

Abandoning the Buck Passing Analysis of Final Value¹

[Penultimate Draft, Please Do Cite Without Permission]

Introduction

In the decade since the buck passing analysis of *good* (BPA) was (re)introduced by T.M. Scanlon in his book, *What We Owe to Each Other*,² there has been a great deal of optimism about the view as an analysis of good or final value. This optimism is not well founded, and so I shall argue that it is time to abandon the BPA.³ This paper will be concerned with the BPA as an analysis of *good* in the sense of *final value* rather than in any attributive use.⁴ My suggestion is not that the BPA cannot be made to work for one narrow technical reason or another. This may or may not be so. Rather, I shall argue that the BPA is unable to deliver on its supposed advantages and that in the end it lacks plausibility as an analysis of final value.

To make this case, it will be necessary to consider what the BPA is and what benefits are supposed to accrue from adopting it. The BPA and its cousin, the *fitting attitude analysis*, are not presented in uniform ways throughout the literature. Scanlon's original presentation of the

¹ This paper was originally written for the Workshop on Kinds of Value and Kinds of Value Beaters at ECAP 6. I would like to thank the organisers, Wlodek Rabinowicz and Kevin Mulligan for giving me the opportunity to present this paper in Poland. I have received valuable feedback from those who attended presentations of the first version of this paper at McGill University's Faculty and Graduate Student Work in Progress Seminar and at the workshop at ECAP 6. Comments on a revised version of the paper from presentations at the Boston University Philosophy Department and the University of Uppsala Philosophy Department were also extremely valuable. I have not been able to remember everyone who has made helpful suggestions, as I have received a great many of them. Necessarily, then, the list of people I should thank individually is longer than the list of those whom I shall thank individually. With apologies to those who belong on this list, but who have been left off by forgetfulness on my part, I would like to thank for their very valuable comments and criticisms: Johan Brännmark, John Broome, Roger Crisp, Louis DeRosset, Jamie Dreier, Iwao Hirose, Kent Hurtig, Kevin Mulligan, Calvin Normore, Jonas Olson, Jan Österberg, Jessica Pepp, Christian Piller, Wlodek Rabinowicz, Toni Rønnow-Rasmussen, John Skorupski, Daniel Star, Asbjørn Steglich-Petersen, Folke Tersman, and Michael Zimmerman.

² See Scanlon (1998) for the original account.

³ In fact, the argument is aimed at both the BPA and closely related views such as the *fitting attitude analysis* of good or final value.

⁴ E.g. 'A good car' or 'a good person'.

BPA is most likely about properties,⁵ whereas other accounts are about language and concepts.⁶ I shall focus on the conceptual reading in this paper, with a tacit assumption that any plausibility that the metaphysical version of the BPA has derives from inferring a view about properties from a view about concepts.⁷

The strategy in this paper is twofold. The first part of the strategy is to show that two types of objections to the BPA, the wrong kind of reasons problem (WKR) and what I shall call the ‘inaccuracy objection’, highlight the central role that value plays in the epistemology of normativity. The most promising solutions to the WKR gain their currency from appeals to prior, independent intuitions about value. And, the inaccuracy objection itself only makes sense if we have intuitions about value that are independent of those about reasons: our very ability to frame the inaccuracy objection at all suggests that there is reason to doubt the BPA.⁸

The second part of the strategy is to show that the BPA, even if it overcomes the WKR and inaccuracy objection, fails to deliver on its proposed attractions. In short, even if it is possible to provide a version of the BPA that is immune to direct counterexample, the BPA is an unconvincing account of final or intrinsic value, and it fails to deliver on what its proponents take to be its principal attractions.

Buck-Passing

⁵ This is Roger Crisp’s reading (Crisp 2005); Gerald Lang (2008) disagrees and thinks that while the BPA is attractive as a view about properties, it was presented as a view about concepts.

⁶ C.f. Danielsson and Olson (2007), Lang (2008), Olson (2004), Rabinowicz and Rønnow-Rasmussen (2004), and Skorupski (2007).

⁷ Lang (2008) cryptically remarks that the BPA may fail as a conceptual analysis but still be a good account of value properties.

⁸ The mere presence of prior, independent intuitions about an *analysandum* does not alone demonstrate that it is not a good target for reduction to its *analysans*. However, I shall argue that there are particular reasons for being worried about this set up for the BPA.

In Scanlon's original presentation, the BPA is a metaphysical story⁹ that naturalises reason giving properties.¹⁰ In more recent discussions, the focus has been on the BPA as a conceptual or perhaps semantic analysis. This makes discussing the BPA difficult at times, since properties are one thing, language and concepts two others. Although I shall focus for the most part on the conceptual account, I shall begin by looking at the metaphysical version.

The original version of the BPA is presented as a response to the view G.E. Moore puts forward in *Principia Ethica* or at least Scanlon's reading thereof. According to Scanlon, Moore's view was the following: there are natural properties (e.g. pleasantness) that serve as a supervenience base for evaluative properties (e.g. goodness). In turn, those evaluative properties serve as a supervenience base for normative or deontic properties (e.g. being a reason). Given the transitivity of supervenience, one knows that there is no change in reasons (on this account) without a change in the natural world. Scanlon's suggestion is that we can cut out the middle man, excluding value from any role in determining what reasons we have. The buck is passed from evaluative properties to natural properties as the basis for reasons.

This need not require us to eliminate value from the picture, however. According to Scanlon, value is a different higher order property that supervenes on natural properties. In particular, it supervenes just exactly (and not accidentally) on those natural properties that give us reasons to favour things. Thus, what it is for an object to have the property of being good is for it to have the natural properties that provide us with reasons to favour that object. Reasons play an important role in the analysis of which objects or states of affairs have value, but value is not supposed to supervene directly on reasons in the way Scanlon reads Moore as saying that reasons supervene directly on value.

Let me note briefly that Scanlon's original language is not unambiguously metaphysical:

⁹ See fn. 2 above.

¹⁰ Scanlon's view about buck-passing has changed, and his current view does not have the same naturalising tendencies. See Crisp (2008) for an in-depth discussion of the most recent developments, and their upshot, in Scanlon's view. Also note that Scanlon's view that reason-giving properties are natural does not entail naturalism about reasons themselves.

[T]he claim that [something is] valuable is not a property that provides us with reasons. Rather, to call something valuable is to say that it has other properties that provide reasons for behaving in certain ways with respect to it (Scanlon 1998, p. 96).

