



In defence of object-given reasons

Michael Vollmer¹ 

Accepted: 16 January 2024 / Published online: 6 February 2024
© The Author(s) 2024

Abstract

One recurrent objection to the idea that the right kind of reasons for or against an attitude are object-given reasons for or against that attitude is that object-given reasons for or against belief and disbelief are incapable of explaining certain features of epistemic normativity. Prohibitive balancing, the behaviour of bare statistical evidence, information about future or easily available evidence, pragmatic and moral encroachment, as well as higher-order defeaters, are all said to be inexplicable in terms of those object-given reasons. In this paper, I provide a rebuttal to all these challenges by drawing attention to the object-given reasons for and against the third doxastic state, namely the suspension of judgement. First, I introduce an original picture of how suspension relates to belief and disbelief – as ways of being decided – which yields a novel weighing model of the corresponding reasons. Second, I demonstrate that this new take on doxastic deliberation can accommodate all five features of epistemic normativity. Finally, I extend the weighing model to cover practical deliberation and the suspension of intentions. The resulting theory provides an explanation of why epistemic and practical normativity differ in certain respects in terms of a significant difference between doxastic and conative suspension.

Keywords Right kind of reasons · Object-given reasons · Epistemic Normativity · Suspension of judgement · Decidedness · Suspending intentions

✉ Michael Vollmer
Michael.Vollmer@uni-heidelberg.de

¹ University of Heidelberg, Heidelberg, Germany

1 Introduction

In the literature on the normativity of attitudes, it is commonplace to distinguish the *right* from the *wrong kind of reasons*. The right kind of reasons are said to bear a special relation to the *correctness* of the particular attitude; are said to be the only or at least the prime *motivating reasons* for those attitudes; and are commonly considered to play a crucial role in assessing the *rationality* of one's attitude formation. According to the paradigmatic take on these right kind reasons, they are the kinds of considerations that are properly connected to the properties of the object at which those attitudes are directed, that is, they are *object-given reasons* of those attitudes.¹ However, several authors have expressed doubts about this latter claim: *epistemic reasons*, so the idea goes, are clearly the right kind of reasons of beliefs. Nevertheless, there are many features of *epistemic normativity* that cannot be accounted for in a theory which only makes use of object-given reasons of belief and disbelief.² In this paper, I rebut these challenges by closely examining the object-given reasons of the third doxastic stance – namely the *suspension of judgement*.³ By introducing a new way of weighing reasons in doxastic deliberation, my account of suspension provides a neat picture of several features of epistemic normativity which are supposed to be in tension with such a framework of object-given reasons. Roughly put, there is a twofold weighing of reasons: on the one hand, reasons that favour belief are weighed against reasons that favour disbelief, while reasons that favour suspension are weighed against reasons that favour *decidedness* where belief and disbelief are both *ways of being decided* on an issue.

I start by outlining the distinction between the right and the wrong kinds of reasons of attitudes, as well as the idea of accounting for this divide in terms of object-given reasons (Sect. 2). Afterwards, I introduce five puzzles from the literature that plague the application of object-given reasons to an analysis of epistemic normativity (Sect. 3). Following these preliminary sections, I develop my rebuttal of these challenges. First, I provide an account of the third doxastic stance – the suspension of judgement – as well as an elaborate model of how the object-given reasons of suspension relate to the object-given reasons of belief and disbelief (Sect. 4). Second, I highlight that this picture of our doxastic options yields the necessary resources to account for the five puzzles of epistemic normativity in terms of object-given reasons (Sect. 5). However, before closing, I turn towards a remaining worry about the preceding account. While it is able to accommodate the potentially problematic features of epistemic normativity, an explanation is still required for why some of these features do not extend to *practical normativity* (Sect. 6). I argue that *suspend-*

¹ I use 'reasons of an attitude' as short for 'reasons which either favour or disfavour an attitude'.

² Versions of this rationale have been explicitly proposed by Schroeder (2012b, 2021), Booth (2014), and Lee (2022). However, it might also be seen as an implicit motive in works which attempt to identify non-evidential epistemic reasons.

³ Other authors have used different labels for this third doxastic stance which is supposed to come with a form of neutrality concerning the question at hand. McGrath (2021a) and Wagner (2022), for instance, have reserved 'suspension' for certain mental actions. In the course of this paper, however, I use it interchangeably with 'indifference', 'withholding', or 'being agnostic'.

ing intentions exemplifies a quite different normative profile than the suspension of judgement.

2 The right kind of reasons as object-given reasons

The distinction between the *right kind of reasons* (RKR) and the *wrong kind of reasons* (WKR) has its origin in the debate about *buck-passing analyses of value properties*.⁴ According to this family of theories, one category of normative properties – namely values – can be reduced to a different dimension of normativity – namely normative reasons. Thus, roughly put, an object O has a valuable property P iff there are normative reasons for a subject to take up a P-specific form of attitude ϕ towards O. For instance, the thick value property of being admirable can be analysed as follows: being admirable is nothing but the fact that there is a normative reason for a subject to admire the object in question. However, simple buck-passing accounts are riddled with one decisive shortcoming. After all, *prima facie* there can be all sorts of normative reasons which speak in favour of admiring, and only a fraction of them has anything to do with admirability. For instance, a monetary incentive might speak in favour of admiring a given person, but this should not make them admirable. To circumvent this objection, buck-passers introduced the RKR–WKR distinction – where a WKR is any sort of consideration which gives rise to this *WKR's problem*. Thus, an object O has a valuable property P iff there are RKR for a subject to take up a P-specific form of attitude ϕ towards O.

