

What is Good Reasoning?

CONOR MCHUGH AND JONATHAN WAY

University of Southampton

What makes the difference between good and bad reasoning? In this paper we defend a novel account of good reasoning—both theoretical and practical—according to which it preserves *fittingness* or *correctness*: good reasoning is reasoning which is such as to take you from fitting attitudes to further fitting attitudes, other things equal. This account, we argue, is preferable to two others that feature in the recent literature. The first, which has been made prominent by John Broome, holds that the standards of good reasoning derive from rational requirements. The second holds that these standards derive from reasons. We argue that these accounts face serious difficulties in correctly distinguishing good from bad reasoning, and in explaining what's worthwhile about good reasoning. We then propose our alternative account and argue that it performs better on these counts. In the final section, we develop certain elements of the account in response to some possible objections.

Suppose that you are wondering whether the match will go ahead today. You look out the window, see that it's sunny, and reason as follows:

It's sunny

If it's sunny, the match will go ahead

So, the match will go ahead.

Here you begin with two beliefs, and reason to a further belief. The argument expresses your reasoning. The first two sentences express beliefs that you reason from, the third sentence expresses the belief you reason to.

This is a paradigmatic example of theoretical reasoning—reasoning towards belief. There is also practical reasoning—reasoning towards intention or, perhaps, action. For instance, you might reason:

I shall go to the match

In order to go to the match, I must catch the 37 bus

So, I shall catch the 37 bus.

Here you begin with an intention and a belief and reason to a further intention. Again, the argument expresses the reasoning—the first two sentences express the intention and belief you reason from, the third sentence expresses the intention you reason to.

The reasoning in both these cases is good. But not all reasoning is good. For instance, the reasoning you might express by saying ‘the coin came up heads last time, so, it will come up tails next time’ is bad reasoning.

What makes the difference between good and bad reasoning? This is an important question. For instance, the quality of one’s reasoning makes a difference to whether one’s beliefs and intentions are well-grounded, and thereby to further important statuses, such as whether one’s beliefs count as knowledge, or whether one’s intentions deserve credit. Good reasoning is plausibly also connected in important ways to reasons (Setiya 2014, Way forthcoming 1), and thus to what we ought to think and do, and whatever else reasons bear on. The answer to the question of what makes for good reasoning will thus have implications for these further central topics in ethics and epistemology.

Why suppose there is a unified answer to the question of what makes for good reasoning—an answer that applies to both practical and theoretical reasoning? Well, despite their differences, practical and theoretical reasoning are two species of the same kind of thing: reasoning. It would be surprising if there were nothing in common, at any level, between what makes for reasoning’s being good in the two cases. This seems enough to defeasibly justify the search for a unified account. If a promising one can be found, that will provide further evidence for the unity of good reasoning.

In the recent literature, there have been two main suggestions as to what makes for good reasoning.¹ The first, which has been made prominent by John Broome, holds that the standards of good reasoning derive from rational requirements. As a rough first pass, good reasoning is reasoning which leads you to be rational. The second, suggested in passing by several authors but rarely developed in detail, holds that the standards of good reasoning derive from reasons. As a rough first pass, good reasoning is reasoning which leads you to do what there is reason to do. The aim of this paper is to argue against these accounts and propose an alternative. As a rough first pass, we suggest that good reasoning preserves *fittingness* or *correctness*: good reasoning is reasoning which is such as to take you from fitting attitudes to further fitting attitudes.

1. Preliminaries

In this section we make some preliminary points which will serve both to further clarify our question and to establish some desiderata for answers to it. We begin with some further points about reasoning.

Reasoning, we have seen, is a kind of transition between attitudes and perhaps other responses. In our paradigm cases, you move from a set of attitudes to a further attitude. We call the attitudes you begin with the *premise-responses* of your reasoning. The attitude you reason to is the *conclusion-response*.

Reasoning does not always lead to new attitudes. Sometimes reasoning leads you to reaffirm or to drop an attitude. However, to keep things manageable, we will focus on reasoning which leads to new attitudes. We will also focus only on reasoning to and from beliefs and intentions, leaving consideration of reasoning involving, for example, desires or suppositions for another occasion.

¹ There are other accounts that we lack the space to discuss here. Harman (1999) thinks that good reasoning increases the explanatory coherence of your attitudes. Boghossian (2008: chs. 11 and 12) and Enoch and Schechter (2008) give accounts of good theoretical reasoning which might be extended to good reasoning in general.

Reasoning is connected to basing and to agent's reasons—the considerations in light of which an agent responds. If you reason from some premise-responses to a conclusion-response, then, at least other things equal, you come to base the conclusion-response on the premise-responses. The contents of the beliefs from which you reason are reasons for which you respond. For instance, if you reason from an intention to drink gin and a belief that there is gin in your glass to an intention to take a sip, then a reason for which you intend to take a sip—a consideration in light of which you form this intention—is that there is gin in your glass. We assume that, because of these connections, there will also be close connections between good reasoning and good basing. Some such connections will come up in what follows.

We now turn to further clarifications of our notion of good reasoning.

First, when we ask what is good reasoning, we want an account of what, in general, *makes* reasoning good, when it is. That is, we want a statement of the *standard* for good reasoning. Specifying this standard may not be the same as giving a metaphysical analysis of good reasoning. For, plausibly, being good reasoning is being good *as* reasoning. A metaphysical analysis of good reasoning should thus draw on a general account of attributive goodness (goodness *qua* thing of a certain kind). But we should be able to specify the standard for good reasoning while remaining neutral on the analysis of attributive goodness. (Compare: we can say what makes for good cars while remaining neutral on the analysis of attributive goodness.)

Second, in the sense we intend, good reasoning can have bad starting points: it can begin from mistaken or unjustified attitudes. Consider Bernard Williams's (1979) famous example in which you intend to drink gin and believe that there is gin in your glass. It is good reasoning to move from these attitudes to intending to take a sip from your glass even if in fact there is petrol in your glass, and even if you ought not to be drinking gin. To say that some reasoning is good, in our sense, is to say something about the *transition* between attitudes, rather than something about the attitudes you begin from.

Third, it is natural to suppose that good reasoning falls into patterns and that those patterns explain why their instances are good. For example, it is plausible that the pieces of reasoning we began with are good because they exemplify the *modus ponens* and *means-end* patterns respectively. We can represent these patterns schematically as follows:

Belief that p , Belief that if p , q $>$ Belief that q

Intention to E , Belief that in order to E , you must M $>$ Intention to M

Given these assumptions, we can investigate what makes for good reasoning by focusing on what makes for a good pattern of reasoning.² The accounts we consider below should be understood accordingly, as specifying conditions on good patterns of reasoning, rather than on instances.

Fourth, to say that it is good reasoning to move from some premise-responses to a conclusion-response is not by itself to say anything about whether any token piece of reasoning was done well or, as we shall say, competently. This is an analogue of a familiar point. Doing what is in fact the right thing does not ensure that you have acted well—for

² While we accept these assumptions, they are not trivial. Moral particularists (e.g. Dancy 2004), would deny that there are any, or many, good patterns of reasoning. However, we cannot engage with particularist arguments here. We note a point where the assumption is dialectically significant below (n. 13). For defence of the claim that good reasoning falls into patterns see Pollock (1987), Boghossian (2008: ch. 5).

instance, if you have done the right thing for the wrong reasons. Believing what there is justification to believe does not ensure that you believe justifiably—for instance, if your belief is not based on the justification.³ Similarly, reasoning in a way that *accords* with a good pattern of reasoning does not guarantee that you have reasoned competently. For instance, the example with which we began accords with the modus ponens pattern. But suppose that in carrying out this reasoning you are guided not by the modus ponens pattern but by the *modus profusus* pattern (Turri 2010):

Belief that p , Belief that q > Belief that r

In that case, although your reasoning accords with a good pattern, you would not be reasoning competently. At the least, competent reasoning requires that you *follow* a good pattern.

