

For the definitive version of this article, see:

Gregory, A. 2014. "A Very Good Reason To Reject The Buck Passing Account" in *The*

Australasian Journal of Philosophy 92:2 pp.287-303

<http://dx.doi.org/10.1080/00048402.2013.813055>

A Very Good Reason To Reject The Buck Passing Account

Alex Gregory

Abstract: In this paper I present a new objection to the buck passing account of value. I distinguish the buck passing account of predicative value from the buck passing account of attributive value. According to the buck passing account of attributive value, facts about attributive value reduce to facts about reasons and their weights. But since facts about reasons' weights are themselves facts about attributive value, this account presupposes what it is supposed to explain. As part of this argument, I also argue against Mark Schroeder's recent account of the weights of reasons, which purports to explain the weights of reasons in terms of further reasons without generating a vicious regress. I then argue that if we abandon the buck passing account of attributive value, it would be ad hoc and unjustifiable to continue to endorse the buck passing account of predicative value. In short, there seems to be little hope for the buck passing account in either form. The paper ends by briefly sketching a novel alternative theory according to which reasons are analysed in terms of the attributive value of motives. I suggest that a normative reason to ϕ is something that would be a good motive for ϕ ing. At least at first glance, this view has numerous merits and few problems.

Keywords: buck passing, reasons, normative reasons, value, attributive value, weights of reasons.

The buck passing account of value claims that facts about goodness can be reduced to facts about reasons.¹ This view has many supporters (e.g. Parfit 2011a: 38, Scanlon 1998: 94-100, Skorupski 2010, Stratton-Lake and Hooker 2006). In this paper I present a new objection to this view. The basic problem is that reasons can themselves be evaluated. It follows that facts about goodness cannot be reduced to facts about reasons, since the latter presuppose the former. After making this objection, I sketch and defend an opposing view, according to which we should reduce facts about reasons to facts about the value of motives.

The paper is split into five sections. In the first section, I outline two distinct buck passing accounts, and give an overview of my argument against them. In the second section, I defend the claim that reasons' weights are attributive value properties. In the third section, I defend the claim that this undermines the buck passing account of attributive goodness. In the fourth section, I defend the claim that if the buck passing account of attributive goodness is false, so too is the buck passing account of predicative goodness. The fifth section of the paper sketches my positive view according to which facts about reasons are reducible to facts about value.

(1) Buck Passing

Normative facts can be divided into two kinds: deontic facts, and evaluative facts. The former category encompasses, amongst other things, facts about oughts, shoulds, and reasons, where

¹ It might instead reduce *concepts* of value to *concepts* of reasons. In this paper I address the metaphysical version of the view, on the assumption that its falsity undermines the conceptual version as well.

a reason is a fact that counts in favour of some course of action.² The latter category encompasses, amongst other things, facts about goodness, badness, betterness, and worseness. This division seems intuitive, even if hard to pin down. The buck passing account is notable for trying to reduce one of these categories to the other: it tries to reduce evaluative facts to deontic facts, and in particular, to facts about reasons. As such it promises to give us a unified account of all normativity, and to explain the relationship between the most fundamental normative facts. Perhaps the most striking feature of the buck passing account is the way in which it entails that value does no real justificatory work. According to the buck passing account, facts about value are justificatorily epiphenomenal.

The buck passing account is normally held primarily as an account of predicative value: of what it is for some object to be good. In this form, the account says:

BPP: For some entity to be good is for some people to have reason to respond to it positively.³

For example, those who endorse BPP might say that for pleasure to be good is for some people to have reason to promote it. We could make BPP more precise in various ways. We could specify what the relevant reasons are. We could also specify whether things get to be good only if everyone has reason to respond to them positively, or whether they get to be good even if only some people have such reasons. We could also specify what it is to respond to something positively. We could also modify the view to account for pragmatic reasons to respond positively to bad things. No doubt there are still other ways in which BPP needs finessing. But for my purposes the account as defined above will be precise enough: my

² Throughout this paper, unless I indicate otherwise, “reason” means “normative reason”.

³ This view and all of those below can be extended in obvious ways for badness.

argument aims to undermine all forms of the view, no matter what they say about these questions.

BPP is an account of predicative goodness – an account of what it is for something to be good. But there are also facts about goodness of other kinds. Most relevantly for my purposes here, there are facts about attributive goodness, facts about what is *good as* what (see Geach 1956). I might say that Sally is a good assassin, meaning that she is good as an assassin, or I might say that this rock is a good doorstep, meaning that it is good as a doorstep. A natural thought is that if BPP is true, then some parallel account is true of attributive goodness. That is, it seems as though those who endorse BPP should also endorse an account such as:

BPA: For some entity to be good as an F is for some people to have reason to respond to it positively insofar as they want an F (cf. Scanlon 2011, Schroeder 2010: 48-9, Skorupski 2010: 83)

For example, one might think that a good assassin is one you have reason to intend to hire, insofar as you want an assassin. Again, one might quibble with the exact formulation of this view, but the details will not be important for my purposes.⁴ With these preliminaries out of the way, here is the argument I shall press against buck passing accounts:

P1) Reasons' weights are attributive value properties.

P2) If reasons' weights are attributive value properties, then BPA is false.

C1) BPA is false.

P3) If BPA is false, BPP is false.

⁴ BPA raises some new issues as well. For example, are the relevant reasons narrow scope or wide scope (see Broome 2002)? We might think that if you want an assassin, you have reason to hire Sally, or we might instead think that you have reason to [if you want an assassin, hire Sally]. One might also wonder whether relevant reasons are best thought of as conditional on our *wants*, rather than on something else. Again, I shall set these issues aside.

