

Reply to Horta: Spectrum Arguments, the “Unhelpfulness” of Rejecting Transitivity, and Implications for Moral Realism

LARRY TEMKIN
Rutgers University

Abstract

This article responds to Oscar Horta’s article “In Defense of the Internal Aspects View: Person-Affecting Reasons, Spectrum Arguments and Inconsistent Intuitions”. I begin by noting various points of agreement with Horta. I agree that the “better than relation” is asymmetric, and point out that this will be so on an Essentially Comparative View as well as on an Internal Aspects View. I also agree that there are various possible *Person-Affecting Principles*, other than the one my book focuses on, that people might find plausible, and that in some circumstances, at least, these might have deontological, rather than axiological significance. In particular, I grant that Horta’s *Actuality-Dependent Person-Affecting Principle*, his *Time-Dependent Person-Affecting Principle*, and his *Identity-Dependent Person-Affecting Principle*, might each be relevant to what we ought to *do*, without necessarily being relevant to which of two outcomes is *better*. But I reject Horta’s claim that essentially comparative principles don’t apply in Spectrum Arguments. I also argue against Horta’s view that the two Standard Views that underlie our intuitions in Spectrum Arguments are contradictory. I question Horta’s (seeming) position that there is no point in rejecting the transitivity of the “better than” relation on the basis of Spectrum Arguments, on the grounds that doing so won’t *solve* the predicament that Spectrum Arguments pose. Finally, I conclude my paper by challenging Horta’s interesting contention that my views about nontransitivity support an anti-realist metaethics, and are incompatible with the sort of realist approach to metaethics that I favor.

Keywords: Transitivity, Spectrum Arguments, Person-Affecting Principles, Internal Aspects View, Better than, Essentially Comparative View, Realism, Anti-realism, Sophie’s Choice, Moral Dilemmas.

Let me begin by acknowledging my gratitude to Oscar Horta for his thoughtful and sensitive comments in his article “In Defense of the Internal Aspects View: Person-Affecting Reasons, Spectrum Arguments and Inconsistent Intuitions” (Horta, 2014), and also for the generous spirit he displayed in presenting them. I will divide my responses into two main parts. In part I, I will note some points where I agree with Horta. In part II, I will note some points about which we disagree.¹

1.

In section 3, Horta discusses different possible types of Person-Affecting Principles, and the question of whether better than can be non-asymmetric. As Horta recognizes, the points he makes in this section are not in tension with my book’s claims. But I agree that they offer useful lessons to bear in mind as we try to determine what needs to be said about the “better than” relation and other analogous relations.

One of Horta’s main claims in section 3 is that the “better than” relation is asymmetric: so if, in any given context, A is better than B, *all things considered*, then it *can’t* also be the case that, in that very same context, B is better than A, all things considered. Even I, who am open to rejecting the Axiom of Transitivity for the “better than” relation, don’t reject the fact that the better than relation is asymmetric! Similarly, while I am open to rejecting the Axiom of Transitivity for the “equally as good as” relation, I accept the standard view that the “equally as good as” relation is symmetric: so if, in any given context, A is equally as good as B, *all things considered*, then it must

1. This article was originally written in response to the talk that Horta presented at the LEAP symposium on my book, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Temkin 2012), at Pompeu Fabra University in Fall 2012, and to an early draft of his article based on that talk. Unfortunately, shortly before this journal was to go to press, I received the final, revised, version of Horta’s article, and I was somewhat surprised to see that he had substantially revised his article, both in terms of adding new material that I hadn’t previously seen, and deemphasizing, or removing, some key claims or passages to which my original article had objected. I have, where possible, adjusted my article in light of Horta’s final changes. However, given the journal’s time constraints, there were some important aspects of the final version of Horta’s article to which I was unable to respond. In addition, in some cases I have thought it worthwhile to retain points that I raised with respect to his original talk and draft, since others may be attracted to views similar to those he previously held, even if Horta, himself, has now changed his mind on the matters. Thus, I acknowledge, here, that some of the points that I will be making are less relevant, or even not relevant, to his published article, and that, in some cases, when I highlight a supposed disagreement between us, the disagreement between us may have lessened, or disappeared altogether, since Horta first presented his views on these topics. To aid the reader, I try to make it plain in the text where I am mainly responding to his talk, as opposed to his published article.

also be the case that, in that very same context, B is equally as good as A, all things considered.