It is notoriously difficult to say what it is to follow a pattern (or ‘rule’) of reasoning. We will not consider this question here, but we shall assume that the notion is in good order. And our focus will be on what makes for good patterns of reasoning, not what makes for competent reasoning.

A final important point is that good reasoning can be *defeasible*. It can be good reasoning to move from some premise-responses to a conclusion-response even if it is not good reasoning to make that move given some further premise-responses. For instance, it might be good reasoning to move from the belief that someone said that the store is open to the belief that the store is open, even if it is not good reasoning to make this move given the further belief that the person who said this is a liar. And it might be good reasoning to move from the belief that you promised to meet your friend to an intention to do so, even if it is not good reasoning to make this move given the further belief that by meeting your friend, you will leave someone to die.⁴

2. The Rational Requirements View

According to the rational requirements view, standards of good reasoning derive from rational requirements. This view has been made prominent by groundbreaking work by John Broome. And the view has significant initial appeal. Consider the following plausible rational requirements:

(Modus Ponens Requirement)

Rationality requires of you that [if you believe that p and believe that if p then q , you believe that q].

(Means-End Requirement)

Rationality requires of you that [if you intend to E and believe that M -ing is necessary for E -ing, you intend to M].

³ In the jargon, believing what there is propositional justification for does not ensure doxastic justification. For the need for a distinction of this sort in the case of reasoning, see also Jenkins (2014).

⁴ On this way of formulating it, ‘good reasoning’ applies to: making this move *given the defeating belief* (Way forthcoming 1). Considered on its own, the move from the belief that you promised to the intention is still good reasoning, in the sense that it instantiates a good pattern.

(Enkrasia)

Rationality requires of you that [if you believe that you ought to F , you intend to F].⁵

Each of these requirements corresponds to a clear case of good reasoning. The first two correspond to the modus ponens and means-end patterns. The third corresponds to the *enkratic pattern*:

$$\text{Belief that you ought to } F > \text{Intention to } F$$

It is therefore a natural hypothesis that these good patterns of reasoning hold *because* of the corresponding rational requirements. More generally, we might think that what makes a pattern of reasoning good is that it can bring you to satisfy a rational requirement (Broome 1999).

Despite this appeal, we think the rational requirements view is unsatisfactory. We take it to face three serious challenges.

First, there is the *symmetry problem*. Rational requirements are symmetrical in that they do not discriminate between ways in which they might be satisfied. For instance, you satisfy Enkrasia if *either* you believe that you ought to F and intend to F *or* you do not believe that you ought to F . But standards of good reasoning are *asymmetrical*. Not all reasoning which would lead you to satisfy a rational requirement is good reasoning. For instance, suppose that you reason ‘I am not going to F , so it is okay not to F ’. This reasoning brings you to believe that it is okay not to F , which, all going well, will bring you to not believe that you ought to F , and so satisfy Enkrasia. That does not make it good reasoning (cf. Schroeder 2004, Kolodny 2005).

This observation undermines the simple version of the rational requirements view suggested above. Partly in recognition of this point, Broome’s more recent work develops the rational requirements view in a different way. Broome’s view is now that standards of good reasoning derive from *basing permissions*. A basing permission is a permission of rationality to base certain attitudes on other attitudes. Broome suggests that reasoning is good just in case it follows a basing permission (2013: 247). So, for example, rationality permits you to base an intention to F on a belief that you ought to F (290), and that is why it is good reasoning to move from the belief that you ought to F to intending to F .

This is still a version of the rational requirements view. For, while the standards for good reasoning derive most immediately from basing permissions, Broome proposes that basing permissions derive from rational requirements. As he says of good instrumental reasoning, ‘it is correct to reason according to the permission just because this is a way to satisfy the requirement’ (2013: 258).

However, this approach merely relocates the symmetry problem. Just as not all reasoning leading to the satisfaction of a rational requirement is good reasoning, not all basing that satisfies such requirements is permitted. For instance, it is not rationally permissible to base the lack of a belief that you ought to F on the lack of an intention to F . Broome calls this a ‘basing prohibition’ (2013: 140–1). But without an account of where these

⁵ Cf. Broome (2013). These requirements are stated roughly. The details are controversial and won’t matter for the points we want to make.

prohibitions come from, the symmetry problem remains. What's more, it seems clear that these prohibitions *can't* derive from rational requirements such as Enkrasia. After all, if you drop your belief that you ought to F on the basis of not intending to F , you do not thereby violate Enkrasia. You come to satisfy it.⁶

The second problem for the rational requirements view is that it does not accommodate the defeasibility of good reasoning. It is good reasoning to move from the belief that someone said that p to the belief that p , and from the belief that you promised to F to an intention to F .⁷ But there is no rational requirement to [believe that p , if you believe that someone said that p] or to [intend to F , if you believe that you promised to F]. For instance, it is possible to be rational while believing that someone said that p without believing that p —after all, you might also believe that the person who said that p is a liar. So the rational requirements view does not explain how there can be good but defeasible reasoning.

This problem turns on Broome's assumption that if you violate a rational requirement, you are thereby irrational. One option for the rational requirements view would be to give up this assumption and so make room for *pro tanto* rational requirements. However, this assumption is an important part of Broome's overall approach. For one thing, it is needed for Broome's account of the property of rationality in terms of rational requirements (2013: 117–8). More centrally still, it underpins Broome's methodology of testing putative rational requirements by considering whether it is possible to be rational while violating them—see esp. (2013: chs. 9 and 10). Finally, the idea that good defeasible reasoning might turn on *pro tanto* rational requirements moves us away from Broome's guiding idea that rationality is a matter of *coherence*. Not all failures to engage in good but defeasible reasoning seem like failures of coherence. So, while it would be interesting to see how the idea that there are *pro tanto* rational requirements might be developed, it is not straightforward for those sympathetic to Broome's general approach to allow this idea.

A different option, which does not require giving up Broome's assumption, is to say that good defeasible reasoning corresponds to (overall) rational requirements which mention 'defeating attitudes'. For instance, perhaps there is a rational requirement to [believe p , if you believe that someone said that p and lack any defeating belief]. The challenge for this approach is to say what a defeating attitude is. On the face of it, a defeating attitude for a piece of reasoning is just one such that it is not good reasoning to make that move given that attitude. But this would make the approach look circular. To avoid this, one might suggest that a defeating attitude is one in the presence of which something that would otherwise be rationally required is not so. But this seems too strong. For instance, reasoning from an intention to E to a means M you believe sufficient for E looks like good but defeasible reasoning. But suppose you also believe that M' is a sufficient, and no worse, means to E . In that case, you need not be irrational if you fail to intend to M —you might instead intend to M' . But it is still good reasoning to move to an intention

⁶ One might hold that good reasoning is a matter of conforming to basing permissions of rationality, but reject Broome's further claim that these permissions derive from rational requirements. This would avoid the symmetry problem. But it would not be a satisfactory account of good reasoning. Since basing and reasoning are closely related, it is near trivial to say that reasoning is good when it is permitted by standards of basing.