C2) BPP is false.

The inferences from P1 and P2 to C1, and from C1 and P3 to C5, are clearly valid. The real work is establishing P1, P2 and P3, which I shall now defend.

(2) P1: Reasons' weights

P1 says that reasons' weights are attributive value properties. Why think this is true?

We often talk about how good reasons are. Someone might say that Labour's education policy is a good reason for voting for them, whereas their foreign policy is a less good reason for voting for them, and go on to claim that they can't think of a single good reason for voting Conservative. At any rate, they might add, there are better reasons for voting Liberal Democrat than Conservative. Claims like these are common enough that it would be a disaster if they turned out to be universally confused. Our philosophical theories about reasons had better permit that claims such as these are true, or at least coherent. So how should we interpret them?

Here is the obvious option: when someone says that something is a good reason, they are making a claim about attributive goodness. Labour's education policy is a particularly good reason for voting for them because it is particularly good *as* a reason for voting for them: it is a good instance of something that favours voting Labour. This option is intuitive in part because claims such as "their education policy is a good reason for voting Labour" seem grammatically similar to claims about attributive value such as "potash is a good fertiliser for potatoes" and "red is a good colour for stop signs"

This view is also attractive because claims about good reasons employ the word “good”, and it’s hard to explain why this would be true unless such claims report facts about value. It would be a bizarre coincidence that we use this word in this context if we aren’t making claims about value; certainly we should prefer a view that explains why we talk of goodness in this context to one that can’t explain this fact.

We might yet try to interpret claims about the goodness of reasons as claims about predicative- rather than attributive- value. We might say that a good reason is something that is good and also a reason.⁵ But this view is highly implausible: something that is a good reason to do something needn’t itself be a good thing. For example, Susan’s pain might be a good reason to help her without her pain being good, and Labour’s foreign policy might be a good reason to vote against them without their foreign policy being good. So this view cannot be sustained.

So it seems as though claims about good reasons are claims about attributive value. Further, it is highly plausible to understand these facts about attributive value as the facts philosophers pick out when they talk about the *weights* of reasons. Where philosophers talk about the weights of reasons, most people just talk about how good those reasons are.⁶ If this is right, then P1 is established. According to this view, to be a weighty reason is just to be something that is good as a reason.

⁵ Here I rely on Geach’s test for the distinction between predicative and attributive adjectives (Geach 1956: 33). Geach also claims that there are no predicative uses of “good” at all (cf. Skorupski 2010: 84, Thomson 2008). Here I stay neutral on that claim. If it were true, P3 of my argument would be redundant.

⁶ If a good reason is a weighty reason, what is a bad reason? A bad reason is one lacking in weight. That is, a bad reason is one that provides little or no justification for the relevant action, just as a bad knife is one that provides little or no assistance in cutting things. Note that if some bad reason provides no justification at all for the relevant action, then this bad reason is not really a reason at all, in just the same way that a sufficiently bad knife ceases to really be a knife.

There are at least three arguments for the claim that good reasons are weighty reasons. First, facts about the weights of reasons are extremely important for decision making. To make a rational choice, one needs to know not merely what one has reason to do, but how much reason one has to do those things. It would therefore be incredible if everyday English had no word for the weight of a reason. Given that talk of “weights” is idiosyncratic to philosophers, a tempting thought is that we normally make claims about what philosophers call weights in some other way, and the obvious possibility is that we do so by talking about how good the relevant reasons are.

The second argument that good reasons are weighty reasons is this. It is natural to think that what you ought to do is determined by what you have weightiest reason to do, rather than by what you have lightest reason to do.⁷ Why is this? The obvious answer is that weightier reasons are *better* reasons. Unless we appeal to this claim, it is hard to explain why one ought to do what one has weightiest, rather than lightest, reason to do.⁸ The same point obviously applies if we instead talk of the “strengths” of reasons rather than their “weights”. Indeed, presumably we are happy to use “strong” and “weighty” interchangeably just because both suggest a positive evaluation.

There is also a third argument for the claim that good reasons are weighty reasons. It is an argument by elimination: there is nothing else we might plausibly identify with the goodness of a reason other than its weight. So far as I know, there is only one other suggestion about how to interpret claims about reasons being good. This is Skorupski’s view, according to which: “a ‘good reason’ in one main use of that phrase, is simply something that is genuinely

⁷ Here I ignore whatever complications supererogation raises.

⁸ I thank Eric Wiland for this point. Also see Kearns and Star 2008: 42-45.

a reason” (2010: 83). But this cannot be correct. It is true that good reasons must be genuine reasons. But it is not as though reasons are either wholly good or not at all good. Rather, their goodness comes in degrees. I might say “Labour’s education policy is a good reason for you to vote for them, but their health policy is a better reason for you to vote for them”. Here I am allowing that both reasons are good, but suggesting that one reason is more good than the other. Skorupski’s view cannot make sense of such claims, since genuineness doesn’t come in degrees.⁹

In summary, it is highly plausible that when we talk about good reasons, we are talking about the attributive value of those reasons, and that facts about the attributive value of reasons are facts about those reasons’ weights, as P1 claims.

There is an objection to address.¹⁰ One might think that “good X” doesn’t always mean “good as an X”, but instead may have some context sensitive meaning (Szabó 2001: 126-135). A good pebble isn’t something that’s good as a pebble, but instead something that is good for some contextually salient purpose. Similarly, we might think that a good reason isn’t necessarily something that’s good as a reason, but may instead be something that is good for some contextually salient purpose. If we thought that, it seems that “good reason” doesn’t necessarily mean “weighty reason”. Might that be a problem for the argument?

