
REASONS AS PREMISES OF GOOD REASONING

BY

JONATHAN WAY

Abstract: Many philosophers have been attracted to the view that reasons are premises of good reasoning – that reasons to ϕ are premises of good reasoning towards ϕ -ing. However, while this *reasoning view* is indeed attractive, it faces a problem accommodating outweighed reasons. In this article, I argue that the standard solution to this problem is unsuccessful and propose an alternative, which draws on the idea that good patterns of reasoning can be defeasible. I conclude by drawing out implications for the debate over pragmatic reasons for belief and other attitudes and for one influential form of reductionism about the normative.

A compelling thought is that there is an intimate connection between normative reasons – reasons *to* act, believe, desire, or otherwise respond – and reasoning. For some authors, this connection is descriptive – a reason for you to ϕ must be something from which you *could* reason towards ϕ -ing. As John Searle puts it, ‘[y]ou have to be able to reason with reasons’ (2001, p. 104). For others, the connection is normative – a reason for you to ϕ must be an appropriate premise for reasoning towards ϕ -ing. As Mark Schroeder puts it, ‘[w]hen an agent is reasoning well, the kinds of things about which he should be thinking are his reasons’ (2007, p. 26).

The descriptive claim, and its proper interpretation, is rightly controversial – it is one of the points at issue in the debate over Bernard Williams’ (1981) ‘internal reasons’ theory. But the normative claim seems near platitudinous. Reasons are meant to guide us to act, believe, desire, or otherwise respond. But to be guided by reasons just is to engage in reasoning, broadly construed. So it is hard to see how reasons could fail to be appropriate premises for reasoning towards ϕ -ing.

The aim of this article is to develop and defend an account of the nature of reasons that begins with this thought. According to what I shall call the *reasoning view*, reasons are premises of good reasoning: what it is for some

The copyright line for this article was changed on 07 June 2016 after original online publication

Pacific Philosophical Quarterly 98 (2017) 251–270 DOI: 10.1111/papq.12135

© 2015 The Author

Pacific Philosophical Quarterly published by University of Southern California and John Wiley & Sons Ltd. This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

consideration to be a reason to ϕ is for it to be a premise of good reasoning towards ϕ -ing.¹ I explain and motivate this view in Sections 1 and 2. In Section 3, I introduce a problem that views of this sort face accommodating outweighed reasons and argue that the standard solution to this problem is unsuccessful. Sections 4 and 5 develop and defend an alternative solution, which draws on the idea that good patterns of reasoning can be defeasible. Section 6 concludes by drawing out some implications of the resulting view for the debate over pragmatic reasons and for one influential form of reductionism about the normative.

1. *The reasoning view*

The aim of this section is to explain how I shall understand the claim that reasons are premises of good reasoning. I begin with the notion of reasoning.

Reasoning is often, and naturally, understood narrowly – restricted, perhaps, to conscious processes, or ones undertaken intentionally, or which involve some kind of calculation. Under any of these restrictions, I may engage in reasoning when I calculate a tip in a restaurant, or try to figure out how to organise my day, but not when I unreflectively form the intention to turn right upon seeing the sign, or realise with a start that it is my brother's birthday on reading the calendar. However, it is also possible to understand reasoning in a broader way, under which these latter transitions do count as reasoning. In the most general sense, any psychological process by which we come to form, revise, or sustain an attitude for a reason – because of or in light of some consideration – counts as reasoning. In the examples above, I intend to turn right because that is the way home, and believe that today is my brother's birthday because today is the 19th. I therefore count as reasoning in this broader sense. Note also that in this sense, we need not be puzzled by the Aristotelian claim that the conclusion of practical reasoning is an action. It is not controversial that actions can be done for reasons.

If reasoning is understood in this broad way, then to say that reasons are premises of good reasoning is to say that a reason is a suitable thing to have as one's reason – as the consideration in light of which one acts, believes, desires, or otherwise responds. This claim has the air of circularity. But the apparent problem here results from the ambiguity of 'reason'. To disambiguate: the claim of the reasoning view is that a *normative* reason is a suitable thing to have as one's *motivating* reason – as the consideration in light of which one responds. This claim is circular only if the property of being that in light of which one responds must be understood in terms of the property of being a reason *to* respond. But we need not accept this. It is possible, indeed common, to act or believe in light of a consideration which is not a normative reason. Nor need you believe that a consideration is a normative reason in order to respond in light of it.²

Reasoning falls into patterns. Some of these patterns, such as those that constitute affirming the consequent, or the gambler's fallacy, are examples of bad or mistaken reasoning. Others are examples of correct or, as I shall say, good reasoning. We can represent two relatively uncontroversial examples of good reasoning as follows:

Belief that p , Belief that if p then $q \Rightarrow$ Belief that q

Intention to ϕ , Belief that ψ -ing is the only way to $\phi \Rightarrow$ Intention to ψ

We can read these claims as saying that it is good reasoning to move from the states mentioned on the left-hand side of the arrow to the state on the right-hand side. When this is true this move is a *good pattern of reasoning*. The states on the left-hand side of such a claim are the *premise-responses* of that pattern; the state on the right-hand side is the *conclusion-response*. The contents of the beliefs among the premise-responses are the *premises* of that pattern – they are the considerations in light of which an agent acts, believes, or otherwise responds.

On the view to be developed here, it is a necessary condition of a consideration's being a reason for an agent to ϕ that there is a good pattern of reasoning from a set of psychological states, including belief in that consideration, to ϕ -ing. This is not, however, a sufficient condition. To start with, patterns of good reasoning are abstract structures. To give an account of what it is for a consideration to be a reason *for an agent*, we must relate the pattern of reasoning which corresponds to that consideration to that agent's psychology. The simplest way to do this would be to require that the agent has the premise-states of that pattern. However, this approach would commit us to the controversial claim that a consideration can be a reason for an agent only if that agent believes that consideration to obtain. To avoid this implication, we can require instead that the agent has any non-doxastic states among the premise-states of the relevant pattern. On this approach, the second of the patterns above corresponds to a reason for you to ψ only if you intend to ϕ . But it is not required that you believe that ψ -ing is the only way to ϕ (cf. Setiya, 2007, pp. 11–13).

