

Normative reasons as good bases

Alex Gregory¹

© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract In this paper, I defend a new theory of normative reasons called reasons as good bases (RGB), according to which a normative reason to ϕ is something that is a good basis for ϕ ing. The idea is that the grounds on which we do things—bases—can be better or worse as things of their kind, and a normative reason—a good reason—is something that is just a good instance of such a ground. After introducing RGB, I clarify what it is to be a good basis, and argue that RGB has various attractive features: it has intuitive implications, makes good sense of the weights of reasons, and attractively explains the relationship between normative reasons and motivating reasons. I then briefly defend the view from objections and compare it to rivals. Finally, I sketch two possible implications of RGB: some kind of constitutivism, according to which the norms that govern us are explained by the nature of agency, and second, the claim that agents who do things for reasons generally do them for good reasons.

Keywords Reasons · Normative reasons · Motivating reasons · Basing · Constitutivism · Normativity · Value · Attributive goodness

If you are debating what to do this evening, you'll consider the facts that speak in favour of, or against, each possible course of action you might pursue. And if you are debating whether Black committed the murder, you'll consider the facts that speak in favour of, or against, each possible belief you might form. In both cases, we ordinarily refer to these favourers and disfavourers as *reasons*. In philosopher's

✉ Alex Gregory
a.m.gregory@soton.ac.uk

¹ Department of Philosophy, University of Southampton, Southampton, Hampshire SO17 1BF, UK

jargon, they are “normative” reasons, since they contribute to justifying our deeds and thoughts. They contrast with “motivating” reasons, which explain some of our deeds and thoughts in a distinctive manner,¹ and with “explanatory” reasons more generally, which may explain just about anything.

Reasons have seen a lot of attention in recent years. Some believe this is a passing fad, but this seems unlikely. Though older philosophers might not have discussed reasons using this label, many nonetheless discussed the same things under other labels (e.g. Ross 2002: 19–20). And to the extent that some other philosophers downplayed the importance of reasons—Kant? Aristotle?—that was a mistake that needs to be rectified. Two other kinds of scepticism about the new reasons program are, I think, more convincing. First, one might think that some discussion of normative reasons is too disconnected from discussion of what actually leads people to do things (cf. Stocker 1976; Williams 1981, 1995). Second, one might think that some treat normative reasons as more fundamental than they are. Most importantly, one might be sceptical about the claim that all facts about goodness should be explained by appeal to facts about reasons (the view is defended most notably in Scanlon 1998: 94–100, see also Scanlon 2011; I present an objection in Gregory 2014, on which this paper builds). These claims seem to me to be closer to the mark. So in this paper I shall defend a theory of reasons that rectifies these mistakes. I will analyse normative reasons partly by appeal to facts about the explanations of what we do, and partly by appeal to goodness. I will argue that a normative reason to ϕ is something that is a good basis for ϕ ing.

I begin in Sect. 1 by describing the theory and its virtues. Section 2 defends it from objections. Section 3 contrasts it with rivals. Finally, in Sect. 4, I sketch two possible implications of the theory: some kind of constitutivism, according to which the norms that govern us are explained by the nature of agency, and the view that agents who do things for reasons generally do them for good reasons.

1 RGB

The theory I shall defend I call *Reasons as Good Bases*, or *RGB* for short. It says:

RGB: A normative reason to ϕ is something that is a good basis for ϕ ing.

I hope that this view will seem intuitive. Some examples may help: imagine that I can save your life by throwing you a rope. The fact that it will save your life is a reason to throw you the rope. According to RGB, this amounts to the claim that the fact that it will save your life is a good basis for throwing you the rope. That sounds right. Or, for another example, imagine that there has been a murder, and the murder weapon is registered to the butler. The fact that the murder weapon is registered to the butler is a reason to believe that the butler did it. According to RGB, this

¹ I say “in a distinctive manner” because your good upbringing might explain your behaviour without being your motivating reason for that behaviour (see, e.g. Darwall 1983: 29). I return to related issues in 2.1.

amounts to the claim that the fact that the murder weapon is registered to the butler is a good basis for believing that the butler did it. Again, that sounds right. Or for a third example, imagine that it will be a hot dry day today. That is not a reason for me to bring my coat with me on my walk. According to RGB, this amounts to the claim that the fact that it is hot and dry today is not a good basis for taking my coat with me on my walk. Yet again, that sounds right. Or, for a fourth and final example, imagine that wishful thinking is never a reason to believe anything. According to RGB, this amounts to the claim that wishful thinking is never a good basis for belief. Yet again, that sounds right. In these and other cases, RGB has very plausible implications.

RGB analyses reasons in terms of bases, and more specifically in terms of good bases. What do I mean by a basis? I hope the notion is intuitive: you might believe the theory of evolution on a certain basis, or take a particular route to work on a certain basis. But we can say a little more: A basis for ϕ ing is what some would call a *motivating reason* for ϕ ing (see, e.g. Smith 1994). Since bases are motivating reasons, we might formulate RGB directly in terms of motivating reasons rather than bases. I might in principle be happy with the rephrased view, but have resisted this formulation for two reasons. First, talk of motivation is most at home in the context of reasons for action, and I want it to be clear that RGB is an account of all reasons, including reasons for attitudes such as reasons for belief. We might use “motivation” in a broader sense to include the grounds on which we believe things, but it is simpler for me to just use the word “basis” and avoid potential confusion. Second, when formulated in terms of motivating reasons, RGB appears to analyse reasons by appeal to reasons. This would not be circular, since the relevant kinds of reason are distinct. But it is nonetheless potentially confusing. So I will stick with talking about bases so far as is possible.

What is the relevant sense of ‘good’ in which bases can be good? ‘Good’ is not here being used in its predicative sense, since things which are normative reasons need not themselves be good things: your pain might be a reason for me to help you despite the fact that your pain is a bad thing, and the presence of clouds might not be a good thing despite the fact that it is a reason to believe that it will shortly rain. So we should instead understand ‘good’ here as being used in its attributive sense. Just as a good knife is not something which is both good and a knife (Geach 1956), but instead something which is good *as* a knife, a good basis for ϕ ing is not something which is both good and a basis for ϕ ing, but instead something which is good *as* a basis for ϕ ing. The idea is that bases for doing things can be better or worse as things of their kind, and a normative reason to ϕ is something that is towards the better end of this spectrum. We might say that according to RGB, normative reasons stand to bases (to motivating reasons) as good knives stand to knives. This fits neatly with the way in which we normally mark the distinction between motivating and normative reasons by distinguishing between ‘reasons’ and ‘good reasons’. We might, for example, ordinarily say that Hitler had his reasons for invading Poland, but that his reasons were (to put it mildly) not very good ones.

