

What is Reasoning?

Conor McHugh and Jonathan Way

Abstract

Reasoning is a certain kind of attitude-revision. What kind? The aim of this paper is to introduce and defend a new answer to this question, based on the idea that reasoning is a goodness-fixing kind. Our central claim is that reasoning is a functional kind: it has a constitutive point or aim that fixes the standards for good reasoning. We claim, further, that this aim is to get fitting attitudes. We start by considering recent accounts of reasoning due to Ralph Wedgwood and John Broome, and argue that, while these accounts contain important insights, they are not satisfactory: Wedgwood's rules out too much, and Broome's too little. We then introduce and defend our alternative account, discuss some of its implications and attractions, and, finally, consider objections.

1. The Question

This paper is about reasoning. Here are some examples of ways you might reason:

If Jane had a beer, then there are none left. Jane had a beer. So, there are none left.

I shall get beer. In order to get beer, I must go to the shop. So, I shall go to the shop.

In these examples, you revise your attitudes in a certain way. You bring some existing attitudes to mind, saying their contents to yourself, and make a kind of transition to a further attitude which you thereby acquire. But attitude-revision doesn't always count as reasoning. Your attitudes might change through brute causation of one attitude by another, or as a result of a bang on the head. In such cases you wouldn't be reasoning. The question we want to address is: what kind of attitude-revision is reasoning?

Our aim is to introduce and defend a new answer to this question, based on the idea that reasoning is a goodness-fixing kind. Our central claim is that reasoning is a functional kind: it has a constitutive point or aim that fixes the standards for good reasoning. We claim, further, that this aim is to get fitting attitudes. We introduce and motivate these claims in section 4, discuss implications and attractions in section 5, and consider objections in section 6. Before presenting our account, we consider recent accounts of reasoning due to Ralph Wedgwood and John Broome, in sections 2 and 3.

We argue that these accounts are not satisfactory: Wedgwood's rules out too much, and Broome's too little. Nonetheless, these accounts contain important insights, which ours builds on.

We begin with some clarifications of our question. First, by 'reasoning' we mean personal-level, conscious, active reasoning. We do not mean subpersonal, unconscious, or automatic information-processing. As Boghossian (2014) puts it, making reference to the well-known distinction between 'System 1' and 'System 2' processes, reasoning in our sense is "System 1.5 and up" (p. 2). Second, although it is convenient to represent bits of reasoning by how you might express them in language, as above, we don't assume that reasoning requires you to engage in speech, inner or outer. Third, as our examples indicate, we will discuss not only theoretical reasoning (reasoning to beliefs) but also practical reasoning (reasoning to intentions). We take there to be good grounds for supposing that reasoning is at some level a unified phenomenon. We are sympathetic to the idea that you can reason to other attitudes besides belief and intention, but we won't assume that here. Fourth, for simplicity we will mostly use examples of reasoning that conclude in the acquisition of a new attitude. Reasoning can also lead to retaining an attitude, or dropping or refraining from forming one (Harman 1986). This will come up once or twice where it is important, but for the most part we will leave it implicit. Lastly, some terminology: following Broome (2013), we will call the attitudes you reason from 'premise-attitudes', and the attitude you reason to the 'conclusion-attitude'. When these attitudes are beliefs, we will call their contents 'premises' and 'conclusions' respectively.

2. Responding to Normative Support

Reasoning has normative significance. It can get you to have new attitudes that are normatively supported by attitudes you already have. So we might think that this relation of normative support is what makes attitude-revision count as reasoning, when it does. Of course, one attitude might cause another that it *happens* to support, without any reasoning taking place, through some non-rational process. By contrast, we might think, in *reasoning* you form a conclusion-attitude *because* it is supported by your premise-attitudes. This sort of account has been defended by Ralph Wedgwood. On Wedgwood's view, to form a belief or intention through reasoning is to "respond directly to the fact that one has come to be in some antecedent mental states or other that rationalize forming [that] belief

or an intention...by forming that very belief or intention” (Wedgwood 2006, 673).¹ In a slogan, reasoning is *attitude-revision in response to rationalization*.

What is rationalization? Wedgwood treats it as *making rational*. Alternatively, it might be *support by reasons*. In that case the view would say that reasoning is a kind of responding to normative reasons.

These would be two possible ways of developing a view of this sort.

Either way, the obvious shortcoming of this view is that it doesn’t account for bad reasoning. Suppose you reason:

If Jane had a beer, then there are none left. There are none left. So, Jane had a beer.

Here, your premise-beliefs do not make your conclusion-belief rational; nor do they give or correspond to reasons for it. But clearly it is possible to reason in this and other bad ways. So responding to rationalization, whatever this amounts to, is not necessary for reasoning.²

Wedgwood is aware of this. He says that his account should be understood as an account of “rational” (good) reasoning, suggesting that it can “cast light on the general nature of reasoning” if “fallacious reasoning is *parasitic* on rational reasoning” (662). However, our question is what reasoning is, not what good reasoning is. So, Wedgwood’s account can help us only if we know in what way bad reasoning is parasitic on good reasoning. Unfortunately, he doesn’t tell us. One thing he says is that even when you reason badly it is “intelligible” that your premise-attitudes could “represent [your] reason for forming” the conclusion-attitude (ibid.). Here, ‘your reason’ means a *motivating* reason - the reason *for which* you form the conclusion-attitude. But, when you reason, your premise-attitudes presumably *do* represent or correspond to your motivating reason(s) for forming the conclusion-attitude. Our question is how this happens. It doesn’t help to be told that its happening must be intelligible.

¹ Wedgwood adds that this response must manifest a certain disposition, if it is to count as reasoning.

² Wedgwood might claim that reasoning involves responding to some *degree* of rationalization, even if not a high degree (cf. Arpaly and Schroeder 2014, 72ff.). But, while some fallacious reasoning might be like that, it seems clear that much of it isn’t; when you make a mistake such as affirming the consequent or the gambler’s fallacy, your premise-attitudes need not provide any rationalization whatsoever for the conclusion-attitude. Admittedly, some transitions seem so transparently bad that it’s hard to make sense of them as reasoning at all. We will provide a better explanation of this phenomenon below.

Nonetheless, this suggests a natural way of generalising Wedgwood's account. When you reason well, he might say, you respond to normative reasons. When you reason badly, you respond to what you mistakenly *take* to be, or *treat* as, normative reasons. It is this *taking* or *treating* that makes a process of attitude-revision count as reasoning, in the bad case.³

The suggestion, in effect, is that reasoning is simply a kind of *basing* - of forming attitudes for motivating reasons. In the good case, these motivating reasons are normative reasons. In the bad case, they are merely taken to be, or treated as, normative reasons. This suggestion has *prima facie* plausibility, since it's clear that reasoning and basing are closely connected. And it preserves a sense in which the good case is prior to the bad one, as Wedgwood claims. Nonetheless, there are serious problems.

First, it doesn't get us very far. Basing is as much in need of explication as reasoning. Many of the problems discussed here regarding reasoning have analogues for basing. For example, what's the difference between an attitude's being based on another and its merely being caused by it? It's not that, in the former case, you form the attitude because you take there to be a normative reason for it: the causation of the attitude might still be of the wrong kind to constitute basing. Nor is it that, in forming the attitude, you manifest a disposition to form attitudes of that kind when you take there to be such a reason for them. This disposition might be a non-rational one; you might have had a brain haemorrhage that makes you compulsively form attitudes in such circumstances.