Scholars have proposed different ideas concerning how exactly one ought to flesh out this general, schematic distinction of RKR and WKR. A natural suggestion, with which I contend in this paper, is that RKR should bear the proper connection to the object of one's attitude and its value properties, or what has come to be known as *object-given reasons* (OGR). After all, WKR are considerations which are not indicative of the existence of such a valuable feature of the object. It must be noted that many diverging theories exist that specify the details of OGR in slightly different ways.⁵ For the purposes of this paper, one can neglect these differences. Since I conceive her account to be highly illustrative, I follow Hieronymi (2005) in specifying OGR in terms of justifying answers to certain *questions*: rationally evaluable attitudes are sensitive to certain *judgements*, which can be understood as *answers to characteristic questions*. Admiration, for instance, is sensitive to the judgement 'X is (very) admirable', which is an answer to 'Is X admirable?' or 'How admirable is X?'. An OGR is a consideration which can be alluded to in justifying an answer to this characteristic question. Paradigmatic WKR like monetary incentives fail this test, as one cannot appeal to them in justifying one's answer.

⁴ For an overview of these theories see Howard (2023). For a more detailed summary of the literature on RKR and WKR see Gertken and Kiesewetter (2017).

⁵ I use 'OGR' in a very broad manner to denote a couple of different theories, which range from standard instances like that of Parfit (2011) to related accounts like that of Hieronymi (2005), but also that of Sharadin (2016b). Furthermore, I regard the more recent line of thought which connects the WKR–RKR distinction to the notion of *fittingness* as a natural successor of the OGR account (McHugh & Way, 2016; Howard, 2019). Thus, fittingness-centred theories can also count as the subject of the discussion to come.

In his seminal work, Schroeder (2012b) proposes a slightly different take on the RKR–WKR distinction. Besides the role that RKR are supposed to play in circumventing the WKR problem, they also exemplify *two further earmarks* that make them interesting in a variety of ways, clearly exceeding the limits of buck-passing accounts. First, and less central to the purposes of this paper, the RKR of an attitude ϕ can be used by a subject as a *basis* for their ϕ -ing. Only RKR are considerations which can also be the content of the *motivating reasons* for the subject. Second, and more important for the discussion to come, only RKR have a bearing on the *ϕ -specific rationality* of the subject’s attitude formation. While sensitivity to monetary incentives might be indicative of a form of *practical* rationality, the rational standards of one’s attitude formation seem to be unmoved by it. The clearest example of this earmark is the idea of *epistemic rationality* of doxastic attitudes. While monetary incentives might yield a reason in favour of a belief, which makes that belief practically rational, such reasons do not seem to have a bearing on its epistemic rationality. This point is, again, in stark contrast to evidence in favour of or against the believed proposition. If one is not sensitive to these kinds of considerations, one fails to be epistemically rational.

Now, the leading question that guides the following discussion is whether OGRs can be made to work as an analysis of RKR *even if* one adopts Schroeder’s more sophisticated notion of RKR. I call the affirmative answer to this question the *right kind as object-given thesis*, or RKOG for short. Many proponents of OGRs as a solution to the WKR problem do seem to endorse RKOG as well.⁶ This is to be expected. After all, RKOG is an attractive thesis that ties one’s reply to the WKR problem to a more general distinction between kinds of reasons for attitudes.

3 Five puzzles from epistemology

One recurrent line of attack on RKOG is its application to *epistemic normativity*. Epistemic normativity, so the idea goes, can be understood in terms of *epistemic reasons*, which, in turn, are nothing but RKR of doxastic attitudes.⁷ Given RKOG, one would expect that all features of epistemic normativity are explained by the OGRs of these attitudes. What are the OGRs of doxastic attitudes? In the case of belief and disbelief, most authors agree that the characteristic question of these attitudes is whether the (dis-)believed proposition is the case. Just like admiration, believ-

⁶ One notable exception appears to be Hieronymi (2013). In her paper, she appears dissatisfied with Schroeder’s extension of the concept of RKR to his earmarks, which is a crucial juncture in his case against RKOG (see Sect. 3).

⁷ It should be noted that whether only epistemic reasons are considerations which can satisfy either of Schroeder’s earmarks is a controversial issue. According to pragmatists, for instance, practical considerations are also able to yield motivating reasons for beliefs (e.g. Schleifer McCormick, 2015; Leary, 2017; Rinard, 2019; Jackson, 2023). If one is convinced by these cases, one should either drop Schroeder’s first earmark or replace it with some weaker claim about typicality. The following discussion should be unmoved by these changes.

ing p is a judgement-sensitive attitude which answers the question ‘Is p the case?’.⁸ Correspondingly, OGRs of belief and disbelief must be considerations which can justify such answers. The paradigmatic instance of such a consideration is *evidence* that supports p or non- p . After all, answering whether p is the case seems to depend on considerations which are *indicative* of p ’s truth or falsity. For the purposes of this paper, I assume a weak form of evidentialism according to which only evidence is an OGR of belief and disbelief.⁹ However, opponents of RKOGR have pointed to five supposed features of epistemic normativity that exceed these OGRs and thus generate counterexamples – in that there must be some RKR at work which are not reducible to OGRs.