⁷ Broome may not accept these examples. However, the worry here also arises for cases of defeasible good reasoning which Broome does accept, such as inductive and abductive reasoning (2013: 191).

to M —this reasoning is not defeated. Variants on this example also undermine the suggestion that defeating attitudes are those which correspond to defeating reasons. Reasoning from an intention to E to an intention to M could be defeated by a belief that there is a clearly better means to E , or by an intention to do something incompatible with M -ing. But these attitudes need not correspond to reasons not to intend to M .⁸

So the rational requirements view cannot easily accommodate defeasible reasoning. It also faces a third problem. We take theories of good reasoning to have implications concerning the *point* or *aim* of reasoning. The thought here is that we should expect theories of good reasoning to help explain why reasoning is worth going in for, or at least makes sense to go in for, when forming and revising your attitudes. If we take theories of good reasoning to tell us what the point or aim of reasoning is, they help to make clear what reasoning can help us achieve, and thus why it is worth going in for.⁹

Given the rational requirements view, the point of reasoning is satisfying rational requirements. We find this claim implausible. Rational requirements tell you to have attitudes which fit together coherently (see e.g. Broome 2013: 152). As Broome and others have persuasively argued, though, there may be no reason to be rational in this sense. Satisfying rational requirements need not be a means to anything else there is reason to do (2013: sec. 11.3). Nor does satisfying rational requirements seem worthwhile for its own sake. Psychic tidiness is not a final value (Kolodny 2007: 251).¹⁰

So if the point of reasoning is to satisfy rational requirements, reasoning seems like a fairly worthless activity. Why go in for this activity which, even when it is going well, can only be relied on to ensure you meet some requirements which are not worth satisfying?¹¹

Overall then, the rational requirements faces three serious problems. While each is worthy of further discussion, they also seem serious enough to lead us to consider alternatives.

3. The Reasons View

According to the reasons view, standards of good reasoning derive from reasons. Roughly, the idea is that good reasoning is reasoning which takes you to responses for which there are reasons. This is an attractive idea. As noted, in reasoning we form and revise attitudes for reasons—in light of certain considerations. It is therefore natural to suppose that reasoning goes well when it leads us to form and revise attitudes for good reasons.¹² Furthermore, the reasons view seems to easily allow for defeasible reasoning. If someone said that p , that is a defeasible reason to think that p . The view thus allows

⁸ Note that a better means to some end might not be a means one has better reason to take.

⁹ For further defence of the connection between good reasoning and the point of reasoning, see McHugh and Way (ms 1).

¹⁰ It might be replied that there are related goals which are worth aiming at for their own sake—perhaps integrity, or authenticity. However, even if this is true, satisfying rational requirements of coherence falls a long way short of achieving these goals.

¹¹ In fact, Broome believes that there is always some reason to satisfy rational requirements (2013: 204). However, as Broome admits, he has no argument for this claim. Furthermore, even if true, this claim is not enough to meet the worry. For the claim that there is some reason to F is a very weak one (cf. Schroeder 2007 and §3 below). An activity which, even when going well, ensures only that you end up doing something there is *some* reason to do still does not seem especially worth going in for.

¹² If you like: reasoning goes well when it ensures that our *motivating* reasons are *normative* reasons.

that it is good but defeasible reasoning to move from the belief that someone said that *p* to the belief that *p*. The view also avoids the symmetry problem for the rational requirements view. For instance, since the fact that you're not going to *F* is not by itself a reason to believe that it's okay not to *F*, the view does not imply that 'I am not going to *F*, so it's okay not to *F*' is good reasoning. Finally, the view also fits well with the recently popular 'reasons-first' approach to normativity. On this approach, all normative and evaluative properties are to be understood in terms of reasons (e.g. Scanlon 1998, Schroeder 2007, Skorupski 2010, Parfit 2011). Those sympathetic to this approach will be sympathetic to understanding good reasoning in terms of reasons.¹³

While several philosophers have made remarks which suggest sympathy for the reasons view, it is hard to find any detailed defence of it.¹⁴ So we will just start with what seems the most natural formulation of the view, and proceed by working through some problems for it, and seeing how it might be developed to avoid these problems. We will not argue that these problems are decisive. But we will argue that solving them requires the reasons view to take on controversial commitments. While we take this to be a serious disadvantage, others might be happy to take these commitments on. Such readers can take our discussion as showing what the best version of the reasons view must look like.

Here is a first version of the view:

(RV₁) The move from $P_1 \dots P_n$ to *C* is a good pattern of reasoning iff, and because, the contents of the beliefs in $P_1 \dots P_n$ are reasons for *C*.

RV₁ has the attractions of the reasons view noted above. But it faces the *problem of bad starting points*. For an example of this problem, consider that while reasons are facts, good reasoning can begin from false beliefs. Consider again Williams's example in which you intend to drink gin and falsely believe that there is gin in your glass. It is good reasoning to move from these attitudes to an intention to drink from your glass. But the consideration that there is gin in your glass is not a reason to intend to take a sip from your glass, since it is not a fact.

A reply to this version of the problem is available. It is common to distinguish between *objective* reasons and *subjective* or *apparent* reasons. Objective reasons are facts which count in favour of (or against) some response. Subjective or apparent reasons are considerations which bear on the rationality of some response and which it is rational to treat as objective reasons.¹⁵ In Williams's case, the consideration that there is gin in your glass is not an objective reason to intend to take a sip but it might be a subjective reason. So perhaps RV₁ should be understood to refer to subjective reasons, rather than objective ones. If so, Williams's case is not a counter-example.

This reply does not solve the underlying problem. Good reasoning can begin from irrational attitudes. But good reasoning from irrational attitudes need not lead you to do

¹³ Two further potential attractions of the reasons view are worth noting. First, it generalises straightforwardly to reasoning which concludes in dropping an attitude. Second, it allows for the possibility of good reasoning without good patterns of reasoning. As noted in §1, we are putting these issues aside here.

¹⁴ Those who seem sympathetic include Dancy (2014), Grice (2001), Harty (2012), Pollock (1987).

¹⁵ A common view is that subjective reasons are considerations which would be objective reasons if true (e.g. Parfit 2011, Schroeder 2007). This view faces a host of problems (Sylvan forthcoming, Way 2012, Whiting 2014, Vogelstein 2012). Rather than relying on this or an alternative theory, we will instead rely on the roles of subjective reasons identified in the text.

what there is subjective reason to do. For instance, suppose you irrationally believe that your plane will crash, despite being aware of all the evidence that flying is safe. It is not rational to treat the consideration that your plane will crash as an objective reason, so this consideration is not a subjective reason. But ‘my plane will crash, so I won’t board’ is still good reasoning. Or suppose that you have some irrational intention—say, out of curiosity you intend to set yourself on fire. It is good reasoning to move from this intention to an intention for what you take to be the necessary means to setting yourself on fire—say, dousing yourself in petrol. But it is not rational to treat the consideration that dousing yourself in petrol is a necessary means to setting yourself on fire as a reason to douse yourself in petrol. So again, good reasoning need not lead you to do what there is subjective reason to do.¹⁶

RV₁ is thus unsatisfactory. But there is an alternative way to develop the reasons view. Even if we allow that good reasoning need not lead you to attitudes for which there are reasons, objective or subjective, we might think that *under favourable conditions* good reasoning will lead you to attitudes for which there are reasons. A straightforward way to spell out this thought is to say that good reasoning is reason-preserving: if it begins from responses for which there is sufficient reason, it leads to a response for which there is a reason.^{17,18} So consider:

(RV₂) The move from $P_1 \dots P_n$ to C is a good pattern of reasoning iff, and because, if there is sufficient reason for $P_1 \dots P_n$, then the contents of the beliefs in $P_1 \dots P_n$ are reasons for C .