Secondly, it is highly controversial whether reasoning necessarily involves taking your premises to be normative reasons for your conclusion. On the natural assumption that 'taking' is just believing, such a condition leads to well-known circularity and regress worries,⁴ and arguably over-intellectualises reasoning. These worries might be avoided by a view which appealed to *treating* as a reason. But it is not clear what treating something as a reason amounts to, if not *reasoning* from it. So such a view would face its own problem of circularity.

³ Grice (2001) defends such a view.

⁴ See Boghossian (2014), Broome (2014), Wright (2014), Setiya (2013) and McHugh and Way (2016 a) for discussion.

Perhaps there is another way of generalising Wedgwood's account, to accommodate both good and bad reasoning. Until this is put forward, however, we cannot accept Wedgwood's view.⁵

3. Rule-Following

Reasoning isn't always good, but it is always constrained. You can't just reason any way you want. This makes it natural to think that reasoning constitutively involves *following rules*. This hypothesis is further supported by the observation that, when we reason, we respond to *general* features that our reasoning instantiates. In the examples given at the start of the paper, you presumably don't reason the way you do just because of the specifics of the situation with Jane and the beer. Rather, you do so because of some more general *rule* or *pattern* that this reasoning falls under. For example, you might follow these rules:

(*Modus Ponens* Rule) From believing *p*, and believing if *p* then *q*, move to believing *q*.

(Means-End Rule) From intending to *E*, and believing in order to *E* you must *M*, move to intending to *M*.

The idea that reasoning is a kind of rule-following also meshes with the plausible thought that *good* reasoning can be understood in terms of good *rules* or patterns of reasoning.⁶ On such a view, it is natural to think that even when you are reasoning incompetently, you are still following rules – just the wrong ones.⁷ So, we might think of reasoning as *rule-following attitude-revision*.

As it stands, this is too weak. There are rules for attitude-revision following which wouldn't count as reasoning, like the rule of forming as many beliefs as possible, or of abandoning all of your intentions.

⁵ For further discussion see Neta (2013). Wedgwood (2012) offers what might be taken as an alternative account of theoretical reasoning, on which you reason from the belief that *p* to the belief that *q* when you believe *p* unconditionally, form a belief in *q conditional* on *p*, and come to believe *q* unconditionally in response to this. However, this account helps itself to a notion of forming one belief 'in response to' another, and is thus no more informative than the account of reasoning in terms of basing.

⁶ Horty (2012), Pollock (1987), McHugh and Way (forthcoming).

⁷ Or perhaps misapplying the right rules. For further motivation of a rule-following account see Boghossian (2014).

Perhaps, then, only certain kinds of rules count. According to John Broome (2013), they must be *rules for operating on contents*. Broome describes reasoning as “a rule-governed operation on the contents of your conscious attitudes” (234). More fully:

“Active reasoning is a particular sort of process by which conscious premise-attitudes cause you to acquire a conclusion-attitude. The process is that you operate on the contents of your premise-attitudes following a rule, to construct the conclusion, which is the content of a new attitude of yours that you acquire in the process” (ibid.).

For example, you might operate on the believed content ‘If Jane had a beer, there are none left’, and the believed content ‘Jane had a beer’, to construct the content ‘There are no beers left’, thereby coming to believe the latter.

Three important clarifications are in order. First, as Broome acknowledges, the rules you follow in reasoning can’t appeal *only* to contents. They must also appeal to attitude-types. There’s a difference between reasoning to a belief from other beliefs, and reasoning to that belief from, say, imaginings or wishes with the same contents. So the suggestion is really that reasoning is following rules for operating on content-attitude *pairs*, or what we are simply calling ‘attitudes’.

Second, rule-following must not be conceived as itself involving reasoning. On one view, following a rule, R, involves intending to conform to R, believing that ϕ ing would constitute conformity to R, and reasoning from these attitudes to the intention to ϕ . If we understand reasoning in terms of rule-following, then this *intentional account* of rule-following is ruled out on pain of circularity (Boghossian 2014). What seems required instead is a *dispositional account*, according to which following a rule is a matter of manifesting a certain sort of disposition. Broome endorses such an account (loc. cit., Ch. 13).

Third, this account faces a crucial question: what is it to ‘construct’ or ‘derive’ an attitude? It would be natural to understand this as *reasoning* to the formation of the attitude. But this leads to circularity: the account would say that reasoning consists in following rules for reasoning. Perhaps, then, we can say that the construction or derivation of a conclusion-attitude is simply its being *caused* by the

premise-attitudes. But in that case Broome's rules will simply tell you to allow certain attitudes to cause certain others. Following rules of this sort does not seem sufficient for reasoning - at least, not given a dispositional account of rule-following.⁸ So there is a difficulty here for Broome's account. We won't press this difficulty, but we do want to emphasise that, when assessing whether Broome's conditions are sufficient for reasoning, it's important not to understand rule-following as intentional and not to read 'construct' or 'derive' as synonyms for 'reason to'. To avoid confusion we will just talk about moving from one attitude to another, as in the formulation of the *Modus Ponens* Rule above. With these clarifications in place, we turn to our main criticism of Broome's account. It is this: following rules for operating on attitudes seems insufficient for reasoning. Consider these rules:

- (R1) For any p, move from p occurring to you to believing p.
- (R2) For any p, move from believing p to intending to murder everyone.
- (R3) Move from any attitude to the proposition that there's a spider around, to dropping that attitude. (An arachnophobe might try to follow this rule in order to avoid phobic episodes.)

It seems clear that you could follow these rules without doing any reasoning.

We expect that Broome would want to resist the claim that one could follow rules R1-R3 without reasoning. He considers another example of a bizarre transition:

- (R4) Move from believing it is raining, and believing if it is raining the snow will melt, to believing you hear trumpets (ibid., 233).

⁸ According to Broome (2014), manifestations of a disposition constitute rule-following when they are also disposed to 'seem right' to you, where a result's seeming right involves the possibility of checking whether it is right, for example by running through the same process again. (In his 2013, 237ff., Broome stops short of claiming that these conditions are sufficient for rule-following, but he does not say what else might be needed.) An attitude can be caused by another in a way that is disposed to occur and to seem right to you in this sense, without any reasoning occurring.

With regard to R4, Broome insists:

“If you derive this conclusion by operating on the premises, following the rule, we should count you as reasoning” (ibid.).

This seems wrong to us. Of course, if rule-following itself involves reasoning, or if ‘derive’ just means ‘reason to’, then, trivially, you could not follow R4 without reasoning. But, as we saw, rule-following must not be construed in this way, and the relevant sense of ‘derive’ must be a process that by itself falls short of reasoning. When we keep these points in mind, it seems clear that following R4 is not sufficient for reasoning. It’s not that in following R4 you *couldn’t* be reasoning. It’s that you *needn’t* be. For example, following R4 seems compatible with knowing that the weather has nothing to do with whether you hear trumpets.⁹ It’s hard to see how attitude-formation of this sort could be sufficient for reasoning. The point is even clearer with R1-R3.¹⁰

Broome might claim that having a disposition to make the transition in R4 *is* a way of implicitly taking facts about the weather to support conclusions about whether you hear trumpets (cf. Broome 2014). But that claim doesn’t seem justified. A disposition to move from one belief to another is just a different thing from taking there to be a relation of evidential support between their contents.

In sum, effecting some rule-governed operation for moving between attitudes isn’t sufficient for reasoning. Broome might reply that his conditions are nonetheless necessary and illuminating. That may be so, and indeed we are happy to accept as much, but we think more can be said. In the following sections, we offer a further centrally important condition on reasoning. Since we accept Broome’s claim that reasoning is a kind of rule-following, we will develop our proposal in the context of this assumption. The assumption is inessential, however: our proposal is independent of the rule-following conception.