1.1 Evaluating bases

I hope that the notion of a good basis is intuitive. It is hardly incomprehensible to claim that, say, wishful thinking is a bad basis for belief. Still, it would be nice to have a little more detail about the nature of bases and how they get to be good. Knives are good to the extent that they are able to cut things. In virtue of what do bases get to be good?

Since RGB aims to explain what it is to be a normative reason, we can't say that bases are good to the extent that they correspond to normative reasons. For obvious reasons, we also shouldn't say that bases are good to the extent that they successfully move us. People are often strongly moved by terrible reasons, and vice versa.

One more attractive possibility might be to say that bases for action are good to the extent that they indicate good consequences of the favoured action, and that bases for belief are good to the extent that they indicate the truth of the favoured belief. This view seems simple and attractive, and might well seem attractive to some. However, it will also seem very unattractive to others, since these claims seem to commit us to some kind of ethical consequentialism and some kind of evidentialism about belief. Such commitments are not obviously false but are also nonetheless controversial.

Another attractive possibility might be to say that bases for action are good to the extent that they are bases on which a virtuous person would act, and that bases for belief are good to the extent that they are bases on which a virtuous person would believe. Again, such a view would be attractive to some, but nonetheless controversial in virtue of the wider commitments it would bring.

I could continue to list other possibilities like these, which combine RGB with particular claims about what makes bases good. For example, we might be able to formulate a Kantian version of RGB. But hopefully this short survey makes clear that adopting a theory about what makes bases good may well commit us to controversial substantive claims in ethics and epistemology. Since RGB is supposed to be an abstract account of reasons that is silent on such substantive issues, it is best to treat RGB itself as modular: as something that can be plugged into various different views about the evaluation of bases, in order to generate a variety of substantive commitments. RGB is attractive independently of substantive claims about the evaluation of bases, and it is best to formulate it in a manner that demonstrates that it can be held in combination with a variety of other views. Even when understood in this ecumenical manner, RGB is still attractive and significant.

Note that this does not rob RGB of content. My claim here is not that RGB employs the phrase "good basis" to refer to some as yet unspecified property. Rather, the point is that RGB can be silent about which *other* properties make a difference to how good something is as a basis for doing something. Here is a comparison: if I claim that what you ought to do is just whatever you have most reason to do, that claim is not robbed of content by my allowing that it is an open question as to which things you have most reason to do. Similarly, when I claim that reasons are good bases, that claim is not robbed of content by my allowing that it is an open question as to which bases are good.

I should note one last possibility that may be at least *somewhat* informative without being too controversial. We might say that bases are good to the extent that they make the thing they favour a good instance of its kind (cf. Finlay 2014: 85–115; Raz 1999: 23). So bases for action are good to the extent that they make the favoured action a good action, and bases for belief are good to the extent that they make the favoured belief a good belief. This is attractive but nonetheless ecumenical since it leaves open the further question of how actions or beliefs are themselves to be evaluated.

There is a related issue that I shall also set aside. Imagine that Red Rum is going to win the Grand National horse race, but the evidence available to you suggests otherwise and your beliefs reflect this evidence. Do you have any reason to bet on Red Rum? *Perspectivists* think that you don't (e.g. Lord 2015). They think of normative reasons in a way that makes them dependent on the agent's beliefs and/or evidence (henceforth: "perspective"). In contrast, *objectivists* think that you do have a reason to bet on Red Rum (e.g. Parfit 2011: 31–32). Objectivists think of normative reasons in a way that makes them independent of the agent's perspective. Still others try to resolve this debate by claiming that "reason" is simply ambiguous between these possibilities (e.g. Schroeder 2008). Sometimes, this ambiguity view is paired with the claim that the relevant distinction is one between the reasons an agent *has*, and the reasons there *are* (see Broome 2013: 65–66 for scepticism). The debate between these positions is beyond the scope of this paper, so I will instead assume that we might develop RGB in either manner, depending on how that debate turns out: we might say that a normative reason is a basis that is good given your perspective, or alternatively, that a normative reason is a basis that is good given the facts. The choice between these options plays no significant role in what follows, with the exception of 4.2, where I mention the relevant complications.

1.2 The weights of reasons

Regardless of how we resolve the issues above, RGB promises to provide a natural account of the weights of reasons. Reasons have weights that come in degrees: they can be stronger or lighter than one another, and they can also be equal, or incommensurable, in weight. Since goodness also comes in degrees, it is easy to see how RGB accommodates the weights of reasons: According to RGB, the weightiness of a normative reason is just how good it is as a basis. I might have a really good basis for ϕ ing but a still better basis for ψ ing. If that is true, then my reason to ψ is weightier than my reason to ϕ .

We should be careful with the claims above. Imagine that A is a better basis for ϕ ing than B. Does it follow that A is a stronger reason to ϕ than B? Surprisingly, it doesn't: perhaps neither A nor B is a reason to ϕ at all, and A is better than B as a basis for ϕ ing only in the minimal sense that it's less bad. For example, perhaps A is "that it will hurt someone" and B is "that it will kill someone". In light of this, we should be clear that according to RGB, a reason to ϕ is something that is a *good* basis for ϕ ing, and that the weight of a reason is given by just how *good* this basis is. Something which is a bad basis for ϕ ing is no reason to ϕ at all, and the weight of a reason is determined by how good it is, never by how bad it is.

There is another worry regarding the weights of reasons that I should address. Imagine that some reason is genuine but very weak. For example, perhaps there is some very weak reason to drink rainwater rather than tap water, since it is very marginally cheaper. Or for another example, perhaps there is some very weak reason to believe P, because you have the tiniest amount of evidence that P. Under such circumstances, we might be hesitant to say that you have a good basis for drinking rainwater, or for (outright) believing P.

Here are two possible responses to this worry. First, we might think that this is just a pragmatic phenomenon (cf. Schroeder 2007: 92–97). Often, we use “good” in a manner that implies that the thing in question is more than *slightly* good. If I say that jogging regularly is good for your health when the impact is only very negligible, that would be misleading. But for all that, we might think that the claim would be true. Equally, we might think that it is true, but misleading, to say of the cases above that you have a good basis for drinking rainwater, and a good basis for believing P. Perhaps such claims are true, but misleading because they imply that the relevant basis is not merely slightly good but instead good to some particular degree.

