



A puzzle about enkratic reasoning

Jonathan Way¹

Accepted: 2 September 2020 / Published online: 15 October 2020
© The Author(s) 2020

Abstract Enkratic reasoning—reasoning from believing that you ought to do something to an intention to do that thing—seems good. But there is a puzzle about how it could be. Good reasoning preserves correctness, other things equal. But enkratic reasoning does not preserve correctness. This is because what you ought to do depends on your epistemic position, but what it is correct to intend does not. In this paper, I motivate these claims and thus show that there is a puzzle. I then argue that the best solution is to deny that correctness is always independent of your epistemic position. As I explain, a notable upshot is that a central epistemic norm directs us to believe, not simply what is true, but what we are in a position to know.

Keywords Reasoning · Enkratic reasoning · Correctness · Objectivism and perspectivism about ought

Suppose you are wondering whether to stop by the deli. You'd like to get bagels but you're also running late for an appointment. After brief consideration, you figure that it won't matter much if you're a little later still and so come to believe that you ought to stop by the deli. On that basis, you form the intention to stop by the deli.

This last step is an example of what Broome (2013) calls enkratic reasoning: moving from a belief that you ought to do something to an intention to do that thing. It seems like good reasoning. Of course, you might make mistakes along the way. You might be wrong about the pros and cons, or how they stack up. You might be

✉ Jonathan Way
j.way@soton.ac.uk

¹ University of Southampton, Southampton, UK

unjustified in these assessments. Still, considered just as a transition of thought, independently of the status of its inputs, enkratic reasoning seems hard to fault.

Indeed, enkratic reasoning seems to play a central role in deliberation. When you deliberate, you try to figure out what to do—to settle on a course of action, by forming an intention. If, during deliberation, you reach the conclusion that you ought to act in a certain way, you can close deliberation, by settling on acting in that way. But this just is to engage in enkratic reasoning: to form an intention on the basis of a belief about what you ought to do. Enkratic reasoning is thus a central way of bringing deliberation to a close.

These remarks concern beliefs about what you ought to do in the ‘all-out’, ‘overall’, or, as I shall say, ‘deliberative’ sense of ‘ought’. In the example above, you might believe that you ought, given only your concern for bagels, to stop by the deli and that you ought, given only your concern for your appointment, to head straight there. These beliefs do not allow you to form intentions. But it is extremely plausible that there is a further question you can ask: given this conflict, what ought I to do? It is answering this question that allows you to move to an intention.¹

So enkratic reasoning seems good—indeed, it seems like it must be good. But there is a puzzle about how it could be. In short, the puzzle is that a highly plausible claim about good reasoning—that good reasoning must preserve correctness—seems to imply that enkratic reasoning is not good. In Sect. 1 I explain and motivate this claim about good reasoning. In Sect. 2, I explain how it seems to imply that enkratic reasoning is not good. In Sect. 3, I discuss three possible responses to the puzzle, and defend my preferred response. In Sect. 4, I defend this response against the objection that it fails to deal with some related puzzles. As I explain there, a notable upshot is that a central epistemic norm directs us to believe, not simply what is true, but what we are in a position to know.

1 Good reasoning

The puzzle begins with the following claim:

(GR) Good reasoning preserves correctness, other things equal.

That is: good reasoning is such as to lead you from correct responses to a further correct response, other things equal. This requires some explanation.²

First, I take reasoning to be a kind of transition between responses. The responses I will be interested in are attitudes, and especially beliefs and intentions. So, for

¹ Although terminology varies, the deliberative ought plays an important role in much normative philosophy. Prominent examples include Broome (2013), Davidson (1980), Gibbard (1992), Korsgaard (1996), Parfit (2011) and Williams (1965) (the source of the term ‘deliberative “ought”’). For useful recent discussions see Kiesewetter (2017), Lord (2015), McPherson (2018) and Schroeder (2011a). Some doubt there is any such ‘ought’ (Tiffany 2007; Baker 2018). Doubts might target either the psychological claim that there is an intelligible concept here or the metaphysical claim that there is a corresponding property. I only presuppose the former.

² For further clarification and defence of GR, see McHugh and Way (2018). Cf. Wedgwood (2007: 99ff), Worsnip (2019).

example, you might reason to the belief that the deli is open from your beliefs that the deli is open until 6 pm and that it's now 5.30 pm. Or you might reason to an intention to go to the deli from your intention to get bagels and a belief that there are bagels at the deli. I will call the responses you reason from your *premise-responses* and the response you reason to your *conclusion-response*. Where these responses are beliefs, their contents are your premises and conclusion, respectively.

Second, in saying that a piece of reasoning is good, I am making a *prospective* assessment of a *transition*. It is *prospective* in that it is an assessment of a type of reasoning, rather than a *retrospective* assessment of any token of that type; that a token piece of reasoning exemplifies a good type of reasoning does not ensure that it is done well, or competently. And it concerns the transition in that it is independent of any assessment of the responses from which the reasoning begins: a piece of reasoning might be good even if it starts from mistaken or unjustified responses.

Third, correctness is a normative status. To say that an attitude is correct is to say that it gets things right; its object merits, is worthy of, or makes fitting, that attitude. This notion shows up in a wide range of normative theorising. For many epistemologists, a central epistemic norm is that it is correct to believe just what is true.³ Many value-theorists are interested in whether the valuable can be analysed in terms of correct valuing—for instance, whether being admirable is being correct to admire.⁴ Moral responsibility theorists are interested in when a person is worthy of praise or blame—that is, when it is correct to praise or blame them.⁵ Plausibly, the same notion of correctness is at play in all of these inquiries. For example, one central mark of this notion is that it is insensitive to incentives. That you would be rewarded for believing, admiring, or blaming does not bear on whether such an attitude would be correct.⁶

Intentions can also be correct or incorrect. Just as correct admiration and correct blame at least correspond to evaluative properties—the admirable and the blameworthy—correct intention corresponds to an evaluative property: the choiceworthy. To choose is to form an intention, and the choiceworthy is what merits choice. Of course, what makes for choiceworthiness is a contentious question. For example, consequentialists will think it has to do with promoting value, deontologists will disagree.⁷ Equally though, it is a contentious question what makes someone admirable, or blameworthy. We will not need to investigate these questions here.

With these clarifications made, we can see why GR is plausible. First, paradigm cases of good reasoning preserve correctness. The clearest example is deductive theoretical reasoning. Deductive theoretical reasoning is guaranteed to preserve truth, and truth is widely taken to be what makes belief correct. But it is also

³ See, e.g. Boghossian (2003), Shah and Velleman (2005), Wedgwood (2007: 100) and Whiting (2013).