To these uncontroversial claims, I would simply point out that the explanation I give for *why* the various Axioms of Transitivity could fail to hold, does not similarly challenge these other standard claims. In challenging the Axioms of Transitivity, I noted that some of the ideals people most value are *Essentially Comparative*. On an *Essentially Comparative View*, the factors that are relevant and significant for assessing an outcome may vary depending on the alternative with which it is compared. This opens up the possibility that, in any given context, the factors that are relevant and significant for comparing A with B, or B with C, may be different from the factors that are relevant and significant for comparing A with C. From this it follows that, in a given context, A might be better than (or equally as good as) B, in terms of *all* of the factors that are relevant and significant for making *that* comparison, and B might be better than (or equally as good as) C, in terms of *all* of the factors that are relevant and significant for making *that* comparison, and yet A might *not* be better than (or equally as good as) C, in terms of *all* of the factors that are relevant and significant for making *that* comparison. Thus, I claimed that on an Essentially Comparative View of ideals, the Axioms of Transitivity regarding the “better than” and “equally as good as” relations may fail, or fail to apply across different sets of alternatives to which we might have thought they should apply.²

However, as indicated, the explanation I give for *why* the Axioms of Transitivity may fail, or fail to apply, offers no reason to doubt the asymmetry of the “better than” relation, or the symmetry of the “equally as good as” relation. This is because whether one accepts an *Internal Aspects View* (according to which how good an outcome is depends solely on the internal features of that outcome), or an Essentially Comparative View of ideals, if, in a given context, A is better than B in terms of *all* of the factors that are relevant and significant for comparing A and B in that context, then it *will* be the case that B is *worse* than (and hence *not* better than!) A in terms *all* of the factors that are relevant and significant for comparing A and B in that

2. Some people believe that on an Essentially Comparative View the Axioms of Transitivity fail to hold, so that they should be rejected. Others insist that there are various ways of preserving the Axiom of Transitivity even on an Essentially Comparative View, so that they never fail, but that the Axioms of Transitivity may *fail to apply* in those cases where Essentially Comparative ideals are relevant for assessing different alternatives. In my book, I use the notion of *non-transitivity* to cover both the cases where we think the Axioms of Transitivity fail, and those where we think they fail to apply across different sets of alternatives to which we might have expected that they should apply. I argue that there are significant practical and theoretical implications of the “all-things-considered better than”, “equally as good as”, and “at least as good as” relations being non-transitive, whether or not this is because the relations fail, or “merely” because they fail to apply in the cases I discuss (Temkin, 2014: 85, note 13). See Temkin 2012: 5, 16-8, 59-60, 66, 163-182, 197-8, 203-214, 223-5, and ch. 13.

context. Thus, “better than” is asymmetric, precisely as Horta has claimed and as I, and everyone else, should readily accept.

Similarly, whether one accepts an Internal Aspects View or an Essentially Comparative View of ideals, if, in a given context, A is equally as good as B in terms of *all* of the factors that are relevant and significant for comparing A and B in that context, then it *will* be the case that B is equally as good as A in terms *all* of the factors that are relevant and significant for comparing A and B in that context. Hence, the “equally as good as” relation is symmetric.

In sum, I have offered an account of why the Axioms of Transitivity might fail to hold, but that account does nothing to challenge the uncontroversial claims that “better than” and “equally as good as” are asymmetric and symmetric relations, respectively.

Let me turn next to Horta’s discussion of different Person-Affecting Principles. Horta distinguishes between three different kinds of Person-Affecting Principles:

The Actuality-Dependent Person-Affecting Principle: In assessing possible outcomes, one should focus on the status of those who exist in the actual world, with the aim of wanting them to be as well off as possible, and (2) ignore the status of those who do not exist in the actual world, except that one wants to avoid harming them as much as possible.

The Time-Dependent Person-Affecting Principle: In assessing possible outcomes that haven’t occurred yet, one should focus on those who will exist in the outcome that will occur first, with the aim of wanting them to be as well off as possible, and (2) ignore the status of those who will not exist in the outcome that will occur first, except that one wants to avoid harming them as much as possible.

The Identity-Dependent Person-Affecting Principle: In assessing possible outcomes, one should focus on the status of those whose identity is already determined, with the aim of wanting them to be as well off as possible, and (2) ignore the status of those whose identity is not determined yet, except that one wants to avoid harming them as much as possible (Horta 2014: 96).

I should mention that none of Horta’s versions of Person-Affecting Principles correspond exactly to the two versions I discuss in my book, my *Narrow Person-Affecting View* and my *Wide Person-Affecting View* (Temkin 2012: 416-45), but this need not concern us here.

