

Partiality and Intrinsic Value.

By: Michael J. Zimmerman

[Zimmerman, M.J.](#) "Partiality and Intrinsic Value." *Mind*. 120.478 (2011): 447-483. Print.

This is a pre-copy-editing, author-produced PDF of an article accepted for publication in *Mind* following peer review. The definitive publisher-authenticated version is available online at: <http://mind.oxfordjournals.org/content/120/478/447.abstract?sid=62746f8f-b460-4c44-aa92-be80bf987e8a>

Abstract:

The fitting-attitudes analysis of value, which states that something's being good consists in its being the fitting object of some pro-attitude, has recently been the focus of intense debate. Many objections have been levelled against this analysis. One objection to it concerns the 'challenge from partiality', according to which it can be fitting to display partiality toward objects of equal value. Several responses to the challenge have been proposed. This paper criticizes these and other responses and then offers a response that, it is claimed, solves the challenge.

Keywords: intrinsic value | value analysis | fitting-attitudes value analysis | philosophy

Article:

1. Introduction

'If I am asked', wrote G. E. Moore in *Principia Ethica*, "'What is good?'" my answer is that good is good, and that is the end of the matter. Or if I am asked "How is good to be defined?" my answer is that it cannot be defined, and that is all I have to say about it' (Moore 1903, p. 58). The concept of goodness, he said, is simple, unanalysable, and, lest his readers entertain any doubts on this score, he proffered his open question argument as proof (Moore 1903, p. 67). But this argument is riddled with problems, as Moore himself later admitted (Moore 1959, p. 98), and it seems perfectly reasonable that we should continue to hope to find an acceptable analysis of what is surely one of the key concepts of ethics.

A very attractive analysis had in fact already been in circulation for several years by the time Moore issued his denunciation of the project of seeking an analysis. In 1889 Franz Brentano declared, '[T]he good is that which is worthy of love, that which can be loved with a love that is correct' (Brentano 1969, p. 18). And although Moore paid no heed to this proposal in *Principia*, it is one that many philosophers have since endorsed in one or another version.¹ In recent years the proposal has come to be called the fitting-attitudes analysis (Rabinowicz and Rønnow-Rasmussen 2004, p. 391), in deference to its definitive formulation by A. C. Ewing, who said: 'We may ... define "good" as "fitting object of a pro attitude"' (Ewing 1948, p. 152). Part of

what is so attractive about this proposal is that it seems to accord so well with common sense and common language. After all, we often refer to what is good as desirable, lovable, adorable, estimable, laudable, admirable, venerable, and so on, and here the suffix ‘-able’ clearly expresses the idea that the item in question is worthy of being desired, loved, adored, and so on — that it is fitting to take up one or other of these pro-attitudes in response to it.

My particular concern in this paper is with the value that something has for its own sake. Moore called such value intrinsic value, in part because he took it to supervene solely on the intrinsic properties of its bearers. This supervenience thesis has recently been challenged.² Some opponents of the thesis are willing to continue to use the term ‘intrinsic value’ to refer to the kind of value at issue (e.g. Kagan 1998), while others prefer a more neutral term, such as ‘final value’ (e.g. Rabinowicz and Rønnow-Rasmussen 1999). Since I subscribe to the supervenience thesis in question (Zimmerman 2001, Sect. 3.6), I am happy to continue, and will continue, to use the term ‘intrinsic value’ to refer to what others prefer to call final value. Nothing I say in this paper, however, will presuppose the truth of the supervenience thesis.

The fitting-attitudes analysis easily lends itself to a restriction to intrinsic value. Here is a rather rough formulation of the thesis with which I will be dealing:

(FA1) x is intrinsically good =df it is fitting for anyone who contemplates x to favour it for its own sake³

(Here, ‘favour’ is an umbrella term covering a variety of pro-attitudes toward x .) There is of course a complementary account of intrinsic badness:

(FA2) x is intrinsically bad =df it is fitting for anyone who contemplates x to disfavour it for its own sake

Presumably we should also say:

(FA3) x is intrinsically neutral =df it is fitting for anyone who contemplates x to be indifferent toward it for its own sake

And, finally, it is important to add:

(FA4) x is intrinsically better than y =df it is fitting for anyone who contemplates both x and y to prefer x for its own sake to y for its own sake

(Although (FA1)–(FA3) ensure that the fitting response to value is in the right ‘general area’, it is the addition of (FA4) that ensures that fitting responses track value precisely. To say that one prefers x for its own sake to y for its own sake is to say that one either (1) favours each of x and y for its own sake but favours x more, or (2) favours x for its own sake but is either indifferent toward or disfavours y for its own sake, or (3) disfavours y for its own sake but is either indifferent toward or favours x for its own sake, or (4) disfavours each of y and x for its own sake but disfavours y more.) I will refer to the conjunction of (FA1)–(FA4) simply as (FA).

Several points should be noted. First, there is a seldom-acknowledged but crucial ambiguity in the term ‘fitting’. (What I am about to say is actually too simple. I will return to the matter below.) ‘Fitting’ has both a weak and a strong sense. In the weak sense, it has the force of ‘may’; in the strong sense, it has the force of ‘must’. In contrast, ‘unfitting’ has just one sense, with the force of ‘must not’. I think it is clear that proponents of (FA) have had the strong sense of ‘fitting’ in mind. According to them, one not only may but must favour the good, disfavour the bad, prefer the better, and so on.

Second, the ‘must’ at issue constitutes only a prima facie or pro tanto requirement and not a requirement all things considered. For obvious reasons, it would be absurd to maintain that everyone always has an all-things-considered requirement to favour the good, and so on. Even if x is intrinsically good and I am contemplating it, it may be that I have an all-things-considered requirement to disfavour it for its own sake. If an evil demon will wreak havoc on the world unless I disfavour x , then it is surely plausible to say that I ought all things considered to disfavour it, even though it is good. (This example raises some important questions, to which I will return.)

Finally, my formulation of (FA) explicitly holds that it is those who contemplate x who are required to have a certain attitude toward it. This is in keeping with the way in which some writers have formulated the thesis,⁴ but there is no mention of contemplation in the formulations

provided by Brentano and Ewing. Why include this restriction? Well, it is clear that whether one is required to do something — either to act in some way or simply to have a certain attitude — depends in part on one's own properties. Stones, for example, are under no requirement to favour what is good; nor are dogs and cats. To have such a requirement one must, I think, have some kind of awareness of, some kind of epistemic familiarity with, the object in question (Bykvist 2009, p. 4). I will not try to spell out the nature of such awareness, but it is this that I mean by 'contemplation'.

2. The challenge from partiality

Despite (FA)'s intuitive appeal, it faces several difficulties. The particular problem that I want to address here has to do with the apparent fact that it is often fitting to display partiality in one's responses to value. Brand Blanshard was one of the first to press this problem in a discussion of Ewing's proposal. He writes:

[I]f goodness were the same as fittingness of favour, they could never fall apart, but it seems to me that occasionally they do. Favour and disfavour include ... many kinds and intensities of attitudes. I think Dr Ewing would say that if one of our children were starving, our judgment that this was bad would mean that the strongest aversion toward it, acute grief about it, and energetic efforts to remove it, would be fitting. Now there happens to be another child starving in central China. To feel nothing about this would certainly be unfitting if we knew its plight. On the other hand, perhaps no one would say that the suffering had the same claim on our feelings as that of our own child. There are many feelings, felt with great intensity, that are obviously suitable in the case of our own child, which we could hardly be expected to show about one that was remote and all but unknown. Now if to call anything bad is to say that it is the fitting object of anti-attitudes, then when we call the remote child's starvation bad, we must mean that it is a great deal less bad than our own child's starvation, since the attitudes fitting in this case are less various and less intense. But if it were put to us expressly that this was what we meant, I think we should say 'No, I did not mean that at all; what the Chinese child is going through may be as bad as anything my own child is suffering, or even worse; I cannot feel as keenly about remote evils as I do about those nearer home, and have no sense that I ought to, but I can still recognize quite clearly that those evils are as great as those near by.' In short, the judged badness may be the same while the fitting attitudes differ; the one, therefore, is not the same as the other. (Blanshard 1961, pp. 287–8)

It is not clear to me whether Blanshard is using 'fitting' in the same sense throughout this passage. He presumably understands (FA) in terms of the strong sense of 'fitting' and thus as implying that we must respond equally to situations of equal value. Given that a distant child's suffering is just as bad as one's own child's suffering, then, (FA) appears to imply that one must

prefer neither episode of suffering to the other. But in characterizing one's preferring the distant child's suffering to one's own child's suffering as fitting, does Blanshard mean to say that one must, or only that one may, indulge in such partiality? I am not sure, but there is no need to settle the question here. For on either reading Blanshard has posed a serious challenge to (FA). His argument may be put as follows: intrinsic value is a kind of impersonal value such that, if (FA) were correct, impartiality in one's responses to situations that are of equal intrinsic value would always be required; sometimes, though, partiality in one's responses to situations of equal intrinsic value is at least permitted, if not required; hence (FA) is false.

This challenge from partiality, as I will call it, is worth pursuing only if the sort of partiality — the parental bias — to which Blanshard alludes is indeed permitted. This, of course, is a large question. I will not address it here. I will simply declare that I think he is surely right about this.