⁴ See Jacobson (2011) for an overview.

⁵ See Eshleman (2014) for an overview.

⁶ For more on correctness, or what I elsewhere call 'fittingness', see Howard (2018) and McHugh and Way (2016).

⁷ Such disputes between moral theorists are only one relevant point of contention, given that non-moral considerations are also plausibly relevant to choiceworthiness.

plausible that, other things equal, non-deductive theoretical reasoning preserves truth. For instance, it is a plausible defeasible rule of good reasoning to believe p if someone says p . And it is also plausible that other things equal, if someone says p , then p .

Paradigmatic forms of practical reasoning also preserve correctness. For example, it is good practical reasoning to move from intending an end to intending a means you believe necessary to achieve that end. And if an end is choiceworthy, and a means is necessary for that end, then that means is also choiceworthy. Plausible defeasible rules of practical reasoning also preserve correctness, other things equal. For instance, it is a plausible defeasible rule of good reasoning to move from believing that you promised to ϕ to intending to ϕ . And it is plausible that other things equal, if you promised to ϕ , then ϕ -ing is choiceworthy.⁸

There is also a more theoretical rationale for GR. The standard for good reasoning should fit with the *point* or *aim* of reasoning. Reasoning is a way of forming new attitudes. But the point of reasoning is not just to form any old attitudes; it is to form attitudes with a certain normative status. The most plausible candidate for this status, I suggest, is correctness. In reasoning we are trying to get things right in our beliefs and intentions. We are not just trying to form beliefs and intentions which cohere with our other attitudes, or for which there are good reasons. Consider, for example, that in theoretical reasoning, we aim to form beliefs that are true. But beliefs can cohere with our other attitudes and be supported by good reasons without being true.⁹

If in reasoning we aim to get things right in our beliefs and intentions, then we should expect good reasoning to be well-suited to success in this aim. Of course, reasoning cannot be expected to generate correct responses from nothing: starting points matter. But what reasoning can do is keep us on track. Thus we can expect good reasoning to preserve correctness, at least other things equal.

Finally, it should be stressed that GR is only a necessary condition on good reasoning. There are serious challenges to the idea that correctness-preservation is sufficient for good reasoning: for instance, it doesn't seem like just any reasoning to belief in a necessary truth is good reasoning. In other work, I discuss how to develop GR into a full account of good reasoning, in light of such challenges (McHugh and Way 2018). But the puzzle I want to raise does not depend on whether this can be done.

2 The puzzle

GR implies that enkratic reasoning is good only if it preserves correctness. That is:

- (GR-ER) Other things equal, if it is correct to believe that you ought to ϕ , then it is correct to intend to ϕ .

⁸ There are obviously questions about how to understand this 'other things equal' clause. I flag this issue when relevant. See McHugh and Way (2018) for discussion.

⁹ See McHugh and Way (2018) for further discussion.

(GR-ER) has the air of plausibility. But on reflection it turns out to be problematic.

Assume the plausible and widely held view that it is correct to believe just what is true. Given this, GR-ER implies:

- (*) Other things equal, if you ought to ϕ , it is correct to intend to ϕ .

It is very plausible that what you ought to do—in the deliberative sense—is constrained by your epistemic position. That is, factors that are not epistemically accessible to you do not bear on what you ought to do. Consider Judith Thomson's (1990: 229) well-known example in which, due to an extraordinary and unforeseeable coincidence, flipping the light switch when you get home will cause a small lightning flash in your neighbour's home, and badly burn them. Thomson held that in this case you ought to refrain from flipping the switch. But many people have a different reaction: given that the harm to your neighbour is unforeseeable, it is okay to flip the switch. This reaction presupposes that what you ought to do is constrained by your epistemic position.

By contrast, the correctness of a response—in the sense introduced—does not seem constrained by your epistemic position. I assumed that it is correct to believe just what is true, but what is true does not generally depend on your epistemic position. And the point is not specific to belief. Consider emotions. You get things right in your emotions when they are directed towards objects with certain evaluative properties. For instance, it is correct to admire the admirable and blame the blameworthy. But whether a person is admirable or an act blameworthy does not depend on your epistemic position: someone might be blameworthy for deeds outside your ken. It thus seems a general principle that the correctness of a response is not constrained by your epistemic position.¹⁰

This apparent difference between correctness and 'ought' gives rise to counter-examples to (*). Consider Thomson's example again. If the fact that flipping the switch will harm your neighbour does not bear on what you ought to do, then presumably you ought to flip the switch. The epistemically accessible facts—principally, that flipping the switch will illuminate the room—support doing so. But if corrent intention is not epistemically constrained, then it is not correct to intend to flip the switch. Given all the facts, including that flipping the switch will harm your neighbor, that is not the right choice.

So (*), and thus GR-ER, appears false. But (GR-ER) followed from the assumptions that enkratic reasoning is good and that good reasoning preserves correctness. Since both these assumptions are plausible, we have a puzzle.

Before moving on, it should be emphasised that this result does not turn on conflating senses of 'ought' and 'correct'. No doubt these terms can express different notions. In particular, each can express epistemically constrained and epistemically unconstrained notions. We can make sense of claims about what you ought to do given all the facts, and what you ought to do given only the epistemically accessible facts. And we can make sense of claims about what it is

¹⁰ For more or less explicit endorsement of this principle see e.g. D'Arms and Jacobson (2000: 69), Deonna and Teroni (2012: 6–7), Gibbard (2005: 339–341), Rosen (2015: 70), Svavarsdóttir (2014: 106), Wedgwood (2007: 156) and Worsnip (2019: 146).

correct to intend given all the facts and what it is correct to intend given only the epistemically accessible facts. However, it is a substantive question whether what you *deliberatively* ought to do depends on all the facts or only the epistemically accessible facts. And it is a substantive question whether what it is correct to intend, in the sense in which good reasoning preserves correctness, depends on all the facts or only the epistemically accessible facts. The puzzle arises because the answers to these questions seem to diverge. In the relevant senses, what you ought to do is epistemically constrained, while what it is correct to intend is not.

3 Responses to the puzzle

The puzzle depends on four claims:

- (1) Good reasoning preserves correctness, other things equal (GR).
- (2) Enkratic reasoning is good.
- (3) What you ought to do is constrained by your epistemic position.
- (4) What it is correct to intend is not constrained by your epistemic position.

Responses to the puzzle must either reject one of these claims or show that they are compatible. In the rest of this paper, I consider some of these options. Since I have already explained why I find it plausible that enkratic reasoning is good and that good reasoning preserves correctness, I will not further consider (1) or (2). While the puzzle might show that one of these claims should be rejected, my interest is in whether they can be retained.¹¹ So I will instead consider versions of the three other options: rejecting (3), rejecting (4), and maintaining that (1)–(4) are compatible. To anticipate, I shall argue that the best solution is to reject (4) and thus hold that what it is correct to intend is epistemically constrained.

