and Rabinowicz’s own frameworks to be compatible with situations like the ones they describe above, these situations must involve two objects between which none of ‘better’, ‘worse’, and ‘equally good’ holds. In a standard example of such a case, the comparison of Mozart with Michelangelo, it still seems plausible to claim that Mozart and Michelangelo are, nonetheless, axiologically on a par. Similar points can also be made for the other stock examples. So, given that this kind of case does not rule out also Gert’s and Rabinowicz’s frameworks, we should grant that objects can be axiologically on a par while none of ‘better’, ‘worse’, and ‘equally good’ holds between them.

But then we may account for the situation Gert and Rabinowicz describe, as a case where two objects x and y are axiologically weakly on

cent defence of a Brentano-styled response to the wrong-kind-of-reason problem, see S. Danielsson and J. Olson, ‘Brentano and the Buck-Passers’, *Mind*, 116 (2007), pp. 511–522. Should we, for example, adopt Brentano’s approach, we could say that even though one is rationally required to prefer the saucer it would not be correct to do so. We would then avoid the unwanted conclusion that the saucer is better. Given this kind of response to the wrong-kind-of-reason problem, my proposal no longer implies anything about what is rationally required and hence it is immune to Gert’s objection.

³⁶ That is, we assume that it expresses a *rational* requirement, rather than some other type of requirement.

a par but none of ‘better’, ‘worse’, and ‘equally good’ holds between them. Since neither ‘better’ nor ‘worse’ holds, one is not rationally required to prefer x to y and one is not rationally required to prefer y to x . Yet one might be both rationally permitted to prefer x to y and rationally permitted to prefer y to x since this is compatible with having the rationally required preference relation, that is, to hold x and y preferentially weakly on a par. The upshot is that we may thus handle this kind of case without giving up the analysis of value relations just in terms of rationally required preference relations.

In summation, Gert’s and Rabinowicz’s frameworks are asymmetric in that they have more value relations than preference relations. This asymmetry seems untenable since the arguments in support of non-standard value relations, which are needed to motivate their extended frameworks, seem to work equally well, *mutatis mutandis*, in support of non-standard preference relations. The general fitting-attitude analysis presented in this paper satisfies value-preference symmetry and handles values and preferences congruently. Furthermore, it does not deem irrational what I have called ‘axiological constrainedness’. It also offers an explanation of why a certain value relation rather than another holds when a certain preference relation is fitting. Finally, my framework can also handle the cases with multiple rationally permitted preferences, which seem to have motivated Gert’s and Rabinowicz’s analyses in terms of a strong and a weak level of fitting.

Thanks to Gustaf Arrhenius, John Cantwell, Marc Fleurbaey, Sven Ove Hansson, Martin Peterson, Wlodek Rabinowicz, the participants of the Paris Seminar on Economics and Philosophy, and two anonymous referees for this journal for valuable comments.