The challenge is also worth pursuing only if the thesis that is challenged is plausible apart from the challenge. I have said that (FA) has strong intuitive appeal, but I also mentioned that it faces several difficulties, the challenge from partiality being only one among many. Although there is no time to investigate any of these other difficulties in detail here, I think I should say something briefly about two of them that have received a fair bit of attention recently.⁵

3. Two further challenges to (FA)

3.1 Buck-passing

(FA) is an integral part of what has come to be known, thanks to Thomas Scanlon, as the buck-passing account of value (Scanlon 1998, pp. 95–7). According to this account, it is not x's being good that renders it fitting of a favourable response; rather, it is those properties in virtue of which it is good that do so. Many philosophers have endorsed this claim, but some have demurred. Blanshard himself is among the demurrers. He insists that the reason why it is fitting to favour what is good is the fact that it is good (Blanshard 1961, p. 284). This indicates a priority of goodness over fittingness of favour that is inconsistent with (FA)'s contention that the former may be analysed in terms of the latter.

I am inclined to side with Blanshard on this issue, although the matter is not an easy one (Zimmerman 2007). Thus I am inclined to reject (FA). But even if (FA) is to be denied for this reason, the challenge from partiality is none the less worth pursuing. This is because (FA) can be

weakened in such a way that the objection just mentioned is sidestepped but the challenge from partiality still applies. The weaker thesis I have in mind is this:

(FA1*) Necessarily, x is intrinsically good if and only if it is fitting for anyone who contemplates x to favour it for its own sake

Similar modifications can of course be made to each of (FA2)–(FA4). I will refer to their conjunction as (FA*).

On (FA*), as on (FA), there is always, and of necessity, a perfect coincidence between intrinsic value, on the one hand, and fittingness of certain attitudes, on the other, in a way that would seem to rule out the sort of parental bias in Blanshard's example. But, unlike (FA), (FA*) is perfectly compatible with the contention that fittingness of favour supervenes on goodness. So, even if this contention is correct and (FA) is therefore false, the challenge from partiality still has a solid target in (FA*), which preserves the same plausible strict connection between what is good, on the one hand, and what is desirable, lovable, adorable, and so on, on the other.

3.2 The wrong kind of reason

The second objection to (FA) that I want briefly to consider, however, also constitutes an objection to (FA*). It has to do with what has come to be called the wrong kind of reason (Rabinowicz and Rønnow-Rasmussen 2004, p. 393). It seems that cases can be given in which there is reason to respond favourably toward something for its own sake even though that thing is not intrinsically good. If this is so, then it is not only (FA) that is thrown into doubt; so, too, is (FA*). What kind of case might achieve this? Well, suppose that some powerful demon threatens to inflict severe suffering on you unless you favour another person's suffering for its own sake.⁶ Surely you have very good reason to comply with the demon's wishes; but, equally surely, this fact does not render the other person's suffering intrinsically good.

In section one, I mentioned this sort of case in order to emphasize the fact that the requirement that, according to (FA) (and (FA*)), we have to favour the good, and so on, is merely pro tanto and not (necessarily) all-things-considered. It may therefore seem that all we need do to defuse the current objection is point out that it is perfectly consistent to say that, although you are required (pro tanto) to comply with the demon's wishes and favour the person's suffering, you are

also required (pro tanto), as (FA) says, to disfavour his suffering. No contradiction, therefore no problem.

But of course this reply will not do. The objection does not target that part of (FA) that implies that if *x* is intrinsically good, then it is fitting to favour it for its own sake. It targets that part that implies that if it is fitting to favour *x* for its own sake, then it is intrinsically good.

Many attempts have been made to handle the objection.⁷ One response that I believe has been unduly neglected is simply to note, first of all, that the wrong-kind-of-reason objection of course concerns the wrong kind of reason. It applies directly to a version of the fitting-attitudes analysis that is couched in terms of reasons — for example:

(FA1r) *x* is intrinsically good =df anyone who contemplates *x* has a reason to favour it for its own sake

But this is not how (FA) was originally formulated. Even if it is granted that you have a reason to favour the other person's suffering for its own sake, it is less clear that the demon's threat makes it fitting for you to favour it (Rabinowicz and Rønnow-Rasmussen 2004, pp. 422–3). Indeed, if we move from 'fitting' to some other terms that are sometimes used in this context — 'worthy', for example (a term used by Brentano himself, who pioneered this sort of analysis), or 'deserving' — the wrong-kind-of-reason objection loses its foothold. What makes something that one is contemplating worthy or deserving of favour is the nature of the thing contemplated, not some threat about what will happen if one fails to contemplate it.⁸ No matter how powerful the demon may be, his threat is impotent, in as much as it can do nothing to render the other person's suffering worthy or deserving of favour. (FA), understood in terms of worthiness or desert, is therefore immune to the objection. A proponent of (FA), so understood, can agree that there is always reason to favour that which is good, but he is not committed thereby to agreeing that whatever there is reason to favour is good. On the contrary, it is only when one has a reason of the right kind to favour something — a reason that is constituted by the thing's being worthy or deserving of favour — that it follows that the thing in question is good.

4. Responses to the challenge from partiality

Let us assume, then, that (FA) (properly understood, but perhaps only when put in the weakened form of (FA*)) survives all objections other than that posed by the challenge from partiality, and let us assume that this challenge correctly presumes that partiality in response to episodes of intrinsic value can be fitting, as in Blanshard's example of parental bias. The challenge is then worth pursuing. How might a proponent of (FA) try to handle it? In this section, I will discuss a baker's dozen of responses to the challenge, rejecting all but the last.

4.1 Personal vs. impersonal value

One response is to say that it is a mistake to conceive of intrinsic value as a kind of impersonal value. Contrary to Blanshard's assertion, it might be said, my own child's suffering is intrinsically worse — worse for me, that is — than the suffering of some distant child who is unknown to me. That distant child's suffering is of course intrinsically worse for its father than my child's suffering is, but there is no contradiction as long as we recognize the personal nature of intrinsic value. And, of course, conceiving of intrinsic value as being personal in this way fits perfectly with the fitting-attitudes analysis, which would then condone, rather than condemn, the sort of parental bias that Blanshard finds fitting.

Moore is well known for arguing that the claim that value can be personal does in fact lead to contradiction (Moore 1903, pp. 150–1). I think he is clearly wrong about this. It is perfectly consistent to say that one and the same episode may be good for me but bad for you. For example, my winning the prize that we both covet may be good for me but bad for you. Perhaps there are various ways in which this might be so. One way is this: what is good for me is what is in my interests, and what is bad for you is what runs counter to your interests.

None the less, I am convinced that Moore is right to declare intrinsic value impersonal. Or rather, to put the point more carefully: he is right to insist that the sort of value with which he and I are concerned, and to which we both refer by the label 'intrinsic value', is impersonal. Let me be clear about this. I am not suggesting that it is because intrinsic value supervenes solely on the intrinsic properties of its bearers that it is an impersonal kind of value, for I have already noted that I will not be presupposing the truth of this supervenience thesis in this paper. Nor am I suggesting that it is because intrinsic value is final or non-derivative that it is an impersonal kind of value (as opposed, for example, to the sort of non-final or derivative value that often goes by the name of 'instrumental value').⁹ After all, the distinction between what is of final or non-derivative value and what is merely of non-final or derivative value can be drawn within the personal kind of value that has to do with what is in one's interests (Thomson 1992, p. 103). (It is plausible to contend, for example, that one's being happy is non-derivatively personally good —

finally good, or good for its own sake, as far as one's interests are concerned. If so, then whatever is conducive to one's happiness will be derivatively personally good.) Rather, the distinction between intrinsic value (the sort of value to which Moore and I — and a great many others — refer by 'intrinsic value') and other kinds of non-derivative value lies elsewhere. In my view, which I have found that few share with me,¹⁰ intrinsic value is a kind of non-derivative moral or ethical value, and instrumental value is derivative value of this same kind (Zimmerman 2001, pp. 23–4). But there is no need to insist on this here (although it is a point to which I will return later). All that needs noting is that intrinsic value is not the same kind of value as the personal value that has to do with what is in one's interests.

Of course, even if I am right about this, it could still be that intrinsic value is a personal value of some other kind. I confess that I have no argument to give against this claim. Instead, let me simply say that, just as Blanshard is surely right to say that the sort of parental bias he describes is fitting (at least in the sense that it is permitted), so too he is surely right to say that the distant child's suffering is, in some very important way, just as bad as my own child's suffering. We all understand what way this is; it is this way of being bad that I express by 'being intrinsically bad'. Thus I reject this first response to the challenge from partiality.

4.2 Personal value revisited

It might be said that there is no need to insist that intrinsic value is itself a kind of personal value. All that we need say is that there is a kind of personal value distinct from intrinsic value which accounts for the propriety of partiality. Given the fact that my own child's suffering is no worse, intrinsically, than the distant child's suffering, (FA) implies, correctly, that I, like everyone who contemplates these states, have a *pro tanto* requirement to respond to them with equal disfavour. However, saying this is perfectly consistent with also saying that, as a parent, I have another *pro tanto* requirement to disfavour my own child's suffering more than that of the distant child. Or, to put the point in terms that are currently popular: while I, like everyone else who contemplates these two states, have an agent-neutral reason to respond equally to them, I also have an agent-relative reason that not everyone else has to prefer one of them to the other. This last reason signals the fact that my own child's suffering is worse for me than the distant child's suffering, even though the latter is just as bad, intrinsically, as the former (Olson 2009, pp. 369 ff.).

As will become clear below, I think that a version of this response is correct but that, unelaborated, it is unsatisfactory. One problem with it (which can perhaps be managed) is that it fails to make clear what sort of personal value is allegedly at issue. I do not think it can be the sort that I have just discussed, namely, that which has to do with what is in one's interests. For it

might be that my own child's suffering does not run counter to my interests (I may be the sort of person who takes pleasure in other people's misery, even that of my own child), yet it remains true that it would be fitting for me to prefer the distant child's suffering to that of my own child. A second problem runs deeper: even though only *pro tanto*, and not *all-things-considered*, requirement is at issue, it cannot be that contemplation of my own child's suffering both does and does not require a particular degree of disfavour on my part. If contemplation of both it and the distant child's suffering requires that I disfavour them equally, it cannot also be the case that such contemplation does not require that I disfavour them equally, let alone its being the case that it does require that I prefer one to the other (Chisholm 1974, p. 7).