3.1 Can we accept each of (1)–(4)?

(1)–(4) are not strictly incompatible. What (1)—(GR)—requires is that good reasoning preserve correctness *other things equal*. Thus it does not imply that there cannot be cases in which it is correct to believe that you ought to ϕ but not correct to intend to ϕ . It only implies that any such cases must be ones in which other things are not equal. So an initially attractive response to our puzzle says that in the cases which generate it, other things aren't equal. If this response can be defended, we can resolve the puzzle without rejecting any of (1)–(4).

A final assessment of this response requires detailed consideration of how GR's 'other things equal' clause is to be understood. However, we need not get into this to see reason for scepticism.

¹¹ The only writer I know of to reject (2) is Southwood (2016). See Broome 2016 for a response. While (1) is more contentious, alternative views of good reasoning face their own issues about enkratic reasoning. For instance, the view that good reasoning must preserve justification forms the basis of an argument that justified beliefs about what you ought to do must be true (Way and Whiting 2016). Since this is a surprising conclusion, this view also does not straightforwardly vindicate enkratic reasoning.

We can distinguish between defeasible and indefeasible good reasoning. Roughly, a good piece of reasoning is defeasible when there are responses that can coherently be added to its premise-responses so that the reasoning is no longer good. Many good pieces of reasoning are defeasible in this way. For instance, the reasoning you could express by saying ‘the deli’s standard hours are 8 am–6 pm; it’s now 5.30 pm, so the deli will be open’ is good. But if you add the premise that the deli’s Facebook page says that it’s closing early today, it is no longer good. Other pieces of reasoning seem indefeasible. If you reason instead from the premise that the deli *is* open from 8 am–6 pm today, then it’s hard to see what could coherently be added to spoil the reasoning. Of course, we can think of considerations that could lead you to give up one of the premises—such as the Facebook announcement. But that is for one of the premises to be undermined, rather than for the reasoning to be defeated.

The point of the ‘other things equal’ clause in GR is to allow for defeasibility. So when a piece of reasoning is indefeasibly good, we might expect a stricter version of GR to apply: indefeasibly good reasoning should preserve correctness in all cases.

Enkratic reasoning seems indefeasible. It is hard to think of attitudes that can coherently be added to its premise-responses which would spoil it. For example, if you believe that you ought to stop by the deli, learning of new reasons against going to the deli will not stop it being good reasoning to move to an intention to go to the deli. New reasons might lead you to drop the belief that you ought to stop by the deli. But again, that is for the premise of the reasoning to be undermined, not for the reasoning to be defeated (cf. Way and Whiting 2016: 1877–1878).

If this is right, we can restate the claims on which the puzzle rests:

- (1′) Indefeasibly good reasoning preserves correctness in all cases.
- (2′) Enkratic reasoning is indefeasibly good.
- (3) What you ought to do is constrained by your epistemic position.
- (4) What it is correct to intend is not constrained by your epistemic position.

These claims are incompatible. (1) and (2) imply

(GR-ER′) In all cases, if it is correct to believe that you ought to ϕ , then it is correct to intend to ϕ . But if (3) and (4) are true, then this is false. In Thomson’s case, it is correct to believe that you ought to flip the switch, but not correct to intend to flip the switch.

At this point, some resist (2′). Indeed, some philosophers hold that there are no cases of indefeasibly good reasoning. This view is motivated by cases of higher-order evidence: evidence which bears on whether you are in a position to respond appropriately to normatively relevant considerations.¹² For instance, if you believe that your reasoning capacities have been impaired, it might not seem good reasoning to perform even a paradigmatic form of good deductive reasoning, such as *modus ponens* reasoning. In the same circumstances, enkratic reasoning might not seem good.

¹² For general discussion see, e.g. Bradley (2019), Christensen (2010) and Lasonen-Aarnio (2014).

Such cases are far from clear-cut; they raise a host of contentious issues. However, even if they do show that there are no cases of indefeasibly good reasoning, it remains plausible that there are types of reasoning which are immune from certain kinds of defeat. Many good pieces of reasoning can be defeated by countervailing considerations or by considerations that undermine the support that the premise appeared to provide for the conclusion-response.¹³ Some good pieces of reasoning—such as deductive theoretical reasoning—cannot be defeated in this way. Furthermore, it remains plausible that reasoning which is immune from defeat in this way preserves correctness in all cases. Even if higher-order evidence can spoil a deductively good piece of reasoning, this is not because it shows that the premise-responses can be correct while the conclusion-response is not. If such reasoning is defeasible, this must be because of some further condition on good reasoning.¹⁴

Enkratic reasoning seems to fall into this category. Even it can be defeated by higher-order evidence, it does not seem defeasible by countervailing or undermining considerations. If this is right, the puzzle remains—we just need to read ‘indefeasibly’ in (1′) and (2′) as meaning ‘indefeasible by countervailing or undermining considerations’.

I conclude that that the puzzle cannot easily be resolved by appeal to GR’s ‘other things equal’ clause. If we are to maintain GR and that enkratic reasoning is good, we must reject either (3) or (4). I now consider these options.

3.2 Rejecting (3): objectivism about ought

There is an important debate between *objectivists* and *perspectivists* about what you ought to do. Perspectivists hold that only epistemically accessible facts bear on what you ought to do. Objectivists deny this; all facts are in principle relevant. Objectivists hold that in Thomson’s case, the fact that flipping the switch will harm your neighbour is relevant to whether you ought to flip the switch, and thus that you ought to refrain from flipping the switch.¹⁵

Objectivists will thus deny (3). In this way, our puzzle might seem to rest on a contentious view about what you ought to do. Indeed, the puzzle might be thought grounds for preferring objectivism to perspectivism.

I do not mean to rule out objectivism. However, if we are interested in vindicating enkratic reasoning, objectivism is not the way to go. As other well-known cases show, if objectivism is true, then enkratic reasoning is not generally

¹³ In our example, the announcement that the deli is closing early today is a countervailing consideration for the premise that the deli’s standard hours are 8 am–6 pm. An example of an undermining consideration is that the deli is not running its standard hours today. For discussion see, e.g. Pollock (1987), Dancy (2004: ch. 3), Schroeder (2011b) and Horta (2013: chs. 2 and 5).

¹⁴ These claims are clearly true if the standard of correctness for belief is truth. Things are more complicated if, as I consider below, the standard of correctness is being in a position to know—it will depend on how this notion is cashed out.