I am not entirely certain that the meaning of “good reason” could really be context sensitive in the way that the meaning of “good pebble” is. It’s hard to think of plausible examples in

⁹ Might we say that a good reason is something that is *probably* a genuine reason, and a better reason is something that is more probably a genuine reason? No: It’s clear that one might think that both R1 and R2 are *certainly* genuine reasons to ϕ but that R1 is nonetheless a better reason to ϕ .

¹⁰ I thank a referee for this journal for this objection.

which “good reason” means anything other than “weighty reason”. Still, for the sake of argument, let’s imagine that there are some such contexts, and that the meaning of “good reason” is therefore context sensitive. Then the apparent threat to the argument would be that it might rest on an equivocation between “good reason” in P1 and P2. Perhaps P1 is true only when “good reason” has the meaning it does in one context, and P2 is true only when “good reason” has the meaning it does in some other context. Then we could render P1 and P2 true only at the cost of making the argument invalid.

However, “good reason” definitely *standardly* means “good as a reason”, and in this sense it means a weighty reason. So P1 is to be read as though “good reason” has the meaning it does in standard contexts. If that is true, then my argument is safe from the objection so long as my defence of P2 doesn’t appeal to the meaning of the phrase “good reason” in some non-standard context. In what follows, it should be clear that this is highly plausible. So I shall set this complication aside: hereafter, whenever I use the phrase “good reason”, I mean it in the standard sense in which it means a weighty reason.

This completes my case for P1. Good reasons are things that are good as reasons, and this is the same as their being weighty reasons. If that is true, then so is P1: reasons’ weights are attributive value properties.

(3) P2: The Regress

P2 states that if reasons’ weights are attributive value properties, then BPA is false. At this stage I should clarify my target. One can imagine a view like BPA which did not aim to reduce facts about value to facts about reasons, but instead merely asserted the presence of a correlation between facts about value and facts about reasons. This non-reductive view is not

my target. The reductive version of the buck passing account is the most plausible. Buck passers of all kinds want to say that facts about value do not provide us with reasons to do things. Reductive buck passers explain why facts about value nonetheless matter because facts about value are reducible to facts about reasons. But a non-reductive buck passing account would involve the denial of this claim. On a non-reductive buck passing account, facts about value neither provide reasons, nor are reducible to facts about reasons. On such an account, it would therefore be a complete mystery why we should care about value at all (cf. p18 below). So it seems that buck passers ought to be reductive buck passers, and claim that facts about values reduce to facts about reasons.¹¹

So in its best form, the buck passing account tries to explain facts about value in terms of facts about reasons. Further, it must explain facts about value by appeal not only to facts about *which* reasons there are, but also to facts about how *weighty* those reasons are. On the most straightforward view, buck passers should say that an object is good to the extent that the reasons for responding positively towards it are weighty.¹² It would be implausible to hold a view on which values reduce to reasons but the weights of those reasons are irrelevant to value.

Summarising, the buck passing account tries to explain facts about value in terms of facts about reasons, including facts about those reasons' weights. But above I defended P1, the claim that reasons' weights are attributive value properties. These two claims are unhappy partners. The buck passing account says that all facts about value reduce to facts about

¹¹ So far as I am aware, only Wedgwood 2009 defends a non-reductive buck passing account.

¹² Actually, the best buck passing view will be far more complex than this, given that (a) there are different ways in which an object can be good, (b) there are different ways in which you can respond positively to an object, and (c) there are different degrees to which you can have some particular positive response. But despite these additional complications, the basic point stands that the weights of the relevant reasons will be relevant to the object's value. I thank Philip Stratton-Lake for helping me to see the complexity here.

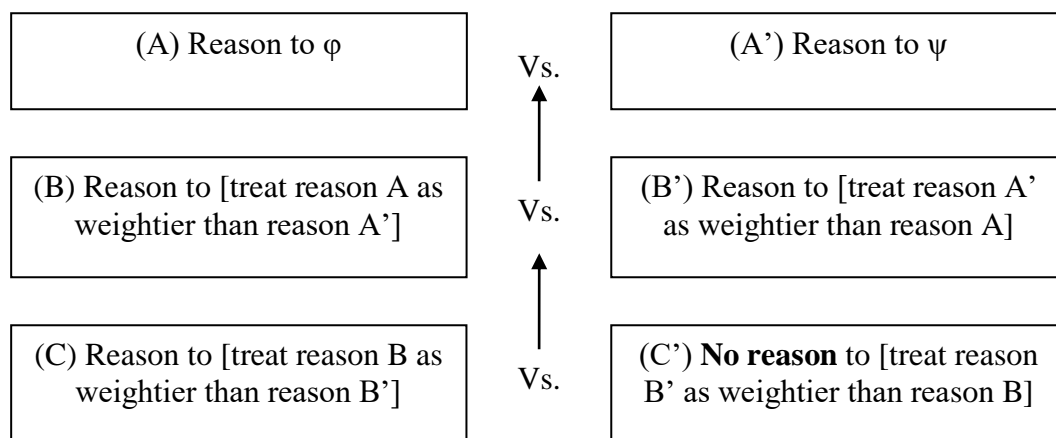
reasons and their weights. That cannot be true if reasons' weights are attributive value properties, as P1 claims. It would generate a vicious regress if we tried to reduce every fact about value to other facts about value. That is, BPA tells us that we should explain facts about how attributively good things are by appeal to facts about how weighty certain reasons are. But since the latter facts are themselves facts about how attributively good things are, we will have to explain them in the very same way, and this process generates an infinite regress: we will never really get an explanation of the original value fact at all. So P2 looks highly plausible: if reasons' weights are attributive value properties, then BPA is false.