However, even if an agent does have the premise-states of a good pattern of reasoning which concludes in ϕ -ing, it does not follow that there is a reason for that agent to ϕ . To say that there is a good pattern of reasoning from a set of psychological states to some response is to say something about the *transition* from those states to that response. If the starting points are bad, good reasoning need not lead you to do what there is reason to do. The clearest examples here involve reasoning which turns on false belief. Consider Bernard Williams' famous character who decides to mix petrol with tonic, because he believes the petrol to be gin (1981, p. 102). Though this

character reasons impeccably, he decides to drink in light of a consideration which is not a reason to do so. (More precisely, it is not an *objective* reason to do – I say more about this below.) To develop our necessary condition into a sufficient condition, we thus need to add some constraints on the psychological states involved in an agent's reasoning.

The simplest move here would be to require that the beliefs involved in an agent's reasoning are true (Setiya, 2007, p. 12). However, we might think that further constraints are also needed. For example, we might think that good reasoning from intentions to do what you ought not to do, or desires for bad outcomes, can also lead us to do what there is no reason to do, even if such reasoning involves no false beliefs.³ If so, then we should suggest a more general constraint: all of the attitudes involved in the agent's reasoning must be correct. On the assumption that beliefs are correct only if true, this more general constraint subsumes the initial one.

Putting these points together, we can state the view to be defended in this article as follows:

- (R1) For the fact that p to be a reason for S to ϕ is for there to be a good pattern of reasoning from the belief that p , perhaps together with other correct attitudes which S has, to ϕ -ing.

Two further clarifications are necessary. First, as anticipated above, R1 concerns *objective* reasons. Objective reasons are facts that support actions or other responses. R1 does not concern *subjective* reasons – believed considerations which bear on the rationality of some response. While it should be possible to extend the reasoning view to subjective reasons, this is a task for another time.⁴

Second, R1 is intended as a *constitutive account* of reasons – an account of what it is for some consideration to be a normative reason.⁵ As such, it takes the correctness of attitudes and the goodness of reasoning to be prior to reasons. One kind of question this raises is how the correctness of attitudes and the goodness of reasoning are related to each other. There are various options here; this article is neutral on them.⁶ But the view is committed to rejecting the suggestion, due to Schroeder (2007, pp. 133–4), that the correctness of a response is just a matter of there being sufficient reason for that response. While I cannot engage with this view in detail, let me note that it is at least not obvious that correctness should be understood in terms of reasons. For instance, it is highly plausible that beliefs are correct just when true but it is not at all clear that there is sufficient reason to believe just what is true. More generally, the best way to assess the claim that correctness might be prior to reasons is to develop and assess accounts of reasons in terms of correctness. This is the task I begin in this article.⁷

2. *The attractions of the reasoning view*

I take the reasoning view, as developed along these lines, to have three central attractions. The first central attraction is that the view promises a unified account of normative reasons – an account which applies to reasons to act, believe, desire, and so forth. Just as there are good patterns of reasoning which conclude in action, there are good patterns of reasoning which conclude in belief, desire, intention, and many other kinds of response. The reasoning view can thus say that reasons for action, belief, and desire correspond to each of these different forms of reasoning – reasons to act are premises of good reasoning which concludes in action, reasons to believe are premises of good reasoning which concludes in belief, and so on.

A unified theory of this sort is something we should want.⁸ Reasons to act, believe, and desire show many similarities. For instance, in each case, we see a structure in which there are a variety of considerations counting for or against a response, which come together to determine whether you ought to make that response. In each case, reasons, when known about, bear on whether it is rational or reasonable to make a response, and on whether an agent should be praised or blamed for the responses they make. There are further similarities too. For instance, parallel issues arise in each domain. To give two examples: in each domain, we can ask whether your reasons depend only on the facts of your situation or also on your epistemic position, and we can ask about whether reasons have to conform to principles. What's more it is plausible – at least as a default presumption – that we should expect the same answers in each case.⁹

All of these points are naturally explained by the supposition that reasons to act, believe, and desire have something important in common – that considerations which are reasons to act, believe, or desire stand in the *same relation* to different responses. I take the aim of a constitutive account of reasons to be to say what it is that reasons to act, believe, and desire have in common – to reveal the nature of this relation. A constitutive account of reasons should thus allow that the reason-relation can have each of these different responses as one of its relata, and it should have broadly plausible results when applied to these different responses. The reasoning view gives us a way to do this.

At the same time, the reasoning view is compatible with very different *substantive* accounts of reasons. Whereas constitutive accounts tell us about the nature of the reason-relation, substantive accounts tell us what considerations stand in this relation – they tell us which considerations are reasons for action, belief, and other responses. The reasoning view allows for very different substantive accounts of reasons because good reasoning towards action may be very different from good reasoning towards belief, or other responses. For instance, it might be good reasoning to move from the utility of ϕ -ing to ϕ -ing but not good reasoning to move from the utility of believing

p to believing p . If so, the reasoning view will tell us that utility provides reasons for action but not belief. So while the reasoning view offers unity at one level, it is quite compatible with diversity at another level.

The second central attraction of the reasoning view is that it captures with ease the very plausible thought with which we began – that reasons to ϕ must be appropriate premises of reasoning towards ϕ -ing. On the reasoning view, this thought is a simple consequence of the nature of reasons.

It is instructive here to contrast the reasoning view with an alternative. Consider the ‘primitivist’ view of Jonathan Dancy (2004), Derek Parfit (2011a, 2011b), and T.M. Scanlon (1998), on which reasons to ϕ are simply considerations which count in favour of ϕ -ing. The *prima facie* problem for this view is that not all considerations which count in favour of some response seem to be suitable premises for reasoning towards that response. Consider so-called *pragmatic* or ‘wrong kind’ reasons: that believing in God would make you happy, or that intending to drink an unpleasant but otherwise harmless toxin will earn you \$1 m. Since the prospect of happiness clearly counts in favour of believing that God exists, and the \$1 m prize clearly counts in favour of intending to drink the toxin, the primitivist should count these considerations as reasons to believe that God exists and to intend to drink the toxin.¹⁰ But it is equally clear that these considerations are not appropriate premises of reasoning to the belief that God exists or to the intention to drink the toxin. The transitions of thought you would express to yourself by saying ‘believing that God exists would make me happy, so, God exists’ or ‘intending to drink the toxin would make me rich, so, I shall drink the toxin’ are bad pieces of reasoning, if they count as reasoning at all. The primitivist view is thus in *prima facie* tension with the plausible thought that reasons to ϕ must be appropriate premises of reasoning towards ϕ -ing. By contrast, the reasoning view supports and explains this claim.

The third central attraction of the reasoning view is that it offers a simple explanation of a notable *restriction* on reasons. While there are normative reasons for action, belief, and desire, there are no normative reasons for perceptual experiences, or for your hair to be turquoise, or for you to be a little bit taller (cf. Scanlon, 1998, p. 20). More generally, there are normative reasons only for things which can be done or had for reasons. This fact is a straightforward consequence of the reasoning view. For while actions, beliefs, and desires can be conclusions of reasoning, in our broad sense, perceptual experiences, the colour of your hair, and your height, cannot.

