

Value and Reasons to Favour¹

Jonathan Way

According to *fitting-attitudes* accounts of value (FA accounts), what it is for something to be good or of value is for it to be a *fitting* object of a pro-attitude. Some FA accounts take this notion of fittingness as primitive; others understand it in terms of normative reasons. Here I shall be concerned with this latter sort of account – accounts which hold that for something to be good is for there to be reasons to favour that thing.

A central attraction of FA accounts, so understood, is that they promise straightforward explanations of the various connections between reasons and value. If what it is for something to be of value is for there to be reasons to favour that thing, then the connections between reasons and value may be a consequence of the nature of value.² In the first two sections of this paper, I motivate and defend one such connection, and show how it is nicely explained by a FA account.

Of course, this form of argument for FA accounts is highly defeasible. If the connections between reasons and value are better explained by an alternative account, these connections do not give us strong reason to accept an FA account. An adequate

¹ Versions of this paper were presented at: the 8th Annual Metaethics Workshop at the University of Wisconsin, Madison; the Theoretical and Practical Rationality Workshop at the Institute for Advanced Study, Hebrew University of Jerusalem; the Aims and Norms: Emotions Workshop at the University of Southampton; and the University of Stirling. I would like to thank the audiences on each of these occasions for very helpful discussion. For written comments or discussion I am especially grateful to John Brunero, David Enoch, Guy Fletcher, Alex Gregory, Errol Lord, Conor McHugh, Jonas Olson, Richard Rowland, Jussi Suikkanen, Andrew Reisner, Pekka Väyrynen, Daniel Whiting, and two anonymous referees for this volume.

² See, for example, Danielsson and Olson, “Brentano and the Buck Passers”, p. 520; Rabinowicz and Rønnow-Rasmussen, “The Strike of the Demon”, pp. 391-2; and Suikkanen, “Reasons and Value”, p. 529. Scanlon’s much discussed ‘redundancy argument’ for FA accounts (*What We Owe To Each Other*, p. 97) should also, I think, be understood as being of this form (cf. Hooker and Stratton-Lake, “Scanlon versus Moore on Goodness”, p. 161). For general discussions of motivations for FA accounts, see Olson, “Buck-Passing Accounts”; Reisner, “Abandoning the Buck-Passing Analysis of Final Value”; and Suikkanen “Buck-Passing Accounts of Value”.

assessment of FA accounts thus requires us to pay close attention to the alternatives to FA accounts. In the third section of this paper, I thus consider whether such alternatives can adequately explain the connection I defend.

The principal class of alternatives to FA accounts I will be interested in are *value-first* views. Value-first views analyse reasons in terms of value, rather than the other way round. A well-known example of this sort of view, although concerning rightness rather than reasons, was held by G.E. Moore.³ Moore held that what it is for an act to be right is for it to maximize value. Just so, a value-first account of reasons might hold that what it is for there to be a reason for you to do something is for your doing that thing to be of some value. Several writers have taken this Moorean view to be one of the central alternatives to FA accounts.⁴

In one way it is surprising that this Moorean view is so often taken as the paradigm alternative to an FA account. The Moorean view is an account of reasons for *action*. But the connections between reasons and value which motivate FA accounts do not hold only, or even primarily, between value and reasons for action; they hold between value and reasons to *favour* – for instance, to desire, choose, hope or wish for, take pleasure in, be glad of, and all the rest of the so-called ‘pro-attitudes’. Thus to assess whether FA accounts provide the best explanation of these connections, we should focus, not on value-first accounts of reasons for action, but on value-first accounts of reasons to favour. Surprisingly though, there has been little discussion of such views.

A central aim of this paper is thus to begin to explore what a value-first view of reasons to favour might look like. I will outline two options that value-first theorists might want to pursue, and bring out some of the commitments these options

³ Moore, *Principia Ethica*.

⁴ See Scanlon *What We Owe To Each Other*, chap. 2; Dancy “Should We Pass the Buck?”; Suikkanen “Reasons and Value”; and Hooker and Stratton-Lake, “Scanlon versus Moore on Goodness”.

involve. However, my aims are not purely exploratory – I also hope to provide some support for FA accounts. I will argue that an FA account better explains the connection between reasons and value which I defend than the two value-first theories I consider. While this certainly does not show that value-first theorists cannot adequately explain this connection, it does suggest that value-first theorists cannot explain it as easily as proponents of FA accounts. This connection between reasons and value fits better with an FA account than with a value-first account.

Three preliminary points are in order. First, I will primarily be concerned with the goodness *simpliciter* of outcomes, or states of affairs. Goodness *simpliciter* contrasts with what is good *for* some person or group. For instance, it might be good (*simpliciter*) if happiness was distributed in proportion to desert, but this would not be good for those who do not deserve to be happy.⁵ While proponents of FA accounts often offer accounts of good-for, and of the other “ways of being good”, I shall not focus on such accounts here.

Second, the FA accounts I will be interested in are intended as *constitutive accounts* of the property of goodness. The aim of a constitutive account is to tell us *what it is* for something to be good – to reveal the *nature* of goodness. Such an account is to be contrasted both with substantive accounts of goodness, which aim to tell us what things are good, as well as with accounts of the concept GOODNESS, or the meaning of the word ‘good’. The value-first views I will be interested in have parallel aims – they aim to tell us what it is for there to be a reason to favour something, rather than to tell us about the concept REASON TO FAVOUR, the meaning of ‘reason to favour’, or the things we have reasons to favour.⁶

⁵ ‘Goodness-for’ is also sometimes called welfare, well-being, prudential value, and various other things.

⁶ For discussion of the notion of a constitutive account see, e.g. Schroeder, *Slaves of the Passions*, pp. 61-72 and Wedgwood, *The Nature of Normativity*, pp. 136-44.

Third, I should emphasize that I do not take value-first views, so understood, to be the only alternative to FA accounts. There are many others. For instance, there are ‘no-priority’ views on which neither reasons nor value are analysed in terms of the other. On one such view, reasons and value simply form distinct, primitive domains. On others, reasons and value are explained in terms of some third factor, such as desires, or rationality. There are also possible views on which value is grounded in, but not reducible to, reasons (or vice-versa).⁷ At most then, the argument of this paper only takes the first steps towards a full defence of an FA account. Nonetheless, since value-first views are often taken to be the natural alternatives to FA accounts, they are important steps to take.