In his original talk, Horta contended that while positions like the Actuality-Dependent, Time-Dependent, and Identity-Dependent Person-Affecting Principles might be relevant, in certain circumstances, to our assessment of what we ought or ought not to *do* in choosing between two alternatives, it doesn’t follow that such judgments are tracking which of the two alternatives, considered just by *themselves*, is *better*, *per se*. Here, too, I agree with Horta.

Consider, for example, Diagram One.




Diagram One

As drawn, Diagram One represents two possible outcomes, Q and R. Each outcome contains two groups, a better-off group and a worse-off group. There is no difference between the levels or number of people in the two better off groups, and similarly, for the two worse-off groups. The only difference concerns the *identities* of the better- and worse-off groups in the two outcomes. In Q, the A people occupy the better-off group and the B people occupy the worse-off group; in R, the reverse is true.

Assuming there are no morally relevant differences between the A and B people—so, for example, each person is equally talented, hardworking, deserving, and so on—it seems clear that, considered just by *themselves*, Q and R are *equally good*. But now, suppose that there were a presently existing *actual* outcome, with people whose identities were already determined, and it looked like P in Diagram Two.




Diagram Two

In P, a presently existing actual outcome, the A people already exist, and they are at a level between that of the better- and worse-off groups in possible

future outcomes Q and R. Suppose, next, that we could transform the P outcome into one like Q or R. That is, we could either raise the A people up, and bring another group of people, B, into existence at a lower level, so as to produce an outcome like Q, or, alternatively, we could bring another group of people, B, into existence at a high level, but at the cost of lowering the A people, so as to produce an outcome like R.

In accordance with Horta's Actuality-Dependent, Time-Dependent, and Identity-Dependent Person-Affecting Principles, it would be permissible, desirable, and perhaps even *obligatory* to bring about Q, and impermissible, undesirable, and prohibited to bring about R. But even if this showed that P would be *improved* by being changed into Q and *worsened* by being changed into R, that *doesn't* entail that Q, considered just by *itself*, is a worse outcome than R, considered just by *itself*. This should be evident, if one considers the fact that instead of the third alternative P, being a presently existing actual outcome, it *might* have been a fourth alternative, O, that was a presently existing actual outcome, where O was *just like* P except that instead of the A people existing at a level between the better- and worse-off groups in Q and R, the B people existed at that level. In that case, it would have been true that on all three of Horta's Dependent Person-Affecting Principles, O would be *improved* by being changed into R, and it would be *worsened* by being changed into Q. Thus, we would have to abandon the non-asymmetry of the "better than" relation *if* we thought that we *could* infer that Q was in *itself* better than R, simply from the fact that it *would* be desirable to transform P into Q, but undesirable to transform P into R; because the same reasoning would then entail that R was in *itself* better than Q, since it would *also* be desirable to transform O into R, but undesirable to transform O into Q.

A fortiori, as Horta contended in his talk, even if there are cases where in accordance with any of his Dependent Person-Affecting Principles it would be obligatory to bring about one outcome, O_1 , rather than another outcome, O_2 , it doesn't follow that the judgments yielded by such principles support the conclusion that, considered just by *themselves*, O_1 is better than O_2 . However, let me add that Horta's position is not merely compatible with my claims in *Rethinking the Good*, it follows directly from the fact that Horta's three Person-Affecting Views are Essentially Comparative as I characterized that notion. As noted previously, on an Essentially Comparative View, different factors can be relevant and significant for assessing the relative goodness of outcomes like Q or R, depending on the alternatives with which they are compared. Thus, Q might have one value in comparison with R when those are the only alternatives, but a different value in comparison with R when each is itself an alternative to some third outcome P, and a still different value in comparison with R when each is itself an alternative to some fourth outcome O.

2.

Let me turn next to some points of disagreement with Horta, some of which are fairly minor, but others of which are not.

First, in discussing my Spectrum Arguments, Horta titles his subsection 4.1 “*No essentially comparative principle applies in Spectrum Arguments*”, and he calls special attention to the claim “that appeals to person-affecting reasons play no role here” (Horta, 2014: 99). I agree with Horta’s claim that Person-Affecting Views do not underlie my Spectrum Arguments.³ However, I take exception to the title of his subsection, which, I believe, is either misleading or mistaken.