Scanlon could, in the spirit of J.L. Mackie, be giving an account of what moral language refers, or tries to refer, to. This would make Scanlon's view consistent with error theory. I do not think that this is his intention, but I do not intend to follow up this exegetical question further and shall read Scanlon metaphysically.¹¹ The language in other versions of the BPA makes it clear that it is intended as a conceptual analysis.¹²

The Allure of the Buck-Passing Analysis

I turn now to the question of what we could gain philosophically, should the BPA be correct. There is at least one very straightforward answer to this question. Philosophers want to know the truth about a variety of things, and so learning that value can be reduced to reasons would constitute a gain in our knowledge. However, the BPA clearly has an allure beyond that.

Proposals from people who have defended the BPA show what advantages it is thought to confer. Gerald Lang writes:

Perhaps the leading virtue of the buck-passing account is that, in Wlodek Rabinowicz's and Toni Rønnow-Rasmussen's phrase, it 'demystifies value'. Value properties are indirectly reduced to reason-providing properties, in that the property of being good is held by the buck-passing account, to be simply the higher-order,

¹¹ I follow Crisp (2005) in doing so.

¹² See Danielsson and Olson (2007) and Skorupski (2007) for clear examples.

non-reason-providing property of having other, lower-order reason-providing properties. As a result, the buck-passing account usefully provides for a conceptual or internal connection between the realm of the evaluative and the realm of the deontic, which means that we can pour all of our philosophical energy into identifying reasons, rather than into identifying values which those reasons supposedly, or hopefully, track.

Accordingly, the buck-passing account supposedly carries a twofold advantage: first, it obviates the need to explain the practical or action-guiding significance of value facts; and second, it obviates or at least mitigates the metaphysical and epistemological difficulties about identifying value facts. (Note that Scanlon's passage states a semantic or conceptual claim about value. Even if this conceptual claim about value cannot be sustained, the buck-passing account might just work as a metaphysical account of value) (Lang 2008, pp. 472-472).

Rabinowicz and Rønnow-Rasmussen write:

The virtues of the FA analysis [‘fitting attitude analysis’, their name for buck-passing] are considerable: it demystifies value and explains why we are justified in our concern for valuable objects. The justification is immediately forthcoming if value is nothing but the existence of reasons for such a concern. Another virtue of the FA analysis is its neutrality on the difficult issues concerning the nature of value judgments. Reducing such judgments to judgments about reasons for pro-attitudes does not beg the question against cognitivist or non-cognitivist theories of evaluative discourse. It is consistent with both, since the interpretation of judgments about reasons is left open... (Rabinowicz and Rønnow-Rasmussen 2004, p. 400)

And finally, Danielsson and Olson:

One feature of the Brentano-style approach that we find particularly attractive is that a kind of internalism will be included in the bargain. This will be a kind of internalism that establishes a necessary link between values and attitudes: necessarily, to claim that an object is valuable is to claim that a pro-attitude towards that object is (would be) correct (Danielsson and Olson 2007, p. 520).

We can tease out four main points of attraction for the buck-passing analysis:

1. It demystifies the link between value and reasons
 - a. By (analytically) explaining why we are justified in being concerned (in the literal sense of feeling concern for) with value.
 - b. By Obviating the need to explain why value is action guiding.
2. It demystifies value by reducing questions about the metaphysics and epistemology of value to questions about the epistemology and metaphysics of reasons.
3. It offers up a kind of internalism about value and pro-attitudes.
4. It may provide the basis for understanding what the distinction is between the descriptive (or descriptive discourse) and the normative, i.e. the non-descriptive (or normative discourse).¹³

A fifth point should be added implicitly: that the buck-passing analysis is correct (which is not my view, but of course that would be an attraction).

Of these four attractions of the BPA, 1) and 4) are the most important. I shall focus on them in the final section. I should like to say something about 2 and 3 before moving forward. 2) is suggested by Lang. I do not share his view that this is a potential attraction for the simple reason that it is very hard to see how the deontic or normative (in its narrow sense) has a clearer metaphysics or epistemology than value. It may be that it does, but the metaphysics and

¹³ This attraction is developed by Skorupski (2007). See the following section for more discussion of his view.

epistemology of non-descriptive properties and concepts are sufficiently difficult that a good deal of work needs to be done to show that one class of them (the deontic) is somehow more tractable than another (the evaluative). In the next section of this paper, when I discuss the WKR, I shall in fact argue that if there is any particular category of the normative which is easier to grasp, it is the evaluative, although I make that claim with all due hesitancy.

3) may be a genuine gain, and thus it is a genuine attraction of the BPA. However, the BPA may not deliver judgement internalism in as satisfactory a fashion as one might wish. It should be noted that there are a great many paths to judgement internalism, and some further argument is required as to why we should be especially happy with the way the BPA delivers it to us. Nonetheless, I shall concede 3) to the BPA for sake of economy.¹⁴ For the moment, I shall leave consideration of the attractions of the BPA behind and look at one of the two most serious obstacles to the BPA, the wrong kind of reason problem (WKR). I shall reconsider whether the BPA can deliver on its attractions in the final section of the paper.

The WKR and What We Can Learn from Two Proposed Solutions

The WKR arises when we notice that there are reasons to favour things that manifestly are not intrinsically good. An easy way to generate examples of reasons to favour (a generic term I shall use for pro-attitudes) something that itself is not good is to imagine incentives offered by an evil demon. The evil demon may ask you to admire cruel actions, on pain of your being tormented, should you fail to do so. Here, the evil demon has given you a reason to admire cruel actions, although it is clearly not the case that cruel actions are finally good.

Attempts to solve the WKR problem have typically relied on one of two strategies. The first is to distinguish between two types of reasons: *object-given reasons* and *state-given reasons*,

¹⁴ Louis DeRosset pressed me to withdraw this concession, because it is doubtful that the BPA in any interesting way *explains* judgement internalism. Rather, the BPA is an account that has judgement internalism as part of its description: what it is for one to believe something to be valuable is for one to believe that one has a reason to favour it. That is (rational) judgement internalism of a kind, and so the BPA has not explained why judgement internalism is true. Rather, it is an account on which it could not be false.

identifying the latter as the wrong kind of reason (for the buck-passing analysis of intrinsic value).¹⁵ The second is to claim that putative examples of something's being the wrong kind of reason are not, in fact, reasons to favour the thing in question at all. Rather, they are reasons for higher-order attitudes towards favouring (e.g. that you desire to favour *x*) or are reasons to do something (e.g. bring it about that you favour *x*).¹⁶ These accounts have not proved promising for reasons discussed elsewhere.¹⁷

In this section I want to look at one solution of each kind. Danielsson and Olson have recently proposed a solution (the DO proposal) that has so far avoided falling prey to the standard kind of counterexample that plagues solutions of the first kind. I shall also look at a proposal that they criticise, John Skorupski's. In both cases, a careful examination of each proposal, I shall argue, should lead us to consider abandoning the BPA. Although I shall voice worries about both proposals, my aim is not to argue that either is unsuccessful *per se*, but that to whatever extent they are successful, they are so in ways that make the BPA unattractive. I shall begin by looking at the DO proposal.