First, and probably the least divisive feature on this list, epistemic normativity is said to *balance* in a *prohibitive* manner: if the evidence is *tied* – or even *close to being tied* – on some issue, neither belief nor disbelief should be rational options – both are prohibited.¹⁰ As we will see in Sect. 6, this property of epistemic normativity is in stark contrast to the way in which practical normativity appears to behave. After all, in the case of ties between practical considerations, both options appear to be permitted. Consider, for instance, Buridan’s ass: a donkey stands between two bales of hay and has no reason to favour one bale over the other. Intuitively, if the donkey were to eat from either of the bales, this would count as rational behaviour. Thus, any account of epistemic normativity in terms of RKR must account for this special feature of prohibitive balancing. However, given RKOGR, all we seem to be left with is a tie between two options which, on the face of it, looks just like the tie in the example of Buridan’s ass. Therefore, we miss some sort of RKR against believing and disbelieving in the case of tied options, which would explain why they are prohibited.

Second, Schroeder (2021, p. 130 f.) points to *bare statistical evidence* and *lottery propositions*. One might argue that any theory of epistemic normativity has to accommodate the following two features of statistical evidence: statistical evidence is evidence and thus counts as an RKR in favour of belief and disbelief,¹¹ yet bare statistical evidence does *not suffice* to rationalise beliefs. At least in certain cases,

⁸ This thesis is often regarded as intimately connected to the so-called *transparency* of doxastic deliberation, which was made prominent by the works of Shah (2003) and Shah and Velleman (2005).

⁹ I use ‘evidence’ in a narrow sense as an indicator of the truth or falsity of the proposition in question. For the discussion to come, I only treat those kinds of considerations as proper justifiers of answers to ‘Is p the case?’. However, one might resist this assumption. According to Schmidt (2023), certain *facts about my first-order evidence* – that it bears an *inductive* or *explanatory* support relation to p – can be epistemic reasons in favour of believing p and disbelieving p . While being a contentious assumption, which might also be incompatible with some ways of spelling out the notion of OGR, the position defended in this paper is compatible with these ideas of Schmidt. Allowing for these additional OGRs in favour of belief and disbelief does not alter the set of puzzles presented in this section nor the solution I propose later in the paper. I thank one of the reviewers of *Philosophical Studies* for bringing this issue to my attention.

¹⁰ For work on this feature of epistemic normativity see Harman (2004), Schroeder (2012a, 2015, 2021), Cohen (2016), Roerber (2016), Snedegar (2017), Berker (2018), Brunero (2022), and Tucker (forth.).

¹¹ It should be noted that this first point is already contentious. According to several authors, some bare statistical evidence is not even an RKR that favours belief in the first place. According to Smith (2016, Chap. 2), information about statistics can fail a *normality* constrain which he attributes to our concept of evidence. Similarly, on Whiting’s (2022, p. 190) *modal* analysis of reasons, statistical evidence might equally fail to exemplify the necessary modal features of reasons.

beliefs that are only based on statistical evidence are not rational. Consider the following example which is prominent in legal epistemology.

Blue Bus

A car has been accidentally hit by a bus. 90% of the buses driving through the area are operated by the Blue Bus Company; the remaining 10% belong to the Red Bus Company. An inquiring observer learns these facts and comes to believe that the bus involved in the accident was run by the Blue Bus Company.

The inquiring observer in Blue Bus, according to the reaction of many scholars, is not rational in forming their belief.¹² However, given that reasons in favour of an option φ count as sufficient iff they are at least as weighty as the reasons against φ -ing, the proponent of RKOg is faced with the following challenge: given that the statistical evidence in Blue Bus is an RKR in favour of belief, there must be some RKR against believing which explains the absence of sufficient reasons. By construction, the only piece of evidence in favour of disbelieving is, again, the statistical evidence. However, it would be absurd to assume that the weight of the RKR based on the 90% chance that it was a Blue Bus Company vehicle is on a par with the weight of the RKR based on the 10% chance that it was a Red Bus Company one. Therefore, the RKRs in favour of disbelief are, in principle, insufficient for explaining the lack of sufficiency. Thus, again, proponents of RKOg seem to be short of an RKR against believing.

Third, Schroeder (2012b) argues that learning about *forthcoming evidence* or *easily available evidence* can function as an RKR against believing or disbelieving. Consider the following case. Detective D is investigating a murder. D has plenty of evidence that gardener G committed the crime. However, there are still some DNA test results that D is yet to learn about, which will be made available to D the next day. Now, according to the supposed intuition, learning that such critical evidence exists and is forthcoming appears to make believing that G is the murderer irrational.¹³ Similarly, it appears to be irrational to form a belief about the weather based on yesterday's forecast; instead, one ought to get out of bed and look out the window. Thus, the fact that one knows that some evidence would be really easy to access appears, again, to constitute an RKR against believing. Yet, neither the information about the forthcoming results, nor the fact about availability is itself evidence concerning the proposition in question. After all, the fact that DNA-results will be imminently available does not make G any more or less likely to be the murderer.¹⁴

¹² This assessment is shared, for instance, by Buchak (2014), Smith (2016, p. 36 f.), and Jackson (2020), among other scholars.