The main problem here seems to be that P1 combined with BPA commits us to analysing the weights of reasons in terms of the weights of further reasons, and this seems like a recipe for generating a vicious regress. But Mark Schroeder has defended a view of exactly this kind (Schroeder 2007a: 123-145).¹³ He claims that the resulting view doesn't generate a vicious regress, but is instead merely committed to a *recursive* theory of the weights of reasons. So in the remainder of this section I argue against Schroeder's view. In principle, there might be some other recursive view like Schroeder's, a view that might enable us to hold onto both P1 and BPA, but which isn't subject to the same objections that Schroeder's view is. But Schroeder's theory is striking precisely because it has this recursive structure, and it's hard to imagine another plausible view of a similar kind. So if Schroeder's view fails, we should take the apparent conflict between P1 and BPA at face value, as P2 claims.

Schroeder offers his account in the broader context of his Humean theory of reasons, but I shall set that aside and focus solely on Schroeder's account of the weights of reasons. His account is a finessed version of the claim that the weights of reasons are analysed in terms of

¹³ Indeed, one wonders if his view is partly motivated by worries along the lines of those in this paper.

further reasons to weigh the initial reasons in a certain way during deliberation. I shall explain the account with reference to the following diagrammatic example:



Here we are wondering whether some reason to ϕ outweighs some reason to ψ : comparing the reasons in the top two boxes. Schroeder claims that to work out the answer, we need to know whether we have more reason to treat the reason to ϕ as weightier than the reason to ψ , or vice versa. So we now must compare the reasons in the middle two boxes. To do that, we need to know whether we have more reason to treat the reason to [treat the reason to ϕ as weightier than the reason to ψ] as weightier than the reason to [treat the reason to ψ as weightier than the reason to ϕ]. So we now must compare the reason in the bottom two boxes. At each stage of this process, we compare the weights of reasons by asking which reason we have higher-order reason to weigh more heavily.¹⁴ In general, the weights of reasons are determined by the reasons we have to weigh them in certain ways. Since the higher-order reasons to weigh lower-order reasons in a certain manner might be thought of as *fans* of the lower-order reasons, I shall call such higher-order reasons *fan-reasons*. In the diagram above, B, B', and C are all fan-reasons.

¹⁴ Note that in the diagram the reasons become higher-order as you go *down* the diagram.

This process looks like it will be endless. But Schroeder's thought is that at some stage there will be no fan-reason at all to weigh one reason more heavily than the other. This is true in my example, in which there are no reasons C' : no reasons at all to treat reason B' as weightier than reason B . At this stage, we need not attend to any further fan-reasons. If there are some reasons to treat reason B as weightier than reason B' , and no reason at all to do the reverse, then we should treat reason B as weightier than reason B' . If reason B is weightier than reason B' , then we should weigh reason A as weightier than reason A' . And if that is so, then we have more reason to ϕ than to ψ .

Summarising, Schroeder's account of the weights of reasons consists in two major claims:

- (I) One reason is stronger than another by there being more fan-reason to weigh the first as stronger than the second rather than vice versa, and this can happen either (a) by there being a less weighty fan-reason to weigh the second reason more heavily or (b) by there being no fan-reason to weigh the second reason more heavily (Schroeder 2007a: 138).
- (II) For any pair of reasons whose weight we might compare, at some stage in the recursive process the latter of these possibilities will be realised, ending the regress (Schroeder 2007a: 137-8).

This account is in many respects clever, but I doubt that it could be true. One might question whether Schroeder provides sufficient grounds for accepting (II). But here I will instead object to the view in other ways. I will try to show that the view is undermotivated and incomplete.

First, it seems that the account is undermotivated in the following manner. Schroeder motivates the view by analogy with epistemic undercutters (Schroeder 2007a: 132-3). The idea is that fan-reasons relate to reasons for action in the same way that epistemic undercutters relate to reasons for belief. For an example of an epistemic undercutter, consider that whilst blood on the knife might be a reason to suspect the butler, the fact that the butler also works as a butcher could undercut this first reason, making the blood on the knife a less good reason for suspecting the butler. This is a bad analogy for Schroeder's account. Epistemic undercutters *change* the weights of other reasons. The fact that the butler is also a butcher makes the blood on the knife a *weaker* reason for suspecting the butler. But this presupposes that that reason previously had a certain weight, which has now been lowered. In contrast, Schroeder isn't trying to explain how the weights of reasons can be changed, but instead what it is for a reason to have a weight in the first place. Knowing that a reason has become stronger or weaker tells us little about its strength unless we know how strong it was to begin with, and that is what Schroeder was trying to explain. In summary, sympathy for the existence of epistemic undercutters shouldn't make us sympathetic to Schroeder's view, since undercutters play a more modest role than fan-reasons do. Undercutters only modify prior weights of reasons, whereas fan-reasons are supposed to wholly determine the weights of reasons.¹⁵

The second problem is that Schroeder's account is lacking in detail. This is so in various ways – for example, he does not explain how reasons can be equal or incomparable in weight. But I suspect that these problems could be dealt with. A more serious incompleteness is a

¹⁵ Perhaps there is *an* analogy of epistemic undercutting in the practical domain, which Dancy calls *attenuation* (Dancy 2004: 41-2). For example, perhaps the presence of other people attenuates the strength of my reasons to help a drowning stranger. But again, we need more than attenuators to explain the weights of practical reasons. We need to know how strong the reasons were before they were attenuated. So attenuators are distinct from fan-reasons.

lack of plausible examples of fan-reasons. If I am weighing up the reasons to choose curry against those to choose pizza, what might be an example of a fan-reason to treat the reasons to eat curry as weightier? Two obvious possibilities spring to mind, but neither is available to Schroeder. First, one reason to treat the reason to eat curry as weightier is just because it is a weightier reason. But this would be circular: this is what we were hoping to explain. Second, one might have some *instrumental* reason to treat the reason to eat curry as weightier: there might be some external benefit to thinking in this manner. But Schroeder plausibly argues that such instrumental reasons cannot be fan-reasons (Schroeder 2007a: 133-8). This seems correct: if an eccentric billionaire offers you a million pounds to treat your reason to eat curry as weightier, this surely can't affect whether you in fact have most reason to eat curry.