Standard ways of defeating attempts to solve the WKR, which acknowledge that reasons of the wrong kind are genuinely reasons, usually involve finding a reason of the wrong kind that fits the criteria that only reasons of the right kind are supposed to meet. The DO proposal avoids this problem by distinguishing between reasons that count towards the 'correctness' of pro-attitudes for which they are reasons and those that do not. *Correctness* is a purely schematic concept; it is undefined and unanalysed, identified only by its role within the DO proposal (although a gesture to its nature is made by an analogy with truth). As will presently become apparent, having an undefined normative concept like *correctness* doing the work in one's

¹⁵ See Piller (2006) for more on the distinction.

¹⁶ See Parfit (2001) and Skorupski (2007) for versions of the higher-order attitude and bringing-it-about accounts of state-given reasons.

¹⁷ See Danielsson and Olson (2007), Lang (2008), and Rabinowicz and Ronn ow-Rasmussen (2004) for thoughtful criticisms of past attempts at solving the WKR. In addition, recent work on reasons for belief suggests problems for the two main solution strategies for the WKR. See Steglich-Petersen (2006).

solution may make the solution resistant to counterexamples, but it ultimately helps to show why we should consider abandoning the BPA.

The DO proposal offers a novel division of the types of reasons. The two types of reason in the DO proposal are *holding-reasons* and *content-reasons*. The DO proposal accepts that both types of reasons are genuinely reasons for pro-attitudes. This distinguishes the DO proposal from many accounts of reasons that deny that something like holding-reasons that are not content reasons, too, are genuine reasons for an agent to favour something.¹⁸

These two types of reasons reflect two different ways in which a fact can be a reason for an attitude. Holding-reasons are reasons, as the authors put it, for ‘having the attitude’, whereas content-reasons are reasons for ‘the correctness of the attitude’. Content-reasons give rise to holding reasons on the DO proposal, but holding reasons do not give rise to content-reasons. The distinction between the two types of reason is easier to understand by considering the analogous case of reasons for belief.

According to Danielsson and Olson, we can divide reasons for belief into strongly analogous categories to those assigned to reasons for pro-attitudes. Some reasons for belief are, to put things a bit loosely, reasons to believe true. These truth-reasons, so to speak, are evidential reasons for belief. Other reasons for belief, those given to us by the personal or social benefits of believing something (pragmatic reasons for belief), are reasons for having a belief. According to the DO proposal, that something is good for you to believe *is* a reason for you to believe it. But, that something is good for you to believe does not count towards its being true. Danielsson and Olson claim that under ordinary circumstances, the truth-reasons for belief determine what one ought to believe. But, those reasons may be defeated in some circumstances by the goodness-reasons, those one gets from its being good to believe something.¹⁹

Correctness-reasons, so to speak, are supposed to be like truth-reasons. Some pro-

¹⁸ See for example Parfit (2001) and Skorupski (2007).

¹⁹ See Reisner (2008) for a fuller account of the ways in which goodness-reasons can defeat truth-reasons.

attitudes are correctly held towards certain objects (or propositions). On the DO proposal, correctness for pro-attitudes is the analogue of truth for belief. Danielsson and Olson call the reasons given by correctness ‘content reasons’. As with belief, there may be reasons to have a certain pro-attitude towards an object (or proposition) that are given by the fact that it is beneficial to have that attitude. These are holding-reasons, but not content-reasons.

As may already be apparent, the important work in the DO proposal is done by introducing the concept of *correctness*. Correctness is a normative concept that serves a parallel role for reasons for pro-attitudes as truth does for reasons for belief.²⁰ Like evidential reasons for belief, content-reasons for favouring generally determine what one ought to do, but they can be defeated under certain circumstances by holding-reasons not derived from content-reasons. Holding-reasons that are also content-reasons are the right kind of reasons for buck-passing, according to the DO proposal, whereas holding reasons that are not also content-reasons are the wrong kind.

Although impressive in a variety of ways, the DO proposal is ultimately unsatisfying (I do not say ‘unsuccessful’, as I am not quite sure what is required for a solution to the WKR to be counted as successful) because the central notion, *correctness*, is undefined and only understandable schematically by the role it plays in their theory. The authors do not offer a specific account of correctness. Instead, they write:

We may well find it useful to talk about truth without being precise about the nature of the truth predicate, that is, without committing to the correspondence theory or to minimalism, or to any other theory. Analogously, we may find it useful to talk about correctness of attitudes without being precise about the nature of this notion, that is, without committing to the idea that to say that some attitude is correct is to concur with

²⁰ Jessica Pepp and Folke Tersman have both point out to me a worry about the analogy between correctness and truth. While both correctness and truth may play regulatory roles for propositional attitudes, truth is a property of the content of the attitude, as well as the attitude itself, in the case of belief. Correctness is not a property of the contents of pro-attitudes.

some possible expression of the attitude — which would be the ‘minimalist’ or ‘expressivist’ alternative — or to say something with more metaphysically robust implications (Danielsson and Olson 2007, p. 517).

In order for the DO proposal to be a persuasive solution to the WKR, it cannot merely put forward that there is some property, *correctness*, which does the fundamental work in distinguishing the right from the wrong kind of reasons, without saying something about the nature of *correctness*. Although there are competing theories of truth, it is a concept so widely used in philosophy that a variety of roles, functions, and constraints can be assumed when working with *truth*, regardless of which theory of *truth* is on offer. This is not so with *correctness*, which has been introduced to provide a solution to the WKR. Without elaborating on *correctness*, we have a purely schematic, and philosophically uninformative, account of what distinguishes the right kind of reasons from the wrong kind of reasons: namely that what distinguishes them is just that property which distinguishes them.