4.3 Supervenience

A third response relies on the thesis that intrinsic value supervenes solely on the intrinsic properties of its bearers. According to this response, it is the intrinsic features of my own child's suffering and of the distant child's suffering that ground a *pro tanto* requirement to disfavour these states equally; it is certain extrinsic features of my own child's suffering that permit, or even require, that I disfavour it more than the suffering of the distant child. Thus (FA) need only be revised as follows in order to meet the challenge from partiality:

(FA4i) x is intrinsically better than y =df it is fitting for anyone who contemplates the intrinsic features of both x and y to prefer x for its own sake to y for its own sake

As I have said, I have no quarrel with the supervenience thesis in question. But others do, and so for this reason alone it would be a pity to have to rely on it in order to rebut the challenge from partiality. A further problem is that, even if the supervenience thesis is true, it may not do the job required of it here. Suppose we accept that there is a sense in which my own child's suffering is worse for me (though no worse intrinsically) than the distant child's suffering. It is not clear that this fact must depend on non-intrinsic features of either state. In order to facilitate the discussion, let us take the bearers of intrinsic value to be concrete, Kim-like states that involve an individual, x , exemplifying a property, P , at a time, t (Kim 1976). I will designate states by means of expressions of the form ' $[x, P, t]$ ' and assume that ' $[x, P, t]$ ' is identical to ' $[y, Q, t']$ ' if and only if x is identical to y , P is identical to Q , and t is identical to t' .¹¹ (We may call x , P , and t the constituents of the former state and y , Q , and t' the constituents of the latter. Note that, since properties are to be finely individuated, so, too, are states.) Suppose now that my daughter Sarah is suffering to degree 10. What is 'the' state at issue? Well, two candidates are: [Sarah, suffering to degree 10, now] and [Sarah, being Michael's daughter and suffering to degree 10, now]. Consider the first of these. Does it call for special disfavour on my part and, if so, does it do so in

virtue of some of its non-intrinsic features? The answer is not clear. Sarah is the state's constituent individual and her suffering to degree 10 is its constituent property, and these constituents are intrinsic to it. It might be argued that the propriety of my disfavoured the state especially strongly depends on these features alone; if so, no non-intrinsic feature of the state plays a role in accounting for this propriety. Or it might be argued that the propriety of special disfavor on my part does not depend on these constituents alone and that some further feature, namely, the constituent individual's being such that she is my daughter, plays a role. What of the second state? If it is this state, and not the first, that calls for special disfavor on my part, there would seem to be no reason to say that it does so in virtue of any of its non-intrinsic features; for Sarah's being my daughter is part of the state's constituent property. This shows that there is at least one way in which one might accept the supervenience thesis at issue while maintaining that the propriety of special disfavor on my part for my daughter's suffering depends solely on the intrinsic features of her suffering.

It could be, of course, that the correct account of the nature of states would imply that the present response succeeds. But unless and until it is settled what the correct account is, we cannot assume that this response provides an acceptable solution to the challenge from partiality. Indeed, given my allegiance to the Kim-inspired account just sketched, I think we should assume that the response is likely to prove unacceptable.

4.4 Ideal observers

Another response involves the well-known idea of an ideal observer, a person who is ideally situated, in terms of both his qualities and his circumstances, to observe what is to be observed. It might be contended that among the relevant qualities is the trait of impartiality — ideal observers just do not take things personally — so that all that is needed to reconcile (FA) with the challenge from partiality is to qualify it in some such way as this:

(FA4io) x is intrinsically better than y =df if one were an ideal observer contemplating both x and y , one would prefer x for its own sake to y for its own sake

(Corresponding amendments would be made to (FA1)–(FA3).) This would leave open the possibility that it would be fitting for someone who is not ideally situated with respect to both x and y to prefer x to y even when the former is not intrinsically better than the latter. In this way the fittingness of a parent's preferring a distant child's suffering to his own child's suffering, even

though the former is not intrinsically better than the latter, could be attributed to the parent's failing to be ideally situated with respect to both instances of suffering.

I am not a fan of ideal observer theories. In this case, as in others, they seem to me to be subject to a dilemma concerning the determination of the features that supposedly make an observer ideal. Either these features are specified independently of the theoretical purpose that they are intended to serve, or they are not. If they are, I cannot see how it can be guaranteed that they will successfully serve this purpose. (On what grounds are we to assume that an ideal observer will be impartial? Is this supposed to be an ideal feature in a parent, or must an ideal observer be childless? In this connection, consider what Judith Thomson has to say:

Is it plausible to think that what has intrinsic goodness just is what a person (all people?) would value for its own sake if he or she were fully informed, free of neuroses, and assessing the matter in a cool hour? No, unless we can show that people really would not love the nasty under this constraint. [Thomson 1992, p. 108])

If, however, the features in question are not specified independently of the theoretical purpose that they are intended to serve, then perhaps success can be guaranteed, but only at the cost of circularity.

4.5 Disinterested observers

One way to circumvent the dilemma just noted would be to think of an 'ideal' observer in purely descriptive terms, as the sort of person who in fact has none of the properties (such as that of being a parent) that would permit or require partiality in response to items of equal intrinsic value. Let us call such an observer 'disinterested' rather than 'ideal'. Then it might seem that all we need do to meet the challenge from partiality is to revise (FA) as follows:

(FA4do) x is intrinsically better than y =df it is fitting for any disinterested observer who contemplates both x and y to prefer x for its own sake to y for its own sake

(Again, corresponding amendments would be made to (FA1)–(FA3).)

I have no quarrel with (FA4do) as such but, as a response to the challenge from partiality, it is unsatisfactory. One problem is that, as it stands, it leaves entirely unspecified what it is to be a

disinterested observer. Surely it is not only parents who fail to be disinterested, for special relations that seem to warrant partiality arise in many different contexts. Perhaps, too, some parents qualify as being disinterested, for it is doubtful that a special relation (of the relevant sort) is grounded purely in the biological fact that one person is the progenitor of another. Another problem is this: if, as I have been assuming, I do not qualify as being disinterested with respect to my own child's suffering, then the present proposal does not apply to me. It certainly does not imply that I am required to disfavour her suffering no more than that of a distant child. Yet there is strong pressure to say just that, given that my child's suffering is indeed no worse, intrinsically, than the distant child's. The present proposal simply does not address this aspect of the challenge from partiality.

4.6 Modes of fittingness

In section 3, I discussed the wrong-kind-of-reason objection to (FA). My response was to say that although the demon's threat surely gives you a reason to favour the other person's suffering, it is less clear that it renders such favour fitting, and it seems quite clear that it does not make the suffering worthy or deserving of favour. Jonas Olson has suggested that we say something similar about the challenge from partiality: although I surely have a reason to prefer the distant child's suffering to my own child's suffering, it does not follow that such preference is fitting. Hence (FA), when couched in terms of fittingness in particular rather than reasons in general, is unscathed (Olson 2009, pp. 373 ff.).

As it stands, this proposal seems unsatisfactory. While I concede that my having a reason to disfavour my own child's suffering more strongly than the distant child's suffering may not entail that it is fitting that I do so and does not entail that the former deserves stronger disfavour on my part, none the less it seems that it is indeed fitting that parents display such partiality and that children deserve it.

Still, one might try to press the current line of thinking by claiming that the challenge from partiality presents not a wrong-kind-of-reason situation in general but a wrong-kind-of-fittingness (or -desert) situation in particular. To do this, one must of course distinguish modes of fittingness (or desert). One could then say that, although it is in some way fitting for me to prefer a distant child's suffering to my own child's suffering, this is irrelevant to the determination of intrinsic value. In so far as such determination is concerned, the way in which it is fitting to disfavour the distant child's suffering is such that it is not fitting for me to disfavour my own child's suffering more.

The success of this response of course hinges on the alleged distinction between modes of fittingness. How might such a distinction be drawn? Olson, following Ewing, suggests that the requirement to prefer the distant child's suffering is a moral one, whereas the requirement not to do so is non-moral. (Like Ewing, he reserves the term 'fittingness' for the latter requirement. I have just suggested that this is a mistake, but the issue is of no great moment. It is the distinction, not the terminology that is used to draw it, that matters.) Now, I grant that the requirement to prefer the distant child's suffering is a moral one; I would say, more particularly, that it is morally fitting that I do so. But I reject the present proposal since, as I said earlier, I also think that, in light of the equal intrinsic badness of the distant child's suffering and my own child's suffering, it is morally fitting that I disfavour them equally. Olson of course rejects this view, offering two objections.¹²

The first objection is that it is implausible to think that in general what is intrinsically good is such that there is a moral requirement to favour it. He cites New Zealand's fiord land and Da Vinci's *The Last Supper* as plausible examples of things that are intrinsically good but such that the failure to favour them does not constitute a moral shortcoming.

I grant that the failure to appreciate the beauty (or, perhaps, those properties that ground the beauty) of the items in question does not constitute a moral shortcoming. It is, rather, an aesthetic shortcoming. But if (a fairly big 'if', I think) these items are not only beautiful but also intrinsically good, then I would contend that the failure to appreciate their intrinsic value (or, perhaps, the properties that ground this value) does indeed constitute a moral shortcoming. (If it is precisely the same properties that ground both the beauty and the intrinsic value of the items, then this reply may require refinement.) I say this because, even if we grant (as I am sure we should) that the items in question are not morally good, still the requirement to favour them would appear to be a moral one — something that is reflected in the common view that we have a *pro tanto* moral duty to promote intrinsic value, wherever it may reside and regardless of whether what has intrinsic value is itself morally good.