1. Reasons and Value: The Linking Principle

In this section, I present and defend the connection between reasons and value I will be interested in. We can begin with the familiar observation that reasons admit of a distinction between the contributory (or ‘pro tanto’) and the overall. Famously, for example, that you have promised to meet a friend may be *a* reason to meet your friend, even if there is *most reason* for you not to meet your friend – say, because meeting your friend would prevent you from saving someone’s life. Much the same goes for reasons to believe and to favour. For example, the fact that the local news reports that the sun will shine tomorrow may be a reason to believe that the sun will shine tomorrow, even if it not a sufficient reason to believe this – say, because there is some reason to doubt what the local news says on such matter. And the fact that

⁷ For the first kind of ‘no priority’ view, see Dancy, “Should We Pass the Buck?”. Smith (*The Moral Problem*) might be taken to defend the second kind of view. For discussions of grounding see, e.g. Fine, “The Question of Realism”; Rosen, “Metaphysical Dependence”; and Schaffer, “On What Grounds What”.

smoking can be pleasant may be a reason to want to smoke, even if it is not a sufficient reason to want to smoke, because of the bad effects on your health.

We see something very similar in the evaluative realm. Outcomes are often pro tanto good, or good in some respects, without being good overall. And when an outcome is pro tanto good, we can at least typically say in what respect it is good. For instance, it might be good in some respects if wealth was distributed equally – say, because then everyone would have enough to meet basic needs – even if this would not be good overall – say, because it meant that the poorest would be poorer than they otherwise could be. And it might be good in some respects if I took up smoking – for smoking would give me some pleasure – even if this would not be good overall, because of the bad effects on my health.

This is an interesting structural similarity between the normative and evaluative domains. But what is more striking, I think, is that these structures are intimately connected. To start with, it seems plausible that if an outcome is good in some respect, there is some reason to favour that outcome. But I think that something stronger is true. I suggest that the following *linking principle* is true:

If some consideration is a respect in which an outcome is good, that consideration is also a reason to favour that outcome.

This principle certainly seems to hold in the examples I have given so far. If a respect in which it would be good if wealth was distributed equally is that everyone would have enough to meet basic needs, then the fact that everyone would have enough to meet basic needs is a reason to want wealth to be distributed equally. And if a respect in which it would be good if I took up smoking is that smoking would give me

pleasure, then the fact that smoking would give me pleasure is a reason to want me to take up smoking.

The linking principle helps to shed light on several parallels between the evaluative and the normative. For example, in both domains, we can distinguish between the final and the instrumental. One respect in which an outcome may be good is if it facilitates a further good outcome. Another respect in which an outcome may be good is if it contains something of final value – something that is good for its own sake. In much the same way, we can distinguish between reasons to favour an outcome for its own sake, and reasons to favour an outcome because of what it will facilitate. I take this to be a difference in the contents of reasons to favour – a difference in the considerations which constitute such reasons. And it is a difference we should expect, if respects in which a thing is good are also reasons to favour that thing.⁸

The linking principle also helps to explain why the theory of reasons and the theory of value face so many of the same issues. Consider, for example, the recent debates between atomists and holists in the theory of reasons. Atomists hold that if some consideration is a reason to do something in one circumstance, then that consideration must be a reason to do that thing in any circumstance in which it obtains. Holists reject atomism, claiming that a consideration can be a reason to do something in one circumstance, even if it is not a reason to do so in others. An exactly parallel issue arises about value – if something would be good in one circumstance, does it follow that it is good in all others? What is more, much the same examples that

⁸ There is an interesting related parallel. When we are dealing with something that is instrumentally good, or which there is instrumental reason to want, it is easy to say what the reason to want it is, and to specify the respect in which it is good. By contrast, when something is of final value, or something which there is final reason to want, it can be difficult to say – in any informative way – what the reason to want it is, or to specify the respect in which it is good. It can sound strained, or at least uninformative, to say that a respect in which someone's pleasure is good, and a reason to favour it, is that it is a pleasure.

are used to motivate holism about reasons can also be used to motivate holism about value. Cases of sadistic pleasure – the torturer who enjoys his work, and the like – are often held up as cases in which the fact that someone would enjoy doing something is not a reason for them to do it, and also as cases in which pleasure is not good.⁹ It thus seems highly plausible that atomism about reasons and atomism about value stand or fall together. The linking principle can help to explain this. For example, if pleasure is always good, then the linking principle tells us that the fact that doing something would be pleasant is always a reason to favour it. Thus given the linking principle, atomism about value supports atomism about reasons, and so holism about reasons to favour supports holism about value.

The linking principle thus seems to be intuitively plausible, and helps to explain some notable parallels between reasons and value. However, there is a possible worry about the linking principle which we should acknowledge. As several writers have observed, reasons seem to be linked to agents in a way that goodness is not.¹⁰ This might seem to make room for counter-examples to the linking principle. Since reasons are always reasons for agents, a feature of an outcome is a reason to favour that outcome only if some agent is suitably related to it. But on the face of it, an outcome can be good whether or not there are any agents suitably related to it. And if so, there can be outcomes which are pro tanto good but which no-one has any reason to favour.

A simple amendment to the linking principle goes most of the way to meeting this worry. Rather than saying that respects in which an outcome is good are also reasons to favour that outcome, we can say that (necessarily) respects in which an

⁹ For holism about both reasons and value, see Dancy *Ethics without Principles*. For discussion of holism about value see, e.g. Fletcher, “Rejecting Well-Being Invariabilism” and Hurka “Two Kinds of Organic Unity”.

¹⁰ Dancy “Should We Pass the Buck?”, p. 170; Suikkanen, “Reasons and Value”, pp. 531-4; Väyrynen, “Resisting the Buck-Passing Account of Value”, n.18.

outcome is good are reasons for any suitably related agents to favour that outcome. This principle is not undermined by the fact that there may be some good outcomes to which no agents are in fact suitably related. However, some might think that this response does not entirely meet the worry. For we might think that there can be good outcomes to which agents *cannot* be suitably related. When applied to such outcomes, the linking principle would be trivially true at best.

Consider an example of Krister Bykvist's.¹¹ Bykvist imagines a state of affairs in which there are happy egrets but no past, present, or future agents. This state of affairs is good in some respect – the egrets are happy. The amended linking principle thus implies that the fact that the egrets are happy is a reason for any suitably related agents to favour this outcome. But we might think that this is true, if at all, only because there can be no agents suitably related to this outcome – for if there were any agents around to do any favouring, we would have a different state of affairs.