As I have presented and analyzed them, Spectrum Arguments *do* arise because of the *Essentially Comparative View* of ideals. Consider, for example, my Spectrum Argument where the first member of the Spectrum involves a very long life with 15 mosquito bites per month and two years of torture, and the last member involves a very long life with 16 mosquito bites per month but no torture. I claimed that two distinct views guided our thinking in making different comparisons along the Spectrum. The *First Standard View* reflects an *Additive-Aggregationist Approach* and the *Second Standard View* reflects an *Anti-Additive-Aggregationist Approach*. Both Views are limited in scope, in that they seem relevant and significant for making certain comparisons but not others. In particular, I pointed out that where the differences in the intensity of pains between two alternatives was very small, the First Standard View seemed plausible and appropriate for comparing those alternatives, so, in particular, it seemed appropriate for comparing my Spectrum’s first alternative with the second, the second with the third, the third with the fourth, and so on. On the other hand, where the differences in the intensity of pains between two alternatives was very large, the Second Standard View seemed plausible and appropriate for comparing *those* alternatives, so, in particular, it seemed appropriate for comparing my Spectrum’s first few alternatives with its last few alternatives. Thus, the factors that seemed relevant and significant for assessing the Spectrum’s first alternative were *different* depending on whether it was being compared with the Spectrum’s second alternative or its last alternative. This is in keeping with the *Essentially Comparative View* of ideals, and is at odds with the *Internal Aspects View* of ideals (see Temkin 2014: section 3, 71; and Temkin 2012: 62-6, 229-31, 369-74).

In sum, while it may be true that there is no *single* Essentially Comparative ideal like a Person-Affecting View that underlies my Spectrum Arguments, I think it *is* true that an Essentially Comparative approach best explains

3. However, like the Spectrum Arguments, Person-Affecting Views challenge the Axioms of Transitivity. See Temkin 2014: sec. 5; and Temkin 2012: ch. 12.

what is going on in my Spectrum Arguments, and why they ultimately put pressure on the Axioms of Transitivity.

Let me turn to a second point. Horta suggests that my two Standard Views are *contradictory*.⁴ I deny this.

The First and Second Standard Views would be *contradictory* if there were any particular judgments which the First Standard View made which were denied by the Second, or vice versa. But this is not, I think, the case. For example, where the First Standard View yields the judgment that the first member of my Spectrum is better than the second, the Second Standard View doesn't *deny* this judgment, rather it is silent. Specifically, as I characterized it in my book, the Second Standard View simply doesn't *apply* for comparisons involving such alternatives. Likewise, where the Second Standard View yields the judgment that the Spectrum's first member is worse than the last, the First Standard View doesn't *deny* this judgment, rather it is silent. Again, as I characterized it, the First Standard View simply doesn't *apply* for comparisons involving such alternatives.

Note, there would be nothing *contradictory* about the claims that John is in love with Mary and that John isn't in love with Tim, *even if* Mary is in love with Tim. There would also be nothing problematic about such claims, and this for the simple reason that "is in love with" isn't a transitive relation! Likewise, there will be nothing *contradictory* or even problematic about the claim that Spectrum's first outcome is better than the second, but the first is *not* better than the last, *even if*, for each pair of adjacent outcomes n and $n + 1$ along the Spectrum, n is better than $n + 1$, *as long as* "all-things-considered better than" *isn't* a transitive relation.

But, of course, *if*, as many believe, the First Standard View *is* relevant and significant for comparing adjacent outcomes along my Spectrum, but doesn't apply for comparing the first and last outcomes, and *if*, as many also believe, the Second Standard View *is* relevant and significant for comparing my Spectrum's first and last outcomes, but doesn't apply for comparing my Spectrum's adjacent outcomes, then there is good reason to believe that "all-things-considered better than" *isn't* a transitive relation, in which case the First and Second Standard Views won't be contradictory, and neither will their respective judgments that the Spectrum's first outcome is better than the second, but not better than the last.