A second way to respond to the worry is to modify RGB.² We might say that:

RGB-modified: A normative reason to ϕ is something that is a good basis for ϕ ing, or else something that is part of a good basis for ϕ ing.

On this view, a weak reason to ϕ might attain its status in virtue of the fact that it, in conjunction with some other facts, would be a good basis for ϕ ing. This would have plausible implications for the cases above: that it’s cheaper might not itself be a good basis for drinking rainwater, but the conjunctive fact that it’s cheaper and you really need to save the money might be a good basis for drinking rainwater. Equally, your very weak evidence that P might not be a good basis for believing P, but that same evidence in conjunction with further corroborating evidence might be.

In what follows, I assume that the first response above is the best one: I tend to think on reflection that it is equally acceptable but more parsimonious. But readers might prefer to think of RGB in the above modified manner, and that does little to affect my claims in what follows.

1.3 Unifying reasons

One virtue of RGB—beyond those noted above—is that it unifies the notion of a reason. There are three apparently different senses of the word ‘reason’ illustrated by the differences between the claims ‘the volcano erupted for a reason’, ‘she murdered him for a reason’, and ‘there are good reasons to study philosophy’. The first ‘reason’ seems to simply be an explanation—an explanatory reason, the second ‘reason’ seems to be a motive—a motivating reason, and the third ‘reason’ seems to be a pro tanto justification—a normative reason.

² Thanks to an anonymous referee for this suggestion.

Normative reasons are distinct from motivating and explanatory reasons. There are normative reasons that don't motivate or explain anything, such as the normative reasons Scrooge has to give more of his money away (well, at least at the start of the story). Vice versa, there are motivating and explanatory reasons that do nothing to justify anything, such as Hitler's motivating reasons for starting the war, and the explanatory reason why clouds are white.

However, whilst the above claims seem correct, it also seems plausible that these senses of the word 'reason' are somehow related. It doesn't seem like an accident that we use the same word in all three cases: this isn't like the ambiguity present in the word 'bank'. It is very easy to see that explanatory and motivating reasons are related: a motivating reason is just one particular kind of explanatory reason. But this leaves a task for a theory of normative reasons, which is to explain how they are connected with explanatory and motivating reasons (cf. Broome 2004: 34; 2013: 50).

RGB succeeds in this task. According to RGB, a normative reason is something that is a good motivating reason. Again, the idea is that the relationship between motivating reasons and normative reasons is the same as that between knives and good knives. That ties the two notions closely together in a manner that makes sense of why we use the same word for both.

So though RGB respects the distinction between normative reasons and other reasons—P might be a good basis for your ϕ ing though you never ϕ on the basis of P, and you might ϕ on the basis that P even though P is not a good basis for ϕ ing—it also unifies normative reasons and other reasons in an attractive manner.³

2 Objections

Despite its appeal, perhaps some objection shows that RGB should be rejected. In this section I address two possible objections.

2.1 Reduction

Someone might agree that there is a perfect one-to-one correspondence between good bases and normative reasons, just as I suggest. But they might explain this correspondence in the opposite manner. That is, they might say that good bases are themselves best analysed in terms of normative reasons. On such a view, RGB—understood as *reducing* reasons to good bases—is false, but one could still accept the plausible claim that R is a normative reason to ϕ if and only if R is a good basis for ϕ ing. But this suggestion is implausible. Once we allow that there is a perfect correspondence between good bases and normative reasons, it is much more tempting to reduce normative reasons to good bases rather than vice versa. This is so for two reasons.

³ In this respect, RGB may fit well with the possibility that other words such as “principle”, “rule”, and “norm” also have systematically related normative and descriptive uses.

First, compare the case with other cases where objects of some kind can be evaluated as items of that kind, and especially good objects of that kind bear their own name. For example, perhaps thinkers can be evaluated in terms of how good they are, as thinkers, and especially good thinkers bear their own name: geniuses. (I'm sure that this is not exactly right, but the example will serve for the point being made.) RGB is parallel to the claim that we should reductively analyse geniuses as good thinkers. The objector's position is parallel to the claim that we should instead reductively analyse good thinkers as geniuses. I hope it is clear that the former is more plausible, and in turn that the same thing is true with respect to RGB: we should reduce normative reasons to good bases, rather than good bases to normative reasons.

Second, compare the case with the reduction of water to H_2O . Why do we think that water reduces to H_2O rather than vice versa? A very tempting answer to this question appeals to the fact that hydrogen and oxygen can each exist by themselves—they can each exist in places where they are not constituents of water, such as in the air (cf. Schroeder 2007: 61–72). Similar reasoning applies in the case at hand, and favours RGB over the reverse reduction. Certainly, there can be attributive values that are not constituents of normative reasons, such as the goodness of toasters, of cars, of eyes, and so on. And certainly there are bases that are not constituents of normative reasons—just take any case in which someone acts on some wholly mistaken basis, such as when I go to the fridge for a beer when there is none left. Since attributive values, and bases, can each exist independently of normative reasons, but not vice versa, it is more plausible to reduce normative reasons to good bases than vice versa.

Before I move on, I should respond to a related worry. I have said that RGB reduces normative reasons to motivating reasons of a particular kind. But we might worry that this reduction cannot be complete, because the notion of a motivating reason itself needs to be explained with reference to normative reasons.⁴ Imagine that I fail a student paper. There might be various ways to explain why I did so. Some of those explanatory reasons will be motivating reasons, but some will not. For example, that I was angry at the time might be an explanatory reason as to why I failed the paper, but not a motivating reason. In contrast, that the paper seemed to show no knowledge of the course would not merely be an explanatory reason but also a motivating reason as to why I failed the paper. How should we distinguish motivating reasons from other explanatory reasons? One possible answer is that a motivating reason is something that the agent believed to be a normative reason (e.g. Schroeder 2007: 14). In the example above, we might think that my anger was not my motivating reason because—unlike the ignorance the paper showed—it was not something that I believed to be a normative reason for failing the paper. But if this is the right way to understand motivating reasons, it might seem as though it would be circular to reductively analyse normative reasons in terms of motivating reasons, as RGB does. We can't analyse normative reasons in terms of motivating reasons if we also want to analyse motivating reasons by appeal to normative

⁴ Thanks to an anonymous referee and Garrett Cullity for this objection.

reasons. There are two broad ways to handle this objection, which I here remain neutral between.