Again, it is useful here to contrast the reasoning view with a rival. According to *value-based views*, reasons are considerations which make a response *good* in some way.¹¹ This view offers no simple explanation of the restriction on reasons (cf. Gibbons, 2010, pp. 346–7). For after all, it might well be good to be a little bit taller, or to have turquoise hair, or to have certain perceptual experiences. In order to avoid the implication that there can be reasons for

these things, the value-based view must be restricted in some way. But it is hard to see how to restrict the view in a principled way. It is not enough, for instance, to say that reasons are considerations which make some causally possible response good, since perceptual experiences and the rest can also be causally possible. The value-based view is thus left saying that reasons are considerations which make states or events which can be had or done for reasons good. But this is to stipulate that only things which can be had or done for reasons are subject to reasons, rather than explain why this restriction applies. The reasoning view thus illuminates a fact about normative reasons which the value-based view does not.

3. *The problem of outweighed reasons and two solutions*

The reasoning view, as developed in R1, thus has several important attractions. However, it also faces a serious problem.

A central fact about normative reasons is that they can be outweighed. To say that you have some reason to ϕ does not imply that you have *most* or even *sufficient* reason to ϕ . It implies only that there is something to be said for ϕ -ing – something which may be outweighed, or perhaps defeated in other ways, by competing considerations.

This point is perhaps most widely acknowledged in the literature on reasons to act. Thus to adapt a famous example from Ross (1930, p. 18), the fact that you have promised to meet a friend for lunch may be a reason to meet your friend even if, because you could save a life by breaking your promise, you ought not meet your friend. But the same holds for reasons for belief, desire, and other attitudes as well. For example, if Betty says that Billy spent the whole evening at home with her, then that is some reason to think that he did so. But if others claim to have seen Billy at the scene of the crime, then we might lack sufficient reason to think that Billy spent the night at home. And if some policy will promote equality, then, if egalitarians are right, there is some reason to want this policy to pass. But if the way in which this policy promotes equality is by making everyone as badly off as the worst-off – if it is a ‘levelling down’ policy – then there is not – even egalitarians should admit – sufficient reason to want this policy to pass.

R1 seems not to accommodate these observations. If p is a reason to ϕ , then R1 tells us that it is good reasoning to move from the belief that p , perhaps together with other relevant psychological states, to ϕ -ing. However, when a reason to ϕ is outweighed, it seems *not* to be good reasoning to move from that reason to ϕ -ing, given the stronger conflicting reasons. In the examples above, it would be a mistake to move from the belief that you promised to keeping your promise, or from Betty’s testimony to the belief that Billy stayed home that evening. (Equivalently, these considerations

are not here suitable things to have as your reasons for keeping your promise, or believing that Billy stayed home that evening.) These examples thus suggest is that it is good reasoning to move from belief that p to ϕ -ing only if p is a *sufficient* reason to ϕ . If that is right, then R1 fails as an account of reasons as such.

Broadly speaking, there are two ways for the reasoning view to respond to this problem. The first is to reject the counter-examples. A natural way to do this is to distinguish between good reasoning from certain of your attitudes and good reasoning given some further attitudes. Thus we might say that while it is indeed good reasoning to move from your belief that you promised to keeping your promise, it is not good reasoning to make this move *given* your further belief that by breaking your promise you can save a life. The thought here is that, just as reasons can be outweighed, good patterns of reasoning can be defeated. If we allow for good but *defeasible* patterns of reasoning, R1 can accommodate outweighed reasons.¹²

A second response to the problem is to revise the reasoning view. This is the standard way in which proponents of the reasoning view, and related views, have allowed for outweighed reasons. Consider Bernard Williams' (1981, 1995, 2001) famous 'internal reasons' theory of practical reasons. According to a fairly standard formulation of this view, there is a reason for an agent to ϕ only if there is a 'sound deliberative route' from that agent's desires to a *motivation* to ϕ . This view allows for outweighed reasons because it is possible to be motivated to ϕ while being more strongly motivated to do something else – thus in Ross's case, sound deliberation might lead you to be motivated to keep your promise while also being more strongly motivated to save the life.¹³ Or consider Michael Smith's 'ideal advisor' account of practical reasons. On Smith's view, what we have reason to do is what we would *want* ourselves to do, if we were fully informed and had deliberated correctly (1994, ch. 5). Since it is possible to have conflicting desires about what to do – even given full information and correct deliberation – this view also has the structure to allow for outweighed reasons.¹⁴

It is straightforward to incorporate these ideas into the formulation of the reasoning view presented here. Rather than saying that reasons to ϕ are premises of good reasoning which concludes in ϕ -ing, we can say, following Smith, that reasons to ϕ are premises of good reasoning which concludes in a desire to ϕ or, following Williams, a motivation to ϕ . More generally, the *standard solution* to the problem of outweighed reasons takes reasons to ϕ to correspond to good reasoning which concludes in something weaker than, but importantly related to, ϕ -ing. Versions of the standard solution vary in what they take this weaker thing to be.

Despite its prominence, I doubt that the standard solution succeeds. In the next sub-sections, I consider in more detail the versions of it just sketched and argue that neither retains all of the central attractions of the reasoning

view. While other versions of the standard solution are possible, I take these problems to motivate consideration of the first kind of solution to the problem of outweighed reasons. In Sections 4 and 5, I argue that this solution is viable and, since it allows us to retain the reasoning view's central attractions, should be preferred to the standard solution.

3.1. THE WEAKER CONCLUSION AS DESIRE

The first version of the standard solution holds that reasons to ϕ are premises of good reasoning which concludes in a desire to ϕ :

- (R2) For the fact that p to be a reason for S to ϕ is for there to be a good pattern of reasoning from the belief that p , perhaps together with other correct attitudes which S has, to a desire to ϕ .

Talk of desires can be understood broadly or narrowly. On the narrow construal, all desires are desires to act. Putative desires for the weather to improve, for there to be life elsewhere in the universe, or for a certain number to be prime, are better understood as wishes or hopes. On the broad construal, no such restriction is in place – any sort of thing can be the object of desire. I shall argue that on either construal, proponents of the reasoning view should be reluctant to adopt R2.

The problem with the narrow construal is clear. With desires understood in this way, R2 offers only an account of reasons to act, rather than a unified account of reasons. It thus lacks the first central attraction of the reasoning view.