Danielsson and Olson try to pin down correctness without actually giving an account of it by giving us a picture first of reasons for belief. According to them, we can imagine two sorts of reasons for belief. There are content-reasons for belief (which will also be holding-reasons for belief) and non-content-derived holding-reasons for belief. Content-reasons for belief bear directly on whether it is true that p . Holding-reasons for belief not arising from content-reasons for belief, on the other hand, do not bear directly on the truth of p . Rather, they bear on (for example) whether it is good for you to believe that p is true.²¹ One could explain the truth of p by appealing to content-reasons, but one could not explain the truth of p by appealing to holding-reasons.

In the case of reasons for pro-attitudes, content-reasons for favouring do not bear directly on the truth of (the contents of) the attitude, but rather on its correctness. Holding-reasons

²¹ Danielsson and Olson (2007) pp. 516ff. Although the details of the accounts differ, there is much in common between their view and Hieronymi (2005).

not derived from content-reasons do not bear on the attitude's correctness. It is proposed that pro-attitudes can be correct or incorrect in the way that beliefs can be true or false. This is the main substance of what we are told about the property of correctness: it plays the same role for reasons for pro-attitudes as truth does for reasons for belief. We are also told that correctness is a normative property. Once we have the property of correctness, we can finish constructing the analogy: content-reasons for favouring can explain why a particular pro-attitude is correct, whereas holding-reasons not derived from content-reasons for favouring cannot.

There are difficulties with this analogy. To see them, it is useful to recall that the WKR arose because of threats to the buck-passing analysis of intrinsic value. On the buck-passing analysis, to claim that something is good or intrinsically valuable is to claim that there is a reason to have a pro-attitude towards it. The WKR arises because there are some reasons to favour things which are not themselves good.

Typically (although not exclusively), one is said to have the wrong kind of reason when one is given incentives for holding a pro-attitude, independently of considerations concerning the conceptual relationship between the attitude and its content. An eccentric millionaire might offer you all of his gold if you desire to count all of the holes in the acoustic ceiling tiles of a large room. There is nothing intrinsically good about counting all the holes in an acoustically tiled ceiling, and yet there at least seems to be some reason to desire to do the counting, that reason being the prospect of gaining the millionaire's gold just for wanting to count the holes (indeed, one need not ever do the counting in this example). Here one is said to have the wrong kind of reason to desire something. What is wrong with it? The object of the desire is not thought to be intrinsically valuable.

Here it is important to note the method for testing a solution to the WKR. The motivation for claiming that such-and-such is the wrong kind of reason is that we lack an *intuition*. The intuition is that the object of the pro-attitude is intrinsically valuable. In the above example, we are supposed to have the intuition that there is nothing intrinsically valuable about counting (or knowing) how many holes there are in the acoustic ceiling tiles of a large

room. Solutions to the WKR will successfully sort the reasons that count for purposes of buck-passing from those that do not, and that success will be measured by whether there are only the right kind of reasons for pro-attitudes having objects that we intuit as being of intrinsic value, and whether there are the right kind of reasons for all pro-attitudes having objects that we intuit as being of intrinsic value.

This point of method allows us to see what is wrong with the analogy between correctness and truth. There are independent accounts of truth and attendant theories of evidence. We need not be distracted by the diversity of accounts of the former or the latter. The point is that when assembling an account of the right and wrong kind of reasons for belief, one can use three concepts, *reason*, *truth*, and *belief*, all of which have some prior analysis, or at least some agreed upon features. To claim that the aim of belief is truth is informative, because we have prior notions of *belief* and of *truth*.²² To say that a content-reason for belief is a reason that counts towards the truth of its contents is informative in the same kind of way, all the more so because what is generally thought to count towards the truth of a belief is evidence. And, we have independent theories of evidence. If content-reasons are the right kind of reasons, and if holding-reasons of the wrong kind are taken as their negation (that is, holding-reasons that are not also content-reasons), then we have an analysis that is informative.²³ To see the informativeness point more clearly in the case of claims that belief aims at truth, it is helpful to note that there are alternative accounts of the aim of belief. For example, one alternative view is that beliefs aim at knowledge. It is not sufficient to identify truth as being that which belief aims at, precisely because beliefs might aim at something else.²⁴

In the case of reasons for pro-attitudes, the concepts we have available on the DO proposal are *reason*, *correctness*, and *pro-attitude*. We lack a prior analysis for correctness, and it is not readily apparent that there is a concept that for pro-attitudes corresponds in its role to

²² At least insofar as one thinks belief has an aim.

²³ Although it may well be implausible. See Steglich-Petersen (2006).

²⁴ See Williamson (2000) for an account of knowledge as the aim of belief.

that played by belief in the case of truth. Even if there is such a concept, if the analogy between belief and pro-attitudes is meant to be strict, then it may turn out that correctness is not the right concept to employ in dividing the right from the wrong kind of reasons. We have been given no informative analyses of correctness, and so it is at minimum difficult to tell what the prospects are for the DO proposal, which relies on *correctness* to do the important work.²⁵

This last point requires some explanation, for it raises a central challenge for both the DO proposal and for solutions that rely on schematically given, unanalysed concepts to do the work. The point is epistemological, but the epistemology bears on the metaphysics. *Correctness* in the DO proposal is a concept that we are invited to grasp based on its schematic role in separating holding-reasons that are also content-reasons from holding-reasons that are not also content-reasons. A consideration that is a holding-reason and also a content-reason for a pro-attitude bears on the correctness of holding a pro-attitude towards its object.

In order to judge whether a consideration is a content-reason, one must judge whether the object of the pro-attitude for which there is a putative content-reason is intrinsically valuable. For example, I want to know whether a firefighter's bravery in the face of a burning building is a content-reason to admire her actions. In order to do this, I must first decide whether bravery in the face of a burning building is intrinsically valuable, for if it is not, there can be no content-reason to admire it. Supposing I judge that it is intrinsically valuable, I now know that bravery in the face of a burning building is a content-reason to admire the firefighter's actions. Now that I know that there is a content-reason to admire the firefighter's actions (as an example of bravery in the face of a burning building), I am in a position to judge that the consideration that is the content-reason contributes to the correctness of the object of the pro-attitude in the same way as evidence contributes to the truth of a belief. The

²⁵ Asbjørn Steglich-Petersen observed to me that there may be another problem with treating correctness as analogous to truth. Because correctness is supposed to be a normative notion, one that matches content-reasons with pro-attitudes, it is tempting to think that truth just is correctness for belief. If this is so, then we have an analysis of correctness for belief (as truth) and might be led to wonder why, if an analysis can be given for belief, one cannot be given for pro-attitudes.

epistemology of determining whether some considerations contribute to the correctness of the content of a pro-attitude thus works back from an initial intuition concerning what is intrinsically valuable.