Olson's second objection is that, given that one's own pleasure is intrinsically good, my view would imply that we are morally required to favour it; yet that seems doubtful. (W. D. Ross once made a similar point [Ross 1939, p. 282].) There is certainly a difficulty here, but I cannot see that it favours Olson's approach over my own. (Like me, Ross took the requirement implicit in intrinsic value to be a moral one.) That is because, just as it seems plausible to deny that I am morally required to prefer my own pleasure to the lesser pleasure of a stranger, so too it seems

plausible to deny that it is fitting that I do so. (This example seems to me to constitute a sort of inverse challenge from partiality. I will attend to it later.)

In the end, I do not need to insist that the requirement to favour that which is intrinsically good is a moral one. All that I need say is that I do not yet see what reason there is to hold that the way in which it is fitting to respond equally to one's own child's suffering and to a distant child's suffering differs from the way in which it is fitting to respond unequally to them. Hence I think we need to look elsewhere for a solution to the challenge from partiality.

4.7 Direct vs. indirect fittingness

Another response to the challenge is this. We should distinguish between what is directly fitting or unfitting, on the one hand, and what is indirectly fitting or unfitting, on the other. Consider, for example, someone who works with Médecins Sans Frontières. It is fitting for such a person to show compassion in the face of the suffering that confronts him. It is also fitting that such a person relieve such suffering efficiently. Indulging one's compassion, though, can compromise the efficiency with which one relieves suffering; in some cases, it can be totally disabling. Those who are in the business of relieving suffering on a large scale thus frequently find that they must curb their compassion if they are to be effective. Often this curbing occurs involuntarily — a phenomenon known as compassion fatigue. People find themselves after a while no longer capable of feeling the degree of compassion that directly fits the suffering to which they are witness. Such fatigue, though directly unfitting, is none the less indirectly fitting, in so far as it enables them to get on with the job of relieving the suffering.

Or consider the phenomenon of agent-regret. If Bernard Williams's lorry driver was really not at fault in running over the child, then, arguably, whatever degree of regret it is fitting for him to feel about what has taken place is no greater than the degree of regret it is fitting for a spectator to feel — as a matter of direct fittingness, that is (Williams 1981, p. 28). However, if the driver were in fact to show no greater regret than a spectator, we would be justifiably suspicious, taking it as a sign that he is disposed not to show greater regret when it is called for — for example, when he is at fault; for it is unlikely that anyone can fine-tune his emotional responses in such a manner. If so, then his actually showing greater regret than a spectator, though strictly an overreaction, would be indirectly fitting, in that it would bespeak a general uprightness of character, a general sensitivity to value, in a way in which his showing no greater regret would fail to do.

We might try saying something similar about the example that Blanshard has given. The suggestion would be that the degree of anguish that it is directly fitting to display in response to one's own child's suffering to a certain extent is no greater, and no smaller, than the degree of anguish that it is directly fitting to display in response to a distant child's suffering to the same extent. However, if a parent were in fact to show no greater anguish over his own child's suffering than over a distant child's suffering, that would be indirectly unfitting, in that it would bespeak a general baseness of character, a general insensitivity to value, in a way in which his overreacting to his child's suffering would not.

I do not accept this response. We should surely acknowledge the distinction between direct and indirect fittingness, and I think it is probably correct to say that the lack of compassion that characterizes compassion fatigue is directly unfitting but indirectly fitting, for the reasons given. The matter of agent-regret is more complicated, although I am inclined to think that it is indeed correct to say that, 'in principle', the driver should feel no worse about what has taken place than a spectator should. But I do not think we ought to extend this diagnosis to the case of parental bias. It is not indirectly but directly fitting that one feel worse about one's own child's suffering than about a distant child's suffering, at least in the sense that one is permitted to feel worse and possibly also in the sense that one is required to feel worse. The special relation between parent and child gives rise to special permissions, I believe, and possibly also to special duties, permissions and duties having to do not only with actions but also with attitudes. I am not sure how best to defend this claim,¹³ but, given that I do subscribe to it, I believe that we must once again look elsewhere for an acceptable response to the challenge from partiality.

4.8 Kinds of attitude

Noah Lemos proposes another response. He says:

[E]ven if we concede that it is more appropriate to have a more intense feeling of grief [or] sadness toward the suffering of one's own child than toward the suffering of a total stranger, this concession does not imply that we cannot explicate intrinsic value in terms of required love, hate, and preference. This is so simply because grief, sadness, and melancholy are not the same attitudes as love, hate, and preference simpliciter. It is not at all obvious that one's contemplation of just the states of affairs my child's suffering and an unknown child's suffering requires that one prefer the latter as such to the former. (Lemos 1994, p. 18)

Lemos is here relying on there being a distinction between the kind of pro- and con-attitudes that are fitting responses to intrinsic value and other kinds of pro- and con-attitudes that one might take toward episodes in which intrinsic value happens to be implicated. Hate, he says, is distinct from grief and sadness. But this is helpful only if we are told just what the distinction is

supposed to be and why it is relevant. Even if hate (the kind of hate that Lemos has in mind) is distinct from sadness, why is not sadness just as fitting as hate when responding to the intrinsic badness of some situation? Similarly, when considering fitting responses to intrinsic goodness, why restrict our attention to love, whatever precisely that is supposed to be? What, for instance, of the kinds of responses that others have declared appropriate to value: desire, satisfaction, appreciation, admiration, joy, and so on? I know of no reason to rule them out as fitting responses to value. Moreover, even if we restricted our attention to love, hate, and preference, as Lemos recommends, I cannot see that doing so helps rebut the challenge from partiality. It seems perfectly fitting that I should hate my child's suffering more than a distant child's suffering, perfectly fitting that I should prefer the latter to the former. (Lemos contends that I am not required to prefer the latter to the former. Perhaps he is right about this, but what his response needs is something stronger if it is to be successful. It needs to be the case that I am required not to prefer the latter to the former, and this is what strikes me as quite dubious.)

4.9 Action vs. attitude

Lemos offers a further response. He says:

[C]laiming that there is no requirement that one prefer in itself the latter to the former is compatible with holding that there is a requirement to alleviate the one rather than the other ... [W]hich case of suffering one is required to alleviate depends on considerations other than the contemplation of just those states of affairs. These other considerations might include the nearness of the children, the probability of success in alleviating their pain, the costs of doing so, and even considerations of loyalty to one's children. But even if, in light of these other factors, there is a requirement to choose to alleviate the suffering of one's own child, it does not follow that one is required to prefer simpliciter the suffering of the unknown child. (Lemos 1994, p. 18)

This response is similar to one that Ewing himself proposed as a solution to the challenge from partiality. According to Ewing, even if one's own child's suffering is itself no worse, intrinsically, than a distant child's suffering, so that it is fitting to prefer neither to the other, none the less it may be better, intrinsically, for one to alleviate the former rather than the latter (Ewing 1948, p. 192).

Let us consider Ewing's proposal first. As Olson has pointed out, it fails to meet the challenge successfully. If it is better for me to alleviate my own child's suffering rather than a distant child's, presumably it is also better, and to the same extent, for the parent of the distant child to do the reverse. If so, (FA) implies that I am required not to prefer my alleviating my child's suffering to the other parent's alleviating his child's suffering. But, to the extent that the challenge

from partiality has intuitive force in the first place, it would seem fitting for me to prefer the former to the latter (Olson 2009, p. 372). This is so, moreover, even if the parent's alleviating his child's suffering were better than my alleviating my child's suffering, because his child's suffering is greater than that of mine.

These objections do not pertain to Lemos's response, however, since his response is couched in terms of there being a special duty, rather than value, associated with alleviating one's own child's suffering. But there are problems with the response none the less. It is of course true that, once we distinguish between attitudes and actions, we may also distinguish between duties of attitude and duties of action. And, if we take seriously the title 'fitting-attitudes analysis', then (FA) explicitly concerns only duties (or permissions) of attitude and not duties of action. Thus it is arguably the case that a parent has a stronger duty to alleviate his own child's suffering than to alleviate that of a distant child, even though he is required to adopt the same attitude toward both. But I would like to see the argument. Lemos mentions some possible considerations: the nearness of the children, the probability of success in attempting to alleviate their suffering, the costs of doing so, and so on. But of course we can imagine cases in which these considerations do not apply, in that the children are equally close, the probability of success is the same in each case, and so on. Lemos also mentions the possibility that loyalty to one's children serves to generate a duty to help one's own child that is stronger than the duty to help a distant child. But he gives no reason to think that this does not also ground a duty to hate one's own child's suffering more strongly than one hates the distant child's. Once again, then, I think we have made no headway in meeting the challenge from partiality.

4.10 Distance

Another response has been provided by Graham Oddie, who draws an analogy between partiality and visual perspective. 'The moon', he says, 'should look bigger than the sun from the vantage point of Earth, and it is no defect of the observer that it appears that way to him', even though the sun is of course bigger than the moon. He continues:

So, even though the facts about size are observer-independent, experiences of size are appropriately observer-dependent, since they depend on the ... perspective of the observer. Analogously, even though facts about value are valuer-neutral, experiences of value can and should be valuer-relative, since they depend on exactly where the valuer is situated with respect to the objects of [favour and disfavour]. (Oddie 2005, p. 63)¹⁴

On the basis of this analogy, Oddie contends that the impersonal nature of value — that is, as I understand him, of intrinsic value (in the sense explained earlier) — does not require impartiality in our response to value. It is only when two equally valuable objects are 'equidistant' from us

that we should be impartial between them; otherwise, we should prefer the good that is ‘closer’ or the evil that is ‘farther away’. In general, our response to a valuable object should be dictated not only by how valuable it is but also by how far we are from it; the response should be directly proportional to the former factor but inversely proportional to the latter. It is only when we are at ‘zero distance’ from the object that it is fitting that our response to value be proportional to the magnitude of value alone. Thus Oddie is, in effect, rejecting (FA) as formulated above but endorsing something like the following amendment to it (see Oddie 2005, pp. 222 ff.):

(FA1zd) x is intrinsically good to degree n =df it is fitting for anyone who contemplates x while at zero distance from it to favour it to degree n for its own sake and for anyone who contemplates x while at non-zero distance from it to favour it to a degree (less than n) that is inversely proportional to his or her distance from it

By virtue of this amendment (and corresponding amendments to (FA2)–(FA4)), he claims to be able to reconcile the insight underlying (FA) with the propriety of partiality.