So Schroeder can't appeal to these two obvious possible examples of fan-reasons. He does himself offer two examples of fan-reasons. His first example of a fan-reason to weigh one reason more heavily than another is the fact that that reason is agent-neutral (Schroeder 2007a: 142). There are two problems with this suggestion. First, Schroeder also claims that all fan-reasons must be agent-neutral (Schroeder 2007a: 142). If the fact that some reason is agent-neutral is a fan-reason to place weight on it, and all fan-reasons are agent-neutral, it is clear that his account will in fact generate a regress, because there will *always* be some further fan-reason to place weight on every fan-reason. That is, this suggestion undermines Schroeder claim (II), above. The second problem with this example is simply that it is hard to generalise. It would be problematic for the account if this were the only kind of fan-reason there is.

Schroeder's second example of a fan-reason to weigh one reason more heavily than another is the fact that that reason favours taking enormously costly means to fantastically frivolous

ends (Schroeder 2007a: 143). But it is again hard to generalise from this example. This fan-reason will only be relevant for instrumental reasons, but we need weightings for final reasons as well as instrumental ones. Worse, one would think that a *frivolous* end is one that one has little reason to pursue, so that this fan-reason presupposes some facts about the weights of reasons.

In summary, Schroeder's account is problematic. It is undermotivated because it relies on a bad analogy between fan-reasons and epistemic undercutters, and it is incomplete because there are not enough plausible examples of fan-reasons (if any). Unless we accept an account like Schroeder's, it is deeply unclear how to reconcile BPA with P1, the claim that reasons' weights are attributive value properties. So it seems that we should accept P2: if reasons' weights are attributive value properties, then BPA is false.

(4) P3: The Scope Of The Buck Passing Account

So far I have defended the following argument:

P1) Reasons' weights are attributive value properties.

P2) If reasons' weights are attributive value properties, then BPA is false.

C1) BPA is false.

That BPA is false is interesting: BPA might have looked like an attractive view. Indeed, some authors, such as Scanlon (2011), and Schroeder (2010: 48-9) endorse it, and if the argument above is correct, this is a mistake. But though this conclusion is itself interesting, I want to go one step further and also defend:

P3) If BPA is false, BPP is false.

P3 states that the falsity of BPA undermines BPP, the buck passing account of predicative value. In this section I shall defend P3, entitling us to the overall conclusion that both BPA and BPP are false: the buck passing account is in general unsustainable. Those who deny P3 will claim that the buck passing account is best held as an account of predicative goodness but not of attributive goodness (e.g. Parfit 2011a: 38-9). But there are at least two reasons to reject this hybrid view.

First, it is ad hoc. It is implausible to claim that “good” is ambiguous between completely unrelated things in the same way that “bank” is. Whilst it seems correct to distinguish predicative goodness and attributive goodness, it nonetheless seems that these are both kinds of one single thing: goodness. It follows that the correct analysis of either kind of goodness should provide insight into the correct analysis of the other (cf. Foot 1961: 59, Hursthouse 1999: 195-6). If the buck passing account were true of predicative value, one would think that a similar account would be true of attributive goodness.

Second, this hybrid view is difficult to justify. The main arguments for BPP are also arguments for BPA. So either these arguments fail and both BPP and BPA are unjustified, or else they succeed and we ought to endorse both views. I now explain the four main arguments for buck passing accounts, and show that they support BPA as much as they support BPP.

The first and most persuasive argument for BPP is the redundancy argument (Scanlon 1998: 97, Stratton-Lake and Hooker 2006: 154-6). If I tell you that you have reason to go to the dentist both because it would stop the pain, and because it would be good to go, it seems that the latter claim is redundant. It would be good to go just because it will stop the pain. Since

facts about predicative value do not seem to contribute anything to facts about reasons, an appealing thought is that they simply restate facts about reasons: that BPP is true. If the redundancy argument justifies BPP, it also justifies BPA. If you want to hire an assassin, and I tell you that you have reason to hire Sue, both because she will kill her target without arousing suspicion, and also because she is a good assassin, it seems that the latter claim is redundant (cf. Nowell-Smith 1969: 161). She is a good assassin just because she will kill her target without arousing suspicion. Since facts about attributive value do not seem to contribute anything to facts about reasons, an appealing thought is that they simply restate facts about reasons: that BPA is true.

The second argument for BPP is that it explains the relevance of value for choice – it demystifies value (Rabinowicz and Rønnow-Rasmussen 2004: 400, Stratton-Lake and Hooker 2006: 162-3). It seems as though it is a mistake to ask whether one should respond positively to good things. BPP explains why this is a mistake: for something to be good just is for one to have reasons to respond to it positively. If this argument justifies BPP, it also justifies BPA. It also seems as though it is a mistake to ask whether, insofar as one wants an assassin, one should respond positively to good assassins. BPA explains why this is a mistake: for someone to be a good assassin is just for you to have reasons to respond to them positively insofar as you want an assassin.