This view helps with some versions of the problem of bad starting points. For instance, if there is sufficient reason to intend to set fire to yourself, then the fact that dousing

¹⁶ These counterexamples assume that rationality is not merely a matter of coherence. Since those who distinguish objective and subjective reasons also make this assumption, this does not beg any questions. Nonetheless, some readers may feel inclined to resist our verdicts on the examples and insist that it is rational to, e.g. treat means to irrational ends as reasons. One way to defend this claim is to note that something is going right if you decide to take such means. However, this is more directly and plausibly accounted for by observing that you’re engaging in good reasoning than by saying that this decision is rational. For what seems to be going right here is just the transition between attitudes, rather than the attitude you end up with. Another way to defend this claim is to appeal to a distinction between local and global rationality. It might be said that the decision to take the means is rational relative to the intention to take the (irrational) end, even if it is not rational relative to the rest of one’s psychology. However, if subjective reasons are taken to correspond to local rationality in this way, it becomes harder to see how they contribute to determining what’s rational overall, as subjective reasons are supposed to do—for presumably what matters for overall rationality are the subjective reasons one has relative to one’s overall psychology.

While several accounts of subjective reasons do allow that irrational attitudes can provide subjective reasons (Parfit 2011, Schroeder 2007, Vogelstein 2012), this is often taken to be an objection to these accounts. Schroeder (2011) defends such views against this charge. However, his defence requires us to give up the claim that subjective reasons are considerations it is rational to treat as objective reasons, as well as the claim that subjective reasons contribute to the rationality of further attitudes. Without either of these claims, we arguably lose our grip on the notion of a subjective reason.

¹⁷ The objections we are about to outline would apply equally—indeed, more so—to a version of the view which appealed only to there being *some* reason for the premise-responses.

¹⁸ A different way to develop this thought would be to say that good reasoning leads to a response for which there is a reason, if the premise beliefs are true. However, this proposal shares many of the problems faced by the view that subjective reasons are considerations that would be objective reasons if true (see n. 15 above). Moreover, it does not handle cases where agents have true means-end beliefs but mistaken intentions, such as the intention to set themselves on fire.

yourself in petrol is necessary for doing so is a reason to intend to douse yourself in petrol. So RV_2 implies that it is good reasoning to move from the former intention to the latter.

However, RV_2 remains vulnerable to some versions of the problem. Consider Williams's example again. On the assumption that there can be sufficient reason to believe falsehoods, there may be sufficient reason to believe that there's gin in your glass even though there is not—and so even though the consideration that there is gin in your glass is not a reason to intend to take a sip. Thus RV_2 does not imply that it is good reasoning to move from the belief that there's gin in your glass to the intention to take a sip.

There are a couple of ways to avoid this problem. First, the proponent of RV_2 might deny that there can be sufficient reason to believe falsehoods. Second, the proponent of RV_2 might appeal to subjective reasons. If you have sufficient *subjective* reason to intend to drink gin and to believe that there's gin in your glass, then that there is gin in your glass is plausibly a subjective reason to intend to take a sip.¹⁹

However, both these responses involve commitments which many will want to avoid. The claim that there can never be sufficient reason to believe what's false—in effect, the claim that there can be no justified false beliefs—is highly, and rightly, controversial. Evidence provides reason to believe, and even very strong evidence can be misleading. Short of thinking that sufficient evidence must be *entailing*—a view that threatens inductive scepticism—it's hard to see why very strong misleading evidence couldn't justify belief.²⁰ Solutions to the problem of bad starting points with such implications are to be avoided if possible.

The appeal to subjective reasons raises a worry about circularity. As we saw, subjective reasons are determined by rational attitudes—irrational attitudes do not generate subjective reasons. Proponents of this version of RV_2 are thus committed to explaining what it takes for token attitudes to be rational without appealing to claims about good reasoning. But when an attitude is formed or sustained through reasoning, the quality of this reasoning seems highly relevant to its rationality. Again, it would be preferable to avoid

¹⁹ A third move would be to remove reference to the *contents* of the premise-beliefs in RV_2 , and say that reasoning is good when, if there is sufficient reason for the premise-responses, *there is* a reason for the conclusion-response. That there is gin in your glass is not a reason to take a sip, but perhaps the fact that your glass was filled from the gin bottle is. A major worry for this view is that it makes good patterns of reasoning too easy to come by, since there will often be some reason for the conclusion-response, whether or not there is sufficient reason for the premise-responses. (We will return to a version of this issue in §5.2 below.) This view would also face the further problems for the reasons view outlined below.

²⁰ Note that the controversy over whether there can be sufficient reason to believe falsehoods persists even when we are clear that we are talking about objective reasons. One can think that only truths can be reasons for belief while denying that only truths can be justified (cf. Williamson 2000). Nor can the controversy be dissolved by distinguishing the reasons there are from the reasons an agent has: some authors hold that the (epistemic) reasons there are must be in some way accessible (e.g. Skorupski 2010). More generally, claims about what there is sufficient reason to believe plausibly correspond to claims about what one *may* believe or what there is *justification* to believe. But if such permissions or justifications always depend on the truth of the target proposition, it is unclear how they can guide us in the many cases in which the truth is inaccessible to us (cf. Gibbons 2013, Lord 2015, McHugh and Way ms 2). These worries notwithstanding, our central point here is just that such infallibilism about justified belief is controversial and so a cost of this version of the reasons view. For defences of such infallibilism see Littlejohn (2012), Schroeder (2015), Whiting (2013).

commitment to the task of explaining what makes for rational attitudes without appealing to good reasoning.²¹

So none of the solutions to the problem of bad starting points seems fully satisfactory. At best, solutions to it require strong and controversial commitments. Furthermore, all versions of the reasons view face two further problems.

First, the view faces a problem familiar to many attempts to analyse other properties in terms of reasons—the *wrong kind of reason problem*. Suppose that it would make you happy to believe in God, or that an eccentric billionaire offers a \$1m prize for intending to drink an unpleasant but otherwise harmless toxin.²² Since the prospect of happiness clearly counts in favour of believing in God, and the billionaire’s offer clearly counts in favour of intending to drink the toxin, it is tempting to think that you here have reasons to believe in God and intend to drink the toxin. But the reasoning you would express to yourself by saying ‘believing in God would make me happy, so, God exists’ or ‘intending to drink the toxin would get me \$1m, so, I shall drink the toxin’ is clearly bad reasoning. This suggests that not all reasons correspond to good reasoning.

There are various strategies available for responding to this objection, which parallel the responses to the wrong kind of reason problem for other reasons-based analyses. In particular, the reasons view might distinguish between the *right* and *wrong* kind of reasons, and claim that good reasoning only corresponds to the right kind of reasons. Alternatively, it might be denied that the incentives in these examples give you reason to believe in God or to intend to drink the toxin. Since these strategies have been discussed in detail elsewhere, we will not consider them here.²³ Instead, we will just make three simple points. First, while several philosophers have tried to solve the wrong kind of reason problem, no solution has yet been widely accepted. It is an open question whether an adequate solution is available. Second, it is an especially hard problem for proponents of the reasons-first approach to solve. By analysing other normative notions in terms of reasons, they make it hard to appeal to those notions in order to solve the problem.²⁴ Third, even if a solution compatible with this approach is possible, the literature thus far suggests that it will be complicated. This makes reasons-first analyses correspondingly less attractive.

The second problem for the reasons view turns on its implications for the point of reasoning. Since the reasons view says that good reasoning takes you to attitudes for which there is some reason, at least when you begin from well-supported attitudes, it implies that the point of reasoning is to reach attitudes for which there is some reason.