First, one might reject the proposed analysis of motivating reasons, and adopt some rival view instead. There are various arguments against the proposed analysis of motivating reasons. For example, some claim that it is too intellectualised (e.g. Arpaly and Schroeder 2013, though cf. Enoch 2011: 225–230, Gregory forthcoming). Others claim that it rules out some possible kinds of akrasia (e.g. Stocker 1979, though cf. Gregory forthcoming). Setiya (2007) claims that the proposed view is implausible because rival theories are better placed to explain our non-inferential knowledge of our own actions. In light of such objections, we might analyse motivating reasons in some other manner, such as in terms of causation of the right kind by other kinds of mental state (e.g. Setiya 2007; Smith 1994: 92; Turri 2011).

Alternatively, if we were to adopt the proposed analysis of motivating reasons, we might continue to hold RGB by denying that the resulting combination of views is problematically circular. Whereas RGB analyses normative reasons in terms of motivating reasons, the view above analyses motivating reasons in terms of *beliefs* about normative reasons. Since these analyses make reference to different things—one to normative reasons themselves, one to beliefs about such reasons—their combination is not strictly circular. Their combination might yet be indirectly circular if we could only explain what it is to have a belief about a normative reason by referring to normative reasons themselves. But we might deny this. To explain what it is to have a belief about a normative reason requires that we analyse the *concept* of a normative reason.⁵ But we might endorse RGB as a metaphysical theory of the nature of reasons themselves, rather than as a theory of the relevant concept. By doing this, we could agree with the proposed analysis that the concept of a normative reason plays a role in the analysis of motivating reasons, but nonetheless maintain RGB by claiming that normative reasons themselves are metaphysically analysed in terms of motivating reasons. The resulting view might bear some resemblance to response-dependent analyses of other properties, which define properties partially in terms of our attitudes towards those very properties (for discussion, see, e.g. McDowell 1998; Street 2008, esp. 239–242; Wright 1988).

2.2 Inappropriate motives

Another objection is that it is sometimes bad to respond appropriately to our reasons, and vice versa, sometimes it is good to respond to things which aren't reasons. For example, even if I have a good reason to believe I'll lose at tennis this weekend, it may be bad for me to respond to this reason since doing so might make my chances still worse (cf. Railton 1984). Vice versa, though a millionaire might incentivise me to intend to drink poison, rendering that intention a good one to have, it's not clear that this incentive gives me a reason for forming that intention (Kavka

⁵ Perhaps even less than this, if we think that a belief need not employ the exact concept but only some approximation thereof.

1983).⁶ Such cases might appear to be counterexamples to RGB, which says that reasons are good bases.

But this objection rests on a misunderstanding of RGB (here I more briefly repeat what I say in Gregory 2014: 300–302). RGB analyses reasons in terms of facts about which things are good as bases, not in terms of facts about which bases it is (predicatively) good to base your responses on. So RGB does not say that reasons are things it would be good to base your responses on. That something is a good basis for ϕ ing is consistent with its being bad that you ϕ on that basis. So in the tennis case above, though it might be bad for you to believe that you'll lose, the evidence might nonetheless be a good basis for that belief. Vice versa, that something is not a good basis for ϕ ing is consistent with its being good that you ϕ on that basis. In the poison case above, though it might be good to intend to drink the poison, it might nonetheless be true that the incentive is not a good basis for this intention.

3 Rivals

Hopefully, at this stage it is clear what RGB says, that it has various attractive features, and that it can overcome some of the most obvious objections. In this section I will briefly compare RGB with rivals.

3.1 Setiya

Kieran Setiya offers an analysis of reasons that may be close to RGB. He claims that:

“ p is a reason for A to ϕ just in case A has a collection of psychological states, C , such that the disposition to be moved to ϕ by C -and-the-belief-that- p is a good disposition of practical thought, and C contains no false beliefs.” (2007: 12, 2014).

Despite some similarities, there are various differences between Setiya's view and RGB. I shall focus on two.⁷ First, Setiya treats the claim above as a true biconditional, but not a reductive account of what reasons are, unlike RGB (Setiya 2007: 13, 2014: 223). RGB is in this respect more explanatorily powerful since it gives a complete analysis of what a reason is.

⁶ This second case might be more controversial than the first: it relies on the claim that there are no so-called “wrong kinds of reasons” (see, e.g. Rabinowicz and Rønnow-Rasmussen 2004: 411–414). If one disagreed with that assumption, all the better for RGB.

⁷ One other difference that I set aside is that as stated, RGB said that normative reasons are good bases. In contrast, according to Setiya's view, whereas normative reasons are propositions, we are moved by psychological states. In this way, Setiya's view rules out the possibility that normative reasons are things we are moved by (cf. Dancy 2000; Setiya 2007: 29–31). This difference is interesting, but beyond the scope of the paper, and at any rate I take it that either view could easily be reformulated to bring it closer to the other in this respect. Most saliently, we could modify RGB so that it said that a normative reason to ϕ is something such that the belief in it is a good basis for ϕ ing.

A second more significant difference is that whereas Setiya's view takes the weight of a reason to be determined by the degree of motivation that good dispositions dispose you to have (2007: 13, 2014: 222, 228–232), RGB claims that the weight of a reason to ϕ corresponds to the degree of goodness it has as a basis for ϕ ing. In this respect, RGB seems to be superior. This is clearest in the case of reasons for belief, since in this context it's not even clear what it would mean to talk about degrees of motivation to believe something (see Way forthcoming), though it does make sense to talk about how good your basis for believing is.⁸

I take it that in the two respects above, RGB is somewhat superior. Further, I take it that RGB has the advantage of parsimony over Setiya's view, in that it is formulated more straightforwardly. Relatedly, I take that RGB has intuitive appeal in that it is common sense to agree that reasons are good bases in a way that it isn't common sense to sign up to Setiya's complex claim. For all these reasons, I take it that we should prefer RGB to Setiya's view.