However, the fact that there can be no agents in this state of affairs is not a good reason to think that there can be no agents suitably related to this state of affairs. (Nor does Bykvist claim otherwise). The notion of a 'suitably related' agent which is relevant here is psychological and perhaps also epistemic. In this sense, for example, we are not suitably related to, and so lack reasons to have attitudes towards, outcomes which we lack the concepts to conceive of. But there is nothing to stop agents being in this sense suitably related to states of affairs that they are not a part of. For example, in Bykvist's case *we* have reasons to have attitudes towards the state of affairs of there being happy egrets but no past, present, or future agents – for instance, reasons to be glad that this state of affairs does not obtain.

¹¹ Bykvist, "No Good Fit".

So the observation that reasons are linked to agents in a way that goodness is not does not by itself undermine a suitably amended version of the linking principle. Nonetheless, Bykvist's example might still seem to cast doubt on the principle. It might just seem implausible that the fact that the egrets are happy is a reason to favour this outcome – for instance, to wish that it obtained. There may seem to be no reason to wish for such a desolate outcome.

There are two points to make about this worry. First, any reasons to favour this outcome would be very weak, since there are so many reasons not to favour this outcome. But this should make us less confident that there is *no* reason to favour this outcome, rather than just a very weak reason. Surely, if anything is a reason to wish that this outcome had obtained, it is the fact that the egrets are happy.¹²

Second, the case would not undermine a comparative version of the linking principle. According to this principle, if some consideration is a respect in which one outcome is better than another, this consideration is also a reason (for suitably related agents) to prefer the one outcome to the other. The fact that the egrets are happy is certainly a respect in which there being happy egrets but no agents is better than there being *unhappy* egrets and no agents. But this fact is also plausibly a reason to prefer the former state of affairs to the latter – consider, for example, that it is a consideration which makes the former *preferable* to the latter.¹³ This comparative

¹² For more on the unreliability of intuitions to the effect that there is no reason for some response, see Schroeder, *Slaves of the Passions*, pp. 92-97.

¹³ Bykvist ("No Good Fit", pp. 10-11) would have us deny that there are reasons to prefer the agent-less happy world to the agent-less unhappy world. He suggests that to prefer one state of affairs to another involves being disposed to choose the former state of affairs over the latter, and claims that we do not have reasons to be disposed to choose the agent-less happy world over the agent-less unhappy world, as this is a choice it is logically impossible to face. However, we might equally take this case to cast doubt on views which tie preferences too closely to dispositions to choose. This move does not seem ad hoc. As Bykvist (p. 11) acknowledges, we should allow that we can have pro-attitudes towards past outcomes, and other outcomes over which we have no control. Preferences seem no exception to this rule – I can prefer the Republicans not to have won the 2000 election; some people prefer never to have been born. It is not clear why logical impossibility, as opposed to historical or causal impossibility, is supposed to make such a difference here.

linking principle could also be used to motivate a FA account, and much of the discussion to follow could be framed in these terms. For brevity though, I will stick to the original formulation (often omitting the qualification about suitably related agents).

2. Explaining the Linking Principle

If the linking principle is true, we should ask why it is true. For it is surely not a brute fact that reasons and value are connected in such a systematic way. I think that FA accounts provide a simple answer. Consider the following FA account:

(FA) For R to be a respect in which an outcome O is good is for R to be a reason for all (suitably related) agents to favour O.

FA is intended as a constitutive account of goodness – an account of what it is for something to be good. Thus if FA is true, the linking principle follows directly from the nature of goodness. Respects in which an outcome is good are also reasons to favour that outcome because what it is for a consideration to be a respect in which an outcome is good is for it to be a reason to favour that outcome.¹⁴

¹⁴ Since the linking principle seems like a relatively fundamental fact about the connection between reasons and value, I find it plausible that an explanation of the linking principle should take this form – that it should appeal to a constitutive account of reasons or value. However, this might be disputed. For instance, it might be suggested that we can explain the linking principle by appealing only to substantive normative principles. For instance, suppose that there is a basic normative principle directing us to favour the good. We might think it follows from this principle that respects in which an outcome is good are reasons to favour that outcome. (Just as it might seem to follow from a basic normative principle directing us to favour happiness that respects in which an outcome would involve happiness are reasons to favour that outcome). However, while this is an interesting suggestion, it seems to me that at this point Scanlon's thoughts about the 'redundancy' of goodness as a source of reasons have some force (*What We Owe To Each Other*, p.97). It is not clear why we would need a basic normative principle to favour the good *in addition to* basic normative principles directing us to favour things which are good (e.g. happiness, or justice).

In the next section, I shall consider whether constitutive accounts of reasons to favour in terms of goodness offer similarly straightforward explanations of the linking principle. But before turning to consider this question, we should note two points about FA.

First, FA analyses goodness simpliciter in terms of considerations which are reasons for *all* (suitably related) agents. This restriction captures the natural thought that goodness simpliciter is impartial or agent-neutral, rather than being relativised to particular agents.¹⁵ It also helps to highlight the normative significance of goodness simpliciter. If what it is for something to be good simpliciter is for it to be the kind of thing which *everyone* has reason to favour, then it is no mystery why we should care about what is good simpliciter.¹⁶

Second, FA faces the famous ‘wrong kind of reason’ objection. According to this objection, not all reasons for everyone to favour an outcome are respects in which that outcome would be good. For example, if an evil demon threatens to destroy the world unless everyone desires another World War, then the fact that the demon has made this threat is a reason for everyone to favour another World War, although it is not a respect in which another World War would be good.¹⁷ I shall not address this problem here. In other work, I argue that proponents of FA accounts should respond by rejecting counter-examples of this kind. They should claim that the demon’s threat

¹⁵ The ‘partiality challenge’ to FA accounts turns on this point. For discussion, see Olson “Fitting Attitudes Analyses of Value and the Partiality Challenge” and Zimmerman, “Partiality and Intrinsic Value”.

¹⁶ This approach also suggests that we might extend the FA account to goodness-for and the other “ways of being good” by restricting the class of agents appealed to. For instance, perhaps what it is for something to be a respect in which an outcome is good-*for* X is for it to be a reason for all suitably related agents who care about X to favour that outcome, because they care about X. For suggestions along these lines, see Darwall, *Welfare and Rational Care* and Schroeder, “Value and the Right Kind of Reason”. For discussion see Fletcher, “Resisting Buck-Passing Accounts of Prudential Value” and Heathwood, “Fitting Attitudes and Welfare”.