¹³ Again, the appeal to future evidence is somewhat contentious. I have met several philosophers who have expressed diverging intuitions concerning such cases. According to them, the subjects are rational in forming or retaining their beliefs in the face of learning about the forthcoming evidence.

¹⁴ As mentioned in footnote 9, Schmidt (2023) allows for certain facts about the first-order evidence to amount to RKRs of belief and disbelief. However, with regard to the availability or forthcomingness of evidence, it seems out of the question that these facts could justify either an affirmative or negative answer to 'Did G do it?' or 'Does it rain?'. Thus – in line with my analysis in subsequent sections – Schmidt instead views these kinds of facts as RKRs in favour of suspending judgment.

Therefore, whatever accounts for the irrationality of believing in these cases cannot be an OGR of belief or disbelief.

Fourth, while being a heavily debated subject, one might endorse *pragmatic* and/or *moral encroachment* as a feature of epistemic normativity.¹⁵ Paradigmatically, defenders of encroachment theses allude to pairs of cases such as the following (DeRose, 1992; Stanley, 2005):

Low Stakes

Hannah is driving home on a Friday afternoon and intends to deposit a paycheque at the bank just because it would be convenient. However, upon noticing long lines at the bank, she considers driving straight home and returning on Saturday morning instead. She remembers the bank being open on a Saturday in the past, and although she thinks to herself that it is possible that the bank changed its hours, she believes that it will be open and drives home.

High Stakes

Hannah is driving home on a Friday afternoon and intends to deposit an important paycheque. If she does not have enough money to pay her mortgage in her bank account by Sunday, her very livelihood is on the line. However, upon noticing long lines at the bank, she considers driving straight home and returning the next morning instead. She remembers the bank being open on a Saturday in the past, and although she thinks to herself that it's possible that the bank changed its hours, she believes that it will be open and drives home.

According to proponents of encroachment, the difference in the stakes involved in these cases creates a diverging assessment of Hannah's rationality. While her belief might be rational in the former case, it fails to be rational in the latter one. Provided that the evidence in both cases remains the same, a proponent of RKO is, again, faced with a puzzle: what could possibly account for this shift in Hannah's rationality? After all, there appears to be a need for some RKR against believing which is not evidence – that is, not an OGR of belief or disbelief.

Finally, following a more recent attack on fittingness-first theories of normativity, Lee (2022) points to the phenomenon of *higher-order defeat* to drive a wedge between RKO and epistemic normativity.¹⁶ Higher-order defeat differs from other forms of defeat by leaving the evidential support relation between the subject's evidence and the believed proposition unaltered.¹⁷ Consider the following case: after

¹⁵ For an overview of most of the debate on pragmatic encroachment see Kim (2017). See Bolinger (2020) for a presentation of moral encroachment theses. Here, I am only concerned with versions of moral encroachment that mirror the pragmatic one (e.g. Fritz, 2020).

¹⁶ In his paper, Lee points to two further features of epistemic normativity. First, certain *facts about the evidence*, such as it being tied, appear to be RKRs against believing. Since these kinds of features are already involved in the discussion of prohibitive balancing, I neglect them. Second, Lee points to *undercutting defeaters*, which one might regard as RKRs against believing. However, I find Schroeder's (2011) account of undercutting defeat in terms of attenuators convincing. Thus, I do not regard an undercutting defeater as an RKR against believing, but rather as a modulating factor that alters the weight of other RKRs of belief.

¹⁷ Therefore, it is neither a *rebutting* nor an *undercutting defeater*. For this point see Christensen (2010), Lasonen-Aarnio (2014), or DiPaolo (2018).

dwelling over some logical derivations, a subject S comes to believe that p entails q. However, S also learns that their drink was laced with a mind-altering drug that inhibits their capacity to perform logical derivations. According to some authors, if S were to continue to believe that p entails q, then S would be irrational. Therefore, one might again expect the higher-order defeater to create some RKR against believing.¹⁸ Since learning about the drug in the water is neither evidence in favour of the fact that p entails q, nor evidence against it, the envisaged RKR cannot be an OGR of belief or disbelief.

To conclude, there are five issues that arise when a thesis like RKOG is extended to account for epistemic normativity: (1) prohibitive balancing; (2) bare statistical evidence; (3) future and easily available evidence; (4) pragmatic and moral encroachment; and (5) higher-order defeat. While issues (2) to (5) rest on contentious assumptions, I take them for granted as correct depictions of epistemic normativity. In the next section, I introduce a feature of the doxastic option space which the opponents of RKOG have neglected – namely the OGRs of the suspension of judgment, the third doxastic option – before returning to these five puzzles in Sect. 5.¹⁹

4 On suspension and decidedness

In this section, I present the core tenets of my defence of RKOG against the aforementioned set of challenges. First, I argue that the third doxastic state of *suspension* is characterised by a different deliberative question than belief and disbelief (4.1). Therefore, it can be sensitive to different considerations, which provides a proponent of RKOG with the resources to rebut the five puzzles. This latter task is the subject of the subsequent section. However, before turning to this endeavour, my proposal must explain how RKRs of suspension, which are supposed to answer a different question than belief, can nonetheless count as RKRs against believing (4.2). My solution to this issue relies on a novel conception of the doxastic option space. I argue that belief and disbelief are *ways of being decided* on an issue, and therefore, are sensitive to a secondary deliberative question (4.3). The resulting picture is a two-scale weighing model of epistemic reasons, on which each scale is concerned with one of the two deliberative questions (4.4).