3.2 Reasons primitivism

Reasons primitivism says that the notion of a reason is irreducible and cannot be analysed at all (Parfit 2011: 31; Scanlon 1998: 17). Reasons primitivism is hard to undermine. It is not as though the theory might be internally inconsistent or extensionally incorrect. But nonetheless, it also doesn't have much in its favour. Certainly, reasons primitivism has little explanatory power, since it must take many truths about reasons as inexplicable. We might adopt the view if we could not find any plausible and informative analysis of what a reason is, but if we can find some such analysis, there is no reason to adopt reasons primitivism rather than the proposed reduction. This is really to say that I need no argument against reasons primitivism. If RGB is plausible, I take that as enough to show that it is preferable to reasons primitivism.

3.3 Reasons as explanations of oughts

We might try to analyse reasons in terms of oughts. John Broome analyses reasons as the things that play a certain role in the explanation of what one ought to do (Broome 2004, 2013: 46–57). In particular, he divides reasons into *perfect reasons*, and *pro tanto reasons*. Perfect reasons are complete explanations of what you ought to do (2004: 34–36, 2013: 49–51). Pro tanto reasons, in contrast, are considerations that play the “for- ϕ ing role” in a “weighing explanation” of what you ought to do (2004: 36–39, 2013: 51–57).

We might combine this view with RGB, and say that a basis is good to the extent that it makes a certain kind of difference to what you ought to do. Such a view

⁸ Way's own view says that reasons should be defined in terms of the role they play in good reasoning (forthcoming). Again, this view may be relatively similar to RGB. However, this view seems distinct from RGB since not all basing is done via reasoning. Way's view seems to have the implication that perception is reason providing because of its connection with reasoning, rather than basing, and this seems implausible to me.

would entitle us to agree with everything I have said in favour of RGB, and everything Broome says in favour of his view. But I think it is more attractive to adopt RGB without combining it with Broome's view. A general worry is just that Broome's view rules out the attractive view that what you ought to do should be understood in terms of what you have most reason to do. If we find this view attractive, we should reject Broome's view which pursues the reverse analysis. But since this objection may be thought to beg the question, here is another (familiar) worry (see Brunero 2012: 812–816). Broome's view seems circular. Whilst his account of perfect reasons may be acceptable, his account of pro tanto reasons appeals to the notion of a weighing explanation. But since this is obviously not *literally* a weighing explanation, the notion of a weighing explanation must only be an analogy. But when we ask what it is an analogy for, the only viable answer seems to be: for the way in which reasons for and against an option combine to determine what you ought to do. But if this is the model that Broome's view appeals to, it appeals to the very thing that needs explaining.⁹

3.4 Reasons as evidence of oughts

Another view that aims to analyse reasons in terms of oughts is Daniel Star and Stephen Kearns' view, according to which reasons are analysed as *evidence* of what one ought to do (Kearns and Star 2008, 2009, also Thomson 2008: 130, 146).

Again, we might combine this view with RGB, and say that a basis for ϕ ing is good to the extent that it is evidence that you ought to ϕ . Such a view would entitle us to agree with everything I have said in favour of RGB, and everything Kearns and Star say in favour of their view. But again, I think it is more attractive to adopt RGB without combining it with their view.

This is because their view is unable to capture all of the normative force of reasons. Imagine that you face a choice between voting for Kate Conflicted and Mary Mediocre. Kate's policies are almost universally better: she has a better health policy, a better defence policy, a better economic policy, and so on. Sadly, she has one *terrible* policy: perhaps she intends to imprison all homosexuals, start an unjust war, or whatever else. The terrible policy all by itself is conclusive evidence that you ought vote Mary rather than Kate. Nonetheless, you might feel a sense of conflict in this choice: if it weren't for this single policy, Kate would be a better candidate. Such a sense of conflict makes sense because defeated reasons still have normative force. But it's hard to explain that fact on Kearns and Star's view. If, taking all your present evidence into account, you have conclusive evidence that P, your present attitude towards your present evidence that $\neg P$ ought to be dismissive:

⁹ Broome responds to a circularity worry in his 2013 (pp. 54–55), but there he explicitly takes for granted that the objector allows that we can understand the notion of a weighing explanation independently of understanding what reasons are. As such, Broome's claims there are probably best understood as addressing a different circularity worry: that the for- ϕ ing role in a weighing explanation cannot be identified independently of our grasp on reasons (see e.g. Kearns and Star 2008: 42–44). This may be true, but the objection in the main text is that one cannot even identify the relevant sense of "weighing explanation" itself without relying on a prior understanding of reasons.

such evidence is just misleading (of course, your attitude might reasonably change if you acquire new evidence: see Harman 1973: 148–149).¹⁰ But that shouldn't be one's attitude towards the reasons to vote for Kate: they aren't merely red herrings that should be ignored, but instead a real cause for regret. Conclusively defeated evidence has no normative force, whereas conclusively defeated reasons do have normative force. So reasons cannot simply be evidence of what one ought to do.

This point is clearest when taken to its extreme. Imagine an omniscient agent. They know everything, including what they ought to do. Such an agent, I take it, would never feel any sense of conflict when faced with apparently conflicting evidence. They would know which pieces of evidence were misleading, and discount them. But such an agent would not be free of all conflict. They might still face moral conflicts, for example. If such an agent should feel no conflict between opposing evidence, but should feel conflict between opposing reasons, Kearns and Star's view must be false.

4 Two possible implications of RGB

Having briefly addressed rival views, in this final section, I tentatively sketch two possible implications of RGB. I do so to highlight the significance of RGB and to indicate avenues of inquiry worth exploring further.

4.1 Constitutivism

A first possible implication of RGB is some kind of *constitutivism* (Korsgaard 2008, 2009; Velleman 2000). By “constitutivism”, I mean the view that the non-normative nature of agency somehow explains or generates the norms governing us. As Velleman describes his ambitions for his version of the view: “my purpose in trying to identify a constitutive aim of action is to find a non-normative foundation for our norms of practical reasoning” (2004: 287, see also Enoch 2006: 169; Korsgaard 2008: 61).

RGB might support some form of constitutivism. Knives can be evaluated as better or worse, and plausibly this is because of the sorts of things knives are: it is in the nature of knives that they can be so evaluated. So the norms that govern knives are explained by the nature of knives. RGB tells us that normative reasons stand to bases as good knives stand to knives. So in turn, according to RGB, bases can be evaluated as better or worse, and plausibly this is because of the sorts of things bases are: it is in the nature of bases that they can be evaluated as better or worse. So according to RGB, the norms that govern bases are explained by the nature of bases. That sounds very much like the constitutivist claim above that the non-normative nature of agency—the nature of basing—explains the norms governing us. To this extent, RGB promises to provide an explanation of why we are subject to norms: because we do things for reasons. It is because we do things for reasons that our

¹⁰ We might phrase this as: Conclusively defeated evidence is thereby undermined.

behaviour is governed by certain standards it is supposed to live up to. This explanatory claim seems attractive.