¹⁷ For discussion see e.g. Rabinowicz and Rønnow-Rasmussen, “The Strike of the Demon”; Danielsson and Olson, “Brentano and the Buck-Passers”; Schroeder, “Value and the Right Kind of Reason”; and Way, “Transmission and the Wrong Kind of Reason”.

is not a reason to desire another World War but only a reason to desire that everyone desires another World War and to bring it about that this is so.¹⁸ I shall here assume that this response is successful. If I am wrong about this, FA and the discussion to follow would have to be modified accordingly.

3. What are the Alternatives to FA Accounts?

A natural reaction to the discussion thus far is to grant that FA accounts do indeed explain why reasons and value are connected in the way stated by the linking principle, but to insist that this connection can be just as easily explained by a value-first account – a constitutive account of reasons in terms of value, rather than vice-versa.¹⁹ The aim of the rest of this paper is to show that things are not quite so straightforward.

How might we give an account of reasons to favour in terms of value? A tempting initial thought is that all that value-first theorists need to do is to reverse the direction of priority of the FA account. I shall begin by considering this idea. I shall argue that while an account of this form would indeed explain the linking principle, it would not be a promising account of reasons to favour. I shall then turn to another natural idea – that we can see how to develop a value-first account of reasons to favour by looking to more familiar value-first accounts of reasons to act. While I agree that this is a good strategy for value-first theorists to pursue, I shall argue that this approach struggles to explain the linking principle, and that it faces an independent problem.

¹⁸ Way, “Transmission and the Wrong Kind of Reason”. See also Parfit, *On What Matters*, appendix A, and Skorupski, *The Domain of Reasons*, chap 4.

¹⁹ This sort of reaction is on view in the discussion of the motivations for FA accounts in Olson, “Buck-Passing Accounts” and Suikkanen, “Buck-Passing Accounts”. It is often encountered in conversation.

3.1. The Object-Based View

FA says that for some consideration to be a respect in which an outcome is good is for that consideration to be a reason for everyone to favour that outcome. The value-first theorist might suggest that we begin by reversing the order of analysis here:

(OB₁) For R to be a reason for all agents to favour an outcome O is for R to be a respect in which O is good.

Since this view takes reasons to favour to be respects in which the *object* of favouring is good, I shall call this the *object-based view*.²⁰

The object-based view certainly explains the linking principle. On this view, respects in which an outcome is good are also reasons to favour that outcome because what it is for a consideration to be a reason to favour an outcome is for it to be a respect in which that outcome is good. The linking principle follows from the nature of reasons.

However, OB₁ raises two immediate worries. First, it is natural to suppose that for some consideration to be a reason for *all* agents is just for it to be a reason for *each* agent. But if that is right, value-first theorists should begin by giving an account of reasons for an agent, rather than giving an account of reasons for all agents.

Second, it is difficult to see how OB₁ could count as an account of the reason-relation – the relation which holds between a consideration, agents, and a pro-attitude, when that consideration is a reason for those agents to have that pro-attitude. Since its right-

²⁰ The object-based view is suggested in passing by Suikkanen, “Buck-Passing Accounts”, p. 770.

hand side mentions neither agents nor a pro-attitude, it does not look to be of the right form to tell us what it is for a consideration to stand in a relation to agents and a pro-attitude.²¹

The object-based view can be revised so as to avoid these worries. Consider:

(OB₂) For R to be a reason for S to favour an outcome O is for R to be a respect in which the content of S's favouring O is good.

OB₂ is an account of what it is for a consideration to be a reason for an agent and it is of the right form to be an account of the reason-relation. And OB₂ still plausibly explains the linking principle. If R is a respect in which an outcome is good, then R is a respect in which the content of anyone's favouring O is good, and so a reason for anyone to favour O.

However, this points towards a first serious problem for the object-based view. If R is a respect in which the content of S's favouring O is good, then R is a respect in which the content of anyone's favouring O is good. So OB₂ entails that if some consideration is a reason for S to favour an outcome, then it is a reason for anyone to favour that outcome. OB₂ thus seems to rule out the possibility of agent-relative reasons to favour – considerations which are reasons for just some people to favour an outcome. But it certainly seems as if there can be such reasons. For instance, if we learn that only one person survived some disaster, then the fact that Sally is your daughter may be a reason for you to hope that she survived, even if it is not a reason for everyone to hope for this. And if you face a choice between killing one and letting

²¹ Compare: an account of what it is for x to be taller than y must mention both x and y on its right-hand side.

five die, the fact that it would be you killing the one may be a reason for you to prefer that the five die, even if it is not a reason for everyone to prefer that the five die.²²

Now the claim that there are agent-relative reasons of this sort is controversial. So some value-first theorists might be happy to accept the consequence that there are no such reasons. But it is hard to deny that a substantive commitment of this sort is a cost of OB₂. Other things equal, we should prefer a constitutive account of reasons – and an explanation of the linking principle – to leave open contested issues in first-order normative theory.

Value-first theorists might thus revise OB₂ again. They might suggest that there are *further* ways in which a consideration can be a reason to favour, other than by being a respect in which an outcome is good simpliciter. For example, in order to accommodate your reason to hope that Sally survived, the value-first theorist might suggest that respects in which an outcome would be good *for* someone to whom you stand in a special relationship are also reasons (for you) to favour that outcome. And in order to accommodate your reasons to prefer the outcome in which the five die, the value-first theorist might suggest that the outcome in which the five are allowed to die is *better-relative-to-you* than the outcome in which you kill the one, and that respects in which an outcome is *good-relative-to-you* are reasons for you to favour that outcome. Value-first theorists might thus suggest:

(OB₃) For R to be a reason for S to favour outcome O is for R to be a respect in which the content of S's favouring O is good (simpliciter) *or* good for someone to whom S stands in a special relationship *or* good-relative-to-S.

²² For a similar reason, OB₂ seems to rule out the possibility of time-relative reasons – considerations which are reasons to favour an outcome at some times but not others. For possible examples of such reasons, and discussion of the challenges they pose to FA accounts, see Bykvist, “No Good Fit”, p.16; Heathwood, “Fitting Attitudes and Welfare”; and Schroeder, “Value and the Right Kind of Reason”, pp. 45-8.

However, while OB₃ makes room for agent-relative reasons, it seems problematic in at least two ways.