¹⁸ It should be noted that there are several diverging theories of the effects of higher-order defeat and only some of them can be made to put pressure on RKOG. For instance, according to Silva (2016) and Staffel (2021), higher-order defeaters only dissolve the *doxastic justification* of the subject's beliefs while leaving their *propositional justification* unaltered. By contrast, Whiting (2019) argues that higher-order defeaters are, in fact, WKRs against believing. Neither view on higher-order defeat would create the supposed challenge to RKOG.

¹⁹ Additional phenomena which are recurrently put forward against evidentialism and, by extension, against RKOG are *self-defeating* and *self-fulfilling beliefs* (Sharadin, 2016a; Raleigh, 2017; Antill, 2019; Silva & Tal, 2021). These attitudes are supposed to instantiate, respectively, instances of OGRs in favour of belief which are not RKRs and RKRs in favour of belief which are not OGRs. Since these two issues break with the pattern of missing RKRs against believing, I will neglect them in the discussion to come. Roughly put, a response to these self-affecting attitudes requires a shift in our understanding of OGRs. In terms of the Hieronymian framework, one has to rephrase the deliberative question of beliefs as 'If I were to believe p, would p be the case?'. Thereby, the equivalence claim of RKOG is saved.

4.1 The suspension of judgement and its deliberative question

As Schroeder (2012b) already notes as a common reply to his cases, a defender of RKOg might point to the OGRs of the third doxastic stance, namely the *suspension of judgement*. Since suspension is also a rationally evaluable attitude, it should also be judgement-sensitive like belief.²⁰ Given that suspension and (dis-)belief are in competition with each other – an assumption that I investigate further in the following subsections – an RKR in favour of suspension counts as an RKR against believing and disbelieving. Now, if suspension is connected to a judgement that answers a *different deliberative question* than belief and disbelief, a proponent of RKOg has some further resources to draw upon in handling the five puzzles: the missing RKRs against belief and disbelief stem from the OGRs of suspension.^{21,22}

Why should one believe that suspension is sensitive to a different question than belief and disbelief? Here, I provide a list of *prima facie* reasons to doubt that suspending on whether *p* is the case is connected to a judgement that answers ‘Is *p* the case?’. First, according to a widespread view, suspension is not directed at propositions but rather at *questions*. It is an *interrogative attitude*, rather than a propositional one (Friedman, 2013; Booth, 2014). One does not suspend *that p*; rather, one suspends on *whether p*. Hence, one might suspect that the characteristic question of suspension also differs.²³ Second, and relatedly, most analyses of the third doxastic stance entail that suspension should come with a different characteristic question. Consider, for instance, Friedman’s (2017, 2019) seminal theory according to which suspension is an attitude that combines open-mindedness on some question *Q* with a pro-attitude towards *inquiring* into whether *Q*. Such an attitude should be an answer to a question like ‘Is *Q* an interesting subject of inquiry?’. On the other hand, *higher-order accounts* of suspension regard the third doxastic stance to be linked to a higher-

²⁰ Here, following the vast majority of the literature, I assume that such a third attitude exists in the doxastic option space. However, opposition to this assumption should not go unnoticed (Crawford, 2022).

²¹ Note that given RKOg, one cannot simply appeal to RKRs in favour of *not having an attitude at all*. As far as this paper goes, I assume that the state of not having any attitude is not connected to a judgement that answers to some question *la Hieronymi*. Therefore, I do not regard not having any attitude as a proper member of the doxastic option space. This assessment is in strong contrast to that of Lilly (2019), who provides one of the only discussions of the OGRs of suspension. Consequently, in my discussion to come, I am not faced with the challenge of keeping the reasons in favour of suspension and in favour of not having any attitude apart.

²² Schroeder (2012b, p. 477) dismisses suspension as being able to handle his counterexamples. His argument rests on a contentious premise regarding the connection of the third doxastic stance and one’s *deliberation*: according to Schroeder, one can only suspend judgment *in deliberation*. Hence, for any non-deliberating subject, the OGRs of suspension should not play a role. Therefore, the five puzzles still emerge if one specifies that the subjects in question are not deliberating. However, I deny this premise of Schroeder’s. If one can assess the rationality of belief and disbelief in non-deliberating subjects, I deem it only natural that one can also assess the rationality of their suspension in those cases. Hence, suspension does not require deliberation.