“Constitutivism” is sometimes taken to have other commitments beyond the explanatory claim above. For example, constitutivist views are sometimes thought to say that agents have to conform to certain norms if they are going to qualify as agents at all (Korsgaard 2009; Velleman 2000: 15–20, also see Enoch 2006). This claim generates a problem: it seems to entail that no agent can possibly violate the norms in question (see e.g. Railton 1997: 70, cf. Korsgaard 2009: 160). Not only is that implausibly optimistic, but it also undermines the status of those norms as norms: a norm that cannot be violated is no norm at all. But RGB is not committed to the claim that we have to conform to certain norms in virtue of the fact that we do things for reasons. Rather, it says that, just as certain norms *apply* to knives in virtue of their status as knives, certain norms *apply* to us in virtue of the fact that we do things for reasons (cf. the understanding of constitutive norms in Williamson 2000: 240). So RGB does not entail that wrongful choice is impossible for agents, and in turn there is no problem of this kind for RGB. I tend to think of this as showing that RGB is a version of constitutivism that is immune to one major objection to such views, but we need not fuss about the label, and I am happy if readers would prefer to think of RGB as a view that is not a version of constitutivism but instead a rival view which nonetheless shares the attractive explanatory power that constitutivist views have.

4.2 The normativity of mind

Above I argued that RGB commits us to a form of constitutivism which explains why certain norms apply to us, but which rightly permits that we can fail to abide by those norms. But one might think that there is a more indirect route from RGB to the problematic claim that we necessarily abide by certain norms. I shall begin by treating this as a possible objection to RGB, but eventually we shall see that the relevant implication of RGB is not problematic but instead potentially attractive.

The basic line of reasoning behind the objection to RGB is this. RGB treats the relationship between motivating and normative reasons as the same as the relationship between knives and good knives. And we might think that in order to qualify as a *knife* at all, an object must be at least somewhat good as a knife. It appears to follow that to count as a basis, a consideration must be at least somewhat good as a basis. If a good basis is a normative reason, as RGB claims, it follows that people must always do things for normative reasons. RGB permits that the reason in question might be very weak, but this is nonetheless still implausible: people can do things for reasons that are not at all good.¹¹

¹¹ Note that this is likely to remain true even if we are *perspectivists* about normative reasons (see 1.1 above). Perspectivists claim that the normative reasons you have are in some sense relative to your perspective, but this view nonetheless still permits that people do things for bad reasons. For example, you might fail to respond properly to the evidence available from your perspective.

The argument is easier to assess if we spell it out as follows:

P1) Normative reasons are good bases.

P2) If K's are evaluable as items of their kind, and X is of kind K, X must be at least somewhat good as a K.

So, C) If someone is ϕ ing on some basis, they must be ϕ ing for a normative reason.

P1 is just RGB, and at this stage merits little further comment. And it should be clear that C is implausible: if RGB entails C, that would be good grounds for rejecting RGB. And the argument is clearly valid. So if we want to defend RGB, we must reject P2.

P2 says that membership of a functional kind requires some competence at the relevant function. For example, anything that qualifies as a knife must be at least somewhat able to cut things, anything that qualifies as a car must be at least somewhat able to transport people, anything that qualifies as an eye must be at least somewhat able to see, and anything that qualifies as a muscle must be at least somewhat able to contract. There is certainly some appeal to P2. It does seem impossible that a soap bubble could ever qualify as a knife, or that a toenail could ever qualify as an eye, and a tempting thought is that such things are impossible because these objects would be completely useless at performing the tasks definitive of knives and eyes. Similarly, we might think that anything that qualifies as a basis must be somewhat good as a basis—must be a normative reason.

However, things are not so straightforward. For it also seems clear that there are blunt knives, broken cars, sightless eyes, and torn muscles, and such possibilities appear to undermine P2. One possible reply is that such things must still be *somewhat* good as instances of their kind, even if the extent to which they are good is very limited indeed. But this seems forced. A broken car might be completely immobile: it's hard to see in what sense this object is even somewhat good as a car.

Aristotle offers a marginally more persuasive defence of P2 when he claims that 'the eye is the matter of sight, so that when sight leaves it it is no longer an eye except homonymously, in the way of a stone or painted eye' (*De Anima* 412b). The idea seems to be that phrases like 'sightless eye' or 'broken car' are similar to phrases like 'false beard' and 'alleged communist'. Just as it would be a mistake to think that a false beard is a kind of beard, or an alleged communist a kind of communist, Aristotle claims that a sightless eye is not a kind of eye, and a broken car is not a kind of car.

However, Aristotle's defence of P2 is also implausible. One way to see this is to think about the relevant inferential entitlements. If Clive is a communist, and Alan is merely an alleged communist, we cannot infer that there are two communists. We cannot make this inference precisely because an alleged communist is not a kind of communist. But things are different with broken functional items. If you own a working *Ford Fiesta*, and a broken *Volkswagen Golf*, we can infer that you own two cars. A natural thought is that we can make this inference because—*pace* Aristotle—a broken car is a kind of car.

In summary, it seems that P2 is indefensible. It follows that the above argument is unsound, and RGB doesn't require that agents always do things for good reasons.

This is good news for defenders of RGB. But wasn't there something attractive about the original line of reasoning? We might think that there is a more modest way to formulate P2, and in turn a more modest implication of RGB. I consider two possibilities.

First, we might modify P2 so that it states that objects of evaluable kinds are necessarily *intended* to be good. The idea is that though knives can be blunt and cars immobile, such objects were at least created with the intention that they would be able to cut things and transport people. This claim, in conjunction with RGB, might seem to support the conclusion that agents always do things on the basis of considerations that *they intend* to be normative reasons. This claim is strikingly similar to the view sometimes labelled "the guise of the good" (see, e.g. Gregory 2013; Setiya 2007; Velleman 1992, as well as Sect. 2.1 above). We might formulate this argument as follows:

P1) Normative reasons are good bases.

P2*) Necessarily: If K's are evaluable as items of their kind, and X is of kind K, X is *intended to be* at least somewhat good as a K.

So, C*) Necessarily: If someone is ϕ ing on some basis, they *intend to be* ϕ ing for a normative reason.