The third argument for BPP is that it explains what all valuable things have in common (Scanlon 1998: 97-8). What is there in common between the goodness of happiness and the goodness of justice? BPP provides a natural answer: they both have properties that provide reasons for positive responses. If this argument justifies BPP, it also justifies BPA. What is there in common between the fact that Sue is a good assassin and the fact that *War and Peace*

is a good book? BPA provides a natural answer: they both have properties that provide reasons for positive responses insofar as you want something of that kind.

The fourth argument for BPP is that it promises theoretical parsimony. It tells us that facts about predicative value reduce to facts about reasons, and it thereby reduces the number of normative entities that we have to posit. Indeed, BPP is often taken to be a central part of a larger project, *reasons fundamentalism*, according to which *all* normative properties can be understood in terms of normative reasons (e.g. Parfit 2011b: 268, Scanlon 2011: 443, Schroeder 2007a: 81, Skorupski 2010: 77). Reasons fundamentalism is highly parsimonious, and for that reason, attractive. But it is far from clear that reasons fundamentalism could be true unless BPP is. If this line of argument justifies BPP, it also justifies BPA. BPA tells us that facts about attributive value reduce to facts about reasons, and it thereby also reduces the number of normative entities that we have to posit. Further, reasons fundamentalism could only be true if facts about attributive value are reducible to facts about reasons, and it seems that BPA is the only possible view that entitles us to this conclusion.

In summary, if the main arguments for BPP are compelling, they also compel us to accept BPA. So if we reject BPA, we have to reject these arguments. And if we reject these arguments, BPP is unjustified. So if we reject BPA, we must think that BPP is unjustified.

Since I have argued against BPA, I am committed to the claim that the above arguments fail. Unfortunately, it would be beyond the scope of this paper to explain exactly *how* they fail. Still, we can know that they do fail, because they apparently support BPA, and BPA is false, as I argued above.

In summary, it would be ad hoc to endorse BPP but reject BPA. Predicative value and attributive value are both kinds of value, and so seem to require similar analyses. It would also be hard to justify endorsing BPP but rejecting BPA: the arguments for the views seem to be the same. These facts together tell us that P3 is true: if BPA is false, BPP is false as well.

There is an objection to P3 that I ought to address. One might think that attributive value should be analysed in purely descriptive terms, and that this provides a rationale for rejecting BPA even if one accepts BPP.¹⁶ For example, one might think that a good car is simply a car that is comfortable, fast, fuel efficient, and so on. I agree that *if* attributive value were purely descriptive, that would threaten P3. But we should not understand attributive value as purely descriptive, for at least two reasons.

First, as I argued above, this view is ad hoc. It commits us to the view that predicative goodness and attributive goodness are completely unrelated, which is independently implausible and renders it a mystery as to why we use the word “good” in both cases (cf. Hare 1970: 79-93). Second, I argued in section 2 that facts about the weights of reasons are facts about attributive goodness. On the present suggestion, facts about attributive goodness are purely descriptive facts. Together, these claims entail that facts about the weights of reasons are purely descriptive facts. But this is highly implausible. If there are any normative facts at all, facts about the weights of reasons are among them.

(5) A Positive Suggestion

¹⁶ Here, “purely descriptive” means “descriptive and not normative”. The truth of metaethical naturalism is irrelevant to the present issue.

I have defended the three premises of this argument:

P1) Reasons' weights are attributive value properties.

P2) If reasons' weights are attributive value properties, then BPA is false.

C1) BPA is false .

P3) If BPA is false, BPP is false.

C2) BPP is false.

Given that this argument is valid and the premises are true, BPA and BPP must each be false.

Buck passing accounts tell us that facts about value are reducible to facts about reasons, but this must be mistaken since facts about reasons presuppose facts about value.

If facts about reasons are characterised partly in terms of facts about value, perhaps we should go the whole way and reduce facts about reasons to facts about value. Such a view reverses the buck passing analysis, analysing reasons in terms of goodness rather than goodness in terms of reasons. Comprehensively defending a view of this kind would be a large task, and I shall have to leave the details for another time. Here I shall just sketch a view of this kind to make clear that it has various merits and no obvious faults. In short, I shall argue that the buck passing account has a serious rival that merits further investigation.

Here is the view, which I shall call *reasons as good bases*, or *RGB*:

RGB: A reason for ϕ ing is a fact that would be good as a basis for ϕ ing.

Or slightly more precisely:

RGB: A reason for ϕ ing that is good to degree D is a fact that would be good to degree D as a basis for ϕ ing.

For example, according to RGB, if Labour's education policy is a very good reason for voting for them, this amounts to the fact that Labour's education policy is very good as a basis for voting Labour. A *basis* is something that explains why an agent performed some action or changed or retained some state of mind. In the context of action, bases are often referred to as *motivating reasons*. But we could also talk about motivating reasons for states of mind: I can say that the reason why you believe that Micronesia is a country is that they said so on *The West Wing*. This belief is not exactly *motivated*, as such, but it is nonetheless held for a reason. So we might equivalently define RGB in terms of motivating reasons rather than bases, so long as we remember we are using "motivating reason" in this broader sense that includes our reasons for holding attitudes. I say more about this below.

RGB has at least six merits.

First, like the buck passing account, RGB delivers normative parsimony. It reduces facts about reasons to facts about value, and thereby reduces the number of normative entities we need to posit. Indeed, RGB allows for the possibility of some kind of *value fundamentalism* according to which all normative facts can be understood in terms of value facts (cf. Moore 1903). If we can reduce all normative facts other than value facts to facts about reasons, and facts about reasons themselves reduce to value facts, then we might be able to reduce all normative facts to value facts.