⁹ It might seem that this is ruled out by Broome’s ‘seeming right’ condition on rule-following (n. 8). But, as Broome makes clear, the relevant notion is seeming right *relative to a rule* - in effect, seeming to conform to a rule. Clearly, a transition could seem to you to conform to R4, while you know that the weather has nothing to do with whether you hear trumpets.

¹⁰ The account we offer below will explain when following R4 or other such rules would count as reasoning and when it wouldn’t. For now, the crucial point is just that following rules for operating on attitudes is not, as Broome claims, sufficient for reasoning.

4. Reasoning as a Goodness-Fixing Kind

It is tempting to say that the problem with R1-R4 is that they are obviously *bad* rules. We think there is something right about this. But, as we saw with Wedgwood's view, we can't say that reasoning is a matter of following the *right* rules, on pain of failing to account for bad reasoning. In this section we will outline and provide some initial motivation for a view of reasoning that explains why following obviously bad rules is typically not reasoning. In the following sections we will indicate some of the view's attractions and address objections. We will not provide a decisive argument for the view, but we hope to do enough to show that it has serious promise.

4.1 Goodness-Fixing Kinds and Constitutive Points

It is a platitude - but, we think, a significant one - that there is such a thing as *good* reasoning. This looks like an attributive rather than predicative use of 'good' (Geach 1956). For a transition to be good reasoning is not for it to be both good *simpliciter* and reasoning.

When we talk of 'a good *F*' in the attributive sense, this is sometimes an abbreviation: we might mean that it is good *for* such and such. For example, 'That's a good pebble!', might mean that it is a good pebble for skimming. Other times, it is not an abbreviation: by 'a good *F*', we mean something that is good *qua F*.¹¹ The standard for being good *qua F* is fixed by what it *is* to be an *F*. When there is such a standard, *F* is a 'goodness-fixing kind' (Thomson 2008). The ordinary use of 'good reasoning' clearly falls into this latter category. To say that someone is engaged in good reasoning is not to say that their reasoning is good for this or that contingent end, but that it is good *qua* reasoning. And it is highly plausible that what counts as good reasoning has to do with what reasoning *is*. That is, reasoning is a goodness-fixing kind.

Many goodness-fixing kinds are functional kinds: they have a constitutive function, point or aim, that helps fix the standard for being good of that kind. For example, there's such a thing as a good knife, and what counts as a good knife is determined by the constitutive function of knives, namely to cut.

¹¹ For some other uses of 'a good *F*', not relevant here, see Szabó (2001).

Activity-kinds can be like this too. Building a house has a constitutive aim: to provide shelter (Korsgaard 2009, 29).¹² It is this aim that helps fix the constitutive standards for building a house well. Are there any goodness-fixing kinds that are not functional kinds in this sense? Thomson (2008) proposes some examples: beefsteak tomato, human being, and tiger. We don't find it obvious that there's such a thing as being good *qua* any of these things - as opposed to being, say, a good beefsteak tomato *for eating*, or a *morally* good human being. But in any case, if these are goodness-fixing kinds, they seem different to the case of reasoning. They are all kinds of organisms, or parts thereof. The standards for being good of these kinds, if there are any, are plausibly fixed by their nature as organisms - by what it takes for them to flourish, or the like. Reasoning isn't an organism and can't flourish.

A natural hypothesis, then, is that reasoning, like knives and house-building, is a functional kind. On this hypothesis, it is in the nature of reasoning that it has a certain point or aim. This point or aim is what fixes the constitutive standards of good reasoning. In what follows we will develop the hypothesis that reasoning is a functional kind. Further support for it will come from its fruits.

4.2 What's the Point?

We are proposing that the standards for good reasoning derive from the constitutive point of reasoning. What *is* this point? As we saw earlier, when reasoning is going well it can get you attitudes with normatively significant properties. This suggests that the point of reasoning might be to get attitudes with one of these properties. Which property?

Some possible answers are suggested by the views discussed in previous sections. Perhaps reasoning aims to get you attitudes that are rational, or that are supported by reasons.¹³ There are problems for both of these suggestions, though. Briefly, neither of them seems to explain why reasoning might be a worthwhile way to go about revising your attitudes: satisfying rational requirements does not seem valuable for its own sake and need not be a means to anything else worth doing (Kolodny 2007);

¹² These specifications of the point of knives and house-building are overly simple, but suffice for present purposes.

¹³ Broome (2013) conceives of reasoning as a means to rationality; cf. also Harman (1986). We don't know of an explicit defence of the view that reasoning is for getting support by reasons, but it is a natural idea.

attitudes that are supported by some reason are easy to come by without reasoning. Both views, moreover, have difficulties explaining how the standards for good reasoning might derive from the aim of reasoning. Many transitions that would bring your attitudes into conformity with rational requirements are not good reasoning, and good reasoning can start from premises that are false, and thus that do not give reasons for anything. We have developed these and further arguments in detail elsewhere (McHugh and Way 2015, forthcoming). Here, we will simply motivate and state our own view. Again, the case for the view is partly constituted by its theoretical fruits, some of which will be outlined below. Moreover, much of what we will go on to say depends mainly on our claim that reasoning is a functional kind, rather than on our claim about what the point of reasoning is. Those with alternative views of the point of reasoning can still accept our overall approach.

Consider theoretical reasoning. The rules of good theoretical reasoning clearly have something to do with preserving truth - with leading from true beliefs to further true beliefs. This is especially obvious in the case of good deductive reasoning. But it is almost as obvious in the case of defeasible reasoning. For example, while good reasoning from testimony can sometimes lead you from truth to falsity, it is nonetheless the case that *normally*, or *other things equal*, if you put true beliefs into such reasoning, you will get true beliefs out.

True beliefs have a normatively significant property: they are correct, right, or, as we will say, *fitting*. So, in preserving truth, good theoretical reasoning also preserves fittingness. Very plausibly, this is, in the first instance, what *makes* such reasoning good. In turn, it is plausible that the *point* of theoretical reasoning is to acquire beliefs which have the normative property of fittingness; theoretical reasoning aims at truth because truth is what counts as getting things right in belief.¹⁴

Other attitudes have standards of fittingness too. When you admire Mandela, desire a vacation, or fear an onrushing tiger, 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. The core idea here is that attitudes are associated with standards for their objects. An attitude is fitting when its object meets its standard. ‘Fitting’ is something of a

¹⁴ Not every fittingness-preserving rule is a good rule of reasoning. For example, any rule whose conclusion-response is a necessarily true belief trivially preserves fittingness. We discuss further conditions on good rules of reasoning elsewhere (McHugh and Way forthcoming). The details will not matter here.

term of art. As noted above, we can also talk of *correctness* or *getting it right*; and we can say of an object that it is *fit* to be, say, admired, or *worth* admiring, or an *appropriate object* of admiration.¹⁵ Most pertinently for our purposes, intentions can be fitting or unfitting; you can be getting it right or wrong in having them, depending on whether what you intend is choiceworthy or not. It is a substantive normative question which actions are choiceworthy, and thus which intentions are fitting, but that we can assess intentions in this way shouldn't be controversial. For example, the intention to phone your mother on her birthday is fitting, whereas the intention to murder everyone is unfitting. The suggestion that reasoning aims at fittingness can thus be applied to practical reasoning. And it seems to be borne out by the good rules of practical reasoning: like good rules of theoretical reasoning, these rules preserve fittingness. Consider the Means-End Rule. This preserves fittingness: if it is fitting to intend the end, and fitting to believe (because true) that the means is necessary, then it will be fitting to intend the means. For example, if it's fitting to intend to get beer, and it's true that in order to do so you must go to the shop, then it's fitting to intend to go to the shop. Thus, it is plausible that the point of practical reasoning is to get you fitting intentions. In general then, we suggest that the ultimate point of reasoning is to get fitting attitudes. In other words, it is to *get things right*.¹⁶

4.3 What Sort of Point?

We have argued that reasoning is a goodness-fixing kind, and more specifically a functional kind. We have also argued that the point of reasoning is to get fitting attitudes. But what is involved in reasoning's having this point?