This is an intriguing response to the challenge from partiality, but I fear it faces a number of problems. The talk of ‘distance’ is suggestive — it is reminiscent of Blanshard's observation that he ‘cannot feel as keenly about remote evils as [he does] about those nearer home’ (Blanshard 1961, p. 288) — but I find it obscure. It is not clear to me just what the dimension is supposed to be along which some objects may be ‘close’ to a person while others are ‘farther away’. It cannot simply be a matter of what the person does happen to care about, for there can surely be cases in which one's response to value is not fitting; otherwise there would be no point to an (FA)-type of account at all. And Oddie is obviously aware of this; he explicitly draws a distinction between ‘real’ and ‘perceived’ distance, thereby acknowledging the possibility of unfitting responses to value (Oddie 2005, p. 230). But just as we cannot understand distance in terms of what people do in fact care about, nor can we understand it in terms of what people ought to care about, because then Oddie's amended version of (FA) would be trivial, amounting to no more than the proposal that it is fitting to respond to value in accordance with how fitting it is to respond to it. So we need to understand the notion of distance in some other way. Oddie provides no explication of the notion, though, and I find it rather elusive.

To the extent that I do understand the notion, however, I think that Oddie's proposal has questionable implications. One problem concerns the distinction between one's child's suffering and one's own suffering. The metaphor of distance suggests (to me, at least) that, if one is ever at

zero distance from anything, it is one's own present suffering from which one is at zero distance, rather than that of one's child. But then Oddie's account implies that one ought to disfavour one's own suffering for its own sake more than one disfavors the equal suffering of one's child for its own sake, and that seems wrong. As I noted earlier, Ross has raised a similar point. He says:

[It is incorrect to say] that a man's own pleasures are, from the point of view of any man, good in the same sense in which the pleasures of others are. For while we can see the rightness, the moral suitability, of his taking satisfaction in the latter, we can see no moral suitability in his taking satisfaction in the former. (Ross 1939, p. 282)

Now, I do not think Ross puts his point very well. He seems to suggest that intrinsic value is somehow relative to individuals when he says that one's own pleasures are not good in the same sense in which the pleasures of others are. I reject any such suggestion. But I take it that all he really wants to say is that it is not fitting (not 'morally suitable') to take satisfaction in one's own pleasures to the extent that it is fitting to take satisfaction in those of others. This fits ill with Oddie's metaphor of distance. Others have made similar points. Ewing, for example, says that it is fitting to admire courage in others but not equal courage in oneself, even though the two instances are equally good (Ewing 1948, p. 155). (Of course, not only does this observation threaten Oddie's version of (FA), it threatens Ewing's own.)

Another problem has to do with the ambiguity, noted earlier, of 'fitting' (or 'appropriate', which is the term Oddie favours). When we consider Blanshard's example of partiality toward one's own child, it is very plausible to say that one may prefer the distant child's suffering to one's own child's suffering; that by itself is sufficient to raise the challenge from partiality. But it is another thing to say that one must prefer the distant child's suffering. This may be the case, but I think that whether it is the case is not so clear. Oddie's approach commits him to this stronger, more controversial claim.

A further problem concerns whether it might be fitting sometimes to 'overreact' to goods and evils to which one is very close. Oddie says that, when one is at zero distance from an object, one's response ought to be exactly commensurate with the magnitude of value inherent in it. But might it not be fitting (in either the strong or the weak sense) for a parent to respond to his own child's suffering so strongly that his reaction is strictly disproportionate to the badness of the suffering?

Yet another problem has to do with the possible fittingness of 'under-reacting', or even 'counter-reacting', to goods and evils to which one is very close. Suppose that, whereas a virtuous

person's suffering is intrinsically bad, a vicious person's suffering is intrinsically good. Similarly, suppose that, whereas a virtuous person's enjoying himself is intrinsically good, a vicious person's enjoying himself is intrinsically bad. (These are controversial claims, of course, but they are plausible, and I use them only for purposes of illustration.) According to Oddie, if one were aware of a vicious person's suffering, one would be required to favour it, and if one were aware of a vicious person's enjoying himself, one would be required to disfavour it. But suppose that the vicious person is one's own child. Oddie seems committed to saying that one should favour his suffering and disfavour his enjoyment more than that of a stranger. Surely not. Surely one is permitted to favour the suffering and disfavour the enjoyment less. Indeed, it seems plausible to maintain that one is even permitted to disfavour the suffering and to favour the enjoyment of one's own child, vicious though he may be.¹⁵ Such a verdict is wholly at odds with Oddie's approach, however.

Finally, it should be noted that Oddie's proposal suffers from a problem similar to that mentioned earlier in connection with the proposal about disinterested observers, and that is that it ignores the fact that there is strong pressure to say that even people not placed equidistantly from a given object are required, in light of its intrinsic value, to respond to it equally.

4.11 Degrees of fittingness

Yet another response to the challenge from partiality requires refining our understanding of fittingness. I have said that there are two senses of 'fitting' (one strong, having the force of 'must' or 'required', the other weak, having the force of 'may' or 'permitted') and one sense of 'unfitting' (having the force of 'must not'). But this is not strictly accurate, for fittingness comes in degrees. This is something that Brentano apparently denies. He says that it is 'ridiculous [to say] ... that for each instance of rejoicing [in what is good] only a certain amount of joy is appropriate ... [S]urely it can never be reprehensible to feel the greatest joy possible in what is good' (Brentano 1969, p. 25). But this does not seem right. To favour what is good may be fitting, but if one favours only slightly something that is very good, then one is under-reacting, and if one favours greatly something that is only slightly good, then one is overreacting. Granted, as Brentano says, perhaps we should not say that such under- and overreactions are 'reprehensible' — that is, that they are positively unfitting — but that of course is compatible with saying that they are not as fitting as a reaction whose intensity corresponds precisely with the degree of value in the object that is being contemplated. The upshot is that the relation between what is required, permitted, fitting, and unfitting involves something more complicated than the simple trichotomy between 'must', 'may', and 'must not'. What we should say, I think, is this (Zimmerman 2001, p. 95): one's requirements correspond with that which it is most fitting to do. Strictly speaking, then, one ought to avoid both under- and overreactions. However, this

does not imply that under- and overreactions are unfitting. If they are reactions of the right sort, then they will go at least some way toward meeting one's requirement, and thus be fitting to some extent, even if they do not fully meet one's requirement. Only if they are reactions of the wrong sort (favour, say, where disfavour is required, or disfavour where favour is required) will they be positively unfitting.¹⁶ Thus, although whatever is unfitting is something that one is required to avoid, the reverse does not hold. A response — an under- or overreaction, say — may be such that one is required to avoid it, and yet it may be fitting to some degree.

Given this refined understanding of fittingness, one might try to reconcile (FA) with the challenge from partiality by claiming that it is possible both that it is fitting, to some degree, to react with equal disfavour to one's own child's suffering and to a distant child's suffering and that it is fitting, to some degree, to prefer the latter to the former. In either case one is reacting with disfavour to something that is bad and is thus complying with the demands of (FA). (FA) therefore allows for partiality after all.

But this response will not do, I think. (FA), I have said, is to be understood as implying that one is required not to respond partially to situations of equal intrinsic value. It implies that the most fitting response to value is one whose intensity corresponds precisely to the degree of value at stake. Thus (to use numbers in a wholly ad hoc manner), if someone is suffering to degree 10 and this is intrinsically bad to degree 10, then the most fitting response to such an episode would be to disfavour it to degree 10. Disfavour to degree 5 or disfavour to degree 15 would be fitting, but not as fitting as disfavour to degree 10. But then, given that the distant child's suffering is equally as bad as one's own child's suffering, the most fitting response would be to prefer neither to the other. Thus, even if we allow that it is fitting to disfavour one's own child's suffering more, we must say that doing so is not as fitting as disfavoursing the two episodes of suffering equally. I assume that Blanshard would deny this, and I assume that we should too. Whether partiality is more fitting than impartiality or whether they are equally fitting is an issue that I think we can leave open, but I assume that we should not say that partiality is less fitting than impartiality in the present case. Thus I reject this response, too.

4.12 Grounds of requirement

Given the many problems faced by the foregoing attempts to respond to the challenge from partiality, it may seem that no response can succeed — that we must either deny the propriety of partiality in response to intrinsic value or reject the general idea underlying the fitting-attitudes analysis. But in fact I think this would be throwing in the towel too soon.

Recall a point I raised earlier: it cannot be that contemplation of my own child's suffering both does and does not require a particular degree of disfavour on my part, even when the kind of requirement at issue is merely pro tanto and not all-things-considered. If we are to conquer the challenge from partiality, then, we must divide the requirements involved. There are in principle two ways in which this might be done. Whenever there is a requirement, there is a 'requirer', that is, a ground of the requirement, and a 'required'. We can thus distinguish requirements either by distinguishing between their grounds or by distinguishing between the things required.