²³ It should be noted that this divergence in objects is not necessary for my case. Thus, even if one finds proposition-directed accounts of suspension appealing, as, for instance, Rosa (2020) or Archer (2022) do, one might still endorse the view that suspension answers to a different question than belief and disbelief. After all, suspension might try to get a different property of these propositions right.

order cognitive state about one's epistemic situation.²⁴ Again, on such an account of suspension, it should answer a question like 'Am I in a good epistemic situation concerning p?'. Finally, every account in the literature which investigates the *fittingness* or *correctness* conditions of the third doxastic state seems to assume that these conditions differ from the ones of belief and disbelief, and that they resemble the correctness conditions of certain higher-order beliefs about one's epistemic situation.²⁵ Thus, these views seem to come with a similar upshot as the higher-order accounts of suspension concerning its characteristic question.

However, the strongest reason to opt for an alternative characteristic question for suspension is a much more straightforward one based on a basic platitude about suspension. Whatever suspending on whether p might be, it amounts to a certain degree of *neutrality* as to whether p is the case. Now, being neutral as to whether p hardly counts as an *answer* to a yes-or-no question like 'Is p the case?'.²⁶ Hence, one should not expect the judgement which corresponds to suspension to provide an answer to that question either. Therefore, I follow the lead of the previous pointers and assume that suspension answers a different question. While a full-fledged theory of suspension would go beyond the scope of this paper, I defend a, rough, characterisation of this deliberative question.

I maintain that the deliberative question governing suspension should be concerned with one's *epistemic situation*. It should not be about the truth or falsity of the proposition in question; rather, it should concern some epistemic propriety of my standing towards answering the deliberative question of belief and disbelief. This basic idea strikes me as quite plausible and is backed by most of the considerations that I alluded to above, such as higher-order views of suspension or work on the fittingness of suspension. The only line of resistance to this broad idea that I was able to identify stems from Friedman's *inquisitive attitude* view of suspension. After all, the degree to which a question is interesting and merits being inquired into differs from an assessment of my epistemic standing towards said question. While a complete rebuttal of this concern goes beyond the scope of this paper, I provide a quick sketch of how one should dismantle it.²⁷ First, as has been recurrently noted in response to Friedman's theory, the third doxastic state does not always come with such an openness to further inquiry (e.g., Ferrari & Incurvati, 2022). Sometimes, we remain doxastically neutral towards a question without putting it on our research agenda. Thus, second, if the deliberative question governing suspension should hold for both pro and anti-inquisitive instances of suspension, then it should refer to an issue common to both. Third, the question that is common to either variety of suspension appears to be nothing but the proposal under investigation – that is, some propriety of my epistemic situation. For the purposes of this paper, I assume that something in the

²⁴ Higher-order views of suspension have been proposed by Crawford (2004), Masny (2020), and Raleigh (2021).

²⁵ Compare this point to the ideas entertained by Rosa (2020), McHugh and Way (2022, Chap. 7), and Vollmer (2023).

²⁶ The pragmatics of a situation might dictate that an expression of doxastic neutrality is a proper *response* or *reply* to somebody's raising of a question. Yet, it does not count as an *answer*.

²⁷ I provide a more detailed depiction of this line of thought in Vollmer (Unpublished Manuscript).

ballpark of ‘*Am I in a position to tell whether p?*’ is the correct depiction of the deliberative question of suspension. As I argue in Sect. 5, such a proposal entails OGRs of suspension which bear some independent plausibility.

4.2 A looming problem

Before turning towards the investigation of these OGRs of suspension, however, the following challenge must be answered, which constitutes an apt reply to the strategy of defenders of RKOG as I have laid it out thus far. Given that suspension and (dis-)belief answer to different deliberative questions, why is it that they are even in the appropriate form of competition with each other? An OGR of belief and disbelief justifies answers to ‘Is *p* the case?’, while an OGR of suspension supports an answer to ‘Am I in the position to tell whether *p*?’. According to the envisaged manoeuvre, OGRs of suspension should constitute RKR against believing. Yet why should reasons that justify a negative stance on the latter question have any impact on how one replies to the former question? Consider a different pair of attitudes connected to diverging questions, such as admiring Ghandi and believing that there are 10 pencils on one’s desk. An OGR favouring this admiration does not constitute an RKR against the belief in question – thus, why should things be any different in the case of suspension and belief? Hence, without providing a theory of how the reasons of suspension bear on the rationality of belief, an appeal to a third doxastic state, which happens to be sensitive to different considerations, should strike any opponent of RKOG as an ad hoc manoeuvre to account for the five puzzles.

Therefore, to proceed with the strategy, one has to amend RKOG with some theory of how answers to different deliberative questions can still have a normative bearing on each other. I presume that there are different ways in which one can accommodate this challenge. For instance, one might point to the apparent disanalogies between pairs of suspension and belief and pairs of admiration and belief like the one above. While having both of the latter attitudes does not create any normative conflicts, suspending judgement on whether *p* while believing *p* gives rise to clear *normative tensions* (Friedman, 2017, 2019). An individual who forms both of these attitudes appears to be *structurally irrational*.²⁸ However, in the remainder of this section, I present a different approach to the connection of (dis-)belief and suspension which is able to circumvent the deep waters of structural rationality.²⁹ Instead, my preferred account can be seen as a slight extension of RKOG itself. It provides a somewhat deeper analysis of why suspension and belief are related, which explains why OGRs

²⁸ Such pairs of belief and suspension pass Worsnip’s (2021) test for structural irrationality – a subject to whom it is disclosed that they both believe *p* and suspend on whether *p* will change their doxastic state in due time. A similar result emerges if one turns to Lee’s (2023) recent analysis of structural rationality in terms of commitments – the commitments of believing *p* and suspending on whether *p* are not mutually satisfiable.