Second, relatedly, RGB explains a central normative concept: that of a reason. One popular view is that there can be no explanation of what reasons are (Scanlon 1998: 17, Parfit 2011a:

31). On such views, reasons are unanalysable. But RGB offers something more informative.¹⁷ (Of course, so far as RGB goes, we have no informative theory of what it is for some fact to be a good basis for an action or attitude. We might yet offer some independent theory analysing what good bases are, or else we might take this as primitive. Here I leave it open which of these options we should pursue.)

Third, one worry for views that analyse reasons in terms of values is that they will not be able to accommodate agent-relative reasons (Way, *forthcoming*). This is because a plausible view is that all values are agent-neutral (Schroeder 2007b), and it is hard to see how one could analyse agent-relative reasons in terms of agent-neutral values. But RGB appears to succeed in this task. According to RGB, R is a reason for you to ϕ but not a reason for me to ϕ if and only if R would be a good basis for your ϕ ing but not a good basis for my ϕ ing.

Fourth, RGB explains why reason-implies-can. A natural thought is that one ought to do something only if one can do it, and an equally natural extension of this thought is that one has a normative reason to do something only if one can do it. RGB explains this truth. RGB analyses normative reasons in terms of the features that potential motivating reasons would have: a normative reason is something that *would* be a good motivating reason. So RGB entails that a fact can be a normative reason to ϕ only if one could ϕ for that reason.

Fifth, RGB tells us something informative about the relationship between normative reasons and motivating reasons. As stated, RGB talks about bases, but as I said above, these are really equivalent to motivating reasons, at least when these are understood in the broader sense to

¹⁷ Of course, RGB will need to compete with, or else be rendered consistent with, other reductive analyses of reasons, such as views according to which reasons are analysed in terms of oughts (e.g. Broome 2004, Kearns and Star 2008).

include the explanations of why agents hold certain attitudes as well as the explanations of why they act. On the standard view, normative reasons are distinct from motivating reasons in that the former contribute to justifying, whereas the latter explain, what agents do (Smith 1994: 94-8). One can have motivating reasons which do little to justify what one is doing, and normative reasons which one fails to respond to. But despite this difference, it is nonetheless tempting to think that there are some important connections between normative and motivating reasons (cf. Dancy 2000, Williams 1981b, 1995). RGB explains one sense in which this is so. RGB tells us that normative reasons are defined in terms of motivating reasons. Normative reasons are things which would have some value property if they were motivating reasons.

Sixth, RGB unifies the concept of a reason. We talk not only about normative reasons and motivating reasons, but also about explanatory reasons, as when I say that the reason why the bridge collapsed is because the girders were too short. A natural thought is that motivating reasons are simply one kind of explanatory reason, and this explains why we use the same term for both. If RGB is correct, normative reasons are simply one kind of (possible) motivating reason, and this explains why we use the same term for both. Together, these claims let us know that reasons are, in general, explanations. Understanding normative reasons in terms of motivating reasons therefore allows us to completely unify the concept of a reason.

In summary, RGB has numerous merits. I repeat that more could be said about the details of RGB and these merits – but my goal here is just to note that *prima facie* appeal of a view along the lines of RGB.

There is one objection to RGB that I should address.¹⁸ One might think that sometimes a fact can be a good reason for ϕ ing but nonetheless be a bad basis for ϕ ing: that there can sometimes be a disconnect between a justification and an appropriate motive. For example, imagine that Steve's mother is dying, and he has the opportunity to save her. Steve's mother, it so happens, is a medical researcher on the verge of finding the cure for cancer. That she is going to find this cure seems to be a normative reason for Steve to save her, but that she will find this cure does not seem like a good motive for Steve to save her. One might think that the only good motive for Steve's saving her would be the simple fact she is his mother – other motives would somehow be inappropriate.¹⁹ Other similar examples can be found, though in what follows I stick with this case for simplicity.

I have three responses to this objection. First, we should not overstate the problem. RGB says that to the extent something is good as a reason for ϕ ing, it is good *to that same extent* as a motive for ϕ ing. So even if the fact that his mother will find this cure is *some* reason to save her, that only shows that it is *somewhat* good as a motive for saving her. If the fact that she is his mother is a better reason for Steve to save her, it would also be a better motive for Steve to save her.

Second, RGB does entail that if some fact is a good reason for ϕ ing then that fact is a good motive for ϕ ing. But that something is a good motive for ϕ ing is compatible with it being (predicatively) bad that one ϕ s from that motive (cf. Korsgaard 2009: 28). Perhaps Steve's saving his mother because she will find this cure will undermine how he feels about their relationship, and this might, in the long run, have various bad consequences (cf. Railton

¹⁸ Thanks to Donald Hubin for the objection.

¹⁹ Obviously, the example is inspired by Williams (1981a: 18), though I doubt he would agree with the overall objection being presented against RGB – see, e.g. Williams 1981b and 1995.

1984). So RGB does not commit us to the implausible claim that it is always good to be motivated by normative reasons. Sometimes, things go badly if you act from good motives.

The overall point here may be clearer when illustrated with reasons for belief. The existence of evil might well be a good basis for believing that there is no God, but it doesn't immediately follow that things go best if you believe on this basis that there is no God: perhaps atheism will make you miserable.