The metaphysical lesson that can be taken from the epistemology of judgements about *correctness* is that the concept we are seeking to analyse, intrinsic value, provides us with an understanding of the crucial concepts in its *analysans*: content-reasons and correctness. *Content-reason* is an intermediating concept between *intrinsic value* and *correctness* in the analysis, so intrinsic value's privileged epistemic role in its analysis is not worrying. However, it is worrying when the putative *analysandum* of it is required for understanding a primitive concept that is supposed to be its *analysans*. Indeed, without further independent considerations suggesting that correctness is, in fact, the primitive concept appropriate for grounding an analysis of intrinsic value, it is tempting to follow the epistemological trail, taking *intrinsic value* as the primitive concept that we should use, instead, to provide an analysis of *correctness*.

What we are left with is a set of schematic, and uninformative claims about the right and wrong kind of reasons for favouring something. The right kind of reason for favouring something is a holding-reason that is also a content-reason for favouring it. Something is a content-reason for a pro-attitude just in case it counts in favour of the attitude's being correct. A pro-attitude is correct just in case the object of the pro-attitude is intrinsically valuable. Intrinsic value is the *analysandum* of the buck-passing analysis, so it cannot be used to specify any of the concepts in the *analysans*, and yet it is by appeal to our intuitions about intrinsic value alone that we have a grasp on what attitudes are correct. In short, the DO proposal is not capable of providing criteria whereby one can make judgements distinguishing the right from the wrong kind of reasons.

This point is of central importance to this paper. As noted in the previous section, it is supposed to be helpful (according to Lang) for our epistemology and metaphysics to work with reasons rather than value. But, to do some of the crucial work for this kind of solution to the WKR, work which involves making distinctions concerning reasons, we are epistemically tied

to value. If we are to decide anything about concepts or properties from their epistemology, then it appears that priority, indeed indispensability, rests on the side of value. If the most successful solutions to the WKR find value epistemically (at least) indispensable for working with reasons, we are left to wonder why we should want the BPA in the first place.

Turning to the other approach to solving the WKR, John Skorupski has proposed a very simple solution. Apparent reasons of the wrong kind are instead reasons to cause oneself to have the relevant pro-attitude. Skorupski explains by way of example. He asks us to imagine a violin performance which lacks all musical merit, but the admiration of which will bring about a finally good state of affairs, for the demon will carry out his evil plans unless you admire the performance:

In the case of the [awful] violin performance, the fact that the evil demon has his evil plans is a sufficient reason for me to *do* something — namely, bring it about that I admire the performance, if I can. In the circumstances, that would be a very good thing to bring about. Over and above that uncontroversial point, there is then the question of whether an evaluative reason relation *also* holds. Does the self-same fact about the evil demon stand in that *distinct* reason relation to me and a certain feeling of mine, namely, *admiring the performance*? The two relations are distinct, since their *relata* are distinct. And once they have been distinguished, a case needs to be made for holding that the second relation holds as well as the first (Skorupski 2007, p. 11).

One difficulty with Skorupski's view is what I have called elsewhere *blocked ascent*.²⁶ In blocked ascent, the evil demon will cause trouble for you if you cause yourself to have the relevant pro-attitude. The only way to get the prize, so to speak, is just to have the pro-attitude. If you are lucky enough to already have it, then necessarily you can have it. So, there can be at least some cases in which the Skorupski solution will fail to provide any reason at all.

²⁶ See Reisner (forthcoming) for more on blocked ascent.

Interestingly Wlodek Rabinowicz has raised this point with Skorupski,²⁷ who is willing to bite the bullet.

Olson and Danielsson argue persuasively in my view that this objection to Skorupski's solution needs to be taken seriously.²⁸ Nonetheless, Skorupski's approach to the WKR has several advantages over the DO proposal. Let us now (turn to his account:

The account appeals to those philosophers (I am one of them) who suspect that what makes normative discourse *normative* - in the broad sense of normative that contrasts with 'descriptive' - is precisely its conceptual reducibility to propositions about reasons. Since that is a quite general thesis it clearly propounds a pretty ambitious reductive programme; one can certainly defend buck-passing about goodness without subscribing to it. Nonetheless the buck-passing account would fit neatly into it, and would in turn gain some plausibility from whatever case can be made for the programme overall... (Skorupski 2007, p. 1).

It will be helpful to render this account more precisely. What does it mean to reduce conceptually a discourse to propositions about something? A discourse is subset of a language, or perhaps a language practice, and propositions are not themselves items in a language on most views. The upshot of the reduction is this:

S: Let us understand propositions as being as finely grained as their constituent concepts. Any sentence of a genuinely normative discourse expresses a fine grained proposition partially composed of the normative concept associated with the normative term in the

²⁷ See Skorupski (2007) fn 20. It should be noted that although Skorupski's view requires him to accept one difficult result, it has several advantages over the DO proposal. The difficult result arises from the blocked ascent argument and has been discussed already. The advantage Skorupski's view gains is quite considerable. Danielsson and Olson must explain why something's being both a content reason and a holding reason has different normative significance to something's just being a holding reason (and not a content reason). Because Skorupski denies that there are the latter kind of reasons, there are no concerns for him about why the differing kinds of reason would have differing moral significance.

²⁸ See Danielsson and Olson (2007) and see Reisner (forthcoming) for related arguments.

discourse. For any fine grained normative proposition not expressed in terms of reasons, there is at least one other non-trivial normative proposition with the same truth value, the only constituent normative concepts of which are reasons.

What makes normative discourse normative, according to Skorupski, is that this reduction can take place. Interestingly, Skorupski does not treat reductions as unique to reasons or as being asymmetric with regards to types of normative terms:

Note that the possibility of this reduction does not exclude the possibility of others: for example of reducing reason predicates to value predicates. It might be that you can reduce the normative concepts in more than one way (Skorupski 2007, fn 2)

Given that, some kind of substantive case must be made for the conceptual priority (or other feature suitable for being a base in a reduction) of reasons over value. Skorupski calls this the 'wrong way round' worry. His argument, as I understand it, is the *redundancy argument*: that value is not doing any extra work over and above the natural reason giving properties, and so it should be explained in such a way that it does not play the role of a middle man between nature and reason (or descriptive and deontic).