There are several ways for things to have a point. Knives have the function of cutting in that they are designed to cut. Hearts have the function of pumping blood in that, roughly, they were selected to pump blood. As we saw earlier, activities can also have a point. Sometimes this point is given by an aim that regulates the activity: for example, house-building has the aim of providing shelter in that only activities that are regulated by an aim of providing shelter count as (intentional) house-building.

¹⁵ This notion of fittingness is most commonly associated with the fitting-attitudes theory of value famously proposed by Brentano (1889/2009), Broad (1930), Ewing (1948) and others. More recent discussions of the same notion include D'Arms and Jacobson (2000), Thomson (2008), Chappell (2012), Svavarsdóttir (2014), McHugh (2014), McHugh and Way (2016 b), Howard (ms.).

¹⁶ For further discussion, see McHugh and Way (forthcoming).

Sometimes, however, the point comes from somewhere else. Consider playing chess. In a sense, the point of this activity is to checkmate your opponent; it is this point that fixes the standards for good chess moves. However, playing chess need not be *regulated* by this point. You can play chess without trying to win. Rather, the point of chess comes from the conventions that constitute the game. These conventions stipulate not only how the pieces move, but also the point, i.e. what counts as winning. In going in for chess, you intentionally place yourself under the standards fixed by the point of checkmating your opponent, but you need not aim to satisfy them.

Like house-building and chess, reasoning is a personal-level activity. Unlike chess, the constitutive standards of good reasoning are not fixed by convention. If these standards derive from reasoning's point, this point is not conventional. This suggests that reasoning is like house-building: it has an aim in the strong sense that only activities regulated by this aim count as reasoning. Several observations further support this idea. Firstly, when you reason, you can't be indifferent to the standards of good reasoning. You cannot reason in any way you like. This would be expected if an activity had to be regulated by the aim that generates these standards in order to count as reasoning. Secondly, there are some cases of reasoning for which it seems clearly true that they are regulated by the aim of getting things right. In paradigm cases of active theoretical and practical reasoning, you reason about a certain question – whether p or whether to ϕ .¹⁷ In such cases your reasoning is clearly undertaken with and guided by the aim of reaching a true belief or a correct decision. Thirdly, even in the less paradigmatic cases where you don't formulate a question, reasoning is nonetheless *directed*. Its directedness distinguishes reasoning from other forms of cognition, such as merely drifting in thought; it explains the particular path that your thoughts take, why you stop where you do, and so on.¹⁸ What gives your thought direction is, plausibly, an aim.

We therefore propose that reasoning is indeed a functional kind in the strong sense that only activities regulated by its aim count as reasoning. This hypothesis captures what's plausible in the views we discussed earlier, while avoiding their problems. Wedgwood is right that we can get an insight into what reasoning is by thinking about good reasoning. But our account, unlike his, explains why bad

¹⁷ Cf. Hieronymi (2013). We don't think that reasoning can be adequately characterised in terms of answering questions. You can answer a question without aiming to get things right.

¹⁸ Cf. Shah (2003, 466).

reasoning is nonetheless reasoning: what you're doing can be regulated by an aim even though you're not serving it well. Broome is right that you can reason without doing it well, and he's right that reasoning involves rule-following, but he's wrong to think that following rules for operating on attitudes suffices for reasoning. R1-R4 are rules you could be following without your activity being in any way regulated by the point of reasoning. You might just be having a laugh, or dealing with your arachnophobia. Again, we don't want to say that you *couldn't* be reasoning in following these rules; for R4, at least, perhaps you could. The point is rather that it would depend whether, in following them, you were aiming to get fitting attitudes. As most naturally imagined, you wouldn't be.

4.4 How does the Point Regulate?

The idea that reasoning is aim-directed raises a worry. Sometimes you pursue an aim by intending to attain it, believing that you can do so by ϕ ing, and reasoning from these attitudes to an intention to ϕ . But since this conception of guidance by an aim, like the intentional model of rule-following, presupposes reasoning, we cannot draw on it here.

If our account is to be tenable, it must be possible to be guided by an aim without reasoning. And, in general, this is possible. Whenever you execute an intention, you are guided by an aim, but this need not involve reasoning. However, it might be said that this is only because you do not always need to take a *means* to executing an intention: some actions are basic. According to our proposal, the transitions you make in reasoning *are* made as a means to an end: that of getting fitting attitudes. So we still need to say something about how this sort of aim-directed behaviour can occur without prior reasoning.

We suggest that, although the transitions made in reasoning are made as a means to an end, the way in which they are guided by this end is not via a means-end belief, and thus does not involve reasoning. In general, agents must be able to be sensitive to certain conditions without representing those conditions as obtaining. Our suggestion is that agents can be sensitive to fittingness-preservation in reasoning without representing their reasoning as fittingness-preserving. They thereby count as aiming at fittingness. In particular, we suggest that the point of reasoning guides through the rules of reasoning that you follow. In following the rules that you follow, you manifest your imperfect

sensitivity to what will serve the aim of getting fitting attitudes. For example, in following the *Modus Ponens* Rule, or the Means-End Rule, a reasoner manifests their sensitivity to the fact that these patterns of reasoning preserve fittingness. Such facts are not themselves premises of the reasoning, any more than the rules themselves are. As Lewis Carroll (1895) taught us, not everything that guides you in reasoning is a premise. But if you are not sensitive to the fittingness-preserving character of a given rule, then, in so far as you aim at fitting attitudes, you will not be willing to follow it.

Of course, things can go wrong, even when your premise-attitudes are fitting. You can misapply a good rule: a transition that doesn't in fact accord with the rule can nonetheless trigger your disposition to follow it. In that case you're still following the rule, but incorrectly.¹⁹ Or you can correctly follow a rule that does not preserve fittingness. For example, you might reason by following the rule of affirming the consequent. In that case you're still being guided by the aim of getting fitting attitudes. It's just that your sensitivity to what will serve this aim is imperfect.

We think it is useful to understand your following rules in reasoning as a way of following a more general rule:

(General Rule) In revising your attitudes, get fitting attitudes and not unfitting ones!

The General Rule is not to be understood as enjoining the acquisition of as many fitting attitudes as possible. Rather, it enjoins you, in revising your attitudes, to do so such that the attitudes you get are fitting and not unfitting. Nor does the rule tell you to revise your attitudes in the first place. Rather, it applies to you only when you are revising your attitudes; if you never do this, you neither comply with nor violate the rule.

Note that following a rule that tells you to *F* is a way of aiming to *F*. Satisfying a rule is neither necessary nor sufficient for following it, but when you follow a rule you *try* to satisfy it. Consider Boghossian's (2008):

(Email Rule) Answer any email that calls for an answer immediately upon receipt!

¹⁹ For the point that it is possible to follow a rule without complying with it, see Boghossian (2008, 480).