(There is also the reverse possibility: sometimes things go well if you act on bad motives. For example, imagine that an eccentric billionaire will reward you if you hire someone on the basis of their race. In these circumstances, it might be that things go best if you make the hire on the basis of race, but race is nonetheless not a good reason (/basis) for hiring decisions.)²⁰

Third, once the above qualifications are noted, I am strongly inclined to think that this commitment of RGB is something we should welcome rather than resist. There is something objectionable about the idea that there could be a divergence between what justifies our choices and what would be a good basis for those choices, and this is precisely the thought that RGB aims to capture (cf. Stocker 1976). Normative reasons are *normative* precisely in the sense that they are supposed to guide our decisions. With respect to the example above, it is appealing to think that either the fact that she will find this cure is genuinely a good reason for Steve to save her, in which case it would also be a good motive for his saving her, or else

²⁰ Though the reward doesn't make race a good reason for hiring decisions, it does give you a reason to do something: a reason to [make a hire on the basis of race]. It provides you with a second-order reason to act for a particular reason.

this would be a bad motive for Steve's saving her, in which case it is not a reason for Steve to save her.²¹

In summary, the objection should not persuade us to abandon RGB. RGB only says that reasons are good *to some extent* as bases for doing things, and RGB only says that reasons are good *as bases* for doing things, and thereby leaves open that it might be bad in other ways to act on these bases. Once we acknowledge these points, the remaining implication of RGB is just the tight link it draws between justification and good motives, and this looks appealing.

(6) Conclusion

In this paper I have presented an objection to the buck passing account of value, and sketched an alternative theory. This objection was that since reasons' weights are attributive value properties, any attempt to analyse facts about value in terms of facts about reasons will generate a vicious regress. My alternative suggestion was to characterise normative reasons in terms of good bases, and I argued that this view has numerous merits.²²

Bibliography

Broome, J. 2004. "Reasons" in Wallace, J., Pettit, P., Scheffler, S. and Smith, M. (Ed.)

Reason and Value: Essays on the Moral Philosophy of Joseph Raz. Oxford pp.28-55

Broome, J. 2002. "Practical Reasoning". In Bermudez, J. and Miller, A. (Ed.), *Reason and*

Nature: Essays in the Theory of Rationality. Oxford University Press. pp.85-111

Dancy, J. 2000. *Practical Reality.* Oxford.

Dancy, J. 2004. *Ethics Without Principles.* Oxford.

²¹ Perhaps this reason is simply silenced by the fact that she is his mother (McDowell 1998: 90-3).

²² Amongst too many others, I thank two anonymous referees for this journal, Bart Streumer, and the audience at SLACRR 2012, for helpful comments on this paper.

- Foot, P. "Goodness and Choice" in *Proceedings of the Aristotelian Society, Supplementary Volume 25* (1961) pp.45-80
- Geach, P. 1956. "Good and Evil" in *Analysis* 17:2 pp. 33-42
- Hare, R. 1970. *The Language of Morals*. Oxford.
- Hursthouse, R. *On Virtue Ethics*. Oxford (1999)
- Kearns, S. and Star, D. 2008. "Reasons: Explanations or Evidence?" in *Ethics* 119 pp. 31-56
- Korsgaard, C. 2009. *Self-Constitution: Agency, Identity, and Integrity*. Oxford
- McDowell, J. 'Are Moral Requirements Hypothetical Imperatives?' in his *Mind, Value, and Reality*. Harvard (1998) pp.77-94
- Moore, GE. 1903. *Principia Ethica*. Cambridge University Press.
- Nowell-Smith, P. 1969. *Ethics*. London: Pelican
- Parfit, D. 2011a. *On What Matters, Volume I*. Oxford
- Parfit, D. 2011b. *On What Matters, Volume II*. Oxford
- Rabinowicz, W. and Rønnow-Rasmussen, T. 2004. "The Strike of the Demon: On Fitting Pro-Attitudes and Value' in *Ethics* 114 pp. 391-423
- Scanlon, T. 1998. *What We Owe to Each Other*. Harvard.
- Scanlon, T. 2011. "The Unity of the Normative" in *Philosophical Studies* 154. pp.443-450
- Schroeder, M. 2007a. *Slaves of the Passions*. Oxford.
- Schroeder, M. 2007b. "Teleology, Agent-Relative Value, and 'Good'" in *Ethics* 117 pp.265-295
- Schroeder, M. 2010. "Value and the Right Kind of Reason" in *Oxford Studies in Metaethics* 5. Oxford. pp.25-55
- Skorupski, J. 2010. *The Domain of Reasons*. Oxford.
- Smith, M. 1994. *The Moral Problem*. Blackwell.

- Stratton-Lake, P. and Hooker, B. 2006 “Scanlon versus Moore on Goodness” in Horgan, T. and Timmons, M. (ed) *Metaethics After Moore*. Oxford. pp.149-168
- Stocker, M. “The Schizophrenia of Modern Ethical Theories” in *The Journal of Philosophy* 73:14 (1976) pp.453-466
- Szabó, Z. “Adjectives in Context” in Kenesei, I. and Harnish, R. (Eds) *Perspectives on Semantics, Pragmatics, and Discourse*. John Benjamins (2001) pp. 119 – 146
- Way, J. “Value and Reasons to Favour”, forthcoming in *Oxford Studies in Metaethics, Volume 8*.
- Wedgwood, R. “The ‘Good’ and the ‘Right’ Revisited” in *Philosophical Perspectives* 23 (2009) pp.499-519
- Williams, B. *Moral Luck*. Cambridge (1981)
- Williams, B. “Persons, Character, and Morality” (1981a) in his (1981) pp.1-19
- Williams, B. “Internal and External Reasons” (1981b) in his (1981) pp.101-113
- Williams, B. “Internal Reasons and the Obscurity of Blame” in his *Making Sense of Humanity, and Other Philosophical Papers*. Cambridge (1995) pp.35-45