On the face of it, though, this is not a plausible idea. Almost any attitude has *some* reason for it, even intending to eat your car (Schroeder 2007). If that was all we wanted we might as well form and revise attitudes at random. And there is a salient and clearly more plausible proposal available, according to which reasoning aims at *sufficient* rea-

²¹ Proponents of this view might argue that attitudes are rational when properly based on sufficient reasons. However, this is plausible only if the reasons in question are *subjective* reasons. Thus this suggestion raises clear worries about circularity. Furthermore, insofar as proper basing and good reasoning are intimately connected, it is far from clear that the suggestion avoids the circularity worry in the text.

²² Kavka (1983).

²³ A further possibility would be to claim that it is good reasoning to move from, e.g. the belief that believing in God would make you happy to the belief that God exists. But this claim seems theory-driven. See Way (forthcoming 2) for further discussion.

²⁴ For example, McHugh and Way (2016) defend a solution that the reasons-first could not endorse.

sons. We don't merely want attitudes which have something to be said for them. What we want from reasoning is, more plausibly, attitudes for which there is sufficient support.

The reasons view could be revised so as to support this proposal. For instance, RV_2 could be revised to say that good reasoning is sufficient-reason-preserving—i.e. reasoning which will lead you to a response for which there is sufficient reason, if you begin from responses for which there is sufficient reason. However, revising the view in this way sacrifices the simple and attractive way in which the reasons view accommodates defeasible reasoning. The reasoning you would express by saying ‘someone said that p , so p ’ or ‘I promised to F , so I shall F ’ is not sufficient-reason-preserving. The reasons view thus faces a dilemma: either it offers a straightforward account of defeasible reasoning but thereby makes an implausible claim about the point of reasoning or it makes a plausible claim about the point of reasoning but sacrifices its account of defeasible reasoning.

A proponent of the reasons view might reply that this dilemma can be avoided. The reasons view can say that good reasoning preserves sufficient reasons *other things equal*. The view of good reasoning we will defend below involves such a clause. So we do not object to this move in itself. Nonetheless, it clearly raises questions about how the clause is to be understood. It thus remains true that this approach sacrifices the straightforward way in which other versions of the reasons view handled defeasibility.²⁵

In any event, the claim that reasoning aims at sufficient reasons is only so plausible. A standard and very plausible claim is that theoretical reasoning aims at truth. But even if developed in the way just suggested, the reasons view is in tension with this plausible claim. This is clearly so if the reasons view is formulated in terms of *subjective* reasons. On any account there can be subjective reason for you to believe what's false. If instead the reasons view is formulated in terms of objective reasons, then the reasons view preserves the standard and plausible claim only given the assumption that there can be sufficient objective reason to believe only what is true. But, as we have seen, this is a rightly controversial claim. It would be preferable if an account of good reasoning could avoid such a commitment.

To sum up, the reasons view faces three serious challenges: the problem of bad starting points, the wrong kind of reason problem, and implausible implications about the point of reasoning. We have seen that there are some prospects for avoiding these problems, but they do not come cheaply. Again, while more could be said here, it also seems worth considering an alternative approach.

4. The Fittingness View

The account of good reasoning that we propose begins with the observation that attitudes have standards of *fittingness* or *correctness*. Perhaps the most familiar example here is belief: beliefs are correct, or fitting, just when true.²⁶ But other attitudes have standards of fittingness too. When you admire Mandela, desire a vacation, fear an onrushing tiger, or intend to catch the last train, your attitude can be fitting or unfitting, depending on the features of its object. For example, it seems fitting to admire Mandela and fear the onrushing tiger. It is not fitting to admire Idi Amin or to fear the onrushing kitten.

²⁵ We discuss ‘other things equal’ clauses below. The reasons view might draw on this discussion.

²⁶ Or just when knowledge (Smithies 2012). It doesn't matter here.

‘Fitting’ is here something of a term of art. We could also express the relevant judgments by saying that it is *correct* to admire Mandela, that one would be *getting it right* in admiring him, or that Mandela is *fit* to be admired, or *worth* admiring, or an *appropriate object* of admiration. These expressions can all be used to express to the notion of fittingness. We use ‘fitting’ as our standard way to do so.

The core idea is that attitudes are associated with standards for their objects. An attitude is fitting when its object meets its standard. For instance, belief has truth as its standard of fittingness, and so a belief is fitting when what is believed is true. The standard of fittingness for other attitudes is more controversial. For instance, it is a substantive question in value-theory what makes something worth desiring—that is, desirable, or good. Here we will continue to focus on reasoning to beliefs and intentions. For illustrative purposes, we will adopt Nishi Shah’s (2008) suggestion that the standard of fittingness for intention is permissibility: that it is fitting to intend to *F* just when it is permissible to *F*.

The basic idea of our account is that good reasoning puts you in a position to acquire fitting attitudes—for instance, given the assumptions just stated, to gain more true beliefs, or intentions to do what is permissible. Of course, good reasoning cannot usually generate fitting attitudes from nothing—starting points matter. But if we start from fitting attitudes, we can expect good reasoning to lead us to further fitting attitudes, other things equal. As a first pass, then, we suggest that good reasoning is *fittingness-preserving* reasoning:

(FV) The move from $P_1 \dots P_n$ to *C* is a good pattern of reasoning iff, and because, other things equal, if $P_1 \dots P_n$ are fitting, *C* is fitting too.

Note again that (FV) concerns *patterns* of reasoning, not token-instances of reasoning. So the condition it proposes—that of *ceteris paribus* fittingness-preservation—must hold in general of a pattern of reasoning. It is not enough that, in a particular instance, you avoid going from fitting responses to an unfitting one. Note also that (FV) is *not* an account of *competent* reasoning, and *a fortiori* does not entail a reliabilist account of *justification* for beliefs acquired through reasoning.

To illustrate how the view works, we can begin by noting how it vindicates the paradigmatic examples of reasoning which motivated the rational requirements view. The modus ponens pattern is straightforward. Since modus ponens reasoning preserves truth, and beliefs are fitting just when true, modus ponens reasoning is (necessarily) fittingness-preserving. The enkratic pattern is similarly straightforward. If it is fitting to believe that you ought to *F*, then it is true that you ought to *F*, and so permissible to *F*. It is therefore fitting to intend to *F*. And the means-end pattern holds too. If it is fitting to intend to *E*, then it is permissible to *E*. And if it is permissible to *E* and *M* is a necessary means to *E*, then it is also permissible to *M*. Thus it is fitting to intend to *M*.²⁷

The fittingness view also does well where the reasons view and rational requirements view face problems. First, and as the above examples illustrate, the view avoids the problem of bad starting points. For instance, it implies that it is good reasoning to move from an intention to drink gin and a belief that there’s gin in your glass to an intention to

²⁷ Cf. McHugh and Way (2015), where we show how the fittingness view vindicates refined versions of these patterns.

drink gin, whether or not there is in fact gin in your glass, and whether or not these attitudes are themselves rational or otherwise well-supported. Note that the assumptions needed here—e.g. that fitting beliefs must be true—are far less controversial than those required for the reasons view to vindicate this reasoning. It is not a cost to be committed to the claim that correct beliefs must be true.

Second, since the view only requires that good reasoning preserves fittingness *other things equal*, it makes room for the defeasibility of good reasoning. Other things equal, if someone said that p , then p . So it is good reasoning to move from the belief that someone said that p to the belief that p . And other things equal, if you promised to F , it is permissible to F . So it is good reasoning to move from the belief that you promised to F to an intention to F . (More on this below.)²⁸

Third, the view rules out the kinds of bad reasoning which the alternative views allowed. Unlike the rational requirements view, it does not imply that ‘I am not going to F , so it is okay not to F ’ expresses good reasoning. Even if you are not going to F , it might be that you ought to do so; this reasoning does not preserve fittingness. And, unlike the reasons view, it clearly implies that ‘believing in God would make me happy, so, God exists’ is not good reasoning. It is not the case—even other things equal—that, if believing something would make you happy, then that thing is true.