It is, however, the challenge of the WKR that makes one think that the middle man does, in fact, have a role to play. Skorupski acknowledges that your admiring the awful violin performance is a valuable state of affairs, under the circumstances, it is just that there is no reason to admire it. Before Skorupski and someone who wants to press the WKR have been able to settle their disagreement about reasons, whether there is reason to admire the violin performance, they have already agreed on the final value of the state of affairs in which you admire the performance. If our metaphysics follows our epistemology, then once again it looks as though value may have priority. While a complicated reduction can be performed to make

the middle man of value into a proposition about reasons, I believe the contortions involved are unnatural enough to raise a question about the desirability of an account that requires them.

The Inaccuracy Objection

The final objection to the buck passing analysis comes via Roger Crisp.²⁹ I shall call it the ‘inaccuracy objection’. The inaccuracy objection is that the BPA loses something when reducing from value to reasons. There are many fine hues of value, so to speak, but there are a relatively limited set of pro-attitudes, and there is no guarantee of the one-to-one match that is needed for the BPA to fully capture everything that is in value.³⁰ In particular, it is difficult to articulate how closely related thick values might be distinguished by an account that relies on affective responses, since the same affective response might be appropriate for the two values.

Some examples help make this point clearer. Consider an elegant movement and a graceful movement. One may have a reason (of the right kind) to admire both movements, to praise them, to appreciate them, and so forth. Indeed, one’s affective responses towards each kind of movement might be broadly the same, and yet one might still detect a difference in the nature of the value involved.³¹

To take another example, consider a person who is decent and another who is kind. Responses to decency and kindness might well overlap. Could one find a difference in the appropriate responses? Perhaps, but it seems likely that any analysis of the difference between the two concepts in terms of reasons for affective responses will be less certain than our intuitive grasp of the difference between the concepts themselves.

²⁹ Crisp (2005) raises a series of objections that I take ultimately to trade on this problem, although they would stand independently were one not to accept my more general argument.

³⁰ Mulligan (1998) defends the view that successful buck-passing style analyses must provide a one-to-one match between attitudes and values.

³¹ A buck-passer might argue that the difficulty in finding different appropriate reactive attitudes to match the values of being graceful and being elegant suggests that these are not in fact distinct values. However, there are some differences in the circumstances of their application; one might apply *elegant* to a mathematical proof, but one would be unlikely to apply *graceful* to it (Irinia Meketa suggested this example to me).

The strategy for pressing the inaccuracy objection is by now clear. One finds two thick value concepts that are closely related, and one presses the difficulties with finding distinct affective attitudes sufficiently subtle to distinguish them. This is, as I understand it, the core of Crisp's objection to the BPA (Crisp 2005). It is worth examining a bit further why this objection is so problematic for the BPA.

Interesting analyses are not mere translations. Presumably, more or less anything that one can say in English, one can also say in French. This does not show that somehow French is a correct account, or analysis, of English, much less a proper base for reduction. At best, it shows (with some additional work) that the two languages have the same logical and semantic expressive power. For an analysis to be successful, the vocabulary being analysed cannot be more powerful than the vocabulary into which it is analysed, otherwise the analysis would be unsuccessful. So, the translatability of the vocabulary of the *analysandum* into that of the *analysans* is an important necessary condition for an analysis, but no more than that. From this, we can set as a necessary condition for the success of the BPA that the vocabulary (conceptual or linguistic) of value have no greater expressive power than (and no failure of overlapping with) the vocabulary of reasons.

The inaccuracy objection is immediately worrying, because it raises doubts about whether this condition is met. There is at least *prima facie* evidence that our value vocabulary outstrips our reasons vocabulary. If this is so, then the BPA is a non-starter. The best strategy for defenders of the BPA may be to adopt an axiology on which we would expect some kind of mismatch between the number (and distinctions amongst) reactive attitudes and value *vocabulary*.

An axiology of this kind holds that there is just one kind of value: final value. However, it is multiply realisable; more specifically, final value has a wide variety of types of supervenience bases. When we use thick value terms like 'graceful', 'elegant' and 'courageous',

we are not referring to distinct kinds of final value. Rather, we are saying something about the supervenience base of some particular instance of final value.³²

This way of resisting the inaccuracy objection may not prove entirely successful for two reasons. The first is that it requires buck-passers to commit to a particular axiology, one in which there is just a single kind of final value: thin final value. Thick final value terms do not refer to different kinds of final value, but different supervenience base types for final value. The BPA is a less attractive account of final value if it can only accommodate one axiology.³³ The second and more significant reason is that employing this strategy is only open to those who are defending the BPA as a strictly metaphysical thesis. As a semantic or conceptual thesis, the BPA will fare no better on this revised axiology of final value. The vocabulary (and matched concepts) of final value includes thick value terms (and concepts). These terms, on the revised axiology, contain two types of information: that there is final value and what the supervenience base type is for the final value. The vocabulary and concepts of reactive attitudes available to the buck-passer are still less fine grained than those of value, although here the information lost is about the supervenience base types rather than the value types. The upshot of this second worry is that the BPA still fails one of the standard tests for conceptual analysis: that the analysed vocabulary not be more informative than the vocabulary of analysis.

Further reflection on what constitutes a successful (reductive) analysis leads an additional worry, separate from the question of whether defenders of the BPA should use a revised axiology. Just obtaining a one-to-one match between value terms and reactive attitudes may not be convincing, although this is not certain.³⁴ A classic instance of a successful analysis

³² Toni Rønnow-Rasmussen and Jan Österberg have separately suggested to me that this option is open to the buck-passer.

³³ Perhaps a defender of the BPA would advertise this as an advantage for it, as the truth of the BPA would settle an important outstanding issue in value theory.

³⁴ Mulligan makes the case that a one-to-one match is the right target for a BPA-style account. See Mulligan (1998).

is from the language (or concepts) of chemistry to that of physics. It is an interesting analysis for two reasons. The first is that the analysis is asymmetric. Chemistry can be analysed in terms of physics, but not *vice-versa*. The second is that, because of the greater explanatory power of physics and some substantial views about what is fundamental in the sciences, chemistry is normally seen as reducible to physics. This is an especially interesting analysis, for it tells us that chemistry is a special case of physics.