Here is one way of fleshing out the first alternative. We could say that, although the contemplation of my own child's suffering and of the distant child's suffering requires my disfavouring them equally, none the less the broader situation that consists of my contemplating these states plus my being the father of my own child requires (or permits) my disfavouring my own child's suffering more. Both requirements exist, but the former is overridden by the latter.

This proposal fits well with the observation that partiality is no less fitting than impartiality in the present case. Nevertheless, it seems to me to founder on a crucial fact, one mentioned in section 3. It is the object of contemplation that accounts for the fact that it is fitting to favour or disfavour that which is contemplated. The present proposal, like any proposal that seeks to solve the challenge from partiality by distinguishing between grounds of requirement rather than between things required, ignores this fact. I think, therefore, that any adequate solution to the challenge must pursue the second alternative. I propose to do so by invoking the concept of basic intrinsic value.

4.13 Basic intrinsic value

Apart from providing a solution to the challenge from partiality, the concept of basic intrinsic value is of interest in its own right and of use in other ways. It is particularly helpful in the matter of the computation of intrinsic value. In my view, which I have elaborated elsewhere, invoking basic intrinsic value dispenses with the need to appeal to any principle of organic unities; the intrinsic value of a whole is indeed the sum of the intrinsic values — the basic intrinsic values — of its (suitably identified) parts (Zimmerman 2001, Ch. 5). Let me briefly explain.

For purposes of illustration, let us presuppose the truth of a very simple version of hedonism, which claims that only pleasure accounts for whatever intrinsic goodness there is in the world

and only pain accounts for whatever intrinsic badness there is in it, and that wherever there is pleasure there is something intrinsically good (no matter who experiences the pleasure or what its object is) and wherever there is pain there is something intrinsically bad. And, in order to facilitate the discussion, let us presuppose, as before, that the bearers of intrinsic value are concrete, Kim-like states that involve an individual, *x*, exemplifying a property, *P*, at a time, *t*, and which I will designate by means of expressions of the form '[*x*, *P*, *t*]'

Consider now a certain episode of pleasure: Peter takes considerable pleasure (pleasure to precisely degree 10, say) in eating a large pepperoni pizza at noon on Saturday. Such an episode comprises many states, among which are the following: [Peter, taking pleasure to degree 10 in eating a large pepperoni pizza, Saturday noon], [Peter, taking pleasure to degree 10 in something, Saturday noon], [Peter, taking pleasure to some degree in something, Saturday noon]. Call these states *s*₁, *s*₂, and *s*₃, respectively. Let us assume that the simple hedonism we are working with assigns degrees of value precisely and directly in proportion to degrees of pleasure and pain. Then the hedonist would say that the episode of pleasure in question is intrinsically good to degree 10. But what values are to be assigned to *s*₁, *s*₂, and *s*₃? Should we say that each of them is intrinsically good to degree 10? Would this not threaten to overestimate the value involved in the episode?

Here is my suggestion. We should distinguish between states that are 'evaluatively adequate' in terms of intrinsic value and those that are not. Ask the simple hedonist how good *s*₁ is, intrinsically, and he will say that it is intrinsically good to degree 10. (This answer is of course dictated solely by the fact that, in virtue of its constituent property, *s*₁ is an episode of pleasure to degree 10. This is all that matters, from the perspective of simple hedonism. It is only the constituent properties of states that 'do the work' in accounting for their values. The fact that it is Peter who experiences the pleasure is of course irrelevant, as is the fact that the pleasure occurs at noon on Saturday. Thus the other constituents of the state are 'idle'.) So too for *s*₂: the hedonist will declare it intrinsically good to degree 10. But he will not, or at least should not, say the same for *s*₃; for *s*₃ is, in virtue of its constituent property, such a 'thin' state that there is no saying how good it is, from the perspective of simple hedonism. Therefore, I submit, from this perspective we should not say that it is good at all; indeed, we should say that it is not good. But this does not mean that we should say that it is bad, or even that it is neutral; rather, it is evaluatively inadequate — it has no intrinsic value at all.¹⁷ States *s*₁ and *s*₂ are different, though. Their constituent properties are not too thin, from the perspective of simple hedonism, for an assignment of intrinsic value to be made. However, from this same perspective, *s*₁ suffers from a contrary defect; its constituent property is too 'thick', in that it involves details (having to do with the object of Peter's pleasure) that are once again strictly irrelevant to the assignment of intrinsic value. (Although only the constituent properties of states 'do the work' in accounting

for their value, not all components or aspects of these properties are involved in this work.) Thus s1 is, in terms of intrinsic value, 'evaluatively superfluous'. It is only s2 that is, in terms of intrinsic value, neither evaluatively inadequate nor evaluatively superfluous, and it is in virtue of this fact that s2 and only s2 (among the three states being considered) has basic intrinsic value. If, when computing the intrinsic value of a complex whole, it is only those parts of it with basic intrinsic value that we take into consideration, then we will not overestimate its value.

The thesis that basic intrinsic value is to be ascribed to all and only states that are evaluatively adequate but not evaluatively superfluous is entirely general. On the simple hedonism just discussed, s2 and only s2 has basic intrinsic value. However, a more sophisticated hedonism, according to which the intrinsic value of a state of pleasure depends in part on what the object of the pleasure is, might say that it is s1 that has basic intrinsic value (and hence that s2, like s3, is evaluatively inadequate).

There is a certain relation between the three states mentioned in my example. The constituent property of s1 entails but is not entailed by that of s2, which in turn entails but is not entailed by that of s3. In virtue of this fact (and of the fact that the states have all their other constituents in common), I propose that we say that s3 is a proper part of s2, which in turn is a proper part of s1. In saying that a state (such as s2) has basic intrinsic value just in case it is neither evaluatively inadequate nor evaluatively superfluous, then, I am saying that it has intrinsic value but none of its proper parts have intrinsic value. In saying that a state (such as s1) has non-basic intrinsic value, I am saying both that it has intrinsic value and that some proper part of it has intrinsic value.

How can invoking basic intrinsic value help us meet the challenge from partiality? As follows.¹⁸ Suppose that Sarah and Stella are both suffering to degree 10. Suppose also, in keeping with the intuition that suffering as such is evil, that the states [Sarah, suffering to degree 10, now] and [Stella, suffering to degree 10, now] — call these s4 and s5, respectively — are basically intrinsically bad to degree 10. But now suppose that Sarah is my daughter while Stella is a stranger. It may seem that, given the propriety of parental bias, it is fitting for me to disfavour s4 for its own sake more than I disfavour s5 for its own sake. But I deny this. In so far as s4 and s5 are the same kind of state (in virtue of their constituent properties), I ought to disfavour them equally. However, we should distinguish these states from others involved in the situation that have thicker constituent properties and are therefore, in terms of intrinsic value, evaluatively superfluous. Contrast, then, s4 with [Sarah, being Michael's daughter and suffering to degree 10, now] — call this state s4*; and contrast s5 with [Stella, being a stranger to Michael and suffering to degree 10, now] — call this state s5*. It is s4*, not s4, that involves the fact that Sarah is my

daughter, and it is $s5^*$, not $s5$, that involves the fact that Stella is a stranger to me. And it is perfectly consistent to say that, although I should be impartial between $s4$ and $s5$, I need not be impartial between $s4^*$ and $s5^*$.

My proposal, then, is that we revise (FA) so that it concerns only states with basic intrinsic value, as follows:

(FA1b) x is basically intrinsically good =df

(1) it is fitting for anyone who contemplates x to favour it for its own sake

and

(2) there is no y such that y is a proper part of x and it is fitting for anyone who contemplates y to favour it for its own sake

Corresponding amendments are of course to be made to (FA2) and (FA3). As for intrinsic betterness, we may say:

(FA4b) x is basically intrinsically better than y =df

(1) x and y have basic intrinsic value

and

(2) it is fitting for anyone who contemplates both x and y to prefer x for its own sake to y for its own sake

These revised definitions are intended to capture the leading idea that impartiality is indeed required of all contemplators of states that have intrinsic value, at least with respect to those parts of such states that have basic intrinsic value (the parts whose constituent properties are precisely what is ‘doing the work’ in providing the states with the value that they have); it is not just those contemplators who are, for example, disinterested in or at zero distance from the object of contemplation that are under such a requirement. (As before, it is the definition regarding intrinsic betterness — in this case, (FA4b) rather than (FA4) — that is designed to ensure that fitting responses track value precisely.) But, by restricting the requirement of impartiality to the contemplation of states with basic intrinsic value, the definitions are intended also to be consistent with the other leading idea that partiality may be fitting when contemplating states that have (non-basic) intrinsic value. Thus is the challenge from partiality met.

Let me now attend to some objections.

First, it may be objected that, as it stands, my proposal says nothing about whether it is indeed ever fitting to respond in a partial way to evaluatively superfluous states such as $s4^*$; it implies only that impartiality is required when responding to states with basic intrinsic value. This is of course true. As I have said, I do indeed think that partiality toward certain states can be fitting, but I will not try to develop any formal account of how this is so. I will rest content here simply with pointing out the compatibility between the proposition that my contemplating both $s4$ and $s5$ requires that I disfavour them equally and the proposition that my contemplating both $s4^*$ and $s5^*$ does not require that I disfavour them equally.

This compatibility claim might be challenged, though, on the grounds that, if being my daughter is essential to Sarah, then $s4$ is identical with $s4^*$ and hence contemplation of the ‘one’ necessarily involves contemplation of the ‘other’. But this would be a mistake. The constituent property of $s4^*$ is distinct from that of $s4$, even if being my daughter is essential to Sarah. Furthermore, if I contemplate $s4$, I do not contemplate Sarah's being my daughter, but if I contemplate $s4^*$, I do.