Finally, the fittingness view implies that the point of reasoning is fitting attitudes. This seems a plausible claim. Unlike coherence and attitudes for which there is merely some reason, getting things right is clearly worth aiming at, insofar as you are revising your attitudes. Furthermore, since fitting beliefs are true beliefs, the fittingness view generalizes the familiar and plausible idea that the aim of theoretical reasoning is truth.

We thus take the fittingness view to have serious promise, and to have clear advantages over the reasons and rational requirements views. However, it also faces some problems of its own and requires further development in various ways. In the next section, we note some of these issues and suggest ways to make progress.

5. Developing the Account

5.1. Is the Fittingness View a Genuine Alternative to the Reasons View?

One might wonder whether the fittingness view is a genuine alternative to the reasons view. The upshot of §3 was that the best version of the reasons view says that good reasoning preserves *sufficient right-kind reasons, other things equal*. But one might think that this does not differ significantly from our claim that good reasoning preserves fittingness, other things equal. For, one might suggest, a fitting response is just a response for which there is sufficient right-kind reason.

We do not want to rule out the possibility of analysing fittingness in this way.²⁹ Furthermore, several of the issues to be discussed in the following sections would arise for

²⁸ There are other ways in which a fittingness view might be developed so as to accommodate defeasible reasoning. For example, a view which took fittingness to come in degrees might hold that good patterns of reasoning preserve some degree of fittingness, and explain defeat by appealing to ways in which further premise-attitudes can affect degrees of fittingness. However, while this approach might be appropriate in some cases, it is not clear how it can be applied to theoretical reasoning, where the standard of fittingness is plausibly truth. We thus prefer the approach in the text.

²⁹ But see McHugh and Way (2016).

this reasons view. So, as we have said, our paper might be taken as showing what the best version of a reasons view would have to look like, and exploring some of the issues that it raises.

Nonetheless, the point remains that proponents of the best version of the reasons view takes on commitments which the fittingness view does not. The reasons view requires a solution to the wrong kind of reasons problem and a defence of the claim that there can never be sufficient reason to believe falsehoods. The fittingness view does not involve these commitments. This is an advantage of the fittingness view. Furthermore, the proposed version of the reasons account looks like the result of an attempt to engineer a reasons-based theory that gives the right results. It seems simpler and more natural to suppose that fittingness is what's doing the work in good reasoning.

5.2. *Necessarily Fitting (or Unfitting) Attitudes*

FV may seem to over-generalise. Suppose that you reason, ‘grass is green, so, 79 is prime’. Since 79 is necessarily prime, this reasoning is guaranteed to lead to a true belief, and so preserves fittingness. But it is not at all clear that this makes it good reasoning. Or suppose you reason, ‘77 is prime, so, grass is blue’. Since it’s necessarily false that 77 is prime, this reasoning also preserves fittingness. But again it doesn’t look like good reasoning. Being fittingness-preserving thus seems insufficient for good reasoning.

Our response to this objection has a number of aspects. To begin with, recall the point made earlier, that there is a difference between reasoning which merely accords with a good pattern of reasoning and reasoning which is done *competently*. Competent reasoning requires (at least) following a good pattern.

It seems clear that someone who reasoned in one of the above ways would not be reasoning competently. But why this is will depend on what pattern such a person would be following. For example, suppose someone reasons ‘grass is green, so 79 is prime’ because they follow this pattern:

Belief that $p >$ Belief that the first number you think of is prime

Since this pattern does not preserve fittingness, anyone following it is not reasoning competently. The same applies to many other patterns which might yield the problematic pieces of reasoning. So our account can explain why many instances of such reasoning would not be competent.

Suppose though that you reasoned in the above ways by following these patterns, respectively:

- (T) Belief that $p >$ Belief that q , where q is some necessary truth
- (F) Belief that $p >$ Belief that q , where p is some necessary falsehood

These patterns do preserve fittingness. So, as developed so far, our account implies that these are good patterns. That seems problematic.

The striking thing about these patterns is that the premise-responses may be entirely irrelevant to the conclusion-responses (as they are in our examples). Although the first pattern preserves fittingness, this is just because the conclusion-response is necessarily

fitting. Although the second pattern preserves fittingness, this is just because the premise-response is necessarily unfitting. Intuitively, good patterns of reasoning aren't like this: they preserve fittingness because of some relationship between the premise-responses and the conclusion-response.

This suggests a simple restriction to our account: good patterns of reasoning are not fittingness-preserving only in virtue of the premise- or conclusion-responses. They preserve fittingness (other things equal) in virtue of a relationship between the premise-responses and the conclusion-response. This restriction implies that (T) and (F) are not good patterns.

There are two worries one might have about this reply. First, it might be thought that the restricted account still over-generalises. Imagine that someone reasoned directly—without any intermediate steps—from the Peano Axioms to Fermat's Last Theorem (cf. Broome 2013, Boghossian 2014). Or suppose someone ignorant of the chemical composition of water reasoned 'this is water, so, it's H_2O '. Neither of these seem like cases of good reasoning, but both exemplify patterns which are not fittingness-preserving just in virtue of the premise- or conclusion-responses. However, in these cases, we are happy to allow that the patterns are good ones. It is just that it is hard to imagine these patterns being followed competently in the absence of further steps (in the first case) or the belief that water is H_2O (in the second). That is why the reasoning in these cases strikes us as bad—it is not competent reasoning.³⁰

A second worry is that the restriction is ad hoc. But in fact we can motivate this restriction in a principled way. As we have argued, the point of reasoning is to get fitting attitudes. So good patterns of reasoning must serve this point. But, moreover, good patterns of reasoning must be able to *guide* us. Since we can't get fitting attitudes just like that, we need a way of doing it. The way we do it is by making transitions. The standards of good reasoning can't by themselves ensure that we acquire fitting attitudes, but they can *keep us on track*. They can do that if they are such that, by following them, normally we will get fitting attitudes out if we put fitting attitudes in.

Patterns that preserve fittingness in the way that (F) does, just in virtue of the premise-responses, don't serve the point of reasoning. They have no tendency to lead to fitting attitudes, since the conclusion-response could be anything at all. Patterns like (T), on the other hand, do serve the point of reasoning. But they cannot guide us. In effect, (T) just says: believe necessary truths! But the reason we need standards of good reasoning in the first place is that we can't do things like this just like that.

This suggests that good patterns of reasoning are those which allow us to acquire fitting attitudes by, or in virtue of, reasoning from fitting attitudes. Patterns of reasoning like (T) and (F) are not like this. (F) doesn't allow you to acquire fitting attitudes at all, except by fluke. (T) doesn't allow you to acquire fitting attitudes in virtue of reasoning from fitting attitudes—it leads to fitting attitudes only because of its conclusion. Acquiring fitting attitudes in virtue of reasoning from fitting attitudes requires us to follow

³⁰ So these cases do suggest that competent reasoning requires more than just following a good pattern. It requires some kind of sensitivity to the goodness of the pattern. Again, this is an analogue of a familiar point. Acting well is not simply a matter of doing the right thing for what are in fact the right reasons. Acting well requires that you do the right thing *because* it's right. Since our focus is on what makes for good patterns of reasoning, not on what makes for competent reasoning, we will not consider the nature of the required sensitivity here.

patterns that are fittingness-preserving in virtue of more than just their premise- or conclusion-response. Thus, the proposed restriction on good patterns of reasoning is not ad hoc.