First, the notion of goodness-relative-to which OB₃ appeals to is problematic. As Mark Schroeder has forcefully argued, it is not at all clear what it means to say that an outcome is good-relative-to-you, if this is meant to be distinct from saying, for example, that an outcome is good *for* you, or something you believe to be good (simpliciter).²³ Thus the object-based view accommodates the second kind of case only by relying on a distinction of questionable intelligibility.²⁴

Second, and perhaps more fundamentally, the strategy of avoiding the problem by simply adding disjuncts to the object-based view does not look like a promising way to build a theory. If the aim of a constitutive account is to tell us about the nature of reasons, we should not be satisfied with a disjunctive account. An account of the nature of reasons should tell us what it is that reasons to favour have distinctively in common. But a disjunctive account will not do this – in effect, a disjunctive account tells us that reasons to favour do *not* have anything distinctively in common. OB₃ thus does not look to be of the right form to provide a constitutive account of reasons to favour.²⁵

²³ Schroeder, “Teleology, Agent-Relative Value, and ‘Good’”.

²⁴ Indeed, it seems as if the most attractive way to make sense of talk of agent-relative value is to appeal to a FA account of value (cf. Smith, “Two Kinds of Consequentialism” and Suikkanen, “Consequentialism, Constraints, and the Good-Relative-To”). Given an FA account, we might say, for example, that an outcome is good-relative-to-you if it is an outcome that *you* have reason to favour. Obviously enough though, value-based theorists cannot make this appeal.

²⁵ Might it be replied that what all reasons to favour have in common is that they are respects in which the outcome favoured is in some way good? No: if ‘good’ here means ‘good simpliciter’ then there are reasons to favour outcomes which are not respects in which an outcome favoured is good simpliciter – this is the problem which motivated OB₃. But if ‘good’ means something more general – something which includes all of the “ways of being good” – then there are respects in which an outcome is good which are not reasons to favour. For instance, respects in which an outcome would be good to use as an example in a philosophy paper need not be reasons to favour that outcome. Nor need respects in which an outcome would be good for *you* be reasons for *me* to favour that outcome.

The object-based view thus looks to face a dilemma. OB₂ offers a simple and unified account of reasons to favour but rules out agent-relative reasons to favour. OB₃ accommodates agent-relative reasons but relies on a questionable distinction and offers only a disjunctive account of the nature of reasons. Since neither of these options look attractive, this looks like a serious problem for the object-based view.

The object-based view also faces a second serious problem. We should expect a constitutive account of reasons to satisfy two plausible desiderata. First, the account should apply not only to reasons to favour, but also to reasons to act and believe. After all, when some consideration is a reason for a pro-attitude, and some consideration is a reason to act or believe, it certainly looks as if these considerations stand in the same relation to different responses – to a pro-attitude, action, or belief. A constitutive account of reasons should thus allow that the reason-relation can have each of these different responses as one of its relata, and it should have broadly plausible results when applied to these different responses.²⁶ Second, we should expect a constitutive account of reasons *not* to apply to things for which there cannot be reasons – things like eye colour, feelings of dizziness, or perceptual experiences.

The object-based view does badly on both these counts. Consider the following generalized version of OB₂:

(OB₄) For R to be a reason for S to respond in way W is for R to be a respect in which the content of W is good.

As a general account of reasons this looks problematic. On the face of it, for example, it seems not to apply to actions, since actions are not the sorts of things which have

²⁶ A unified constitutive account is of course compatible with very different *substantive* accounts of reasons for action, belief, and pro-attitudes – very different accounts of the sorts of considerations which favour action, belief, and pro-attitudes.

contents. More seriously still, the view does not provide a plausible account of reasons for belief. What the view tells us is that reasons to believe are respects in which it would be good if what you believe is true. But this seems plainly false. For example, one respect in which it would be good if happiness was distributed in proportion to desert is that the virtuous would be happy. But this is not a reason to believe that happiness *is* distributed in proportion to desert. The object-based view thus seems not to extend naturally to reasons for action or belief.²⁷

Furthermore, OB₄ does seem to allow for some states for which there cannot be reasons. For example, since perceptual experiences have contents, there can be respects in which the content of a perceptual experience is good. And so OB₄ allows that there can be reasons for perceptual experiences. Similarly, if pains have contents, the view will allow that there can be reasons for pains. Of course, the value-first theorist might *stipulate* that the view does not apply to such states. But we might have hoped that an account of reasons would *explain* why it is only some kinds of states, and not others, which are subject to reasons.²⁸

The object-based view thus seems not to satisfy two plausible desiderata on a constitutive account of reasons to favour. This seems to me a second serious objection to the object-based view.

3.2. The Attitude-Based View

²⁷ I say that the second of these objections is more serious because there is a somewhat plausible reply to the first. The object-based theorist might note that actions are intentional under some descriptions, and suggest that we treat the descriptions under which an action is intentional as the ‘contents’ of an action. Understood in this way, the generalized object-based view may be plausible as an account of reasons to act. Whether or not this response succeeds will depend on whether this broader understanding of ‘contents’ can be motivated in a non-ad hoc way.

²⁸ For a related point, see Gibbons, “Things That Make Things Reasonable”, pp. 346-8.

So the tempting first thought – that value-first theorists can match the FA account by simply reversing its direction of priority – turns out not to promise a satisfactory account of reasons to favour. I now turn to a second natural thought – that we can develop a plausible value-first account of reasons to favour by looking to value-first accounts of reasons to act.

Several philosophers have offered such accounts. For example, according to Michael Smith, facts about what we have reason to do are facts about what it is desirable for us to do.²⁹ And according to Joseph Raz, ‘reasons are facts in virtue of which...actions would be good in some respect, and to some degree’.³⁰ More generally, many philosophers have been attracted to a view of reasons to act along the following lines:

For R to be a reason for S to do A is for R to help to explain why it would be good, either instrumentally or for its own sake, if S did A.^{31, 32}

Value-first theorists might thus suggest the parallel account of reasons to favour:

²⁹ Smith, *The Moral Problem*, chap. 5.

³⁰ Raz, *Engaging Reason*, p.23.

³¹ As well as Raz, *Engaging Reason*, see Brunero, “Reasons as Explanations” and Finlay, “The Reasons That Matter”. An slightly different view is that reasons to act are *respects* in which so acting would be good (cf. Wedgwood, “Intrinsic Values and Reasons for Action”). I take the discussion to follow to apply to this view too.

³² It is important that there is some restriction on the “ways of being good” that this analysis appeals to. For not just any way in which an action can be good can provide reasons for it. For example, the facts that help to explain why a certain action would be a good assassination, or a good example of how not to treat people, are not necessarily reasons for that action. However, some might think that it is not only final and instrumental value which provide reasons. For instance, some might think that there is sometimes reason to act because so acting would have symbolic value; others might think that there is sometimes reason to act because so acting would be good *for* you, even if it is not also good simpliciter. As I observe in n.40, these additions would not affect the discussion to follow.