¹⁵ For objectivist views, see Graham (2010), Schroeder (2018a), Thomson (1990: 172–173, 229–242) and Wedgwood (2013). For perspectivist views, see Broome (2013: ch. 3), Lord (2015), McHugh and Way (2017) and Kiesewetter (2017: ch. 8).

good. Consider Frank Jackson's (1991) Doctor case. You are a doctor treating a patient with a minor but non-trivial illness. You could administer drug A, drug B, or drug C. You know that one of drugs A and B will completely cure the patient and the other will kill them. However, you do not know, and cannot find out, which of drugs A and B is the cure and which is the killer. You do know that drug C will partially cure the patient. You are therefore in a position to know that you objectively—given all the facts—ought not give drug C. However, it is clear that it is not good reasoning to move from the belief that you objectively ought not give drug C to the intention not to give drug C. It would be grossly irresponsible to rule out giving drug C. If objectivism is true though, what you ought to do *is* what you objectively ought to do. Thus the objectivist must also deny that it is good reasoning to move from the belief that you ought not give drug C to the intention not to give drug C. So objectivism does not offer a way out of the puzzle that vindicates enkratic reasoning.¹⁶

This may seem too quick. It might be replied that in enkratic reasoning you move from a belief that you ought to *act* to an intention to act. But in the Doctor case, you believe that you ought *not* act in a certain way. So the case doesn't show that objectivists cannot vindicate enkratic reasoning.

This reply shows that we need to be more precise about how we are understanding enkratic reasoning. In enkratic reasoning, you move from believing that you ought to ϕ to intending to ϕ . But what can we substitute for ' ϕ ' here? I think it is plausible that enkratic reasoning is good whenever ϕ -ing is something that you believe to be under your control, in the sense that you will (not) ϕ if and because you intend (not) to (cf. Broome 2013: 290). Omissions can be under your control in this sense, and so beliefs about what you ought not to do can feature in enkratic reasoning. But this seems the right result. To judge that you ought not ϕ , when ϕ -ing is under your control, is to judge that ϕ -ing is not to be done—it is to take ϕ -ing to be ruled out. This seems to commit you to not ϕ -ing, and thus it seems good reasoning to move to an intention not to ϕ . If so then, since, in the Doctor case, you are aware that giving drug C is under your control, objectivism fails to vindicate a plausible form of enkratic reasoning.

It might be suggested that the Doctor case is itself a counter-example to this precisification of enkratic reasoning (cf. Graham 2010; Wedgwood 2007: 30–31; 2013). After all, the objectivist will take the Doctor case to be one in which you can know that you ought not give drug C, and that not giving drug C is under your control, and yet it is not good reasoning to move to an intention not to give drug C. But this description of the case is theory driven: the intuitive verdict on the Doctor case is that you ought to give drug C (cf. Kiesewetter 2017: 210). Nor is it natural to think that the following expresses good reasoning: 'I don't know which of drugs A and B is the full cure; it is clear that I ought not give drug C. Nonetheless, I'll give

¹⁶ Given objectivism, the transition from believing that you objectively ought not give drug C to an intention not to give drug C does preserve correctness. This doesn't mean that objectivists are committed to thinking that this transition is good reasoning: GR is only a necessary condition on good reasoning.

drug C'. The case gives us no independent grounds for rejecting the proposed precisification of enkratic reasoning.^{17,18}

3.3 Rejecting (4): moderate objectivism about correctness

The upshot of the previous two sections is that if we are to maintain that good reasoning preserves correctness, and that enkratic reasoning is good, we must reject (4), the claim that what it is correct to intend is not constrained by your epistemic position. This is my preferred response to the puzzle, and I begin to defend it here. While I won't provide a direct argument against (4), I will show how two natural arguments for (4) can be resisted. In the absence of further arguments for (4), I thus propose to take the plausibility of (1)–(3) as grounds for rejecting (4).

3.3.1 The argument from plausible cases

The first argument says that (4) follows from the way I have characterised the notion of correctness.

An attitude is correct if it gets things right—if its object merits, is worthy of, or makes fitting, that attitude. Plausible cases include that it is correct to believe truths, admire the admirable, and blame the blameworthy. But whether a proposition is true, or a person admirable or blameworthy, does not generally depend on your epistemic position. There are many other examples too: an outcome might be desirable, and thus correct to desire, due to factors we are ignorant of. We should therefore accept *objectivism* about correctness: the correctness of a response is not constrained by your epistemic position. And objectivism implies (4).

In response, we should distinguish *strong objectivism* and *moderate objectivism*. Strong objectivism holds that the correctness of a response is, in all cases, independent of your epistemic position. Moderate objectivism holds that the

¹⁷ Wedgwood (2007: 30–31; 2013: 489–90; cf. Srinivasan 2015) takes Williamson's anti-luminosity argument (2000: 93–106) to show that we are not always in a position to know what we ought to do. He suggests that it follows that there will be Doctor-style counter-examples to the claim that enkratic reasoning is good. There will be cases in which you have three options, ϕ , ψ , and χ , such that (i) you ought to ϕ , (ii) you know only that either you ought to ϕ or you ought to ψ , (iii) you know that whichever of ϕ -ing and ψ -ing is impermissible is very seriously wrong, and (iv) you know that χ -ing is less seriously wrong. However, the anti-luminosity argument does not show that such cases are possible. If successful, it only shows that you can fail to be in a position to know what you ought to do in a *borderline case*—roughly, a case which would be near the border in a series which moves gradually from the clearly permissible to the clearly required.

¹⁸ Another way of rejecting (3) bears mention. Contextualists about 'ought' deny that what you ought to do is constrained by *your* epistemic position. Rather, truths about what you ought to do, in a context of utterance, are constrained by *an* epistemic position determined by that context—this need not be the subject's (cf. Dowell 2013, Björnsson and Finlay 2010). Relativists (e.g. MacFarlane (2014: ch.11) and expressivists (cf. Schroeder (2018a) might make similar claims. However, this view does not resolve the puzzle. If correct intention depends on all the facts, the puzzle arises so long as truths about what a subject ought to do can, as contextualists accept, depend on a less than omniscient epistemic position.

correctness of a response *need not* depend on your epistemic position. Only strong objectivism implies (4): moderate objectivism can allow that there are responses whose correctness depends on your epistemic position. And it is moderate objectivism that is most directly motivated by the examples which the argument from plausible cases begins from. These examples suggest that the correctness of some responses is determined independently of your epistemic position. They do not show that the correctness of all responses is so determined.¹⁹

Indeed, it fairly natural to think that the correctness of some responses is epistemically constrained. Consider *interrogative* attitudes: for example, wondering, or being curious, about a certain question.²⁰ Plausibly, there are questions which merit curiosity, and questions which do not. Whether we have free will, whether Trump will be reelected, and whether the deli is open might all merit curiosity; whether there is an odd number of dust particles on my bookshelf does not. But whether a question merits curiosity seems epistemically constrained. If you know whether the deli is open but I do not, then even if it is fitting for me to be curious about whether the deli is open, it is not fitting for you to be. Or consider credences. Correct credence corresponds to something very much like a value property—credibility. But whether a proposition is credible—and thus whether it is correct to have a certain credence in it—surely depends on your epistemic position.