The broad construal faces three problems. First, when understood in this way, R2 lacks the second central attraction of the reasoning view: it does not capture the plausible idea that reasons to ϕ must be appropriate premises of reasoning towards ϕ -ing. R2 tells us that reasons to ϕ are appropriate premises of reasoning which concludes in a desire to ϕ . But to reason to a desire for a belief, or a desire for an intention, is not necessarily to reason towards that belief or intention. It is not good reasoning to move from a desire for a belief to that belief.

Second, R2 also lacks the third central attraction of the reasoning view – it offers no straightforward explanation of why only things which can be done for reasons are subject to reasons. Since it can be good reasoning to move to a desire (in the broad sense) for a perceptual experience, or for turquoise hair, or to be a little bit taller, the view allows for reasons for such states. To avoid this consequence, R2 must appeal only to desires for things that can be had or done for reasons. However, if it is merely asserted that it is only such desires that are relevant, then the restriction on reasons is stipulated, rather than explained. R2 is thus in much the same position as the value-based view, when it comes to this restriction.¹⁵

Third, it is doubtful that R2 solves the problem for R1. R2 solves the problem of outweighed reasons only if whenever there is an outweighed reason for S to ϕ , it is good reasoning for S to move to a desire to ϕ . This claim is somewhat plausible in the case of reasons for action.¹⁶ However, it is far less plausible in the case of reasons for belief, desire, and other attitudes. Consider our earlier example, in which there is an outweighed reason to believe that Billy stayed home that evening, since Betty testified that he did but others claim that he did not. It is not plausible that an agent who is reasoning well will come to want to believe that Billy stayed home. Since the belief that Billy stayed home is not worth wanting to have – it would not, for instance, constitute knowledge of an interesting truth, and it need not be in any way useful – good reasoning would not lead you to want to believe this.¹⁷ So, R2, when construed broadly, fails to solve the problem of outweighed reasons.

R2 thus fails whether we construe desires broadly or narrowly. The first version of the standard solution is unsatisfactory.

3.2. THE WEAKER CONCLUSION AS MOTIVATION

The second version of the standard solution holds that reasons are premises of good reasoning which concludes in motivation:

- (R3) For the fact that p to be a reason for S to ϕ is for there to be a good pattern of reasoning from the belief that p , perhaps together with other correct attitudes which S has, to a motivation to ϕ .

This view retains two of the central attractions of the reasoning view. It allows that reasons to ϕ are appropriate premises of reasoning towards ϕ -ing, since to be motivated to ϕ is to move towards ϕ -ing. And since perceptual experiences and the colour of your hair cannot be the objects of motivation, the view explains why there can be no reasons for perceptual experiences or the colour of your hair. On the face of it though, the view does not retain the second attraction of the reasoning view. Given the natural assumption that only *actions* can be the objects of motivation, the view offers no account of reasons for belief, desire, and other attitudes.

This may be too quick. There is such a thing as being inclined to believe – when you are deliberating about what to believe, you might feel yourself pulled now one way, now the other. We might take such inclinations to be doxastic motivations. Alternatively, we might follow Setiya's (2014, p. 234) suggestion that reasons to believe are premises of good reasoning which concludes in increased confidence.¹⁸ Insofar as we take degrees of confidence to stand to belief as degrees of motivation stand to action, this may seem a sufficiently unified account.¹⁹

However, neither of these suggestions seem satisfactory. Consider a case in which you learn that p , which is a reason to believe q , but in which you are already warranted in being certain that q , or that q is false. In such a case, good reasoning does not require you to increase your confidence in q , or to feel some inclination to believe q – either because you are already fully confident that q , or because it is plain that the newly acquired reason to believe q is conclusively outweighed. Consider also cases in which you learn that p , which is a reason to believe q , but at the same time acquire stronger, or equally weighty, reasons not to believe q . Again, good reasoning does not require you to increase your confidence in q , or to feel an inclination to believe q .

Replies to these objections are possible. In the first sort of case, it might be suggested that p cannot be a reason to believe q if you are warranted in being certain that not- q . But while this has some plausibility, it cannot be applied to the opposite case – we can clearly gain further reasons for claims we are already right to be certain in. In the second sort of case, it might be said that although learning that p does not require you to increase your confidence that q , given the contrary reasons, it would require you to increase your confidence that q in other circumstances. However, on pain of circularity, defending this response would require specifying these other circumstances without reference to reasons. It is not clear that this could be done. And in any case, the well-known difficulties with explaining reasons in terms of the responses we ought to make, or would be rational to make, in other circumstances, should make us wary of appealing to such counterfactuals in an account of reasons (cf. Dancy, 2004, ch. 2; Smith, 1994, ch. 5; Lord and Maguire, forthcoming).²⁰

Thus, I take it that both of the suggestions for how to develop R3 as an account of reasons to believe face serious problems. This confirms the initial impression that R3 is best suited to be an account of reasons to act, rather than a unified account of reasons.

So neither of the familiar versions of the standard solution solves the problem of outweighed reasons without undermining some of the central attractions of the view. At best then, these approaches come with significant costs. Of course, there are other, less familiar ways in which the standard solution might be developed.²¹ Nonetheless, the problems with R2 and R3 should also lead us to consider alternative ways to solve the problem of outweighed reasons. Above, I noted that there are prospects for arguing that cases of outweighed reasons are not in fact counter-examples to the version of the reasoning view we began with, R1. In the next two sections, I develop and defend this response. If this response succeeds, exploration of further versions of the standard solution will be unnecessary.

4. *Defeasible patterns of reasoning*

In Ross's case, it seems not to be good reasoning to move from the belief that you promised to meet your friend to doing so, when you believe that you will thereby leave someone to die. Nor does it seem good reasoning to move from Betty's testimony to the belief that Billy stayed home that evening, when you know that others dispute this. If you reasoned in one of these ways, you would be making a mistake.

These claims seem to be counter-examples to R1. R1 tells us that when p is a reason to ϕ , there is a good pattern of reasoning from the belief that p (perhaps with other attitudes) to ϕ -ing. There is a good pattern of reasoning from the belief that p to ϕ -ing, I said, when it is good reasoning to move from the belief that p to ϕ -ing. So R1 implies that when p is an outweighed reason to ϕ , it is good reasoning to move from the belief that p to ϕ -ing. The cases seem to show that this is not so.