There is some attraction to this argument, but it also faces various obstacles. One problem is that P2* is not plausible as stated. We tend to think of many biological entities as evaluable as items of their kind—to reuse my examples above, we evaluate eyes and muscles as better or worse as members of their kinds—but such entities were not (I assume) created with any intention. We might try to fix the argument by treating "intention" in some loose sense that allows us to say that eyes and muscles are "intended" *by evolution* to fulfil a certain function. But when understood in this manner, the content of C* is harder to discern, since the sense in which evolution "intends" for eyes and muscles to perform certain functions is highly contested (see, e.g. the essays in Allen et al. 1998). Certainly, when the argument is read in this manner it is not clear that the conclusion amounts to the guise of the good, since on that view the bases on which we act are *literally* thought to be good, and it's clear that evolution doesn't *literally* think about eyes or muscles at all. A second problem for P2* is that it's not clear that such intentions are necessary even in the case of artefacts. I might blindly follow some instructions to create some object, and discover only at the end of the process that the instructions are instructions for creating a bomb. In such circumstances I need not start with the intention to make a good bomb, but a bomb is created nonetheless. So though this form of the argument has some attraction, it would require much further work to show that it can overcome these obstacles.

A second way to modify the original argument is to agree that some functional items are useless as members of their kind, but nonetheless claim that such objects are somehow *abnormal*. In turn, we might think that RGB entails that there is something abnormal about acting for bad reasons. Shortly, I discuss what notion of

normality is in play here. But for now let's stay with an intuitive understanding of the term.¹² This revised argument can be formulated as follows:

P1) Normative reasons are good bases.

P2+) Normally: If K's are evaluable as items of their kind, and X is of kind K, X is at least somewhat good as a K.

So, C+) Normally: If someone is ϕ ing on some basis, they are ϕ ing for a normative reason.

P2+ allows that membership of a functional kind doesn't *require* competence at the relevant function, and that membership of a functional kind doesn't necessarily hinge on any prior intentions. Instead, P2+ says that membership of a functional kind *normally* involves competence at the relevant function. The normal expectation is that knives can cut things, that cars can transport people, that eyes can see, and that muscles can contract. But P2+ permits that sometimes things are abnormal, and when that happens, an item might be wholly incompetent at performing the functions characteristic of its kind.

If we do accept P2+ understood as a claim about normality, we should accept C+, according to which the reasons for which people do things are normally good ones. I take it that C+ represents one way of formulating the claim that the mind is normative (cf. Davidson 2001a, b, c, 2004; Dennett 1987, especially pp. 2–35). It says that people who do things for reasons normally do so for good reasons.

One complication is raised here by the distinction between objectivism and perspectivism that I mentioned in Sect. 1.1, above. As I said there, objectivists think of normative reasons in a manner that makes them insensitive to our beliefs and evidence, whereas perspectivists think of normative reasons in a manner that makes them sensitive to some such facts about our perspective. If we are objectivists, then C+ above *may* seem implausible. For according to objectivists, what you have normative reason to do is independent of the evidence available to you. But according to C+, people normally do things for normative reasons. This combination of claims would tell us that people normally do things in response to reasons that they may well have no access to, and this may be implausible (cf. Lewis 1974: 336). So perhaps we should endorse the argument above only if we also endorse perspectivism. But a lot hangs here on how exactly we understand "normal", since objectivists might claim that circumstances where evidence is misleading just *are* abnormal circumstances.

Whatever we say about this issue we will still need to clarify the sense of "normality" in play in this argument. A natural reading of 'normal' takes it to mean 'most frequent'. But this natural reading would render P2+ false. Good headphones transmit noise to our ears, but this is consistent with the possibility that the majority of headphones in the world are broken ones in waste dumps.

How else might we understand normality here? One way to understand the relevant notion of normality here is in terms of defaults. We might say that the

¹² One might analyse my use of "normal" in terms of Ruth Millikan's very technical sense of "Normal" (1984). I won't argue against that possibility here, but vice versa, it isn't what I have in mind here, which is instead some other more everyday sense of "normal".

default expectation is that members of functional kinds are somewhat good as members of their kind. This allows for exceptions, but tells us that in those cases where there are exceptions, we should expect a positive explanation of the deviation. So for example, we by default expect knives to be sharp, but knives can obviously be blunt so long as there is some explanation of how it came to be that way, such as overuse, or poor manufacturing. Similarly, we might think, the default expectation is that people do things for good reasons, but they can obviously do things for bad reasons so long as there is some explanation of the fault: a heuristic gone wrong, or poor education (cf. Dennett 1987: 18).

Another related possibility is that we should understand ‘normal’ in the above argument as indicating a *generic*: the sense in which cats have four legs, and humans have two eyes (see Leslie 2012 for further references, including on the relationship between generics and defaults). If we understood P2+ as stating a generic, then the conclusion of the argument—that if someone is doing something for a reason, it’s normally a good reason—should also be read as a generic, and this might seem plausible (cf. Setiya 2012: 128–158). That said, there is some uncertainty about how to understand generics, and that uncertainty would be inherited by the claims in this argument if we understood the relevant claims as generics.

Whichever way we go, the argument above might seem attractive, and in turn RGB might be thought to imply one (modest) version of the claim that agents who do things for reasons generally do them for good reasons.

5 Conclusion

In this paper my primary goal has been to defend RGB:

RGB: A reason to ϕ is something that is a good basis for ϕ ing.

I have argued that we have a good basis for believing in RGB, and no good basis for rejecting it. I have argued that we have a good basis for rejecting rivals to RGB. And I have argued that RGB is a good basis for endorsing some kind of constitutivism about the norms governing agents, and is a good basis for believing that agents who do things for reasons generally do them for good reasons.¹³

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Allen, C., Bekoff, M., & Lauder, G. (Eds.). (1998). *Nature’s purposes*. Cambridge, MA: MIT Press.
 Arpaly, N., & Schroeder, T. (2013). *In praise of desire*. Oxford: Oxford University Press.

¹³ Numerous people have given me helpful comments on this paper, but special thanks to Jonathan Way, Conor McHugh, and an anonymous referee for this journal.