Third, Horta suggests that there is a powerful reason to resist being driven to "reject the conclusion that transitivity does not apply to [Spectrum cases, since]... rejecting that betterness is transitive *does not solve the problems*

4. Horta was clearer about this in his talk, than he is in his article, where he seems to deemphasize this claim. However, he still seems to hold the view in question. See, for example, note 9 of Horta 2014: 100, where he writes "Note that the *contradictions* between the applications of the [Spectrum Arguments'] principles sometimes works in different directions... (emphasis added)".

implied by the conflict between different standard views (emphasis added)” (Horta 2014: 107). Later, Horta emphasizes that “*rejecting transitivity fails to facilitate a solution*. In fact, it make it harder, if not impossible, to do so” (Horta 2014: 104). In essence, then, Horta believes that there is no *point* in rejecting transitivity on the basis of Spectrum Arguments, if doing so won’t help us to *solve* such arguments.⁵

To a large extent, I agree with these remarks. Indeed, I emphasize some of these very same sentiments in my book, when I’m discussing the costs and benefits of accepting or rejecting the different positions underlying my impossibility arguments. But, I’m not sure, exactly, what is supposed to *follow* from such observations.

I am a philosopher seeking the truth. We would very much *like* to be able to answer certain questions in a certain way. But what if the sad truth is that the answers we seek are not to be found, or cannot be answered in the way we had thought or hoped. Is it not an important advance in our understanding of the normative realm if we learn that this is so?

I am reminded here of the Socratic claim regarding why the Oracle called him the wisest of all men. It was, Socrates claimed, because whereas most people thought that they knew a lot, and were wrong, he knew that he knew nothing (other than the fact that he knew nothing!).⁶ As Socrates recognized, it can be as important to know what we don’t know, and to learn what we can’t know, as to continue to seek solutions along a path where they can’t be found.

This raises a related point. Many years ago, when I first began thinking about and teaching these issues, Carl Hofer, who was then still an undergraduate, was quite pleased by my results, and the implications he took them to have. Hofer was worried about the dominance of consequentialist reasoning in much of contemporary normative reasoning. At the time, Hofer thought it would be better if moral philosophers spent more time focusing on considerations of character, of the sort championed by Aristotle, or on deontological considerations, of the sort championed by Kant. For Hofer, my results suggested a vindication of sorts for those who thought that a focus on consequences, and in particular on the aim of bringing about the best available outcome by maximizing the good, was the wrong way to proceed in determining how we ought to act, morally.

Now I don’t abandon the appeal to consequences lightly, if at all. Indeed, as I emphasize in my book, I don’t really even see how that could be a possible option, and I am certain that there would be *enormous* costs to such a move. Moreover, and more troubling, I’m not sure that virtue-based or deontological-based reasoning can wholly avoid the sorts of worries that

5. Horta was, I think, even clearer about this in his talk, where he really emphasized this point.

6. The position in question is often attributed to Socrates on the basis of a key passage (21d) of Plato’s *Apology*.

arise in my book from Spectrum Arguments or an Essentially Comparative View of ideals.⁷ But having said all that, the mere fact that abandoning the transitivity of the “all-things-considered better than” relation wouldn’t help us to decide what to do if, for example, we found ourselves facing a Spectrum-type choice and we wanted to bring about the best available outcome, *doesn’t* show us that transitivity *shouldn’t* be abandoned. Perhaps, reluctantly, it should. However intuitively unpalatable, this is an alternative that requires careful consideration. Indeed, perhaps taking such an option seriously will force us to pursue other paths in our exploration of the normative realm that may ultimately prove to be more fruitful than the paths on which most moral philosophers have focused up until now.

Let me conclude my response to Horta with some comments on the topic of section 6 of his article, moral realism.

As I note in my book, Derek Parfit once claimed that *if* my arguments were sound, they amounted to the most skeptical argument against moral realism since David Hume’s arguments.⁸ Since neither Parfit nor I are moral skeptics, that gave us both reasons to *hope* that my arguments aren’t sound—at least, if Parfit were right in his assessment about their implications. But, of course, hoping doesn’t make it so! Moreover, as my previous comment suggests, how devastating my results may prove to be will ultimately turn on whether other fruitful paths in the normative realm might be found that do not fall victim to my arguments.

A key question to be addressed concerns how much of the normative realm depends on our being able to provide a coherent ordering, in the form of a transitive ranking, of outcomes or choices. Even if, in the end, there is no meaningful transitive ranking of outcomes that we can correctly appeal to in our normative deliberations, it is arguable that there may still be many full-blooded realist considerations that would rightly have a bearing on what choices we ought, morally, to make in the living of our lives. Perhaps some of these would be deontic-, caring-, or virtue-based in nature. But, as implied above, perhaps some of these would be along new lines yet to be discovered and developed; lines which, perhaps, will only be discovered if we are forced to look in new directions for navigating the normative realm.