This last claim, however, might be challenged in turn. To say that, when I contemplate $s4^*$, I contemplate Sarah's being my daughter surely presupposes that I am aware that I am Michael and therefore that $s4^*$ involves my daughter's suffering, whereas $s5^*$ does not. But I might not be aware of this. Furthermore, if I were not, surely there would be no reason to say that it is directly

fitting for me to prefer $s4^*$ to $s5^*$. (Perhaps I should repeat here a point that I made earlier. If it is fitting for me to display partiality toward Sarah, this presumably will not be merely because she is, and I am aware that she is, my daughter, but rather because of some special relation between us that has developed partly due to the fact that she is my daughter. The present discussion is therefore somewhat oversimplified.)

In response, let me just say this. Sarah has the property of being Michael's daughter and the property of being my daughter. If these are distinct properties, then I have misidentified the states regarding which I may display partiality. These states are not $s4^*$ and $s5^*$, respectively, but rather the following: [Sarah, being my daughter and suffering to degree 10, now] and [Stella, being a stranger to me and suffering to degree 10, now]. The difficulty is then solved. If, however (as I suspect is the case), the property of being Michael's daughter and the property of being my daughter are not distinct, then we cannot simply say that I may display partiality regarding $s4^*$ and $s5^*$ but must say something more guarded, perhaps along these lines: I may display partiality regarding these states when I conceive of them under this aspect (an aspect that involves my recognizing that I am Michael) even if not under that aspect. I confess that I do not know how best to develop this idea. The problem is an instance of a wider issue, having to do with the opacity of propositional attitudes (among which are the various attitudes of favour and disfavour) and the role that indexicals play in reports of such attitudes. As far as I know, no theory of how to deal with this wider issue has met with general acceptance. I am no expert on the matter and gladly leave it to others to deal with. If and when an acceptable theory is found, I will happily apply it to the present context.

Another objection is that my proposal 'downplays the particularity that ... is involved in the appropriate caring for a person for her own sake' (Bykvist 2009, p. 23).¹⁹ The idea is that what renders the especially strong disfavour of my daughter's suffering fitting is the fact that she is suffering rather than the fact that a daughter of mine is suffering. This is a very difficult issue. There is surely something attractive in the idea that true love involves such a focus on the particular. As Shakespeare tells us, 'Love is not love, Which alters when it alteration finds.'²⁰ Yet the question is not whether my love for my daughter is true but whether it is fitting. As I see it, if an attitude or act is to be fitting, it must satisfy some universalizability condition, and it is not easy to see how doing so is compatible with such an emphasis on particularity. I am inclined therefore to deny that any particularity can be part of what renders love fitting or appropriate. If true love requires such particularity, then that either shows that such love is not fitting or appropriate (which of course is not to say that it is unfitting or inappropriate) or, at least, that such particularity plays no role in its being fitting or appropriate.

Yet another objection is that my proposal does not so much solve the challenge from partiality as dissolve it. According to me, intrinsic value is at bottom an impersonal kind of value that precludes partiality. I have said that it might be fitting to respond partially to states with non-basic intrinsic value but, it may be said, non-basic intrinsic value is not a genuine kind of value at all, since (according to me, at least) it is idle, superfluous: all the value that there is in the world is accounted for by states that have basic intrinsic value. And if non-basic intrinsic value is not a genuine kind of value, then it cannot be this that explains how partiality in response to states could ever be fitting. But, if so, my concession that partiality could be fitting is empty, since I have robbed this possibility of any conceivable rationale.

The claim that non-basic intrinsic value is not a genuine kind of value is, I think, misleading. It is a claim that is sometimes made about derivative value of all sorts, whether non-basic intrinsic value or the various kinds of extrinsic value (such as instrumental value, signatory value, and so on). It is certainly true that only non-derivative value accounts for whatever value there is in the world; for derivative value is purely parasitic on non-derivative value. But, I would say, derivative value is none the less a genuine kind of value. It makes perfect sense, and is perfectly correct, to say that some things are good, others bad, even if only derivatively so. Here, though, we need not try to resolve this issue, because the objection misrepresents my position. It is true that I take it that s4 and s5 require equal disfavour from me in light of (what I am assuming for the sake of illustration to be) their equal basic intrinsic badness. But it is not true that I take it that s4* permits or requires special disfavour from me in light of its non-basic intrinsic badness. I do think that s4* is non-basically intrinsically bad, and I do think that I may, and perhaps even must, respond partially to it, but it is not its non-basic intrinsic badness that renders partiality on my part fitting. It is something else that does so. Perhaps we should say, in a manner reminiscent of the second response above, that it is a certain (non-derivative) personal value that s4* has that renders such a response fitting. (This would be perfectly consistent with saying that s4* is, in terms of intrinsic value, evaluatively superfluous.) In any case, whatever the explanation, it is not to be traced merely to the fact that s4* is non-basically intrinsically bad.

A final objection is that my proposal requires that it be a conceptual truth that basic intrinsic value is an impersonal kind of value and that, while this may be true, it is not conceptually so. Here I am not sure what to say. I am indeed inclined to regard it as a conceptual truth (which is not to say that it is an obvious one). However, I am also quite willing to retreat from (FA1b) to

(FA1b*) Necessarily, x is basically intrinsically good if and only if

(1) it is fitting for anyone who contemplates x to favour it for its own sake

and

(2) there is no y such that y is a proper part of x and it is fitting for anyone who contemplates y to favour it for its own sake

and also, of course, to modify (FA2b)–(FA4b) accordingly. This somewhat weaker combination of theses meets the challenge from partiality equally well and is, besides, not open to Blanshard's objection (with which, as I noted in Sect. 3, I have some sympathy) that it is x's value that grounds the requirement to adopt some attitude toward it.

As an addendum, let me say something briefly about the inverse challenge from partiality having to do with the question whether it is fitting for one to favour one's own pleasures as one favours those of others. The strategy that I applied to the original challenge can be applied here too. A distinction is to be drawn between [Michael, being pleased to degree 10, now], on the one hand, and [Michael, being someone who is Michael and who is pleased to degree 10, now], on the other. If the former state is basically intrinsically good, then my proposal implies that I, like, everyone else who contemplates it, am required to favour it. But if the former state is basically intrinsically good, then the latter state is not, and so my proposal is silent on whether I am required also to favour it. Perhaps, as Ross and others hold, I am not. I will not venture an opinion on the matter here.

5. Concluding remarks

Two quick comments in conclusion.

First, it may be that partiality of a sort different from that exemplified by parental bias is sometimes fitting when responding to value. Time heals all wounds, some say. Others say that, whether or not it does, it ought to or may do so — that it is fitting to disfavour a recent tragedy, say, more strongly than one that occurred a long time ago. Lemos, for example, says:

Different emotions and feelings toward the death of one's child or one's parents seem more appropriate to these events when they are recent than when they have receded into the past. It seems inappropriate to have precisely the same emotions and feelings about their death at different times, for there comes a point when it is appropriate not to feel acute grief and sadness, to let these feelings pass, and to feel in their place a sadness that is less intense. (Lemos 1994, pp. 17–18)

Temporal partiality of the sort Lemos describes is, like the parental partiality with which we have been concerned, perfectly in keeping with my claim that the contemplation of states with basic intrinsic value requires impartiality. Indeed, the door is open for the fittingness of partiality of various sorts. However, that other sorts of partiality are sometimes fitting seems to me not obvious. Or more exactly: while it certainly seems sensible to say, like Lemos, that it is fitting that grief should diminish over time, whether it is directly fitting that it should do so, or merely indirectly fitting (since at some point we need to get on with our lives), is, I think, a difficult question.

Second, I have said very little about just what is involved in ‘contemplating’ something of intrinsic (dis)value, such as one's own child's suffering or a distant child's suffering. It is very likely that, in the former case, one's mental representation of the state will be vivid, graphic, whereas in the latter case it will be relatively pallid and sketchy. Moreover, graphic contemplation of some tragedy is likely to evoke far stronger disfavour than sketchy contemplation will, a fact that some may take to explain both why a father typically does respond more strongly to his own child's suffering than to a distant child's suffering and why it is fitting that he do so. But I do not think this can be quite right, since the graphic–sketchy distinction, however precisely it is to be drawn, does not in fact seem to square with what it is plausible to say about the fittingness of partiality. It is quite possible, after all, for one to contemplate a distant child's suffering in vivid detail (especially if aided by on-site cameras), and it is also possible to contemplate one's own child's suffering in a less vivid way (especially if one has received the news second-hand). In such a case, one may find oneself reacting more strongly to the former than to the latter, and such a disparity in reaction seems to be neither objectionable nor expressive of partiality toward the distant child. What I think we must do, when trying to measure the extent to which someone displays partiality and trying to determine whether such a display is fitting, is compare cases in which the person's contemplation is held constant on the graphic–sketchy scale. Thus the fittingness of the sort of parental partiality we have been discussing should be understood in terms of its being fitting to disfavour one's own child's suffering more strongly than a distant child's suffering, when one contemplates these instances of suffering in equally graphic or sketchy ways. This would allow for the possibility that a parent disfavour a distant child's suffering more strongly than his own child's suffering, without thereby violating any norms of fittingness, if he contemplates the former more graphically than the latter. Similarly, I think we should say that it is possible that one respond unequally to states that have

equal basic intrinsic value, without thereby violating any norms of fittingness, as long as one contemplates one of these states more graphically than the other.²¹

Footnotes

1 See among others: Broad 1930, p. 283; Ross 1939, pp. 275–6; Ewing 1948, p. 152; Chisholm 1986, p. 52; Lemos 1994, pp. 12, 15; Scanlon 1998, p. 96.

2 See, for example, Korsgaard 1983, Kagan 1998, Rabinowicz and Rønnow-Rasmussen 1999.

3 To put the analysis somewhat more precisely: *x* is such that, for all persons *S*, the favouring of it by *S* for its own sake would be fitting if *S* were to contemplate *x*. Such a subjunctive reading is to be applied to all relevant propositions that follow.