However, we need not accept this. What is clear is that in these cases you make a mistake in your reasoning. But we need not take the claim that it is good reasoning to move from the belief that p to ϕ -ing to imply that if you reason from the belief that p to ϕ -ing, you make no mistake in your reasoning. We can distinguish between the claim that it is good reasoning to move from the belief that p to ϕ -ing and the claim that it is good reasoning to move from the belief that p to ϕ -ing given further attitudes you might have. Given this distinction, the fact that it is good reasoning to move from the belief that p to ϕ -ing does not ensure that if you move from the belief that p to ϕ -ing, you make no mistake. For it might not be good reasoning to make this move given further attitudes you have.

The examples above can be taken to illustrate this distinction: although it is good reasoning to move from the belief that you promised to meet your friend to doing so, it is not good reasoning to make this move given the belief that you will thereby leave someone to die. But there are many other examples that also suggest that we need a distinction of this sort. The most familiar cases are induction and abduction. It is plausibly good reasoning to move from the belief that all observed Fs are Gs to the belief that the next F will be G. But it is not good reasoning to make this move given the further belief that conditions for observing Fs have been unrepresentative. Or consider the move from believing that the match was struck to believing that the match lit, or from believing that Tweety is a bird to believing that Tweety flies. These seem like good patterns of reasoning, corresponding to the defeasible generalisations that matches light when struck, and that birds fly. But it is not good reasoning to make these moves given the further beliefs that Tweety is a penguin (in the former case) or that the match was wet (in the latter).²² Practical reasoning affords further examples still. For instance, although it is plausibly good reasoning to move from an intention for an

end to an intention for what you take to be the best means to that end, it is not good reasoning to make this move when you know that taking this means will prevent you achieving other of your ends.

It might be objected that if it is not good reasoning to move from some premise-responses to a conclusion-response given some further attitudes then, if you have those further attitudes, we can infer that it is not good reasoning to move from those premise-responses to the conclusion-response. How could this inference fail? In response, I suggest that claims about good reasoning, and thus good patterns of reasoning, need not only concern transitions between premise- and conclusion-responses. Further relevant attitudes can also fall within the 'scope' of claims about good reasoning – we can make claims of the form: it is not good reasoning to [move from premise-responses P to conclusion-response C given further attitudes A]. It does not follow from a claim of this form that if you have the further attitudes A, then it is not good reasoning to move from P to C. And it is important that we allow claims about good reasoning to take this form. When assessing someone's reasoning we are not only interested in the considerations in light of which they reached their conclusion; we are also interested in the considerations which they may have ignored when doing so.

The distinction between good reasoning from some of your attitudes and good reasoning given further attitudes allows R1 to solve the problem of outweighed reasons. Say that the move from some premise-responses to a conclusion-response is a good but *defeasible* pattern of reasoning when it is good reasoning to move from those premise-responses to that conclusion-response but not good reasoning to make this move given some further possible attitudes. Once we allow for such patterns of reasoning, R1 allows for outweighed reasons. R1 implies that when p is a reason to ϕ , there is a good pattern of reasoning from the belief that p (perhaps with other attitudes) to ϕ -ing. But since patterns of reasoning can be defeasible, it does not follow that it is good reasoning to move from the belief that p to ϕ -ing given other attitudes you may have. The observations which give rise to the problem of outweighed reasons thus do not undermine R1. It doesn't follow from the observation that, given certain other attitudes, it may not be good reasoning to move from the belief that you promised to keeping your promise, or from someone's testimony to accepting that testimony, that the patterns of reasoning here are not good patterns of reasoning. All that follows is that these patterns are defeasible.

5. *The weight of reasons*

R1 allows for outweighed reasons while straightforwardly retaining the central attractions of the reasoning view. R1 thus has clear advantages over

versions of the reasoning view which adopt the standard solution. However, proponents of the standard solution might suggest that their views also have an important advantage over R1.

The standard solution allows for an elegant account of the *weights* of reasons. If reasons correspond to good reasoning which concludes in desire or motivation then, roughly, one reason is stronger than another if the reasoning corresponding to the first concludes in a *stronger* desire or motive than the reasoning corresponding to the second (cf. Smith, 1997, p. 267; Setiya, 2007, p. 13). More generally, if reasons to ϕ are premises of good reasoning which concludes in something which comes in varying strengths, then the strength of a reason to ϕ can be understood in terms of the strength of the response appropriate to ϕ -ing. This simple account of the weight of reasons is a clear attraction of the standard solution.²³

However, the standard solution is not the only way to give an account of the weight of reasons. The notion of a defeasible pattern of reasoning suggests a natural alternative. Schematically, we can say that where p is a reason for S to ϕ and q is a reason for S to ψ , p outweighs q if the pattern of reasoning corresponding to p as a reason to ϕ defeats the pattern of reasoning corresponding to q as a reason to ψ .

What is it for one pattern of reasoning to defeat another? Roughly, I suggest that one pattern of reasoning defeats another if it is good reasoning to move from the premise-states of the one to its conclusion, given the premise-states of the other, but not vice-versa. Slightly more precisely, suppose R_1 is a good pattern of reasoning which concludes in ϕ -ing and R_2 is a good pattern of reasoning which concludes in a conflicting conclusion ψ . For R_1 to defeat R_2 is for it to be good reasoning to move from the premise-responses of R_1 to the conclusion-response of R_1 , given the premise-responses of R_2 , but not good reasoning to move from the premise-responses of R_2 to the conclusion-response of R_2 , given the premise-responses of R_1 .²⁴ For example, in Ross's case, it is not good reasoning to move from the belief that you promised to meet your friend to doing so, given the belief that you will thereby leave someone to die. But it is good reasoning to move from the belief that by staying here you can save a life to staying here, even given the belief that by doing so you will break a promise. This is what it is for the pattern of reasoning which concludes in saving the life to defeat the pattern concluding in keeping your promise, and so for the reason to save the life to outweigh the reason to keep the promise.

This account of what it is for a reason to be outweighed requires further elaboration and development.²⁵ Nonetheless, I take it to be enough to indicate that the standard solution is not the only way for the reasoning view to provide an account of the weights of reasons – the solution to the problem of outweighed reasons that I have recommended naturally suggests an alternative.

6. *Conclusion and implications*

The reasoning view, as developed in R1, offers an attractive account of the nature of reasons. R1 provides a unified account of reasons, vindicates the compelling idea that reasons to ϕ must be appropriate premises for reasoning towards ϕ -ing, and offers a straightforward explanation of why only things which can be done for reasons are subject to reasons. And, as I have now argued, R1 can accommodate outweighed reasons, and suggests a natural view of the weights of reasons.

Of course, much more would need to be done for a full defence of R1. In addition to the issues about the weights of reasons just noted, more would need to be said about the notions of reasoning and correctness which are at the heart of the view, and there are further objections to the view to consider, and alternatives to compare it to.²⁶ Rather than pursue these issues here, I want to conclude by drawing out two ways in which the view bears on further issues in normative theory.