Nor is there anything incoherent in the idea that the same property of correctness could apply in a way that depends on your epistemic position in some cases but not others. Think about it this way. Each response is subject to a standard of correctness—a principle that specifies the conditions under which that response is correct. The standard of correctness for belief might be truth, the standard of correctness for credence might be what is probable given your epistemic position. As these examples illustrate, standards of correctness can vary in whether they specify conditions that depend on your epistemic position. But they are all standards for the same property of correctness.

So the way I have characterised correctness does not force us to accept (4). That there are plausible cases in which the correctness of a response is epistemically unconstrained does not show that the correctness of a response is always epistemically unconstrained.

In response, it might be suggested that correct intention is itself a plausible example of a response whose correctness is epistemically unconstrained. But this is unconvincing. It is correct to intend the choiceworthy. And while it is clear that the admirable and blameworthy are epistemically unconstrained, it is far from clear that the choiceworthy is. Consider the Doctor case. Since there is a good chance that drug A is the killer, it seems too risky to give drug A. I'm inclined to think that this

¹⁹ The literature tends to define objectivism, either about 'ought' or correctness, as I have defined strong objectivism (cf. Kiesewetter 2017: 196, Lord 2015: 26, Worsnip 2019: 146). However, moderate objectivism is the negation of a natural definition of perspectivism, as the view that correctness, or 'ought', must be determined by epistemically accessible considerations. Furthermore, many discussions focus on objectivism about how you ought to act. Here the difference between moderate and strong objectivism is idle.

²⁰ Cf. Friedman (2013).

is relevant to whether giving drug A is choiceworthy, and thus that choiceworthiness is epistemically constrained.²¹ This strikes me as a difference between choiceworthiness and some other evaluative properties: that Jones might be a killer does not make him blameworthy, whereas that drug A might be a killer does make giving drug A unchoiceworthy. But I don't want to put weight on this thought. The important point is negative: epistemic position is not clearly irrelevant to choiceworthiness in the way that it is to admirability and blameworthiness. So we shouldn't take correct intention as a further plausible case of a response whose correctness is epistemically unconstrained.

3.3.2 *The argument from reasoning*

The second argument says that (4) follows from the claim that in reasoning we aim at correctness.

In reasoning, we aim to get things right—to form correct attitudes. In particular, in practical reasoning, or deliberation, we aim to form correct intentions. But in deliberation, we aim to settle on doing what is best—that is, the thing to do given all the facts. And if the aim of the deliberation is to settle on doing what is best, and the aim of deliberation is to form a correct intention, then what it is correct to intend must be what is best. And since what is best is epistemically unconstrained, what it is correct to intend must be epistemically unconstrained.

The crucial premise of this argument is that in deliberation we aim to settle on doing what is best, given all the facts. I do not think we need to accept this premise. I know of two reasons we might accept it, and I do not think either is compelling.

The first appeals to some observations about deliberation. It makes sense for deliberating agents to seek out further information, when possible. And it makes sense for better-informed advisors to rely on their own information, rather than just the subject's information. These observations are neatly explained if in deliberation we aim at what is best (Lord 2015: 44ff).

However, we can also explain these observations by assuming that deliberators are concerned to do what they ought to do. Of course, given perspectivism, what you ought to do depends on your epistemic position. But we need not think that your epistemic position is constituted only by facts you are currently aware of. We can hold instead that your epistemic position is constituted by facts which are available to you by the time of action (Dowell 2013; Kiesewetter 2017: ch. 7; Björnsson and Finlay 2010). While this notion of availability needs spelling out, such a view clearly promises to explain why it makes sense for you to seek out further facts, and

²¹ I assume that epistemic considerations—considerations of what might happen, or what there is a chance of happening—are relevant to choiceworthiness only if choiceworthiness is epistemically constrained. This is because it is hard to see how to weigh epistemic considerations against facts outside your epistemic position. For instance, it is hard to see how to weigh the fact that drug A might be the killer against the fact that it is not the killer. Cf. Kiesewetter (2017: 203).

for better informed advisors to take them into account. So these observations do not strongly support the view that deliberation aims at what is best.²²

The second reason turns on an analogy between practical and theoretical reasoning. In theoretical reasoning, it seems, we aim at truth. For example, you might aim to come to a true belief about whether the deli is open. And if the aim of theoretical reasoning is epistemically unconstrained, we might think that the aim of practical reasoning will be too.

This line of thought is not watertight: the aims of practical and theoretical reasoning might differ in whether they are epistemically constrained. As will emerge though, I think it is important that practical and theoretical reasoning are alike in this regard. So I want to show that the argument fails even granting this point. To do so, we need to reconsider a claim which has thus far been taken for granted: that it is correct to believe just what is true. This claim is plausible and widely held. But it is not trivial or incontestable. Again, to say that a belief is correct is to say that it is merited by, or fitting to, its object. It is a substantive question which objects of belief—which propositions—merit belief.

Although the view that what merits belief is truth is a natural one, it is not the only contender. Another natural view, which can also be found in the literature, is that what merits belief is being in a position to know.²³ This view makes the correctness of belief depend on your epistemic position. But as we have seen, that is compatible with moderate objectivism about correctness.

With this view on the table, we can see where the argument from analogy goes wrong. For it is not at all clear that in theoretical reasoning we aim only at the truth. It is at least equally plausible that in theoretical reasoning, we aim to *figure out* the truth. To figure out the truth is not just to come to a true belief; it is to come to know. And whether you know a truth clearly depends on your epistemic position. The thought that in theoretical reasoning we aim at truth is thus compatible with the view that the aim of theoretical reasoning—and thus what it is correct to believe—is epistemically constrained. So it does not underwrite an argument by analogy that the aim of practical reasoning is epistemically unconstrained.

I conclude that neither of the reasons for accepting that in deliberation we aim at what is best are persuasive, and thus that the argument from reasoning fails. We can accept that reasoning aims at correctness while denying (4).

Before moving on, let me address a worry. As stated, our puzzle assumes that it is correct to believe just what is true. But I have now suggested that we should be open to rejecting this assumption. It might seem that if we do, the puzzle dissolves.