If you follow this rule, then you aim to answer emails that call for an answer immediately. When you overlook one, this is not just a violation of the Email Rule, but a *failure*: you have failed to do what you were aiming to do. In contrast, someone who does not follow the Email Rule will also violate it, but this violation will *not* be a failure on their part, because they weren't trying to satisfy it.²⁰

We can thus say, roughly, that reasoning is attitude-revision that follows the General Rule. In putting things this way, we assume that reasoning involves rule-following, as Broome claims. This assumption is not essential, however: our core claim, that reasoning is a functional kind aiming at fitting attitudes, is independent of the rule-following conception, not a mere refinement of it. And it is important and fruitful in its own right (or so we are arguing). An account that omits it fails to capture the essence of reasoning.

Nonetheless, the formulation in terms of rule-following is useful. One reason it is useful is that, independently of anything argued for here, it's very plausible that reasoning involves rule-following.²¹

So there must be a way of following a rule that does not presuppose reasoning (Boghossian 2014). As noted above, a full account of such rule-following is likely to appeal to certain sorts of dispositions. What we are suggesting, in effect, is a dispositional account of the aim-directedness of reasoning. While we cannot spell out this account fully here, the point that reasoning plausibly involves rule-following gives grounds for optimism about the possibility of doing so without making our account of reasoning circular.

Still, one might have worries. We claim that in following ordinary rules of reasoning, such as the *Modus Ponens* or Means-End Rules, reasoners are also following the higher-level General Rule. Even if this claim does not lead to circularity, one might wonder: what grounds are there for thinking it true? Answering this question will help to fill out and further support the account.

Note that, if you *were* sensitive to fittingness, rules like the *Modus Ponens* and Means-End Rules are just the sorts of rules you would tend to follow (more on this below). Of course, people also sometimes follow bad rules, like the rule of affirming the consequent. But this is because our

²⁰ This is not to endorse Boghossian's claim that anyone following the Email Rule takes the receipt of an email to be a *reason* to answer it immediately.

²¹ Though see Valaris (forthcoming) for an opposing view.

sensitivity to what will preserve fittingness is imperfect.²² And notice that a reasoner who sees that affirming the consequent does not preserve truth will not be indifferent to this fact. They will stop using this rule, at least when they bear this fact in mind. Thus, even bad reasoning dispositions, which are not sensitive to fittingness, are nonetheless typically regulated by higher-order dispositions that are so sensitive. This is not to say that our higher-order dispositions are always good. But there do seem to be limits to the possibility of reasoning in ways that are recognised as not fittingness-preserving. For example, it's hard to make sense of a reasoner continuing to affirm the consequent, aware that she is doing so, while recognising that this pattern of reasoning is fallacious. This does not seem to be genuine reasoning.

Not just our bad reasoning dispositions, but also our good ones – our dispositions to conform to the rules of good reasoning – are regulated by higher-order dispositions that suggest an overarching concern with fittingness. These include higher-order dispositions regarding which rule to follow when two rules conflict, our dispositions *not* to use these ordinary rules when they will lead to unfitting attitudes, and so on. Many sorts of cases illustrate this. Here are two.

Defeat. Many of the rules we follow in reasoning are defeasible. Examples might include:

(Testimony Rule) From believing someone said p , move to believing p .

(Promise Rule) From believing you promised to ϕ , move to intending to ϕ .

You might believe that Joe said p , but also believe that Joe is a liar. In that case, you will refrain from applying the Testimony Rule and coming to believe p . You might believe that you promised to meet Jenny for lunch, but also believe that by doing so you would leave someone to die. In that case, you will refrain from applying the Promise Rule and coming to intend to meet Jenny for lunch. Why? Because these are cases where following these rules will lead to unfitting attitudes.

²² For the point that a dispositional account of rule-following can accommodate dispositions to misapply rules, see Forbes (1983).

Backwards reasoning. We have been focusing on reasoning that concludes in the acquisition of an attitude. But, as noted earlier, reasoning can conclude in other ways. Suppose you believe p and if p then q , but when you consider q it strikes you as very improbable. Then, instead of applying the *Modus Ponens* Rule and coming to believe q , you might instead drop one or both of the other beliefs. Similarly, if you intend to E and think that in order to E you must M , but M -ing strikes you as a very bad idea, you will probably drop your intention to E rather than reasoning to the intention to M . Why, in these cases, do you not apply the *Modus Ponens* and Means-End rules? Because you see that doing so will lead you astray. You see that something has gone wrong with your existing attitudes, and that the way to make your attitudes fitting is to revise these attitudes rather than to follow where they lead.

Our overall pattern of reasoning dispositions thus strongly suggests an overarching, if implicit, concern with getting fitting attitudes. And this concern is not a contingent matter. It's not as though we just happen to reason in these ways. To follow our ordinary rules *wherever* they lead, even to belief in propositions we are antecedently certain are false, or to intentions for courses of action we think disastrous, would not be recognisable as reasoning.

That completes our initial presentation and motivation of our account of reasoning. In a slogan, reasoning is rule-following that aims at fittingness. In the next section, we will briefly discuss some implications of this account, and in doing so add to the case for it. Then we will turn to objections.

5. Implications and Attractions

5.1 Basing and Deviant Chains

We saw earlier that basing and reasoning are closely related. Reasoning is a paradigmatic way of coming to base one attitude on another, and it is plausible that any state of basing corresponds to reasoning that the subject is disposed to do. The nature of reasoning will thus constrain possible basing relations. You don't count as basing one attitude on another unless your holding the attitude is appropriately tied to a disposition to make the relevant transition in a way that is guided by the aim of

getting fitting attitudes. That's why an attitude formed because of a compulsion to form such attitudes in the presence of certain others does not thereby count as based on those other attitudes.²³

5.2 Reasons

Suppose that normative reasons are suitable for reasoning from: they are suitable premises of reasoning.²⁴ Given this, our view of reasoning suggests a plausible account of what it takes to be a normative reason, and of the connection between normative and motivating reasons.

The point of reasoning is to get fitting attitudes. The best way to get a fitting attitude is to start from fitting attitudes, including true beliefs, and take steps that keep you on track – that is, fittingness-preserving steps. Indeed, it's not clear how else we could go about it, given that we lack god-like direct access to which attitudes are fitting. And reasoning has to be something whose standards can guide us; it would be no good if the standards just told us to form fitting attitudes, without telling us how to do that. So, given the nature of the exercise, suitable premises of reasoning to an attitude are truths from which you can make fittingness-preserving transitions to that attitude. So a normative reason for an attitude is a truth from which you can make a fittingness-preserving transition to that attitude.²⁵

The motivating reasons for which you hold attitudes are determined by how those attitudes are based: they are premises of reasoning to those attitudes, reasoning you either do or are disposed to do. So, if the premises of reasoning should serve the point of reasoning, then your motivating reasons should be normative reasons. That is, they should be truths from which you can make fittingness-preserving transitions to the attitudes held for those reasons. On this view, we can say that normative reasons are good motivating reasons.²⁶ This explains why the terms 'normative reason' and 'good reason' can be used almost interchangeably: the latter means 'good *qua* motivating reason'.

5.3 Aims of Attitudes

²³ One might wonder whether beliefs based on perception are counter-examples to these claims. We discuss perceptual belief below.

²⁴ Setiya (2014), Way (forthcoming).

²⁵ We defend such an account of normative reasons further in McHugh and Way (2016 b).

²⁶ Cf. Gregory (forthcoming).

It is sometimes claimed that attitudes like belief have aims or functions corresponding to their fittingness conditions (e.g. Velleman 2000). As is sometimes remarked (Wedgwood 2003), this is *prima facie* puzzling: beliefs do not literally aim at anything. The present account explains what's right in this claim: it's reasoning to these attitudes that has this aim. For example, belief aims at truth in this sense: forming a belief through reasoning is forming it in a way that is regulated by the function of getting fitting, and therefore true, beliefs. This view accommodates the observation that the factors that influence belief outside of reasoning need not be truth-directed.²⁷ And, if basing is closely tied to reasoning, it can help explain why such influence doesn't amount to believing for non-evidential reasons.