Even if we could generate a one-to-one match for the BPA analysis, we might not be satisfied that we had done anything very interesting, vis-à-vis conceptual analysis. We would have shown translatability, but we would not have discovered which vocabulary has conceptual priority. Some independent argument about why reasons are more fundamental than values is required; otherwise we may as well say that we have given a buck-passing analysis of reasons in terms of value.³⁵ Of course, showing that the expressive power of reasons vocabulary outstrips that of value would provide an asymmetry, and asymmetries can be evidence of reductions. However, the inaccuracy objection puts us in just the reverse position. At least as far as the concepts covered by value terms are concerned, the vocabulary of value may well outstrip that of reasons. If the asymmetry goes the other direction, the only kind of buck-passing account one *could* give, if one were not just to abandon the project, would be a buck-passing account of reasons.³⁶

The Failure of Demystification

In the introduction to this paper, I claimed that even if the BPA can be made to work, it would not be desirable. In this section, I shall explain why. Remember that there were four main attractions to the BPA:

³⁵ Skorupski suggests that there may be full translatability in both directions. Symmetry for him will be less of a concern than it may be for other buck-passers, as he is in part aiming to give a general account of normative, i.e. non-descriptive, vocabulary, the mark of which on his view is that it can be translated into reasons vocabulary.

³⁶ Just to be clear, I am not suggesting that doing so would be worthwhile.

1. It demystifies the link between value and reasons
 - a. By (analytically) explaining why we are justified in being concerned (in the literal sense of feeling concern for) with value.
 - b. By obviating the need to explain why value is action guiding.
2. It demystifies value by reducing questions about the metaphysics and epistemology of value to questions about the epistemology and metaphysics of reasons.
3. It offers up a kind of internalism about value and pro-attitudes.
4. It may provide the basis for understanding what the distinction is between the descriptive (or descriptive discourse) and the normative, i.e. the non-descriptive (or normative discourse).

I shall concede points (1a) and (3) to the buck-passers.³⁷ (4) is difficult to assess without considering the details of Skorupski's proposal, but some comments are nonetheless in order. The BPA, as Skorupski observes, cannot deliver the descriptive/non-descriptive distinction on its own. One needs an account of how reasons relate to other deontic concepts (duty, requirement, etc.) to accomplish his task. Skorupski suggests that his broader project may lend credence to (and be bolstered by) the BPA, but buck-passers should reflect carefully on whether they would be happy with this being one of the upshots of their project. Because as presented, the thesis is that all value expressions can be given in the language of reasons, and that the reverse is at least possible.³⁸ If one shows that there is a one-to-one correspondence between value concepts or terms and (rational) affective response concepts or terms, and one takes this

³⁷ Louis DeRosset and Folke Tersman have both pressed me not to concede (3) to the buck passer. They both observe that the kind of internalism offered up by the BPA is not especially strong, and perhaps not especially important. I think that they are right, but as defenders of the BPA are not claiming that they have delivered a very robust form of internalism, it still seems reasonable to allow them (3), even if this only adds slightly to the attraction of the BPA.

³⁸ Skorupski does not think, however, that the reverse would be correct; the important thing to note when there is a one-to-one correspondence in vocabulary is that some further argument is required to show which direction, if any, is the proper one for an analysis.

as a good reason to think that one might reduce to the other, the direction of reduction will depend on further arguments. One-to-one correspondence is insufficient to guarantee that the reduction will go in the direction that the buck-passers expect.

As per above, inter-translatibility does not show very much about the relation between vocabularies or the concepts they represent. But, we might want to reject a candidate translation when one vocabulary clearly fails to pick up important implications of another. This may be one reason to think that we should not aim to reduce evaluative concepts to deontic or reasons concepts (or indeed to reduce deontic or reason concepts to evaluative concepts). Deontic and reasons concepts are explicitly guiding, whether it be of actions, beliefs, or affective states. Indeed, many philosophers, I among them, think that deontic concepts function as propositional operators that are indexed to an agent.³⁹ The propositions governed by these operators must be about — or have the same subject if we are speaking sententially — the same individual to whom the operator is indexed. This kind of view rejects the notion that there are propositions of the form: it ought (or there is a reasons for it) to be the case that *s*, where *s* is a state-of-affairs.⁴⁰ Although *ought* is a controversial case, *duty* is not. Duties clearly apply to agents and their actions/beliefs/affective states, or rather propositions concerning them. Value terms, on the other hand, clearly can apply to states-of-affairs. This clear distinction in the applicability of value terms and concepts on the one hand, and deontic or reasons terms and concepts on the other, is reason for at least mild scepticism about a project that tries to reduce one class of concepts or terms to the other.

³⁹ Consider what I shall call an ‘o-form’ sentence: Jane ought that Jane go to the store. Here we have a proposition, Jane goes to the store, governed by a the propositional operator *ought*, which is indexed to Jane (at the level of propositions, ‘Jane’ is its subject when viewed sententially).

⁴⁰ Deontic logicians often give some account of ‘it ought to be the case that’ oughts. I think this is a mistake, if one is trying to provide a logic of obligation. If one is trying to produce a formal semantics for the English word ‘ought’, then it may be appropriate, but the value of the latter exercise strikes me as limited. I favour reading ‘ought to be’ as expressing an ideal evaluative claim, meaning ‘it would be best that’, or something along those lines.

That being said, the force of this objection should not be overstated, at least as pertains to the reduction of value terms to deontic or reasons terms. A putative translation of value terms into deontic language may not be very worrying. Perhaps a good state of affairs is one that we have reason or a duty to favour. But, the reverse presents a problem. Saying that a state of affairs is good does not say anything about a particular agent (or about agents in general) and is not explicitly guiding. It thereby loses an important aspect of at least some deontic concepts and is therefore a poor candidate as a translation of them.

(1b) and (2), if delivered on, would make make a successful account of the BPA a desirable thing to have. Unfortunately, the BPA does not deliver on either, at least not directly enough to make it attractive. (1b) promises that the BPA will help to explain the link between action and value. Of course, the BPA in most of its current guises is not about action. It is about pro-attitudes. In order to make their account attractive on this score, buck-passers must show some kind of link between value pro-attitudes and action. This is notoriously difficult to do. Difficulties of two kinds arise.

The first is that a great many pro-attitudes give us no clear reasons to act. One might admire a painting, a person, or a country, but the practical upshot of admiring each of these may be different, if it can be discovered at all. Having a reason for a pro-attitude, provided it is the right kind of reason for the BPA, might be a reliable indicator that there is also a reason to act lurking about, and thus that when there is a value, there is a reason for action.⁴¹ This is insufficient for demystifying the link between value and action, because all we now know is that there are reasons for action when something is valuable. We do not yet have an explanation of just exactly why (it will have something to do with there also being reasons for pro-attitudes), and we certainly have no notion yet of what kinds of reasons for actions particular values indicate (or generate).

⁴¹ See Crisp (2008) for an extensive discussion of this issue.