²² Some might take ‘ought’ judgments made in light of information that is not available by the time of action to support the view that deliberation aims at what is best (cf. Lord 2015; MacFarlane 2014: ch.11). However, as Dowell (2013) argues, it is far from clear that such judgments concern the deliberative ought.

²³ Smithies (2012: 283–285). Although he does not speak of ‘correctness’, Williamson’s (forthcoming) claim that the ‘primary norm’ of belief is to believe only what you know is naturally understood as supporting a similar view. Compare his remark that ‘knowledge sets the standard of appropriateness for belief... mere belief is a kind of botched knowing’ (2000: 47).

However, this is a mistake. Here is the puzzle again. We assume that (1), good reasoning preserves correctness, and (2), enkratic reasoning is good. These claims imply:

(GR-ER) Other things equal, if it is correct to believe that you ought to ϕ , then it is correct to intend to ϕ .

If it is correct to believe just what you are in a position to know, GR-ER implies:

(**) Other things equal, if you are in a position to know that you ought to ϕ , it is correct to intend to ϕ .

But if (3), what you ought to do is constrained by your epistemic position, and (4), what it is correct to intend is not, (**) is false. In Thomson's example, you are presumably in a position to know that you ought to flip the switch—the epistemically accessible facts, such as that flipping the switch will illuminate the room, manifestly support doing so. But if what is correct to intend is not epistemically constrained, it is not correct to intend to flip the switch. Thus the puzzle arises whether it is correct to believe just what's true or just what you are in a position to know.²⁴

4 Related puzzles

I have proposed that we solve the puzzle by rejecting (4)—the claim that correct intention is epistemically unconstrained. I've shown that two natural arguments for (4) fail; absent further arguments, we can thus take the plausibility of (1)–(3) as grounds for rejecting (4). In this section, I consider an objection. It might be thought that my preferred solution leads to problems when applied to closely related puzzles. I consider three versions of this concern.

4.1 Quasi-enkratic reasoning

In *quasi-enkratic* reasoning, you move from believing that you ought to have an attitude to that attitude. Examples include reasoning from believing that you ought to believe p to believing p , reasoning from believing that you ought to desire X to desiring X , and reasoning from believing that you ought to admire S to admiring S .

Quasi-enkratic reasoning might seem good. But the puzzle arises for it too. If what you ought to believe, desire, and admire is epistemically constrained, but what it is correct to believe, desire, and admire is not, then such reasoning does not preserve correctness. If we were to respond to this puzzle in the way we responded to the original puzzle, we would be led to say that the correctness of all responses is epistemically constrained. But that is to give up on even moderate objectivism about correctness.

²⁴ The puzzle does require that correct belief is factive. This view is widely held; I defend it in Sect. 4.3.

However, we don't need to respond to this puzzle in the same way. Rather, we should deny that quasi-enkratic reasoning is good. It is therefore not a problem if GR fails to vindicate it.

To begin with, it should be stressed that quasi-enkratic reasoning is not enkratic reasoning. In enkratic reasoning, you move from a belief that you ought to do something—that you take to be under your control—to an intention to do that thing. Quasi-enkratic reasoning, by contrast, moves from a belief that you ought to have an attitude to that attitude. Since this is a different transition, there is no immediate reason to think that because enkratic reasoning is good, quasi-enkratic reasoning is good.

Once we are clear about this difference, it is not obvious why we should think that quasi-enkratic reasoning is good. Consider two examples: reasoning from believing you ought to believe p to believing p and from believing you ought to intend to φ to intending to φ . One useful heuristic for testing a piece of reasoning is to consider how it would be expressed in thought or talk. These pieces of reasoning could be expressed: 'I ought to believe p , so, p ' and 'I ought to intend to φ , so, I'll φ '. These passages of thought do not look cogent: their premises and conclusions concern very different topics.

This disconnect is made vivid by cases where you believe that you ought to have an attitude because it would be beneficial to have it. Suppose you reason: 'believing that God exists will make me happy; if believing that God exists will me happy, I ought to believe in God; so, I ought to believe in God; so, God exists'. Something seems to have gone wrong in one's reasoning here. But the first step is just *modus ponens*. The fault must thus lie with the second, 'quasi-enkratic', step.²⁵

I therefore doubt that quasi-enkratic reasoning is good, and so doubt that the proposed response to the puzzle overgeneralizes as suggested.

4.2 Sufficient evidence reasoning

Consider 'sufficient evidence reasoning': reasoning from the belief that there is sufficient evidence for p to believing p . This looks like a form of good reasoning. Indeed, it looks indefeasibly good—at least, it seems immune to defeat by countervailing or undermining considerations. But it doesn't seem to preserve correctness. Sufficient evidence can still be misleading. Thus whether the standard of correctness for belief is truth or being in a position to know, there can be sufficient evidence for p though it is not correct to believe p .²⁶

This case seems very similar to our original puzzle. It looks like a form of indefeasibly good reasoning which does not preserve correctness in all cases. But

²⁵ This argument does not assume that the benefits of an attitude can make it the case that you ought to have that attitude. It only assumes that someone might believe this.

²⁶ Insofar as evidence provides the "right-kind" of reason for belief, sufficient evidence reasoning exemplifies a more general type of reasoning – roughly, reasoning from the belief that you ought, for right-kind reasons, to have an attitude, to that attitude. It might be thought that this more general type of reasoning also raises the puzzle. I am sceptical, but the suggestion raises too many issues to discuss here. On right-kind reasons, see Gertken and Kiesewetter (2017).

the solution to the original puzzle doesn't help. This puzzle arises even if the standard of correctness for belief is being in a position to know, and thus epistemically constrained.²⁷

One might respond to this puzzle by denying that sufficient evidence can be misleading. This denial is not without precedent and can be motivated in various ways.²⁸ Nonetheless, a commitment to the infallibility of sufficient evidence is a cost. Fortunately, we need not adopt this commitment. Even if we grant that this puzzle about sufficient evidence reasoning shows that indefeasibly good reasoning need not always preserve correctness, this is not enough to solve the puzzle about enkratic reasoning. It thus does not undermine the case for denying (4).

Cases in which sufficient evidence reasoning fails to preserve correctness are *opaque to the subject*. That is, you cannot coherently think that the move from believing that there is sufficient evidence for p to believing p will fail to preserve correctness. That is because you cannot coherently think both that there's sufficient evidence for p and that it is incorrect to believe p . This is clearly true if the standard of correctness for belief is truth—it's incoherent to think that there's sufficient evidence for p although p is not true. But it is also plausibly true if the standard of correctness for belief is being in a position to know. It seems incoherent to think that there is sufficient evidence for p although you are not in a position to know p . Whatever leads you to doubt that you are in a position to know should, it seems, lead you to doubt that the evidence for p is sufficient.