5.4 Fittingness

As we noted earlier, the idea that attitudes have standards of fittingness, or correctness, is a familiar one. Still, one might wish for further elucidation of fittingness. While a full account of it is beyond the scope of the present paper, we think that the account of reasoning proposed here sheds light on it. In particular, the account sheds light on how attitudes get to have standards of fittingness.²⁸

Plausibly, holding an attitude of a particular type involves being disposed to use it as a premise-attitude in specific ways of reasoning. For example, the belief that p involves a disposition to reason towards further beliefs with contents that are supported by p together with other believed contents. The intention to ϕ involves a disposition to reason towards intentions to take believed means to ϕ ing, intentions not to do things believed incompatible with ϕ ing, and so on.

The point of this reasoning is to get fitting attitudes. Whether it will tend to achieve this goal depends systematically on the features of the *object* (or content) of the attitude from which you reason. Which features? It depends on the further attitudes to which this reasoning leads. When you reason in one of these ways from the belief that Jane had a beer to further beliefs, these further beliefs will tend to be true if Jane did indeed have a beer, but not otherwise: if she didn't have a beer, then it would be a lucky accident if the further beliefs reasoned to were true. In general, whether reasoning from the

²⁷ Shah (2003). Elsewhere, Shah (2013) suggests an account of reasoning similar in some respects to ours.

²⁸ See the works cited in n. 15 for further discussion.

belief that p tends to satisfy its point depends on whether p is true. Similarly, when you reason in one of these ways from the intention to go to the shop to further intentions (e.g. the intention to put on your shoes), these further intentions will tend to be fitting if the initial intention is, but not otherwise: if it's not fitting to intend to go to the shop, then it would be a lucky accident if the further intentions reasoned to were fitting. In general, whether reasoning from the intention to ϕ tends to satisfy its point depends on the features of ϕ ing that make intending to ϕ fitting.

So an attitude-type's contribution to reasoning fixes certain features that an object must have, in order for reasoning from an attitude of that type, with that object, to tend to satisfy its point. Since having such an attitude *is* in part being disposed to reason in these ways, this yields a sense in which it is appropriate to have the attitude only when its object has those features. We suggest that this appropriateness just is fittingness, and that the features in question thereby constitute the fittingness condition for that type of attitude.

This account predicts that attitudes whose nature doesn't involve any particular contribution to reasoning won't have fittingness conditions. This prediction seems to be confirmed. For example, fantasizing that p seems to be such an attitude. While it can be fitting to *intend* to fantasize that p , the attitude of fantasizing itself doesn't seem to have a fittingness condition.

The foregoing is not a reductive account of fittingness, since it mentions fittingness. Still, in indicating the role of reasoning in fixing an attitude's standard of fittingness, we take it to illuminate the nature of fittingness.

6. Objections

6.1 Hypothetical Reasoning

You might reason:

Suppose Jane had a beer. Then there are none left. But then Jeremy will be upset. So, if Jane had a beer, then Jeremy will be upset.

Hypothetical reasoning of this sort is ubiquitous, and it might seem to pose a problem for the account of reasoning offered here. There is nothing wrong with using falsehoods as premises of hypothetical reasoning. And it doesn't seem as though the steps of hypothetical reasoning aim at truth. You might even be aiming to derive a contradiction. So how can hypothetical reasoning have the function of getting fitting attitudes?

Properly understood, hypothetical reasoning poses no threat to our account. Crucially, hypothetical reasoning, unlike 'straight' theoretical reasoning, involves attitudes other than belief. A proposition assumed for hypothetical reasoning is not the content of a belief, but of a supposition; it is thus not a premise of your reasoning in our restricted sense. On the other hand, the conclusion-attitude of a piece of hypothetical reasoning *is* typically a belief - e.g. a belief in a conditional, as above, or in the negation of a proposition assumed for *reductio*. And you arrive at this belief in a way that does seem directed at believing the truth.

Still, one might wonder how supposition functions in hypothetical reasoning, and in what sense the individual steps of this reasoning are regulated by the function of getting fitting attitudes. A full answer to this question would require a worked-out account of supposition, something that is beyond the scope of this paper. However, we will indicate what we think is a promising direction for developing such an account. Supposing, we might note, is in some ways like believing and in some ways unlike it.²⁹ It is like believing in that it represents a content as true. It is unlike believing in that it's not necessarily unfitting when its content is false: you're not getting things wrong merely by supposing a falsehood. In this way, unlike belief, supposition isn't answerable to how things actually are. But this doesn't mean that it's answerable to nothing. Rather, when you reason under a supposition, your further suppositional steps are answerable to the situation where the supposition you're reasoning under is true. (A situation here can be understood as a set of propositions – it needn't correspond to a possible world). For example, when reasoning under the supposition that Jane had a beer, given the background assumption that there was only one beer left, it's unfitting – wrong – to suppose that there are still beers in the fridge, and it's fitting to suppose that there are none. That's

²⁹ Broome (2013, 265) and Wright (2014) suggest that supposing that *p* is a kind of make-believe attitude to *p*. What we say here can be seen as a way of developing this thought.

why, when you suppose that Jane had a beer, it would be good reasoning to move to the further supposition that there are no beers left, and bad reasoning to move to the further supposition that there are some beers left. Ultimately, you discharge the initial assumption – you stop operating under the supposition that it's true – and make a transition to a belief, answerable to your actual situation. The rules of conditional proof and *reductio* give us ways of doing this that serve the aim of truth.

Note that the differences in the fittingness-conditions of belief and supposition reflect differences in their functional roles, in accordance with the account sketched in section 5.4. For example, the supposition that p , like the belief that p , will tend to lead through reasoning to further attitudes of the same kind, with contents supported by p ; but it will not tend to lead to intentions to act in ways that would further one's ends if p were true.

The possibility of hypothetical reasoning illustrates a more general point, emphasised by Harman (1986). We must distinguish reasoning, a psychological process, from derivations or arguments, whether formal or informal. Derivations and arguments are abstract objects that can be used to model certain features of reasoning. We can also reason *about* them. But there are many features of reasoning that they don't capture. One such feature is the attitudes involved in reasoning. Derivations and arguments express relations between contents.

6.2 Perceptual Beliefs

Consider the transition from a perceptual experience as of something's being red to the belief that it is red. Is this transition a case of reasoning? It might seem not. Perceptual beliefs often seem to be acquired passively, without our doing anything. And there seem to be important differences between acquiring beliefs by inference from other beliefs, and acquiring them directly based on perception. Yet the account we have offered seems to allow that such transitions could count as reasoning, since they could aim at truth. More specifically, they could manifest the sorts of dispositions that we have suggested can constitute aiming at fitting attitudes. This, then, is a second objection to our account: it mistakenly counts the acquisition of perceptual beliefs as reasoning.

In response, we note, first, that the acquisition of *many* perceptual beliefs will not count as reasoning, on our account. Recall that reasoning in the sense we are considering is conscious and personal-level

(§1). If reasoning is understood in this way, as a conscious transition of thought, then beliefs acquired without any such transition, as many perceptual beliefs are, are not acquired through reasoning. For example, the belief that something is red can be acquired through perception without becoming consciously occurrent. However - moving to the second part of our response - when such a transition from perception to belief does occur, and meets the conditions we have proposed, then we are happy to count it as reasoning. This seems to us the right thing to say about such cases. A conscious transition from perceiving something to be red to believing that it is red, guided by the aim of truth, is not a brute causal process, like acquiring a belief through a blow to the head, nor is it a mere blind reaction to a stimulus. It is a movement of thought manifesting a complex disposition that includes, for example, sensitivity to defeaters: for example, if you believe something is illuminated by red light, you will not make this move. Thinking of this transition as reasoning also fits well with – and so allows us to retain – the straightforward account of the connection between basing and reasoning suggested above, on which all basing corresponds to a disposition to reason. Nor should we be misled by the point that transitions from perception to belief need not be effortful or involve higher order-thought. Many inferential transitions from belief to belief lack these properties. The transitions from perception to belief that meet our conditions do not seem relevantly different from such belief-to-belief transitions.