²⁹ Furthermore, one might suspect that an appeal to structural rationality at this juncture in one’s theory of RKR is a contentious move. After all, according to the prominent works of Kolodny (2005), Kiesewetter (2017), and Lord (2018), the principles of structural rationality should reduce to facts about the subject’s RKR. However, according to the envisaged theory, one requires the concept of structural rationality *prior* to delineating the set of RKR. Thus, this invocation of structural rationality is incompatible with one of the main contenders in an analysis of this variety of rationality.

in favour of the former are RKR against the latter as well as why there is a normative tension in having both attitudes.

4.3 Ways of being decided

To introduce my solution to the abovementioned challenge, it makes sense to draw attention to the following considerations: suspension, as described thus far, is an attitude that is sensitive to a *negative* judgement, answering the question of ‘Am I in a position to tell whether p?’. However, is there also an attitude that is sensitive to an *affirmative* judgement on that question? That is, is there any *antagonistic* attitude opposed to suspension in the way that belief is opposed to disbelief? Given that suspension as the third, neutral doxastic state can be seen as a form of *undecidedness* concerning the deliberative question governing belief and disbelief, it appears only natural to look out for some attitude of *decidedness* with regard to that question, which comes with a commitment to the affirmative judgement of ‘I am in a position to tell whether p.’. Postulating the existence of such a fourth doxastic option of decidedness would neatly fit into the picture of doxastic deliberation that I have envisaged up to this point.

However, assuming such an additional, fourth doxastic attitude of decidedness strikes me as quite contentious. It is clearly not in consonance with the state of the literature in epistemology. There, the consensus appears to be that the doxastic option space only has three members ($\{\text{believe; suspend; disbelief}\}$), rather than four. Hence, even if there might be such a *sui generis* attitude of decidedness, models of doxastic deliberation seem to be able to do without it.³⁰ This first, negative observation leads me to a second, positive one, which can be regarded as the focal point of my new account: instead of requiring any distinct attitude of decidedness, our attitudes of belief and disbelief function as *ways of being decided*. Decidedness on the question of whether p can be regarded as a *disjunctive state*, which is either manifested by believing p or by disbelieving p. Therefore, qua ways of being decided on the matter, belief and disbelief come with a commitment to the deliberative question of suspension in that they are sensitive to an affirmative judgement, answering the question ‘Am I in a position to tell whether p?’. Thus, belief and disbelief play a *dual role* in deliberation. They are sensitive to two kinds of judgement; consequently, they constitute answers to two deliberative questions.

Such a dual role of belief and disbelief would lend a defender of RKOG all the necessary resources to accommodate the puzzles from the previous subsection. First, since belief and disbelief, qua decidedness, commit one to an affirmative judgement on ‘Am I in a position to tell whether p?’, they must be in rational tension with suspension. Just as belief and disbelief are in competition with each other since they are related to opposed judgements on one and the same deliberative question, so too are suspension and (dis-)belief. Second, and relatedly, an OGR in favour of suspension

³⁰ Of course, it is not hard to find some attitude which is linked to a positive answer to the deliberative question of suspension. After all, a higher-order belief about one’s being in a position to tell whether p would fit the bill. And, if higher-order views of suspension are correct, such an affirmative higher-order cognitive state might have any right to count as a distinct doxastic option.

would be an OGR against the disjunctive option of decidedness, and, given RKOg, it should count as an RKR against being decided on the matter. Furthermore, a reason against a disjunctive option ϕ -or- ψ should entail a reason against performing either ϕ or ψ . Thus, the RKR against decidedness explains why an OGR in favour of suspension would constitute an RKR against belief and disbelief. Therefore, both open issues generated by attributing a divergent deliberative question to suspension can be solved by viewing belief and disbelief as ways of being decided.

Before moving on to the details of how to conceive of the weighing of reasons on such a bifurcated picture of doxastic deliberation, some qualms remain that my proposed account faces: how can belief and disbelief be related to two deliberative questions? Relatedly, how can they be sensitive to two diverging sets of OGRs – one related to the question ‘Is p the case?’ and the other to ‘Am I in a position to tell whether p ?’? Stated succinctly, I need to provide some more motivation for how something like the abovementioned dual role of belief and disbelief could be possible.³¹