PRAGMATIC REASONS FOR ATTITUDES

Suppose that believing in God would make you happy or that intending to drink an unpleasant but otherwise harmless toxin would win you \$1 m. As we noted in Section 2, some accounts of reasons support the claim that these considerations are reasons to believe in God, and to intend to drink the toxin. And for some philosophers, it is clear that this is the right verdict. But for others, it is a mistake. Considerations of this sort are not, properly speaking, reasons to believe in God, or to intend to drink the toxin. Rather, they are only reasons to *want* or to *bring about* these attitudes.²⁷ R1 supports an argument for this latter view. A reason to believe in God must be a premise of good reasoning which concludes in the belief that God exists. And while the fact that believing in God would make you happy is a premise of good reasoning which concludes in *wanting* to believe in God and in *bringing about* this attitude, it is not a premise of good reasoning which concludes in belief in God. The reasoning you would express to yourself by saying 'Believing in God would make me happy, so God exists' is plainly bad reasoning. So R1 supports the claim that pragmatic considerations of this sort are not, properly speaking, reasons to believe or intend. Rather, they are reasons to want or bring it about that you believe or intend.²⁸

IDEAL AGENT ACCOUNTS OF REASONS

Something like the reasoning view is often taken to support the idea that normative reasons can be reduced to the responses of an idealised agent (e.g. Smith, 1994, ch. 5). The natural thought here is that if reasons are premises of good reasoning, we will be able to reduce reasons to the

responses of an agent who is reasoning well. However, this thought rests on the assumption that agents who reason well will have some response corresponding to each of their reasons. And while versions of the reasoning view which adopt the standard solution support this assumption – for instance, on Smith's view, agents who are reasoning well have some desire corresponding to each of their reasons – R1 does not.²⁹ If reasons correspond to defeasible patterns of reasoning, rather than requirements to respond in some way, then we should not expect our idealised selves to have some response corresponding to every reason. The view defended here thus undermines a crucial assumption of one influential form of reductionism about reasons.³⁰

Philosophy, Faculty of Humanities
University of Southampton

NOTES

¹ For other authors who have been attracted to this sort of approach, see, for example, Gibbons (2010), Hieronymi (2005), Raz (1978), Setiya (2007, 2014), and, perhaps, Williams (1981).

² While these claims would be accepted by many philosophers, they are nonetheless controversial. As it would take us too far afield to defend them, they must here be treated as assumptions. I plan to discuss these issues elsewhere.

³ For the claim that we lack reason to do what would satisfy desires that we ought not to have, see, e.g. Raz, 1999, and Dancy, 2000.

⁴ For more on the distinction between objective and subjective reasons, see, e.g. Parfit, 2011a; Schroeder, 2008; Sylvan, 2015; Vogelstein, 2012; Way, 2012; Whiting, 2014.

⁵ For discussion of the notion of a constitutive account see, e.g. Fine, 1994; Rosen, 2010, pp. 122–6; Schroeder, 2007, pp. 61–72; and Wedgwood, 2007, pp. 136–44.

⁶ For instance, one might try to understand the goodness of reasoning in terms of the correctness of attitudes. See McHugh and Way, forthcoming, for exploration of this option.

⁷ For others who take correctness to be prior to reasons, see Chappell, 2012; Danielsson and Olson, 2007; and Thomson, 2008. For the claim that correctness is the only basic normative property, see Brentano, 1889/2009; McHugh and Way, forthcoming.

⁸ Cf. Gibbons, 2010; Kearns and Star, 2009, esp. pp. 219–22; Schroeder, 2007, p. 113, n.16, 2008, p. 70; Finlay, 2006.

⁹ I say more about this with respect to the question about principles in Way, 2013a, and with respect to the question about epistemic positions in Way, 2013b.

¹⁰ See Hieronymi, 2005, pp. 442–3, for this point and discussion of such examples in a related context. The second example is, of course, due to Kavka, 1983.

¹¹ Value-based views are most familiar as views of reasons to act. See, for example, Raz (1999, p. 23): 'reasons are facts in virtue of which...actions would be good in some respect, and to some degree'. For a more general defence see Finlay, 2006. Note that value-based views are also in tension with our starting idea that reasons must be appropriate premises of reasoning, since 'wrong kind' considerations clearly make attitudes good.

¹² A variant on this idea is to say that reasons are premises of reasoning which is good *to some extent*. On this view, when a reason is outweighed, the corresponding reasoning is not good *overall*. I am not sure that this suggestion differs substantively from that in the text. The

discussion in §4 and §5 below could be taken as a way of spelling out the difference between *pro tanto* and overall good reasoning.

¹³ For formulations along these lines, see Korsgaard, 1986, p. 19; Millgram, 1996, p. 198; and Parfit, 2011b, p. 269. On other formulations of Williams' view, the conclusion of the relevant reasoning is a desire to ϕ (Johnson, 1999, p. 53; Setiya, 2011, p. 4) or the judgment that there is reason to ϕ (McDowell, 1995, p. 97). Williams' own formulations vary. Sometimes they seem to rule out conflicting reasons: e.g. 'A has a reason to ϕ only if there is a sound deliberative route from A's subjective motivational set...to A's ϕ -ing' (2001, p. 91; cf. 1995, p. 35). However, immediately following this, Williams notes that 'it is natural to [this] condition as implying not just that A has a reason to ϕ , but that he or she has more reason to do that than to do anything else' (2001, p. 91). It is plausible that this is why he sometimes talks instead of sound deliberative routes to a motivation to ϕ (e.g. 1981, p. 110; and note that Williams observes the possibility of conflicting reasons at 1981, p. 104). For a formulation of the reasoning view in this way which is explicitly intended to allow for outweighed reasons, see Setiya, 2014, p. 222.

¹⁴ While he does not emphasise the point, it is clear that Smith develop his view in this way in part in order to allow for conflicting reasons. See for example his remark that 'The analysis of normative reasons I propose tells us...what we have *pro tanto* normative reason to do, not what we have all things considered normative reason to do. This is because...with regard to a particular set of circumstances our fully rational selves could have several conflicting desires about what is to be done' (Smith, 1996, p. 167; cf. 1997, pp. 266–7).

¹⁵ This is not surprising. R2 is very similar to Smith's (1994, ch. 5) account of reasons, and Smith's account of reasons is a version of a value-based theory. For Smith, what we have reason to do is what it is desirable to do, and what it is desirable to do is what our fully informed and rational selves would want our actual selves to do.