To say this is not deny that there are *any* important differences between the two cases. For example, nothing we have said contradicts the claim that perceptual beliefs are epistemically basic, in the sense that they are not based on further beliefs, or that perceptions, unlike beliefs, can justify beliefs without themselves being justified. In certain contexts there may be a point to calling perceptual beliefs ‘non-inferential’, and to withholding the label ‘inference’ from the transitions by means of which they are acquired. But it remains plausible that both belief-to-belief inference and perception-to-belief transitions that meet our conditions fall under a common psychological genus.³⁰

³⁰ Among those who defend the claim that there is reasoning from perception to belief are Pollock and Cruz (1995, 195) and Brewer (1999). Others endorse it in passing (e.g. Pryor 2004). Very many of the epistemologists who reserve the term ‘inference’ for belief-to-belief transitions nonetheless accept that there are rational transitions from perception to belief; for a representative passage see, e.g., Peacocke (1999, 216).

6.3. The Taking Condition

According to Paul Boghossian, theoretical reasoning is subject to the

(Taking Condition) Theoretical reasoning necessarily involves taking your premises to support your conclusion.³¹

Our account does not obviously vindicate the Taking Condition. This might be thought to be grounds for rejecting our account.

To assess this objection, we need to know what it is to “take” your premises to support your conclusion. In discussing Wedgwood’s view, we noted two natural interpretations. On the first, taking your premises to support your conclusion is *believing* that your premises support your conclusion. On the second, taking your premises to support your conclusion is *treating* your premises as supporting your conclusion. The Taking Condition does not undermine our account, on either of these interpretations. On the first, the condition is, as Boghossian argues, too strong to be plausible. On the second, the condition is too weak to pose a problem. Plausibly, treating your premises as supporting your conclusion just is *reasoning* from the former to the latter. More generally, we doubt that there is a way to interpret the condition so that it is substantive, plausible, and problematic for our account. But rather than rehearse the many possible interpretations, we will indicate how our account accommodates some of the motivations for the Taking Condition.

Boghossian (2014, 5) suggests that the Taking Condition is needed to make sense of the *active* character of reasoning. However, the Taking Condition alone does not seem to explain this: taking it that some process is normatively supported does not turn that process into agency. In fact, Boghossian’s suggestion seems to be that the Taking Condition follows from the *aim-directedness* of reasoning, which in turn explains its active character. We agree that what makes reasoning active is that it is aim-directed. If Boghossian is right that this aim-directedness entails the Taking Condition, then our account does after all vindicate the condition. If not, then there is no good motivation for the Taking Condition here. Either way, this motivation generates no problem for our account.

³¹ Boghossian (2014, 5). Boghossian calls theoretical reasoning ‘inference’.

The Taking Condition may also be motivated by the thought that theoretical reasoning *commits* you to thinking that your premises support your conclusion. This can be brought out by the observation that there is something incoherent about assertions of the form:

(Bad) r , so, p ; but r does not support p .

The first half of this assertion expresses reasoning from believing r to believing p . The second half declares that r does not support p . Such assertions seem defective – the second half seems to contradict the first. The Taking Condition offers one explanation of this incoherence. On the assumption that taking is believing, reasoning ‘ r , so, p ’ entails that you believe that r supports p . That is why, when you deny that r supports p , you seem to contradict yourself.³²

But just as we need not explain the oddness of ‘ p but I don’t believe p ’ by saying that believing p entails believing that you believe p , we need not explain the incoherence of (Bad) by appeal to the Taking Condition. Our account offers an alternative explanation. Theoretical reasoning is guided by the aim of acquiring fitting beliefs. If p does not support q , then reasoning from p to q is not a good way to pursue this aim. So, reasoning from p to q while judging that p does not support q amounts to taking what you acknowledge to be an unreliable means to your end. That looks plainly irrational. While more might be said here, this seems enough to give a sense in which reasoning from p to q commits you to thinking that p supports q , and so to explain why assertions of (Bad) seem incoherent. To sum up: while our account fails to vindicate a strong version of the Taking Condition, this version is anyway implausible. Furthermore, our account captures the central motivations for the condition. The condition thus poses no threat to our account.³³

6.4 What About the Science?

Reasoning is a psychological process that can be studied empirically. Psychologists and cognitive scientists have proposed various accounts of how we reason. It might seem surprising that we have

³² Cf. Setiya (2013, 191), Hlobil (2014).

³³ We discuss motivations and problems for the Taking Condition further in McHugh and Way (2016 a).

offered an account of what reasoning is without engaging with empirical findings and theories. One might even wonder whether the account offered here isn't incompatible with empirical findings. For example, we make certain systematic errors in both theoretical and practical reasoning, such as interpreting evidence in ways that are biased towards our existing beliefs, and failing to give sufficient weight to possibilities that we find hard to imagine.³⁴ These sorts of processes, it might be thought, don't aim at fitting attitudes.

In response, we make three points. Firstly, there are different questions one might ask about reasoning. We have been asking the very general question what reasoning, as such, is. Another question is: in what ways do human beings tend to reason? A third is: how is reasoning implemented in human cognitive systems? Correct answers to these questions can constrain each other. But they are distinct questions.

The evidence of systematic errors pertains most directly to the second of these questions, and is compatible with our answer to the first. Nothing in what we have said suggests that human reasoning can't be subject to systematic error, or affected by mechanisms with diverse functions. Our account does suggest that there are limits on the extent to which you can reason in a way that you are aware won't lead to fitting attitudes. But that is compatible with systematic, inadvertent error.

This might seem too quick. Mercier and Sperber (2011) argue, partly based on evidence of systematic error, that the main *function* of reasoning is a social one, to do with the exchange of arguments in order to facilitate advantageous communication. This might seem incompatible with our claim that the function of reasoning is to get fitting attitudes. However, Mercier and Sperber's account need not be seen as a rival to ours. It can be interpreted as making a claim about the evolutionary function of human reasoning mechanisms, rather than about the constitutive function of reasoning as such. At any rate, we think that their claim is most plausible when interpreted in this way. After all, thinkers not prone to systematic error, and with a different evolutionary history to ours, not involving exchange of arguments, could still count as reasoners. And if Mercier and Sperber were right about the social function of reasoning, then it's hard to see why fallacious reasoning which serves this putative function (e.g., as they claim, confirmation bias) should nonetheless count as *bad* reasoning.

³⁴ See Kahneman (2012) for an entertaining survey.

Our second point is that our view is compatible with the possibility of reasoning by engaging in the processes which psychologists argue that reasoning involves – for instance, applying probabilistic heuristics (Chater and Oaksford 1999) or constructing mental models (Johnson-Laird 2008). If you engage in attitude-revision following the General Rule by doing one of these things, your doing it counts as reasoning. Note that the ordinary rules of reasoning we have used as examples in this paper are indicative only. It might be that the rules human thinkers in fact follow are different. It might be, for example, that we often follow domain-specific rules (Cheng and Holyoak 1985).