For R to be a reason for S to favour an outcome O is for R to help to explain why it would be good, either instrumentally or for its own sake, if S favoured O.

Since this view analyses reasons for pro-attitudes by appealing to the value of pro-attitudes, I shall call this the *attitude-based view*.

The attitude-based view may seem to lack the immediate appeal of its analogue for reasons to act. For perhaps the clearest way in which a pro-attitude can be good is by being instrumentally good – as, for example, when admiring an evil demon will prevent the end of the world. But at least according to some philosophers, the instrumental goodness of a pro-attitude does not provide reasons for that attitude, but only reasons to want or bring about that attitude.³³ And even if we do not go this far, it seems plain that many of what we might think of as ordinary or standard reasons for pro-attitudes – the ‘right kind’ of reasons – have little to do with the instrumental value of those attitudes. For there are surely many cases in which we have reasons for pro-attitudes which are not instrumentally good. For example, there are plausibly reasons to want people you have never met to be happy, even if your wanting this has no effect on whether they are happy, or any other good effects.³⁴

If the attitude-based view is to be at all plausible then, we will have to think that pro-attitudes can be non-instrumentally good. Just as epistemologists sometimes claim that beliefs which are true, or which constitute knowledge, are of final value, value-first theorists should say that pro-attitudes towards good outcomes are of final value – of ‘conative value’, as we might say. Fortunately, this claim does seem fairly plausible. It seems to be good to want people to be happy, to hope for a cure for

³³ See n.18.

³⁴ Compare Raz, “Reasons: Practical and Adaptive”, pp. 47-8 and Heuer, “Beyond Wrong Reasons”, pp. 177-8.

cancer, and to be glad when justice is done.³⁵ The goodness of these attitudes does not just seem to be a matter of their good effects. These attitudes seem good for their own sake, insofar as they are directed towards good outcomes. More generally then, we might accept the *conative value claim*:

Pro-attitudes towards good outcomes are good for their own sake.

So long as we grant this claim, there is no quick refutation of the attitude-based view. Ordinary or ‘right kind’ reasons to favour can be explained by appealing to conative value. For example, your reasons to want those you have never met to be happy might be explained by the conative value of this desire, even if it has no instrumental value. However, I shall argue that the attitude-based view faces a difficulty explaining the linking principle, and that, like the object-based view, it faces a dilemma concerning agent-relative reasons.

3.2.1. The Attitude-Based View and the Linking Principle

The linking principle says that respects in which an outcome is good are also reasons to favour that outcome. To explain why this is so, the attitude-based theorist might argue as follows:

- (1) If an outcome is good in some respect, then it is good for its own sake to favour that outcome.

³⁵ Compare Hurka, *Virtue, Vice, and Value*, especially chap.1.

(2) If it is good for its own sake to favour an outcome, then respects in which that outcome is good help to explain why it is good to favour that outcome.

(3) Considerations which help to explain why it is good to favour an outcome are reasons to favour that outcome.

So, (4) Respects in which an outcome is good are reasons to favour that outcome.

The conclusion of this argument is the linking principle and (3) follows directly from the attitude-based view. (2) seems highly plausible. For instance, if it is good to hope for a cure for cancer because a cure for cancer would be good then surely the respects in which this would be so – e.g. that a cure for cancer would save many lives – help to explain why it is good for its own sake to hope for a cure for cancer. The crucial claim for the attitude-based view is thus (1).

(1) may seem to follow straightforwardly from the conative value claim.

However, notice that (1) requires a strong version of this claim to be true:

Pro-attitudes towards outcomes which are good in some respect are good for their own sake.

We can contrast this claim with a weaker version of the conative value claim:

Pro-attitudes towards outcomes which are good overall are good for their own sake.

This weaker claim does not support (1). If an outcome is good in some respect but bad overall, the weak conative value claim does not imply that it is good for its own sake

to favour this outcome. Thus unless the strong conative value claim is true, the attitude-based view does not imply that there are reasons to favour such outcomes, and so does not explain the linking principle.

The problem for the attitude-based view is that the strong conative value claim is much less plausible than the weak conative value claim. One way to see this is by considering examples. Consider an outcome in which wealth is distributed equally, but everyone is extremely poor. If egalitarians are right, this outcome is good in some respect. But even egalitarians can admit that this outcome is not good overall. And it does not seem plausible that favouring this outcome is good for its own sake. It does not seem to be good for its own sake to hope for everyone to be extremely, if equally, badly off, or to be glad if this outcome comes about. Or consider Bykvist's example discussed in section one. Although it would be good in one respect if there were happy egrets but no past, present, or future agents, it does not seem to be good for its own sake to wish that this outcome obtained.

We can reinforce these examples by noting that the appeal of the conative value claim is connected to the attractive idea that it is good to have pro- and con-attitudes which are appropriate to the value of their objects.³⁶ But pro-attitudes towards outcomes which are bad overall are not appropriate to their objects. Other things equal, it is appropriate to favour only what is good overall.³⁷ So this attractive idea seems only to support the claim that it is good for its own sake to favour overall good outcomes, not merely pro tanto good outcomes.

I can think of two ways in which attitude-based theorists might try to defend the strong conative value claim in light of these concerns. First, the attitude-based

³⁶ Hurka, *Virtue, Vice, and Value*, p. 30.

³⁷ Things might not be equal if a bad overall outcome is good for someone you care about, or would involve you doing something you should not do. I discuss some of these complications in the next section.

theorist might distinguish between overall and pro tanto pro-attitudes. Thus it might be suggested that there is a difference between, e.g. hoping for something overall and hoping for it in some respect. It might then be said that although it is not good for its own sake to have *overall* pro-attitudes towards merely pro tanto good outcomes, it is good for its own sake to have *pro tanto* pro-attitudes towards such outcomes. For example, it might be said that while it is not good for its own sake to hope overall for a poor but equal distribution of wealth it is good to hope for this outcome in some respect.

It is not entirely clear to me what it is to favour something pro tanto, or in some respect. However, even if we can make sense of this idea, it does not answer the objection. Even if there is such a thing as favouring something pro tanto, the linking principle is plausibly read as saying that respects in which an outcome is good are reasons to favour that outcome *overall*. Thus if egalitarians are right, the equality of the poor but equal distribution is some (outweighed) reason to hope for this outcome overall.³⁸ The attitude-based view explains this reason only if hoping for this outcome overall is good for its own sake. But the response we are considering grants that this is not so. It merely adds that there may be *another* attitude – hoping for this outcome pro tanto – which is good for its own sake. So this first response does not answer the objection.³⁹

Second, the attitude-based theorist might insist that it *is* good for its own sake to favour merely pro tanto good outcomes. To defend this claim, the attitude-based theorist might suggest that just as we should be suspicious of intuitions that there is no reason to do something, we should be suspicious of intuitions that there is nothing

³⁸ Compare the end of section 1 above.