5.3. *Other Things Equal*

The fittingness view makes crucial appeal to an ‘other things equal’ clause. Such clauses can seem suspicious. Notoriously, some have taken ‘other things equal’ claims to be vacuous. With many others, this seems to us a mistake. While ‘other things equal, birds fly’ and ‘other things equal, promises should be kept’ seem clearly true, ‘other things equal, birds have stripes’ and ‘other things equal, promises should be broken’ seem clearly false. So such claims are not vacuous (cf. Nickel 2010). Nonetheless, it would be nice to say more.

There are various ways to develop the account. One option is to appeal to a probabilistic condition: good reasoning is *likely* to lead from fitting attitudes to fitting attitudes. Another option is to appeal to a ‘safety’-style condition: when reasoning is good, it could not easily happen that the premise-responses are fitting but the conclusion-response is not. However, we do not find either of these suggestions promising. When a piece of reasoning is defeated it may well be probable that it will not lead from fitting responses to a fitting response, and this may be no accident. For example, it will not be probable that what a liar says is true, and it is no accident if it is not.

These observations should not cast doubt on our claim that, other things equal, good reasoning preserves fittingness. It is a common feature of ‘other things equal’ claims that they do not immediately support claims about what is probable or could easily happen. For instance, ‘other things equal, matches light when struck’ and ‘other things equal, fish eggs develop into fish’ both seem true. But it need not be probable that any particular match lights when struck and it can easily happen—it does easily happen—that fish eggs do not develop into fish.³¹

This suggests that we should seek a better understanding of our ‘other things equal’ clause, rather than a replacement for it. One interesting recent proposal is that ‘other things equal’ clauses can be elucidated by appeal to what happens *normally* (Nickel 2010). As a rough first pass:

Other things equal, *F*s are *G*s, just in case all normal *F*s are *G*s.

This requires some clarification and refinement. First, ‘normal’ does not mean ‘usually’. Even if most matches are wet, and so do not light when struck, it remains true that normal matches light when struck. Second, we might distinguish between a normal *F* and normal conditions for an *F*—a wet match is abnormal in the first way, a match in an oxygen free chamber is abnormal in the second way. For brevity, we will take a ‘normal *F*’ to be one which is normal in both these ways. Third, normality comes in respects. A raven may be normal in respect of colour but not in respect of number of wings. To accommodate this the proposal above can be revised:

³¹ As Nickel (2010) points out, ‘other things equal’ claims thus seem very close in meaning to generic claims. The idea of analysing generic claims by appeal to a notion of normality is a popular one. For further references and discussion see Leslie (2012).

Other things equal, *Fs* are *Gs*, just in case all *Fs* which are normal in a respect determined by the predicate ‘*G*’ are *Gs*.³²

Can anything more substantive be said about normality? We note two commonly cited marks of normality. First, normal *Fs* reveal the nature of *Fs*. For example, a normal artifact is one that is functioning properly, in conditions for which it was designed. It thus reveals what the artifact is supposed to do. Similarly, a normal organism is one that is functioning properly in conditions which it, or some feature of it, was selected for (cf. Millikan 1984, Lance and Little 2007). Second, a normal *F* is an *F* for which interfering factors are absent. Thus, what happens to a normal *F* requires no special explanation (cf. Smith 2010, Pietroski and Rey 1995). For example, a normal match’s lighting when struck requires no special explanation. These conditions are connected: because normal conditions lack interfering factors, they reveal natures.

We want to suggest that this understanding of ‘other things equal’ clauses offers a promising way to develop the fittingness view. Understood in this way, the view says that:

(FV) It is good reasoning to move from $P_1 \dots P_n$ to *C* iff, and because, *normally* if $P_1 \dots P_n$ are fitting, *C* is fitting too.

How are we to assess the claims of the form ‘normally if $P_1 \dots P_n$ are fitting, *C* is fitting’? We suggest looking to the corresponding claim about the fittingness condition of the responses involved. For instance, consider:

Normally, if the belief that someone said that *p* is fitting, the belief that *p* is fitting.

Since the fittingness condition for belief is truth, this implies:

Normally, if someone said that *p*, then *p*.

That is:

Normally, testimony is reliable.

This claim is plausible. The claim that testimony which is normal in respect of its truth-value is true might be seen as analogous to the claim that matches normally light when struck. That is, we might see testimony as a social practice with the function of passing on information. Thus, when testimony is functioning as it is supposed to, in the kind of conditions it was designed for, testimony is reliable. If this is right, then the refined fittingness view supports the claim that it is good reasoning to move from the belief that someone said that *p* to believing that *p*. This is a plausible result.

Consider some other patterns of reasoning and the claims about normality they imply:

It is good reasoning to move from the belief that you promised to *F* to intending to *F* only if, normally, promises are permissible to keep.

³² See Nickel (2010) for this and further precisifications.

It is good reasoning to move from the desire that p to the belief that p only if, normally, what you desire is the case.

It is good reasoning to move from the belief that the coin came up heads last time to the belief that it will come up tails next time only if, normally, if the coin came up heads last time, it will come up tails next time.

In the first case, the implication seems plausible. Again, promise-keeping might be seen as a social practice which functions so as to bring about obligations and whose point is, perhaps, to allow us to co-ordinate our behavior in mutually beneficial ways. Viewed in this way, plausible counter-examples to the claim that all promises are permissible to keep—for instance, promises which are coerced or the keeping of which would be disastrous—look like abnormal cases (cf. Lance and Little 2007). So the view implies that it is good reasoning to move from the belief that you promised to F to intending to F .

In the second and third cases, the implications are false. So the view implies that wishful thinking and the gambler's fallacy are both bad patterns of reasoning.

The view thus has plausible implications in core cases. What's more, it offers a straightforward explanation of the defeasibility of good reasoning. The defeasibility of good reasoning is explained by the defeasibility of normal conditions. For instance, suppose that John, who is an inveterate liar, said that p . Although testimony is normally reliable, John's testimony is not normally reliable. The view thus explains why it is good reasoning to move from the belief that someone said that p to believing p , but not good reasoning to make this move when you believe that it was John, an inveterate liar, who said that p .

So, while much more would need to be said, we take the idea that other things equal clauses should be understood by appeal to normality to fit well with the view we are developing here.³³

5.4. Reasoning to the Belief that a Fittingness Condition Obtains

We objected to the first version of Broome's view on the grounds that it licenses bad reasoning, like 'I am not going to F , so it is okay not to F '. It might seem that a similar objection applies to our view. Consider the move from intending to F to believing that it is okay to F . On the assumption that the standard of fittingness for intention is permissibility, this reasoning preserves fittingness. More generally, the fittingness view seems to license reasoning from any attitude to the belief that the fittingness condition for that attitude obtains. Such reasoning might seem clearly bad. It might seem to amount to unacceptable bootstrapping.