¹⁶ Although see Dancy, 2004, pp. 21–2 for doubts.

¹⁷ Parallel claims apply, I suggest, in the case of outweighed reasons to desire.

¹⁸ This is to simplify slightly. Setiya proposes that when p is a reason to believe that q , it is good reasoning to become more confident that q in the light of p and some of your other attitudes than in the light of those attitudes alone. Although Setiya does not elaborate, he might deny that it follows from this that when p is a reason to believe that q , it is good reasoning to move from the belief that p to increased confidence that q . If so, he may avoid the problems noted below. However, it is hard to see how this implication could be denied without distinguishing, along the lines I suggest in Section 4, between good reasoning from certain attitudes and good reasoning given further attitudes. But as I argue below, making this distinction is enough to escape the problem of outweighed reasons. The standard solution would thus be redundant.

¹⁹ Setiya's suggestion could perhaps be generalized to reasons for desire and other attitudes which come in degrees. It is less clear how the first suggestion could be generalized. It is not clear that there is any such thing as being inclined to desire something (as opposed to wanting to desire something, or desiring something to some extent).

²⁰ A different reply to the second case is to suggest that good reasoning permits but does not require you to feel some inclination to believe q . As I note below though (n. 23), proponents of the standard solution must take reasons to correspond to what good reasoning requires, rather than what it merely permits, if they are to give an account of the weights of reasons.

²¹ For instance, it might suggested that reasons to ϕ are premises of good reasoning which concludes in a *disposition* to ϕ , or in *taking, treating, or placing weight on* a consideration as a reason to ϕ . (For the latter notions, see, e.g. Scanlon, 1998; Schlosser, 2012; Bratman, 1996; Gibbard, 1990; pp. 160–4; Schroeder, 2007, ch. 7).

²² See Horty, 2012, from whom I take the standard example of Tweety, and Pollock, 1974, 1987, for extended discussion of this kind of reasoning. I hope to consider the relationship

between the notion of a defeasible pattern of reasoning and Horty's notion of a default rule elsewhere.

²³ Note that this account requires the assumption that if it is good reasoning to move from A to B then, at least insofar as you are reasoning at all, you will be reasoning badly if you do not move from A to B. Otherwise, the account of the weights of reasons above would allow each of two reasons to be stronger than the other (cf. Setiya, 2007, p. 13; 2014, p. 229).

²⁴ This allows for cases in which neither R_1 nor R_2 defeats the other. See Setiya, 2014, pp. 229–31 for discussion of some important implications of this point.

²⁵ One important question is whether R1 can accommodate other ways in which it is often claimed that reasons can interact, besides being outweighed – for instance, being *attenuated* or *disabled* (cf. Dancy, 2004; Horty, 2012; Pollock, 1974, 1987; Schroeder, 2011). A reason to ϕ is attenuated if it is made weaker by a consideration which is not a reason not to ϕ . It is relatively straightforward for R1 to allow for this possibility. For even if p is a reason to ϕ , it may not be good reasoning to move from the belief that p to ϕ -ing given other attitudes, even if those other attitudes do not correspond to reasons not to ϕ . For instance, in Lehrer and Paxson's famous example (1969, p. 228; cf. Schroeder, 2007, p. 93), it is not good reasoning to move from the belief that the person you just saw steal a book looks like Tom Grabbit to the belief that Tom Grabbit stole the book, given the further belief that Tom Grabbit has an identical twin. But it does not follow from this that the fact that Tom Grabbit has an identical twin is a reason not to believe that he stole the book. Disabling is more challenging. A reason to ϕ is disabled when a consideration which would otherwise be a reason to ϕ is not, in this case, a reason to ϕ . R1 may seem to rule this possibility out. If it is good reasoning to move from the belief that p to ϕ -ing, then if p , the fact that p is a reason to ϕ – no further facts can 'disable' this reason. However, R1 can allow for some forms of disabling (cf. Setiya, 2014, pp. 226–8). For instance, where a good pattern of reasoning includes more than one premise, the falsity of one premise will count as a disabler of the truths of the other premises. I grant though that R1 may not accommodate all putative cases of disabling. However, while I will have to leave discussion of this issue for another occasion, it is worth noting that this is a general problem for the reasoning view – it does not reveal any advantage the standard solution has over R1.

²⁶ As well as the primitivist and value-based views noted earlier, there is also the view that reasons are evidence, or explanations, of what you ought to do (see, respectively, Kearns and Star, 2009; Broome, 2004) or of correctness (Thomson, 2008; Chappell, 2012).

²⁷ For the first reaction see, e.g. Rabinowicz and Rønnow-Rasmussen, 2004; Danielsson and Olson, 2007. For the second see, e.g. Gibbard, 1990; Parfit, 2011a; Skorupski, 2010.

²⁸ The argument here draws a conclusion about reasons from a premise about good reasoning. It should be stressed that the reasoning view is not committed to the claim that the best or only way to discover what we have reason to do is by considering good reasoning (cf. Setiya, 2007, p. 117). Nonetheless, we may *sometimes* have clearer intuitions about good reasoning than about reasons. I take this to be an example of such a case.

²⁹ We should also note that one of the arguments in Section 3.1 counts directly against this assumption. If it is not always worth wanting to have attitudes for which there is some reason, then we should not expect our ideal selves to want to have such attitudes.

³⁰ I am grateful to Alex Gregory, Benjamin Kiesewetter, Conor McHugh, Debbie Roberts, Simon Robertson, Kieran Setiya, Justin Snedegar, Daniel Star, Daniel Whiting, and several anonymous referees for written comments or discussion. The article also benefited from discussion with audiences at the Practical Reason and Metaethics Conference at the University of Nebraska, Lincoln; the British Society for Ethical Theory Conference at the University of Stirling; Boston University; the University of Leeds; Ohio State University; and the University of Southampton. Work on this article was supported by the Arts and Humanities Research Council [grant number AH/K008188/1].