Our third point is that cognitive scientists are often interested in information-processing in general, rather than reasoning in our more specific sense. Much of what they describe concerns what we might call System <1.5 processes. These processes are not irrelevant to System ≥ 1.5 reasoning. They clearly subserve it, and so are part of the answer to the third question distinguished above. But again there is no incompatibility with our account here. For example, the output of a System <1.5 process might occur to you as a putative consequence of some things you believe, a consequence you may then endorse, forming a further belief. This endorsement can be a step in reasoning, but the System <1.5 process that made this consequence occur to you is not itself part of your reasoning.

7. Conclusion

This paper has been concerned with reasoning, understood as a certain sort of personal-level, conscious process of attitude revision. Reasoning can be done badly, but it can't be unconstrained. We have argued that there is a constitutive function, getting fitting attitudes, that must guide any episode of thought in order for it count as reasoning, and we have developed this hypothesis in the context of a rule-following conception of reasoning. We have outlined how such a view can illuminate not only the nature of reasoning but also the standards for good reasoning, basing, normative and motivating reasons, and fittingness.³⁵

³⁵ For valuable feedback, thanks to Davor Bodrozic, Jean-Marie Chevalier, Julien Dutant, Tyler Doggett, Fabian Dorsch, Magnus Frei, Alex Gregory, Marie Guillot, Ulf Hlobil, Frank Hofmann, Yair Levy, Errol Lord, Kieran Setiya, Lee Walters, Daniel Whiting, several anonymous referees and editors, and audiences at the University of Barcelona, the University of Cardiff, the Collège de France, Kings College London, the University of Leipzig, the University of Luxembourg and the University of Regensburg. This work was supported by the Arts and Humanities Research Council [grant number AH/K008188/1].

References

- Arpaly, N and Schroeder, T. (2014). *In Praise of Desire*. Oxford: OUP.
- Boghossian, P. (2008). Epistemic Rules. *Journal of Philosophy* 105 (9):472-500.
- Boghossian, P. (2014). What Is Inference? *Philosophical Studies* 169, 1-18.
- Brentano, F. (1889/2009). *The Origin of our Knowledge of Right and Wrong*. Oxford: Routledge.
- Brewer, B. (1999). *Perception and Reason*. Oxford: OUP.
- Broad, C.D. (1930). *Five Types of Ethical Theory*. London: Routledge and Kegan Paul.
- Broome, J. (2013). *Rationality Through Reasoning*. Chichester: Wiley-Blackwell.
- Broome, J. (2014). Comments on Boghossian. *Philosophical Studies* 169, 19-25.
- Carroll, L. (1895). What the tortoise said to Achilles. *Mind* 4 (14):278-280.
- Chappell, R. (2012). Fittingness: The Sole Normative Primitive. *The Philosophical Quarterly* 62, 684-704.
- Chater, N., and Oaksford, M. (1999). The Probability Heuristics Model of Syllogistic Reasoning. *Cognitive Psychology* 38, 191-258.
- Cheng, P.W., and Holyoak, K.J. (1985). Pragmatic Reasoning Schemas. *Cognitive Psychology* 17, 391-416.
- D'Arms, J., and Jacobson, D. (2000). The Moralistic Fallacy: On the 'Appropriateness' of Emotions. *Philosophy and Phenomenological Research* 61, 65-90.
- Ewing, A.C. (1948). *The Definition of Good*. London: Routledge and Kegan Paul.
- Forbes, G.R. (1983). Scepticism and Semantic Knowledge. *Proceedings of the Aristotelian Society* 84, 223-37.
- Geach, P. T. (1956). Good and Evil. *Analysis* 17 (2):33 - 42.
- Grice, P. (2001). *Aspects of Reason* (ed. R. Warner). Oxford: OUP.
- Gregory, A. (forthcoming). Normative Reasons as Good Bases. *Philosophical Studies*.
- Harman, G. (1986). *Change in View: Principles of Reasoning*. Cambridge, Mass.: MIT Press.
- Hieronimi, P. (2013). The Use of Reasons in Thought. *Ethics* 124 (1): 114-27.

- Hlobil, U. (2014). Against Boghossian, Wright and Broome on inference. *Philosophical Studies* 167 (2):419-429.
- Howard, C. (ms.). The Fundamentality of Fit.
- Johnson-Laird, P.N. (2008). Mental Models and Deductive Reasoning. In J.E. Adler and L.J. Rips (eds.), *Reasoning: Studies of Human Inference and Its Foundations*. Cambridge: CUP.
- Kahneman, D. (2012). *Thinking, Fast and Slow*. Penguin.
- Kolodny, N. (2007). How Does Coherence Matter? *Proceedings of the Aristotelian Society* 107, 229-63.
- Korsgaard, C.M. (2009). *Self-Constitution: Agency, Identity, and Integrity*. Oxford: OUP.
- McHugh, C. (2014). Fitting Belief. *Proceedings of the Aristotelian Society* 114, 167-87.
- McHugh, C., and Way, J. (2015). Broome on Reasoning. *Teorema* 34, 131-40.
- McHugh, C., and Way, J. (2016 a). Against the Taking Condition. *Philosophical Issues* 26.
- McHugh, C., and Way, J. (2016 b). Fittingness First. *Ethics* 126, 575-606.
- McHugh, C., and Way, J. (forthcoming). What is Good Reasoning? *Philosophy and Phenomenological Research*.
- Mercier, H., and Sperber, D. (2011). Why Do Humans Reason? Arguments for an Argumentative Theory. *Behavioral and Brain Sciences* 34, 57-74.
- Neta, R. (2013). What is an inference? *Philosophical Issues* 23 (1):388-407.
- Peacocke, C. (1999). *Being Known*. Oxford: OUP.
- Pollock, J. (1987). Defeasible Reasoning. *Cognitive Science* 11 (4):481-518.
- Pollock, J.L., and Cruz, J. (1999). *Contemporary Theories of Knowledge* (2nd ed.). Oxford: Rowman and Littlefield.
- Pryor, J. (2004). What's Wrong with Moore's Argument? *Philosophical Issues* 14, 349-78.
- Setiya, K. (2013). Epistemic Agency: Some Doubts. *Philosophical Issues* 23 (1):179-198.
- Setiya, K. (2014). What is a Reason to Act? *Philosophical Studies* 167 (2):221-235.
- Shah, N. (2003). How Truth Governs Belief. *Philosophical Review* 112 (4):447-482.
- Shah, N. (2013). Why We Reason the Way We Do. *Philosophical Issues* 23 (1):311-325

- Svavarsdóttir, S. (2014). Having Value and Being Worth Valuing. *The Journal of Philosophy* 111, 84-109.
- Szabó, Z.G. (2001). Adjectives in Context. In I. Kenesei and R.M. Harnish, *Perspectives on Semantics, Pragmatics, and Discourse: A Festschrift for Ferenc Kiefer*. Amsterdam: John Benjamins.
- Thomson, J.J. (2008). *Normativity*. Chicago, Ill.: Open Court.
- Valaris, M. (forthcoming). What Reasoning Might Be. *Synthese*.
- Velleman, J.D. (2000). On The Aim of Belief. In his *The Possibility of Practical Reason*. Oxford: OUP.
- Way, J. (forthcoming). Reasons as Premises of Good Reasoning. *Pacific Philosophical Quarterly*.
- Wedgwood, R. (2006). The Normative Force of Reasoning. *Noûs* 40 (4):660–686.
- Wedgwood, R. (2012). Justified Inference. *Synthese* 189 (2):1-23.
- Wright, C. (2014). Comment on Paul Boghossian, "What Is Inference". *Philosophical Studies* 169, 27-37.