This observation explains how sufficient evidence reasoning can be indefeasible despite failing to preserve correctness in all cases. A good piece of reasoning is defeasible if there are responses which can coherently be added to its premise-responses which stop it being good reasoning. But the fact that sufficient evidence reasoning may fail to preserve correctness does not imply that this condition is met. As we have just seen, the kind of considerations that might lead you to think that it is incorrect to believe p cannot be coherently combined with the belief that there is sufficient evidence for p . So they can't provide defeaters for sufficient evidence reasoning.

Thus, even if we deny that indefeasibly good reasoning preserves correctness in all cases, we should accept:

- (1'') Indefeasibly good reasoning can only fail to preserve correctness in ways that are opaque to the subject.

Together with (2'), the claim that enkratic reasoning is indefeasibly good, (1'') implies that any cases in which it is correct to believe that you ought to ϕ but incorrect to intend to ϕ must be opaque to the subject. But if correct intention is not epistemically constrained, this is false. In the Doctor case, you can coherently

²⁷ Indeed, this puzzle also arises if it is correct to believe just what there is sufficient evidence for. That's because, as cases of higher-order evidence illustrate, it appears that there can be sufficient but misleading evidence about what there is sufficient evidence for.

²⁸ For claims in this vicinity see e.g. Littlejohn (2012), Schroeder (2015), Whiting (2013) and Williamson (forthcoming).

believe that you ought to give drug C although it is not correct, given all of the facts, to intend to give drug C.

The upshot is that even if the puzzle about sufficient evidence reasoning shows that good reasoning need not preserve correctness in all cases, we still need to deny (4), that correct intention is not constrained by your epistemic position.²⁹

4.3 Instrumental reasoning

Suppose that you intend to get bagels and believe that you can get bagels at the deli. You might then reason to an intention to stop by the deli. This looks like a form of good (though defeasible) practical reasoning. However, if correct intention is epistemically constrained but correct belief is just true belief, it does not preserve correctness. If ϕ -ing is choiceworthy and you can ϕ by ψ -ing, it doesn't follow that ψ -ing is choiceworthy. If the fact that you can ϕ by ψ -ing is epistemically inaccessible to you, there might be little to be said for ψ -ing.

So if instrumental reasoning is good, and correct intention is epistemically constrained, correct belief must also be epistemically constrained. This is why I was happy to grant that correct intention and belief were alike in this respect in Sect. 3.1.2.

I don't think this is problematic. As discussed, it is a substantive question what the standard of correctness for belief is. The view that it is correct to believe just what is true is a natural view, but so is the view that it is correct to believe just what you are in a position to know. The choice between these views must be made on theoretical grounds. The argument that correct intention is epistemically constrained, together with these observations about instrumental reasoning, contribute to the case for the latter view. I take this to be an important upshot of the preceding discussion.

Of course, one might hold that the standard of correctness for belief is epistemically constrained but deny that it has to do with knowledge. But previous points also tell against this position. In theoretical reasoning we aim at truth—or to figure out the truth. The normative status we seek in theoretical reasoning must thus be factive. And the most natural standard for belief that is both factive and depends on your epistemic position is being in a position to know.³⁰

One might worry, again, that this treatment of instrumental reasoning will generalise to rule out moderate objectivism. But this does not seem to be the case. Consider another form of instrumental reasoning: reasoning from a desire to get bagels, and a belief that you can get bagels at the deli, to a desire to stop by the deli.

²⁹ Of course, if we are to retain GR, and allow for misleading sufficient evidence, cases in which sufficient evidence reasoning fails to preserve correctness must be cases in which other things are not equal. A full defence of GR thus requires an interpretation of its 'other things equal' clause that gets this result. But these issues are to one side, here.

³⁰ Of course, there are other candidate standards that meet these conditions. What makes belief correct might be that it is likely given your epistemic position and *also* true. Or what makes belief correct might simply be that it is knowledge. However, the first of these looks ad hoc. The second does not allow us to distinguish prospective and retrospective assessments of correctness. It is thus vulnerable to the objections of Schroeder 2018b.

If desires can figure in reasoning, this seems like good reasoning. And it is plausible that what is correct to desire—desirable—is not epistemically constrained. However, this reasoning preserves correctness even if the standard of correctness for belief is being in a position to know. If it is desirable to get bagels and you are in a position to know that you can get bagels at the deli, then it is desirable to stop by the deli, at least insofar as it will allow you to get bagels. Of course, stopping at the store might not be desirable overall. So it might not be correct to desire overall to stop by the deli. But the move from a desire to get bagels to an overall desire to stop by the deli is clearly defeasible at best; we should not expect it to preserve correctness in all cases.³¹

5 Conclusion

In this paper, I have presented a puzzle about enkratic reasoning. Each of the following claims is plausible, but they seem incompatible:

- (1) Good reasoning preserves correctness, other things equal (GR).
- (2) Enkratic reasoning is good.
- (3) What you ought to do is constrained by your epistemic position.
- (4) What it is correct to intend is not constrained by your epistemic position.

I have argued that—at least if we are to retain (1) and (2)—the most promising solution is to reject (4). While (1)–(4) are not strictly incompatible, given (1)’s ‘other things equal’ clause, we cannot easily resolve the puzzle by appeal to this clause. And if we reject (3), we have to reject (2) anyway. By contrast, there is no clear case for accepting (4). Nor does this solution run into problems when applied to related puzzles. The plausibility of (1)–(3) should thus be taken as an argument for rejecting (4). And in turn we can take the denial of (4) as the basis for an argument that what it is correct to believe is just what you are in a position to know.