4 See, for example, Broad 1930, Ross 1939, Chisholm 1986, Lemos 1994. See n. 1 above.

5 Other difficulties concern: the possibility of various kinds of neutrality; proportionality between (dis)favourings; and impossible (dis)favourings. On the first two issues, see the discussion in Zimmerman 2001, pp. 115–17 and 117–19. On the third issue, see Bykvist 2009.

6 Cf. Rabinowicz and Rønnow-Rasmussen 2004, p. 407, which borrows from Crisp 2000, p. 459.

7 Among others: Parfit 2001; Rabinowicz and Rønnow-Rasmussen 2004; Olson 2004; Persson 2007; Danielsson and Olson 2007.

Cf. Chisholm 1986, p. 52: ‘the contemplation of just *A* and *B* as such by *x* requires that *x* prefer *A* as such to *B*’. Cf. also Danielsson and Olson (2007), who make a very similar point when they propose solving the wrong-kind-of-reason problem by means of invoking a distinction between ‘content-reasons’ and ‘holding-reasons’.

9 Strictly speaking, it is only basic intrinsic value that is non-derivative. (See Zimmerman 2001, pp. 25, 154 ff.) I will return to the concept of basic intrinsic value below.

10 Exceptions include: Ross 1939, p. 282; Lemos 1994, pp. 12–13; Sumner 1996, pp. 20 ff.

11 This is only a rough account of the identity conditions of states. For refinement, see Zimmerman 2001, pp. 54 ff.

12 Olson 2009, pp. 375–6. Here Olson is criticizing remarks that I made in an earlier unpublished version of this paper.

13 I take it that, although the ‘may’ and ‘ought’ that express permissions and duties of action imply ‘can’ (in some suitably robust sense), those that express permissions and duties of attitude

do not. (This does not, I believe, disqualify permissions and duties of attitude from being moral permissions and duties. Contrast Olson 2009, p. 375.) So the fact that a person finds himself unable to feel compassion for a stranger's suffering, say, does not alter the fact that compassion is none the less fitting.

14 Oddie writes in terms of desire and aversion in particular, rather than favour and disfavour in general.

15 In so saying, I do not mean to suggest that, as a parent, one should not be especially concerned with trying to turn one's child away from a life of vice, but only that, in the meantime, it is fitting that, as a parent, one welcomes his being happy.

16 There may be some exceptions to this claim, although I am inclined to think not. Perhaps on occasion a reaction that is of the 'right sort' can be so disproportionate to its object — so far under or over the mark — as to be positively unfitting.

17 Why not say that s_3 is good, but just not good to any precise degree? This may be a tempting position to take as long as we are dealing only with unsophisticated theories such as the very simple hedonism under discussion, but its shortcomings become apparent as soon as we attempt to apply it to certain more sophisticated theories. Consider, then, a less simple version of hedonism according to which the object of pleasure does affect the value of pleasure: whereas benevolent pleasure is intrinsically good, malicious pleasure is intrinsically bad. Suppose that at midnight Peter takes pleasure to degree 10 in Paul's suffering. This episode comprises the following states, among others: [Peter, taking malicious pleasure to degree 10, midnight] and [Peter, taking pleasure to some degree in something, midnight]. Let us assume that the hedonism we are now working with once again assigns degrees of value precisely in proportion to degrees of pleasure (and pain), but does so in two ways: directly, if the pleasure is not malicious; inversely, if the pleasure is malicious. Then the first of the two states just mentioned will turn out to be intrinsically bad to degree 10. But what about the second? Not only is it too thin for us to assign any precise degree of value to it, it is so thin that we cannot even say whether it is good or bad. Should we say that it none the less has value, even though it is neither good, nor bad, nor neutral? What sense would there be in such a declaration? What could tempt us to make it, short of a blind allegiance to the claim that all states have intrinsic value? Since there is no good reason to assign any intrinsic value to the too-thin state in this case — indeed, there is good reason not to — there is equally good reason not to assign intrinsic value to too-thin states in general; for such a practice affords a simple and uniform approach to the evaluation of all such states.

18 What follows is a solution I first proposed in Zimmerman 2001, pp. 119 ff.

19 Here Bykvist is criticizing my account as I presented it in Zimmerman 2001, pp. 121–3.

20 Sonnet 116. Cf. Frankfurt 1999, p. 166: 'The significance to the lover of what he loves is not that of an exemplar ... but ineluctably particular.'

21 Earlier versions of this paper were presented to audiences at the Universities of Lund, Montreal, Stockholm, and Uppsala, as well as at the conference on Partiality and Impartiality in Ethics held in 2007 at the University of Reading for which it was originally written. For helpful comments on these or other occasions I am very grateful to the Editor and to Per Algander, David Alm, Lori Anthony-Cummings, Gustaf Arrhenius, Johan Brännmark, Krister Bykvist, Erik Carlson, John Cottingham, Sven Danielsson, Dan Egonsson, Karin Enflo, Björn Eriksson, Brian Feltham, Kathrin Glüer-Pagin, Brad Hooker, Noah Lemos, Jonas Olson, Francesco Orsi, Jan Österberg, Peter Pagin, Christian Piller, Wlodek Rabinowicz, Andrew Reisner, Toni Rønnow-Rasmussen, Robert Shaver, Philip Stratton-Lake, Sarah Stroud, Daniel Svensson, Christine Tappolet, Åsa Wikforss, and Fiona Woollard.

References

Blanshard Brand. Reason and Goodness. London: George Allen and Unwin; 1961.

Brand M, Walton D, editors. Action Theory. Dordrecht: Reidel; 1976.

Brentano Franz. The Origin of Our Knowledge of Right and Wrong, tr. R. Chisholm and E. Schneewind. London: Routledge and Kegan Paul; 1969. Originally published in German as Vom Ursprung sittlicher Erkenntnis. Leipzig: Duncker und Humblot, 1889.

Broad C D. Five Types of Ethical Theory. London: Kegan Paul, Trench, Trubner; 1930.

Bykvist Krister. 'No Good Fit: Why the Fitting Attitude Analysis of Value Fails'. *Mind* 2009;118:1-30.

Chisholm Roderick M. 1974. 'Practical Reason and the Logic of Requirement'. In Körner 1974, pp. 1–17.

Chisholm Roderick M. Brentano and Intrinsic Value. Cambridge: Cambridge University Press; 1986.

Crisp Roger. Review of Joel Kupperman, Value ... and What Follows (Oxford: Oxford University Press, 1998). *Philosophy* 2000;75:458-62.

Danielsson Sven, Olson Jonas. 'Brentano and the Buck-Passers'. *Mind* 2007;116:511-22.

Egonsson D, Josefsson J, Petersson B, Rønnow-Rasmussen T, editors. Exploring Practical Philosophy. Aldershot: Ashgate; 2001.

Ewing A C. The Definition of Good. New York: Macmillan; 1948.

Frankfurt Harry G. *Necessity, Volition, and Love*. Cambridge: Cambridge University Press; 1999.

Kagan Shelly. 'Rethinking Intrinsic Value'. *Journal of Ethics* 1998;2:277-97.

Kim Jaegwon. 1976. 'Events as Property Exemplifications'. In Brand and Walton, 1976, pp. 159–77.

Körner S, editor. *Practical Reason*. Oxford: Basil Blackwell; 1974.

Korsgaard Christine M. 'Two Distinctions in Goodness'. *Philosophical Review* 1983;92:169-95.

Lemos Noah. *Intrinsic Value*. Cambridge: Cambridge University Press; 1994.

Moore G E. *Principia Ethica*. Cambridge: Cambridge University Press, rpt 1993; 1903.

Moore G E. *Philosophical Papers*. London: George Allen and Unwin; 1959.

Oddie Graham. *Value, Reality, and Desire*. Oxford: Clarendon Press; 2005.

Olson Jonas. 'Buck-Passing and the Wrong Kind of Reasons'. *Philosophical Quarterly* 2004;54:295-300.

Olson Jonas. 'Fitting Attitude Analyses of Value and the Partiality Challenge'. *Ethical Theory and Moral Practice* 2009;12:365-78.

Parfit Derek. 2001. 'Rationality and Reasons'. In Egonsson, Josefsson, Petersson, and Rønnow-Rasmussen 2001, pp. 17–39.

Paul E F, Miller F D, Paul J, editors. *The Good Life and the Human Good*. Cambridge: Cambridge University Press; 1992.

Persson Ingmar. 2007. 'Primary and Secondary Reasons'.
<http://www.fil.lu.se/HommageaWlodek/site/papper/PerssonIngmar.pdf>.

Rabinowicz Wlodek, Rønnow-Rasmussen Toni. 'A Distinction in Value: Intrinsic and For Its Own Sake'. *Proceedings of the Aristotelian Society* 1999;100:33-52.

Rabinowicz Wlodek, Rønnow-Rasmussen Toni. 'The Strike of the Demon'. *Ethics* 2004;114:391-423.

Ross W D. *Foundations of Ethics*. Oxford: Oxford University Press; 1939.

Scanlon Thomas M. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press; 1998.

Sumner L W. *Welfare, Happiness, and Ethics*. Oxford: Clarendon Press; 1996.

Thomson Judith Jarvis. 1992. 'On Some Ways in Which a Thing Can Be Good'. In Paul, Miller, and Paul 1992, pp. 96–117.

Williams Bernard. *Moral Luck*. Cambridge: Cambridge University Press; 1981.

Zimmerman Michael J. *The Nature of Intrinsic Value*. Lanham: Rowman and Littlefield; 2001.

Zimmerman Michael J. 'The Good and the Right'. *Utilitas* 2007;19:326-53.