The other significant attraction of the BPA, according to some of its proponents,⁴² is that it may help us with the epistemology and metaphysics of value. Here I am sceptical. We now know a very great deal about our options concerning the epistemology and metaphysics of value. These topics have been at the centre of debates about metaethics for the last 105 years of work in analytic ethics. If we are unsure about the metaethics of value, it is because we find values an awkward fit for a fully naturalised conception of the world, and some philosophers are uncomfortable with accepting non-natural properties. J.L Mackie famously remarked on the queerness, metaphysically and epistemologically, of value,⁴³ in no small part because value at least seems to have the remarkable feature of guiding our actions (or at least pulling us towards or pushing us away from particular actions). This queerness problem, however it is best resolved, is inherited by reasons. Indeed, if the BPA is successful, it should pick up the odd characteristics that make value so problematic (otherwise it is less plausible as an analysis), and so we have good reason to think that the same problems will arise in the reasons domain as those that have arisen in the domain of value.

Conclusion

In this paper, I have argued that we have little reason to accept the BPA, both because of the epistemological priority of value over reasons, as is revealed by putative solutions to the WKR and because of the inaccuracy objection. Further the BPA looks unlikely to deliver on its most interesting promises: explaining the link between value and action and demystifying our metaphysics.

The crux of the problem with the BPA rests with the failure of its proponents to motivate it as a plausible analysis. Against the epistemological priority of value, a defender of

⁴² This is suggested by Lang (2008).

⁴³ See Mackie (1946) for his initial presentation of the now famous queerness argument. The most influential version is in Mackie (1977).

the BPA could well say that it is not uncommon when doing a conceptual analysis that we have a better grip on the *analysandum* than the *analysans*.⁴⁴ Perhaps this is even unsurprising, since we are in some sense discovering that the *analysans* fits the *analysandum*. The epistemological situation for the buck-passer is, I think, somewhat worse than that of the ordinary conceptual analyst.

The favourable situation is one in which we have a conceptual (or actual) vocabulary, say that of chemistry, which we begin to think can be analysed in terms of another vocabulary (say that of physics). We may have a better grip on chemistry, and so it may take some time and effort to probe around for the right analyses in physics for particular fragments of chemistry. The most plausible solutions to the WKR show us that the BPA is not in this favourable position. Substantive solutions to the WKR have failed across the board. The very best solutions are schematic. The schematic solutions require an undefined new primitive normative concept (e.g. *correctness*) to make a formal distinction between the right and the wrong kind of reasons. This new primitive, being undefined and unanalysed, just takes on the contours of our intuitions about value. Here we are not discovering how some antecedently accounted for concept(s) turn out to provide a good base for analysis for some other class of concepts about which we have stronger intuitions. Instead, we are taking a concept (value) that we appear to have decent epistemological access to, and we are in effect creating a model for it in a vocabulary that is itself equally mysterious. This is not a favourable circumstance for conducting an analysis.

The inaccuracy objection raises serious doubts about whether the vocabulary of reasons is suitable for analysing the vocabulary of value. These doubts are so serious that it looks very unlikely that there can be a successful BPA for all value terms. But, this paper is called ‘Abandoning the Buck-Passing Analysis of Final Value’. Is the inaccuracy objection fatal to a BPA for final value alone?

⁴⁴ Michael Zimmerman and David Chalmers pointed this out to me.

It may be. Perhaps anything that there is a reason to favour has final value, although what thick kind of value is involved will remain dark as far as the BPA is concerned. But, if the analysis is not strong enough to pick up important aspects of final value, i.e. what sort of final value we have (aesthetic, moral, etc.), it appears that the analysis has failed a crucial test: that information not be lost in the analysis. The arguments presented here do not prove that the BPA is wrong, but they leave us little reason to expect that it is correct.

Works Cited

- Crisp, Roger (2005) 'How to Avoid Passing the Buck', *Analysis* 65, pp. 80-85.
- (2008) 'Goodness and Reasons: Accentuating the Negative', *Mind* 117, pp. 257-265.
- Danielsson, Sven and Olson, Jonas (2007) 'Brentano and the Buck Passers', *Mind* 116, pp. 511-522.
- Hieronymi, Pamela (2005) 'The Wrong Kind of Reason', *The Journal of Philosophy* 102, pp. 437-457.
- Lang, Gerald (2008) 'The Right Kind of Solution to the Wrong Kind of Reason Problem', *Utilitas* 20, pp. 472-489.
- Mackie, J.L (1946) 'A Refutation of Morals', *Australasian Journal of Philosophy and Psychology* 24, pp. 77-90.
- (1977) *Ethics: Inventing Right and Wrong*, London: Pelican Books.
- Mulligan, Kevin (1998) 'From Appropriate Emotions to Values', *The Monist* 'Secondary Qualities Generalised' ed. P. Menzies, pp. 161-188.
- Olson, Jonas (2004) 'Buck-Passing and the Wrong Kind of Reasons', *Philosophical Quarterly* 54, pp. 295-300.

Derek Parfit (2001), 'Reasons and Rationality' in *Exploring Practical Philosophy*, Eds. Dan Egonsson, Jonas Josefsson, Björn Petersson, and Toni Rønnow-Rasmussen, Ashgate Press.

Piller, Christian, (2006) 'Content-Related and Attitude-Related Reasons for Preference', *Philosophy*, Supp. vol. 59, pp. 155-218.

Rabinowicz and Rønnow-Rasmussen (2004) 'The Strike of the Demon: On Fitting Pro-Attitudes and Value', *Ethics* 114, pp. 391-423.

Reisner, Andrew (2008) 'Weighing Pragmatic and Evidential Reasons for Belief', *Philosophical Studies* 138, pp. 17-27.

----- (forthcoming) 'The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reason Problem', *Philosophical Studies*.

Scanlon, T.M. (1998) *What We Owe to Each Other*, Cambridge: Harvard University Press.

Skorupski, John (2007), 'Buck-Passing about Goodness' in *Hommage à Wlodek*, online source: <http://www.fil.lu.se/hommageawlodek/site/papper/SkorupskiJohn.pdf>

Scanlon, T.M. (1998) *What We Owe to Each Other*, Cambridge: Harvard University Press.

Steglich-Petersen, Asbjørn (2006) 'No Norm Needed: On the Aim of Belief', *The Philosophical Quarterly* 56, pp. 599-516.

Williamson, Timothy (2000) *Knowledge and Its Limits*, Oxford: Oxford University Press.