Finally, I am a realist in thinking that *if* we end up abandoning the Axioms of Transitivity, it will be because we are recognizing that there are compelling reasons to accept Essentially Comparative principles like the Narrow Person-

7. Soon after I sent this article off, Morten Dahlback sent me an email suggesting that virtue theorists and deontologists will also face compelling versions of the Spectrum Arguments, together with a sketch of an argument for why this is so. In a few subsequent exchanges, we were able to revise and tighten up his argument, showing that this is, indeed, the case.

8. Actually, Parfit has made this claim to me on multiple occasions over the years during discussions about my work.

Affecting View, the Pareto Principle, person-affecting versions of Utility and Maximin, and combinations of positions like the First and Second Standard Views. That is, I believe that there may be good *reasons* to abandon the Axioms of Transitivity, and that we will only be driven to such a position by the force of such reasons. So, this makes me a *realist* about reasons, as opposed to a skeptical anti-realist who denies that there could be *reasons* to accept or reject *any* particular principles, including principles of consistency like the Axioms of Transitivity.

My position here is similar to that of some people who believe in the possibility of genuine moral dilemmas. Consider one classic example of a so-called moral dilemma, *Sophie's Choice* (Styron 1979). Sophie seemingly faced three choices: she could save her son, in which case the Nazis would murder her daughter, she could save her daughter, in which case the Nazis would murder her son, or she could do nothing, in which case the Nazis would murder both her children.

Anti-realists about reasons believe that in the most fundamental sense it doesn't matter what Sophie chooses. Specifically, they believe that, ultimately, there are *no reasons* of any kind guiding Sophie's choice, or any other choice for that matter. On the anti-realist position, there is no reason either way for Sophie to save her son, save her daughter, allow both to be killed, or, for that matter, to pursue a fourth option of killing both of her children herself, and perhaps a lot of other innocent victims who would otherwise have survived!

Those realists about reason who believe in the possibility of genuine moral dilemmas (many realists do not), believe that there are overwhelmingly compelling reasons for Sophie to save her son, but that there are *also* overwhelmingly compelling reasons for Sophie to save her daughter, and that the nature and structure of the reasons in question are such that they do not, and cannot, balance or cancel each other out. On their view, the fact that she can't save *both* of her children doesn't alter the fact that Sophie *should* save her son, which she *can* do, *and* that she *should* save her daughter, which she *also can* do. Accordingly, for such moral realists, Sophie is facing a moral blind alley, or moral dilemma, in the sense that whether she chooses to save her son or her daughter she will, in a deep and fundamental way, have acted wrongly. On this view, whatever choice Sophie makes, she will have unavoidably acted contrary to compelling reasons for acting otherwise than she did, reasons which were not cancelled out or balanced by the likewise compelling reasons on which she chose to act.

Importantly, such realists would vehemently deny that there are *no* reasons applicable to Sophie's choice. Contrary to the anti-realist, they would insist that Sophie has strong reason not to let *both* children be killed, and even stronger reason not to kill her children herself along with a number of other innocent victims who would otherwise survive. But they believe that the very real and compelling reasons that have a bearing on her situation

put Sophie in a “no win” situation. She *must* choose to save her son or her daughter, but whatever choice she makes will be wrong! On the view in question, moral dilemmas are truly tragic situations from which there is no rational escape. But they only obtain because there genuinely *are* moral reasons whose nature and structure give rise to such dilemmas.

My point in discussing moral dilemmas is not to defend the view that there are such dilemmas, but to illustrate how such a view is consistent with a realist view of reasons. Similar thinking applies, I believe, regarding my Spectrum Arguments, or the other arguments I have given threatening the Axioms of Transitivity. One may be a realist about reasons, but believe that an Essentially Comparative View of ideals is true. *If* the Essentially Comparative View is true, then there may be no transitive ordering of the alternatives in my Spectrum Arguments, or of many of the other sets of alternatives my book discusses. But this won't be because the anti-realists are right. Rather, it will be because of the nature and structure of the genuine reasons that exist, and that bear on the alternatives in question.

Or so I believe, anyway.

BIBLIOGRAPHY

- Horta, O., 2014: “In Defense of the Internal Aspects View: Person-Affecting Reasons, Spectrum Arguments and Inconsistent Intuitions”, *Law, Ethics and Philosophy* 2: 87-107.
- Styron, W., 1979: *Sophie's Choice*, New York: Random House.
- Temkin, L., 2012: *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*, Oxford: Oxford University Press.
- 2014: “Rethinking the Good - A Small Taste”, *Law, Ethics and Philosophy* 2: 58-86.