REFERENCES

- Bratman, M. (1996). 'Identification, Decision, and Treating as a Reason', in his *Faces of Intention*. Cambridge: Cambridge University Press, pp. 185–206.
- Brentano, F. (1889 [2009]). *The Origin of Our Knowledge of Right and Wrong*. Oxford: Routledge.
- Broome, J. (2004). "Reasons," in S. Scheffler, M. Smith, P. Pettit and R. J. Wallace (eds) *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*. Oxford: Oxford University Press, pp. 28–55.
- Chappell, R. (2012). "Fittingness: The Sole Normative Primitive," *Philosophical Quarterly* 62(249), pp. 684–704.
- Dancy, J. (2000). *Practical Reality*. Oxford: Oxford University Press.
- Dancy, J. (2004). *Ethics Without Principles*. Oxford: Oxford University Press.
- Danielsson, S. and Olson, J. (2007). "Brentano and the Buck-Passers," *Mind* 116(463), pp. 511–522.
- Fine, K. (1994). "Essence and Modality," *Philosophical Perspectives* 8, pp. 1–16.
- Finlay, S. (2006). "The Reasons That Matter," *Australasian Journal of Philosophy* 84(1), pp. 1–20.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings*. Oxford: Clarendon Press.
- Gibbons, J. (2010). "Things That Make Things Reasonable," *Philosophy and Phenomenological Research* 81(2), pp. 335–61.
- Hieronymi, P. (2005). "The Wrong Kind of Reason," *Journal of Philosophy* 102(9), pp. 437–57.
- Horty, J. (2012). *Reasons as Defaults*. Oxford: Oxford University Press.
- Johnson, R. (1999). "Internal Reasons and the Conditional Fallacy," *The Philosophical Quarterly* 49(194), pp. 53–71.
- Kavka, G. (1983). "The Toxin Puzzle," *Analysis* 43(1), pp. 33–36.
- Kearns, S. and Star, D. (2009). "Reasons as Evidence," in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics*, Volume 4. Oxford: Oxford University Press, pp. 215–42.
- Korsgaard, C. (1986). "Skepticism about Practical Reason," *Journal of Philosophy* 83(1), pp. 5–25.
- Lehrer, K. and Paxson, T. (1969). "Knowledge: Undefeated Justified True Belief," *Journal of Philosophy* 66(8), pp. 225–37.
- Lord, E. and Maguire, B. (Forthcoming). "An Opinionated Guide To The Weight Of Reasons," in E. Lord and B. Maguire (eds) *Weighing Reasons*. Oxford: Oxford University Press.
- McDowell, J. (1995). 'Might There Be External Reasons?', in his *Mind, Value, and Reality*. Cambridge, MA: Harvard University Press, pp. 95–111.
- McHugh, C. and Way, J. (Forthcoming). 'Fittingness First', *Ethics*.
- Millgram, E. (1996). "Williams' Argument Against External Reasons," *Noûs* 30(2), pp. 197–220.
- Parfit, D. (2011a). *On What Matters*, Volume 1. Oxford: Oxford University Press.
- Parfit, D. (2011b). *On What Matters*, Volume 2. Oxford: Oxford University Press.
- Pollock, J. (1974). *Knowledge and Justification*. Princeton, NJ: Princeton University Press.
- Pollock, J. (1987). "Defeasible Reasoning," *Cognitive Science* 11(4), pp. 481–515.
- Rabinowicz, W. and Rønnow-Rasmussen, T. (2004). "The Strike of the Demon: On Fitting Pro-Attitudes and Value," *Ethics* 114(3), pp. 391–423.
- Raz, J. (1978). "Introduction," in J. Raz (ed.) *Practical Reasoning*. Oxford: Oxford University Press, pp. 1–17.
- Raz, J. (1999). *Engaging Reason*. Oxford: Oxford University Press.
- Rosen, G. (2010). "Metaphysical Dependence: Grounding and Reduction," in B. Hale and A. Hoffmann (eds) *Modality: Metaphysics, Logic, and Epistemology*. Oxford: Oxford University Press, pp. 109–36.
- Ross, W. D. (1930 [2002]). *The Right and the Good*. Oxford: Clarendon Press.

- Scanlon, T. M. (1998). *What We Owe To Each Other*. Cambridge, MA: Harvard University Press.
- Schlosser, M. (2012). "Taking Something as a Reason for Action," *Philosophical Papers* 41(2), pp. 267–304.
- Schroeder, M. (2007). *Slaves of the Passions*. Oxford: Oxford University Press.
- Schroeder, M. (2008). "Having Reasons," *Philosophical Studies* 139(1), pp. 57–71.
- Schroeder, M. (2011). "Holism, Weight, and Undercutting," *Noûs* 45(2), pp. 328–44.
- Setiya, K. (2007). *Reasons without Rationalism*. Princeton, NJ: Princeton University Press.
- Setiya, K. (2011). "Internal Reasons," in K. Setiya and H. Paakkunainen (eds) *Internal Reasons*. Cambridge, MA: MIT Press.
- Setiya, K. (2014). "What is a Reason to Act?," *Philosophical Studies* 167(2), pp. 221–235.
- Searle, J. (2001). *Rationality in Action*. Cambridge, MA: MIT Press.
- Skorupski, J. (2010). *The Domain of Reasons*. Oxford: Oxford University Press.
- Smith, M. (1994). *The Moral Problem*. Oxford: Blackwell.
- Smith, M. (1996). "Normative Reasons and Full Rationality: A Reply to Swanton," *Analysis* 56(3), pp. 160–8.
- Smith, M. (1997). *Ethics and the A Priori*. Cambridge: Cambridge University Press.
- Sylvan, K. (2015). "What Apparent Reasons Appear To Be," *Philosophical Studies* 172, pp. 587–606.
- Thomson, J. (2008). *Normativity*. Chicago, IL: Carus Publishing.
- Vogelstein, E. (2012). "Subjective Reasons," *Ethical Theory and Moral Practice* 15(2), pp. 239–57.
- Way, J. (2012). "Explaining the Instrumental Principle," *Australasian Journal of Philosophy* 90(3), pp. 487–506.
- Way, J. (2013a). "Value and Reasons to Favour," in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics*. Oxford: Oxford University Press, pp. 27–49.
- Way, J. (2013b). "Intentions, Akrasia, and Mere Permissibility," *Organon F* 20(4), pp. 588–611.
- Wedgwood, R. (2007). *The Nature of Normativity*. Oxford: Oxford University Press.
- Whiting, D. (2014). "Keep Things In Perspective: Reasons, Rationality, and the A Priori," *Journal of Ethics and Social Philosophy* 8, pp. 1–22.
- Williams, B. (1981). 'Internal and External Reasons', in his *Moral Luck*. Cambridge: Cambridge University Press, pp. 101–13.
- Williams, B. (1995). 'Internal Reasons and the Obscurity of Blame,' in his *Making Sense of Humanity*. Cambridge: Cambridge University Press, pp. 35–45.
- Williams, B. (2001). "Some Further Notes on Internal and External Reasons," in E. Millgram (ed.) *Varieties of Practical Reasoning*. Cambridge, MA: MIT Press, pp. 91–97.