- Broome, J. (2004). Reasons. In J. Wallace, P. Pettit, S. Scheffler, & M. Smith (Eds.), *Reason and value: Essays on the moral philosophy of Joseph Raz* (pp. 28–55). Oxford: Oxford University Press.
- Broome, J. (2013). *Rationality through reasoning*. Oxford: Blackwell.
- Brunero, J. (2012). Reasons as explanations. *Philosophical Studies*, 165, 805–824.
- Dancy, J. (2000). *Practical reality*. Oxford: Oxford University Press.
- Darwall, S. (1983). *Impartial reason*. Ithaca: Cornell University Press.
- Davidson, D. (2001a). Actions, reasons, and causes. In *Essays on actions and events* (pp. 3–19). Oxford: Oxford University Press.
- Davidson, D. (2001b). Mental events. In *Essays on actions and events* (pp. 207–225). Oxford: Oxford University Press.
- Davidson, D. (2001c). Radical interpretation. In *Problems of rationality* (pp. 125–139). Oxford: Oxford University Press.
- Davidson, D. (2004). Incoherence and irrationality. In *Problems of rationality* (pp. 189–198). Oxford: Oxford University Press.
- Dennett, D. (1987). *The intentional stance*. Denver: Bradford books.
- Enoch, D. (2006). Agency, schmagency: Why normativity won't come from what is constitutive of action. *Philosophical Review*, 115(2), 169–198.
- Enoch, D. (2011). *Taking morality seriously*. Oxford: Oxford University Press.
- Finlay, S. (2014). *Confusion of tongues*. Oxford: Oxford University Press.
- Geach, P. (1956). Good and evil. *Analysis*, 17(2), 33–42.
- Gregory, A. (2013). The guise of reasons. *American Philosophical Quarterly*, 50(1), 63–72.
- Gregory, A. (2014). A very good reason to reject the buck passing account. *Australasian Journal of Philosophy*, 92(2), 287–303.
- Gregory, A. (forthcoming). Might desires be beliefs about normative reasons? In J. Deonna & F. Lauria (Eds.), *The nature of desire*. New York: Oxford University Press.
- Harman, G. (1973). *Thought*. Princeton: Princeton University Press.
- Kavka, G. (1983). The toxin puzzle. *Analysis*, 43(1), 33–36.
- Kearns, S., & Star, D. (2008). Reasons: Explanations or evidence? *Ethics*, 119, 31–56.
- Kearns, S., & Star, D. (2009). Reasons as evidence. In *Oxford studies in metaethics* (Vol. 4, pp. 215–242). Oxford: Oxford University Press.
- Korsgaard, C. (2008). The normativity of instrumental reasoning. In *The constitution of agency* (pp. 27–68). Oxford: Oxford University Press (reprinted).
- Korsgaard, C. (2009). *Self-constitution*. Oxford: Oxford University Press.
- Leslie, S. (2012). Generics. In G. Russell & D. Fara (Eds.), *The Routledge companion to philosophy of language* (pp. 355–367). London: Routledge.
- Lewis, D. (1974). Radical interpretation. *Synthese*, 23, 331–344.
- Lord, E. (2015). Acting for the right reasons, abilities, and obligation. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 10). Oxford: Oxford University Press.
- McDowell, J. (1998). Values and secondary qualities. In *Mind, value, and reality* (pp. 131–150). Cambridge: Harvard University Press.
- Millikan, R. (1984). *Language, thought, and other biological categories*. Cambridge: MIT Press.
- Parfit, D. (2011). *On what matters* (Vol. 1). Oxford: Oxford University Press.
- Rabinowicz, W., & Rønnow-Rasmussen, T. (2004). The strike of the Demon: On fitting pro-attitudes and value. *Ethics*, 114(3), 391–423.
- Railton, P. (1984). Alienation, consequentialism, and the demands of morality. *The Philosophical Review*, 13(2), 134–171.
- Railton, P. (1997). On the hypothetical and non-hypothetical in reasoning about belief and action. In G. Cullity & B. Gaut (Eds.), *Ethics and practical reason* (pp. 53–80). Oxford: Oxford University Press.
- Raz, J. (1999). *Engaging reason*. Oxford: Oxford University Press.
- Ross, D. (2002). *The right and the good*. Oxford: Oxford University Press.
- Scanlon, T. (1998). *What we owe to each other*. Cambridge: Harvard University Press.
- Scanlon, T. M. (2011). The unity of the normative. *Philosophical Studies*, 154(3), 443–450.
- Schroeder, M. (2007). *Slaves of the passions*. Oxford: Oxford University Press.
- Schroeder, M. (2008). Having reasons. *Philosophical Studies*, 139(1), 57–71.
- Setiya, K. (2007). *Reasons without rationalism*. Princeton: Princeton University Press.
- Setiya, K. (2012). *Knowing right from wrong*. Oxford: Oxford University Press.
- Setiya, K. (2014). What is a reason to act? *Philosophical Studies*, 167, 221–235.
- Smith, M. (1994). *The moral problem*. Oxford: Blackwell.

- Stocker, M. (1976). The schizophrenia of modern ethical theories. *The Journal of Philosophy*, 73(14), 453–466.
- Stocker, M. (1979). Desiring the bad: An essay in moral psychology. *The Journal of Philosophy*, 76(12), 738–753.
- Street, S. (2008). Constructivism about reasons. In R. Shafer-Landua (Ed.), *Oxford studies in metaethics* (Vol. 3, pp. 207–245). Oxford: Oxford University Press.
- Thomson, J. (2008). *Normativity*. Chicago: Open Court.
- Turri, J. (2011). Believing for a reason. *Erkenntnis*, 74(3), 383–397.
- Velleman, D. (1992). The guise of the good. *Noûs*, 26, 3–26.
- Velleman, D. (2000). *The possibility of practical reason*. Oxford: Oxford University Press.
- Velleman, D. (2004). Replies to discussion of the possibility of practical reason. *Philosophical Studies*, 121(3), 277–298.
- Way, J. (forthcoming). Reasons as premises of good reasoning. *Pacific Philosophical Quarterly*. doi:[10.1111/papq.12135](https://doi.org/10.1111/papq.12135)
- Williams, B. (1981). Internal and external reasons. In B. Williams (Ed.), *Moral luck* (pp. 101–113). Cambridge: Cambridge University Press.
- Williams, B. (1995). Internal reasons and the obscurity of blame. In B. Williams (Ed.), *Making sense of humanity, other philosophical papers* (pp. 35–45). Cambridge: Cambridge University Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Wright, C. (1988). Moral values, projection, and secondary qualities. *Proceedings of the Aristotelian Society Supplementary*, LXII, 1–26.