Acknowledgements Thanks to Olivia Bailey, Claire Field, Alex Gregory, Wooram Lee, Franziska Poprawe, Keshav Singh, Daniel Whiting, a referee for *Philosophical Studies*, and audiences in Berlin, Chapel Hill, Cologne, Frankfurt, New York, and Southampton for comments and discussion. This paper emerges from collaborative work with Conor McHugh. Special thanks to him for extensive comments and discussion. Thanks also to the Murphy Institute at Tulane University for a fellowship during which I worked on this paper. Finally, thanks to Stein’s Deli in New Orleans for excellent bagels.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended

³¹ Similar points apply to putative examples of good reasoning from belief to other responses whose correctness conditions are epistemically unconstrained, such as admiration and blame.

use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Baker, D. (2018). Skepticism about ought *simpliciter*. *Oxford Studies in Metaethics*, 13, 230–252.
- Boghossian, P. (2003). The normativity of content. *Philosophical Issues*, 13, 31–45.
- Björnsson, G., & Finlay, S. (2010). Metaethical contextualism defended. *Ethics*, 121, 7–36.
- Bradley, D. (2019). Are there infeasible epistemic rules? *Philosophers' Imprint*, 19, 1–19.
- Broome, J. (2013). *Rationality through reasoning*. Chichester: Wiley-Blackwell.
- Broome, J. (2016). Responses. *Philosophical Studies*, 173, 3431–3438.
- Christensen, D. (2010). Higher-order evidence. *Philosophy and Phenomenological Research*, 81, 185–215.
- D'Arms, J., & Jacobson, D. (2000). The Moralistic fallacy: On the 'appropriateness' of emotions. *Philosophy and Phenomenological Research*, 61, 65–90.
- Dancy, J. (2004). *Ethics without principles*. Oxford: Oxford University Press.
- Davidson, D. (1980). How is weakness of will possible? In D. Davidson (Ed.), *Essays on actions and events* (pp. 21–42). Oxford: Clarendon Press.
- Deonna, J., & Teroni, F. (2012). *The emotions*. London: Routledge.
- Dowell, J. (2013). Flexible contextualism about deontic modals: A puzzle about information-sensitivity. *Inquiry*, 56, 149–178.
- Eshleman, A. (2014). Moral responsibility. *The Stanford Encyclopedia of Philosophy*, ed. E. Zalta. <https://plato.stanford.edu/archives/win2016/entries/moral-responsibility/>.
- Friedman, J. (2013). Question-directed attitudes. *Philosophical Perspectives*, 27, 145–174.
- Gertken, J., & Kieseewetter, B. (2017). The right and wrong kinds of reasons. *Philosophy Compass*, 12, e12412.
- Gibbard, A. (1992). *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard University Press.
- Gibbard, A. (2005). Truth and correct belief. *Philosophical Issues*, 15, 338–350.
- Graham, P. (2010). In defence of objectivism about moral obligation. *Ethics*, 121, 88–115.
- Horty, J. (2013). *Reasons as defaults*. Oxford: Oxford University Press.
- Howard, C. (2018). Fittingness. *Philosophy Compass*, 13, e12542.
- Jackson, F. (1991). Decision-theoretic consequentialism and the nearest and dearest objection. *Ethics*, 101, 461–482.
- Jacobson, D. (2011). Fitting attitude theories of value. In E. Zalta (Eds.), *The Stanford encyclopedia of philosophy*. Retrieved from <https://plato.stanford.edu/archives/spr2011/entries/fitting-attitude-theories/>.
- Kieseewetter, B. (2017). *The normativity of rationality*. Oxford: Oxford University Press.
- Korsgaard, C. (1996). *Sources of normativity*. Cambridge: Cambridge University Press.
- Lasonen-Aarnio, M. (2014). Higher-order evidence and the limits of defeat. *Philosophy and Phenomenological Research*, 88, 314–345.
- Littlejohn, C. (2012). *Justification and the truth-connection*. Cambridge: Cambridge University Press.
- Lord, E. (2015). Acting for the right reasons, abilities, and obligation. *Oxford Studies in Metaethics*, 10, 26–52.
- MacFarlane, J. (2014). *Assessment sensitivity*. Oxford: Oxford University Press.
- McHugh, C., & Way, J. (2016). Fittingness first. *Ethics*, 126, 575–606.
- McHugh, C., & Way, J. (2017). Objectivism and perspectivism about the epistemic ought. *Ergo*, 4, 121–145.
- McHugh, C., & Way, J. (2018). What is good reasoning? *Philosophy and Phenomenological Research*, 96, 253–274.
- McPherson, T. (2018). Authoritatively normative concepts. *Oxford Studies in Metaethics*, 13, 253–277.
- Parfit, D. (2011). *On what matters*. Oxford: Oxford University Press.
- Pollock, J. (1987). Defeasible reasoning. *Cognitive Science*, 11, 481–518.

- Rosen, G. (2015). The alethic conception of moral responsibility. In R. Clarke, M. McKenna, & A. M. Smith (Eds.), *The nature of moral responsibility: New essays* (pp. 65–88). Oxford: Oxford University Press.
- Schroeder, M. (2011a). *Ought*, agents, and actions. *The Philosophical Review*, 120, 1–41.
- Schroeder, M. (2011b). Holism, weight, and undercutting. *Noûs*, 45(4), 328–344.
- Schroeder, M. (2015). Knowledge is belief for sufficient (objective and subjective) reason. *Oxford Studies in Epistemology*, 5, 226–252.
- Schroeder, M. (2018a). Getting perspective on objective reasons. *Ethics*, 128, 289–319.
- Schroeder, M. (2018b). Believing Well. In C. McHugh, J. Way, & D. Whiting (Eds.), *Metaepistemology* (pp. 196–212). Oxford: Oxford University Press.
- Shah, N., & Velleman, D. (2005). Doxastic deliberation. *The Philosophical Review*, 114, 497–534.
- Smithies, D. (2012). The normative role of knowledge. *Noûs*, 46, 265–288.
- Srinivasan, A. (2015). normativity without cartesian privilege. *Philosophical Issues*, 25, 273–299.
- Svavarsdóttir, S. (2014). Having value and being worth valuing. *Journal of Philosophy*, 111, 84–109.
- Southwood, N. (2016). The motivation question. *Philosophical Studies*, 173, 3413–3430.
- Thomson, J. J. (1990). *The realm of rights*. Cambridge, MA: Harvard University Press.
- Tiffany, E. (2007). Deflationary normative pluralism. *Canadian Journal of Philosophy*, 37, 231–262.
- Way, J., & Whiting, D. (2016). If you justifiably believe that you ought to ϕ , you ought to ϕ . *Philosophical Studies*, 173, 1873–1895.
- Wedgwood, R. (2007). *The nature of normativity*. Oxford: Oxford University Press.
- Wedgwood, R. (2013). Akrasia and uncertainty. *Organon F*, 20(4), 484–506.
- Whiting, D. (2013). On the norm and aims of belief. In T. Chan (Ed.), *The aim of belief*. Oxford: Oxford University Press.
- Williams, B. (1965). Ethical consistency. *Proceedings of the Aristotelian Society, Supplementary*, 39, 103–124.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Williamson, T. (Forthcoming). Justification, excuses, and sceptical scenarios. In J. Dutant and F. Dorsch (eds), *The new evil demon*. Oxford: Oxford University Press.
- Worsnip, A. (2019). Isolated correct reasoning. In M. B. Jackson & B. B. Jackson (Eds.), *Reasoning* (pp. 129–151). Oxford: OUP.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.