³⁹ This response arguably does confirm the suggestion that it is good for its own sake to have pro-attitudes which are appropriate to their objects, since it is presumably appropriate to have pro tanto pro-attitudes towards pro tanto good outcomes. As we saw though, this suggestion does not support the strong conative value claim.

good about something. When something is in some minor way good, but in far more serious ways bad, we may easily be misled into thinking that this thing is in no way good. It might be thought that this is what is going on in the cases I have described. Since the poor but equal distribution of wealth is very bad overall, it would be bad for its own sake, and to a great extent, to favour this outcome. But this does not mean that it is not also good for its own sake to favour this outcome – it might just be that the goodness of favouring it is heavily outweighed.

This is a fair response but I think that we can go some way to answering it. To start with, the case against the attitude-based view does not rest on the claim that there is nothing good about favouring pro tanto good outcomes. It rests on the more specific claim that it is not good for its own sake to favour such outcomes.⁴⁰ And we can evaluate this claim without relying solely on intuitions about whether favouring some outcome is in any respect good for its own sake. For claims about what is good for its own sake correspond to claims about what is *worth* wanting for its own sake. For example, if outcomes are pro tanto good insofar as they contain happiness (or knowledge, or justice...) then happiness (knowledge, justice) must be worth wanting for its own sake. Thus if pro-attitudes towards merely pro tanto good outcomes are good for their own sake, then it must be worth wanting such attitudes for their own sake.

Now this is not to say that all such attitudes are worth wanting for their own sake. (Just as the claim that happiness is worth wanting for its own sake is not the claim that all instances of happiness are worth wanting for their own sake). So the

⁴⁰ And notice that the attitude-based view really does need the claim that it is good *for its own sake* to favour such outcomes. If this is not so, the attitude-based view explains the linking principle only if respects in which an outcome is good always help to explain why it is instrumentally good to favour that outcome. But that is clearly not so. It need not be instrumentally good to favour a good outcome. And even when it is, the respects in which the outcome is good need not help to explain why. Parallel points apply to versions of the attitude-based view which allow symbolic value or goodness-for to provide reasons (see n.32).

claim in question here is not obviously absurd. Nonetheless, this claim does seem to me to lack independent plausibility. I can see why pro-attitudes towards overall good outcomes might be thought to be worth wanting for their own sake – attitudes of this sort are ordinarily worth having, and so seem worth aiming to have, and worth encouraging in others. But I find it hard to see why we should think the same thing about pro-attitudes towards merely pro tanto good outcomes. Such attitudes are not ordinarily worth having, and so do not seem worth aiming to have, or worth encouraging in others.⁴¹

The strong conative value claim thus continues to seem to me to be doubtful. At the least, it is very unclear whether it is true. It is thus a cost of the attitude-based view's explanation of the linking principle that it is committed to this questionable claim. This is a clear respect in which the FA account's explanation of the linking principle is preferable to that offered by the attitude-based view.

3.2.2. Agent-Relative Reasons and the Sources of Conative Value

In this section, I argue that the appeal to conative value raises a further, independent worry for the attitude-based view.

As we saw in discussing the object-based view, it is plausible that there are agent-relative reasons to favour – considerations which are reasons for just some people to favour an outcome. When Sally is one of the possible survivors of some disaster, the fact that Sally is your daughter is a reason for you, but not for everyone, to hope that she survived. And when you face a choice between killing one and letting

⁴¹ To put the point a slightly different way, it seems strange to think that attitudes of this sort might belong on the value-theorist's list of things of final value, alongside more familiar candidates such as happiness, knowledge, and justice.

five die, the fact that it would be you killing the one is a reason for you, but not for everyone, to prefer that the five die. Can the attitude-based view capture these reasons?

The view certainly faces an initial problem here. (For brevity, I shall focus on the first case). For suppose that it is not good overall if Sally survives. In that case, the conative value claim does not imply that it is of conative value for you to hope that Sally survives. Suppose instead that it is good overall if Sally survives. In that case it is of conative value for you to hope that Sally survives. However, it will also be of conative value for *anyone* to hope that Sally survives. So either the fact that Sally is your daughter will not help to explain why it is of conative value for you to hope that Sally survives, or it will be a reason for everyone to hope that Sally survives. As developed so far then, the attitude-based view struggles to accommodate agent-relative reasons to favour.

There is a solution to this problem. The conative value claim says that it is good for its own sake to favour good outcomes. But it does not say that this is the *only* way in which a pro-attitude can be good for its own sake. And we may well want to deny that this is the case. For example, we might think that it is also good for its own sake for you to hope that Sally survives, even if this outcome would not be good overall. More generally, we might think that it is good for its own sake to have a special concern for those to whom you stand in a special relationship. And if we think this, then the fact that Sally is your daughter will help to explain why it is good for you to hope that she survives – for it will explain why, in hoping for her to survive, you are hoping for an outcome which will be good for someone to whom you stand in a special relationship. In a similar way, we might think that it is good for its own sake

for you to prefer to let the five die, even if things would be better overall if you killed the one.

However, this proliferation of sources of conative value raises an obvious question. Is there anything we can say in general about when a pro-attitude is of conative value? Is there anything which unifies the various cases in which pro-attitudes are good for their own sake? And the answer seems to be that there is. What seems to unify the cases is that in each case the pro-attitudes which are good for their own sake are those which there is sufficient reason for. In the first case, there is sufficient reason for you to hope that Sally survives, even if it would not be good overall if she does. In the second case, there is sufficient reason for you to prefer to let the five die, even though things would be better overall if you killed the one.

The problem for attitude-based theorists is that they cannot accept this answer. Attitude-based theorists cannot explain how conative value arises by appealing to what there is reason to favour, because they aim to explain reasons to favour in terms of value. Attitude-based theorists thus face a dilemma. On the one hand, they could hold that it is of conative value to favour an outcome just in case that outcome is good overall. However, in that case they fail to accommodate considerations which are reasons for just some agents to favour an outcome. On the other hand, they could allow conative value to arise in a variety of ways. But in that case, they are barred from giving the natural answer to the question of just when conative value does arise.