We think that such reasoning is in fact good and does not constitute bootstrapping. First, note that the relevant reasoning for belief—from the belief that p to the belief that p is true—seems like good, if boring, reasoning. So this sort of reasoning is not *generally*

³³ We note one immediate concern. As we've seen, it seems true that, other things equal, fish eggs develop into fish. So on our account, it is good reasoning to move from the belief that this is a fish egg to the belief that this will develop into a fish. This might seem implausible, given how few fish eggs develop into fish. However, we think that this is simply a case of defeasibility. Compare someone who infers of this match that it will light when struck. This is a good piece of reasoning, even if in fact the match is wet. The fish egg reasoning does not seem to differ in kind from this case. The reasoning is defeated by further facts, but the reasoner ignorant of these facts would not be guilty of bad reasoning.

bad. Second, reasoning from an intention to *F* to the belief that it is okay to *F* must be distinguished from reasoning from a *belief* that you will *F* to a belief that it is okay to *F*. In ordinary English, either of these pieces of reasoning could be expressed by saying ‘I shall *F*, so, it’s okay to *F*’. But these pieces of reasoning are very different. If the latter was good it would yield a form of bootstrapping: the fact that you are going to *F* would ensure that *F*-ing was okay. But, since the fact that you are going to *F* doesn’t ensure that *F*-ing is okay, the fittingness view correctly predicts that it’s bad reasoning. (Similar points apply to other related patterns, such as ‘I intend to *F*, so, it’s okay to *F*’.)

We suspect that the former reasoning strikes us as bad when we confuse it with the latter. The former reasoning does not yield bootstrapping. Crudely put: if *F*-ing isn’t okay before you intend it, it won’t be okay afterwards either, so this reasoning will just lead to a false belief. Of course, the reasoning will lead to a true belief if *F*-ing *is* okay. But that’s as we should have expected.

6. Conclusion

We have argued that the fittingness view, according to which good reasoning preserves fittingness, other things equal, is preferable to the rational requirements view and the reasons view. We have also explored some ways in which the fittingness view might be developed—for instance, by understanding the ‘other things equal’ clause in terms of normality. Of course, many further issues remain—for instance, we have not considered how the fittingness view may be extended to reasoning which involves attitudes other than beliefs and intentions, or reasoning which concludes in dropping an attitude.³⁴ Nonetheless, we take the discussion to make a good *prima facie* case for the fittingness view and, thereby, to indicate that a unified theory of good reasoning is possible.³⁵

References

Boghossian, P. 2008. *Content and Justification*. Oxford: Oxford University Press.

—. 2014. What is Inference? *Philosophical Studies* 169, 1–18.

Broome, J. 1999. Normative Requirements. *Ratio* 12, 398–419.

—. 2013. *Rationality Through Reasoning*. Chichester: Wiley-Blackwell.

Dancy, J. 2004. *Ethics Without Principles*. Oxford: Oxford University Press.

—. 2014. From Thought to Action. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, Vol. 9. Oxford: Oxford University Press.

Enoch, D. and Schechter, J. 2008. How Are Basic Belief-Forming Methods Justified? *Philosophy and Phenomenological Research* 76 (3), 547–579.

Gibbons, J. 2013. *The Norm of Belief*. Oxford: Oxford University Press.

Grice, P. 2001. *Aspects of Reason*. Oxford: Oxford University Press.

Harman, G. 1999. *Reasoning, Meaning, and Mind*. Oxford: Clarendon Press.

Harty, J. 2012. *Reasons as Defaults*. Oxford: Oxford University Press.

³⁴ McHugh (2014) discusses reasoning to dropping attitudes.

³⁵ For very helpful comments on earlier versions of this paper, we would like to thank John Brunero, Alex Gregory, Frank Hoffman, Lee Walters, Daniel Whiting, several anonymous referees, and audiences at the Institute of Philosophy at the University of London, New York University Abu Dhabi, the University of Reading, and the University of Southampton. This work was supported by the Arts and Humanities Research Council (grant number AH/K008188/1).

Jenkins, C. 2014. Naturalism and Norms of Inference. In O. Flanagan and A. Fairweather (eds.), *Naturalizing Epistemic Virtue*. Cambridge: Cambridge University Press.

Kavka, G. 1983. The Toxin Puzzle. *Analysis* 43, 33–36.

Kolodny, N. 2005. Why Be Rational? *Mind* 114 (455), 509–563.

—. 2007. How Does Coherence Matter? *Proceedings of the Aristotelian Society* 107, 229–263.

Lance, M. and Little, M. 2007. Where The Laws Are. In R. Shafer-Landau (ed), *Oxford Studies in Metaethics*, Volume 2. Oxford: Oxford University Press.

Leslie, S.J. 2012. Generics. In G. Russell and D.G. Fara (eds.), *The Routledge Companion to Philosophy of Language*. New York: Routledge.

Littlejohn, C. 2012. *Justification and the Truth-Connection*. Cambridge: Cambridge University Press.

Lord, E. 2015. Acting for the Right Reasons, Abilities, and Obligation. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, Volume 10. Oxford: Oxford University Press.

McHugh, C. 2014. Fitting Belief. *Proceedings of the Aristotelian Society* 114, 167–187.

—. and Way, J. 2015. Broome on Reasoning. *Teorema* 34, 131–141.

—. and —. 2016. Fittingness First. *Ethics* 126, 575–606.

—. and —. ms 1. What is Reasoning?

—. and —. ms 2. Objectivism and Perspectivism about the Epistemic Ought.

Millikan, R. 1984. *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.

Nickel, B. 2010. Ceteris Paribus Laws, Generics, and Natural Kinds. *Philosophers' Imprint* 10 (6): 1–25.

Parfit, D. 2011. *On What Matters*. Oxford: Oxford University Press.

Pietroski, P. and Rey, G. 1995. When Other Things Aren't Equal: Saving Ceteris Paribus Laws from Vacuity. *British Journal for the Philosophy of Science* 46 (1): 81–110.

Pollock, J. 1987. Defeasible Reasoning. *Cognitive Science* 11: 481–518.

Scanlon, T.M. 1998. *What We Owe To Each Other*. Cambridge, MA: Harvard University Press.

Schroeder, M. 2004. The Scope of Instrumental Reason. *Philosophical Perspectives* 18, 337–364.

—. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.

—. 2011. What Does It Take To Have A Reason? In A. Reisner and A. Steglich-Peterson (eds.), *Reasons for Belief*. Cambridge: Cambridge University Press.

—. 2015. Knowledge is Belief for Sufficient (Objective and Subjective) Reason. In T. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, Volume 5. Oxford: Oxford University Press.

Setiya, K. 2014. What is a Reason to Act? *Philosophical Studies* 167 (2), 221–235.

Shah, N. 2008. How Action Governs Intention. *Philosophers' Imprint* 8, 1–19.

Skorupski, J. 2010. *The Domain of Reasons*. Oxford: Oxford University Press.

Smith, M. 2010. What Else Justification Could Be. *Noûs* 44, 10–31.

Smithies, D. 2012. The Normative Role of Knowledge. *Noûs* 46, 265–288.

Sylvan, K. forthcoming. What Apparent Reasons Appear To Be. *Philosophical Studies*.

Turri, J. 2010. On The Relationship Between Propositional and Doxastic Justification. *Philosophy and Phenomenological Research* 80 (2), 312–26.

Vogelstein, E. 2012. Subjective Reasons. *Ethical Theory and Moral Practice* 15 (2), 239–257.

Way, J. 2012. Explaining the Instrumental Principle. *Australasian Journal of Philosophy*.
— forthcoming 1. Reasons as Premises of Good Reasoning. *Pacific Philosophical Quarterly*.
— forthcoming 2. Two Arguments for Evidentialism. *The Philosophical Quarterly*.

Whiting, D. 2013. Nothing But The Truth: On The Norms and Aims of Belief.
In T. Chan (ed.) *The Aim of Belief*. Oxford: Oxford University Press.

— 2014. Keep Things in Perspective: Reasons, Rationality, and the A Priori. *Journal of Ethics & Social Philosophy* 8.

Williams, B. 1979. *Internal and External Reasons*. Reprinted in his *Moral Luck*.
Cambridge: Cambridge University Press, 1981.

Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.