4. Conclusion

I have argued that the FA account offers a better explanation of the linking principle than either the object-based view or the attitude-based view. Both of these views face

problems accommodating agent-relative reasons to favour. In addition, the object-based view fails to satisfy two plausible desiderata on an account of reasons, and the attitude-based view struggles to explain the linking principle. Of course, as a case for the FA account, this remains highly defeasible. There remain many possible value-first accounts that I have not been able to consider, not to mention other possible kinds of explanations of the linking principle. Nonetheless, insofar as the two accounts I have considered are two very natural options for value-first theorists to pursue, the problems with these accounts do seem enough to suggest that it is going to be considerably less straightforward for value-first theorists to explain the linking principle than for proponents of FA accounts. This is enough to suggest that the linking principle fits better with FA accounts, and so to lend some support to FA accounts.

References

- Bykvist, Krister. 2009. No Good Fit: Why the Fitting Attitude Analysis of Value Fails. *Mind* 118 (469): 2-30.
- Brunero, John. Forthcoming. Reasons as Explanations. *Philosophical Studies*.
- Dancy, Jonathan. 2000. Should We Pass the Buck? In Anthony O' Hear (ed), *The Good, the True and The Beautiful*. Cambridge: Cambridge University Press, 159-73.
2004. *Ethics without Principles*. Oxford: Oxford University Press.
- Danielsson, Sven and Olson, Jonas. 2007. Brentano and the Buck-Passers. *Mind* 116 (463): 511–522.
- Darwall, Stephen. 2002. *Welfare and Rational Care*. Princeton: Princeton University Press.

- Fine, Kit. 2001. The Question of Realism. *Philosophers' Imprint* 1 (1): 1-30.
- Finlay, Stephen. 2006. The Reasons That Matter. *Australasian Journal of Philosophy* 84 (1): 1-20.
- Fletcher, Guy. 2009. Rejecting Well-Being Invariabilism, *Philosophical Papers* 38 (1): 21-34.
- Forthcoming. Resisting Buck-Passing Accounts of Prudential Value. *Philosophical Studies*.
- Gibbons, John. 2010. Things That Make Things Reasonable. *Philosophy and Phenomenological Research* 81 (2): 335-61.
- Heathwood, Chris. 2008. Fitting Attitudes and Welfare. In Russ Shafer-Landau (ed), *Oxford Studies in Metaethics, Volume 3*. Oxford: Oxford University Press, 47-74.
- Heuer, Ulrike. 2010. Beyond Wrong Reasons: The Buck-Passing Account of Value. In Michael Brady (ed), *New Waves in Metaethics*. Basingstoke: Palgrave, 166-184.
- Hooker, Brad and Stratton-Lake, Phillip. 2006. Scanlon versus Moore on Goodness. In Terry Horgan and Mark Timmons (eds), *Metaethics After Moore*. Oxford: Oxford University Press, 149-68.
- Hurka, Thomas. 1998. Two Kinds of Organic Unity. *Journal of Ethics* 2: 299-320.
2003. *Virtue, Vice, and Value*. Oxford: Oxford University Press.
- Moore, G.E. 1903 *Principia Ethica*. Cambridge: Cambridge University Press.
- Olson, Jonas. 2009. Fitting Attitude Analyses of Value and the Partiality Challenge. *Ethical Theory and Moral Practice* 12: 365-78.
- Forthcoming. Buck-Passing Accounts. In Hugh La Follette (ed), *International Encyclopedia of Ethics*. Oxford: Blackwell.
- Parfit, Derek. 2011. *On What Matters*. Oxford: Oxford University Press.

- Rabinowicz, Wlodek and Rønnow-Rasmussen, Toni. 2004. The Strike of the Demon: On Fitting Pro-Attitudes and Value, *Ethics* 114 (3): 391-423.
- Raz, Joseph. 1999. *Engaging Reason*. Oxford: Oxford University Press.
2009. Reasons: Practical and Adaptive. In David Sobel and Stephen Wall (eds), *Reasons for Action*. Cambridge: Cambridge University Press, 37-57.
- Reisner, Andrew. 2009. Abandoning the Buck-Passing Analysis of Final Value. *Ethical Theory and Moral Practice*, 12 (4): 379-395.
- Rosen, Gideon. 2010. Metaphysical Dependence: Grounding and Reduction. In Bob Hale and Aviv Hoffmann (eds), *Modality: Metaphysics, Logic, and Epistemology*. Oxford: Oxford University Press, 109-36.
- Schaffer, Jonathan. 2011. On What Grounds What. In David Chalmers, David Manley, and Ryan Wasserman (eds), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press, 347-83.
- Schroeder, Mark. 2007. Teleology, Agent-Relative Value, and ‘Good’. *Ethics* 117 (1): 265-295.
2007. *Slaves of the Passions*. Oxford: Oxford University Press.
2010. Value and the Right Kind of Reason. In Russ Shafer-Landau (ed), *Oxford Studies in Metaethics, volume 5*. Oxford: Oxford University Press, 25-55.
- Scanlon, T.M. 1998. *What We Owe To Each Other*. Cambridge, MA: Harvard University Press.
- Skorupski, John 2010. *The Domain of Reasons*. Oxford: Oxford University Press.
- Smith, Michael. 1994. *The Moral Problem*. Oxford: Blackwell.
2009. Two Kinds of Consequentialism. *Philosophical Issues* 19 (1): 257-72.
- Suikkanen, Jussi. 2005. Reasons and Value – A Defence of the Buck-Passing Account, *Ethical Theory and Moral Practice* 7 (5): 513–535.

2009. Buck-Passing Accounts of Value. *Philosophy Compass* 4 (5): 768-779.
2009. Consequentialism, Constraints, and the Good-Relative-to: A Reply to Mark Schroeder. *Journal of Ethics & Social Philosophy*.
- Väyrynen, Pekka. 2006. Resisting the Buck-Passing Account of Value. In Russ Shafer-Landau (ed), *Oxford Studies in Metaethics, Volume 1*. Oxford: Oxford University Press, 295-324.
- Way, Jonathan. 2012. Transmission and the Wrong Kind of Reason. *Ethics* 122 (3): 489-515.
- Wedgwood, Ralph. 2007. *The Nature of Normativity*. Oxford: Oxford University Press.
2009. Intrinsic Values and Reasons for Action. *Philosophical Issues* 19 (1): 342-363.
- Zimmerman, Michael. 2011. Partiality and Intrinsic Value. *Mind* 120 (478): 447-483