While my appeal to a dual role of attitudes might strike one as mysterious, I take it to be the natural consequence of translating a much less mysterious idea in the literature on normative reasons into an RKOg framework. Here, I refer to the idea that at least some normative reasons are *contrastive*: a consideration R might be a normative reason to ϕ *rather than* ψ , but not a reason to ϕ *rather than* χ .³² When applied to epistemology, authors such as Snedegar (2017) and Tucker (forth.) have defended the idea that epistemic reasons must be conceived of in this fashion. A consideration does not simply favour believing or suspending. Instead, it supports believing *rather than* suspending, or believing *rather than* disbelieving. Now, if one conceives of these RKRs in terms of RKOg, one faces a *prima facie* difficulty in accommodating this contrastivist idea. If a consideration supports the judgement on ‘Is p the case?’ to which belief is sensitive, then why should that not count as a reason in favour of belief, regardless of whether one compares it to disbelief or suspension? However, this shortcoming can be lifted if one turns to my dual-role view of some attitudes: a contrastive reason to ϕ rather than ψ must be an OGR that pertains to the deliberative question A to which both ϕ and ψ are sensitive qua being related to incompatible judgements on A . However, this does not entail that the same consideration will also be an OGR that favours ϕ over χ if these two attitudes are related to incompatible judgements on a different deliberative question B . To state this less abstractly, in deliberating about whether to believe *rather than* disbelieve, one seeks to answer one question – namely ‘Is p the case?’ – and in deliberating about whether to believe *rather than* suspend one is answering a different one – namely ‘Am I in a position to tell whether p ?’. To conclude, while it is a substantial commitment of my account,

³¹ I have to thank one of the reviewers of *Philosophical Studies* for pressing me on this issue.

³² For an overview of this concept see the seminal work by Snedegar (2017) who, in his book, defended the thesis that *any* reason must be conceived in a contrastive manner. For a forceful defence of the idea that the weight of *at least some* kinds of reasons merit a contrastive reading see Muñoz (2021).

the mere possibility of having an attitude be sensitive to two diverging deliberative questions and sets of OGRs does not strike me as overly problematic.³³

4.4 The resulting picture

Based on the aforementioned account of doxastic deliberation, I conceive it as being concerned with two distinct questions and two *prima facie* independent sets of OGRs that bear on those questions. Thus, to get a grasp of the upshots of my outlook, one must use a *two-scale* weighing model. Each scale corresponds to one of the two deliberative questions and the corresponding set of OGRs. On *scale A*, concerning the question ‘Is *p* the case?’, the OGRs of belief and disbelief are weighed against each other. Following the assumptions from the previous section, the result on scale *A* is purely determined by the evidence that concerns the proposition in question. Hence, the new components which my theory of suspension and decidedness adds do not alter the result on this scale. However, next to this traditional picture of weighing epistemic reasons, there is also a *scale B*, which is concerned with the question ‘Am I in a position to tell whether *p*?’. Here, the OGRs of suspension and decidedness are the important factors. To put this picture into the idiom of contrastive reasons, scale *A* concerns the reasons to believe *p* rather than disbelieve *p* (and vice versa), while scale *B* concerns the reasons to {believe *p* or disbelieve *p*} rather than suspend on whether *p* (and vice versa).

Now, to derive the *overall status* of either of these options, I propose the following, natural suggestions: suspending on whether *p* is rational iff suspension does not lose on scale *B* – that is, the OGRs to suspend are *at least as weighty as* the OGRs to be decided on whether *p*. Believing *p*, on the other hand, is rational iff believing does not lose on scale *A* and decidedness does not lose on scale *B*. Analogously, disbelieving *p* is rational iff disbelieving does not lose on scale *A* and decidedness does not lose on scale *B*. Therefore, an OGR in favour of suspension can make believing *p* impermissible: it just has to outweigh the reasons in favour of decidedness on scale *B*.³⁴

³³ One might wonder whether this theory can be spelled out on other accounts of OGRs as well. I assume here that it can. Consider, for instance, correctness- or fittingness-centred theories. My dual-role view of belief entails that this attitude comes with two relevant standards of correctness or fittingness, where one is relevant for determining the OGRs in favour of believing rather than disbelieving and the other is relevant for determining the OGRs in favour of believing rather than suspending.

³⁴ Alternatively, one might also put this in a more traditional format in that an option counts as rational iff it *does not lose out* to any of the competing options in a *pairwise comparison*. Thus, if one looks at the doxastic option space {believe; suspend; disbelieve}, one can determine the rationality of either option by examining how well they fare with regard to the two competing ones. Again, following the insights of this section, the set of OGRs relevant for a comparison between belief and disbelief *differs* from those relevant for a comparison between belief and suspension.

5 Revisiting the five puzzles

In this section, I delve into the question of what kind of considerations justify answers to ‘Am I in a position to tell whether p?’. Specifically, I investigate what the OGRs of suspension and decidedness might be and, thus, which factors determine the outcome on scale B. I organise my thoughts on this matter in such a manner that accounting for these OGRs makes solutions to the five puzzles from Sect. 3 apparent. First, I present some general thoughts regarding how my account is able to handle the phenomenon of prohibitive balancing (5.1). Second, I argue that it is an independently plausible position that *live error possibilities* are OGRs of suspension (5.2). This observation grants a proponent of RKO the required resources to handle the behaviour of statistical evidence, cases of future and easily available evidence, as well as pragmatic and moral encroachment. Finally, I propose that higher-order defeat might work quite similarly to undercutting defeat in that it acts as an *enabler* or *modulator* concerning the OGRs relevant on scale B (5.3).

5.1 On the balance of evidence

My novel approach to the analysis of prohibitive balancing rests on the following claim: evidence concerning p – the OGRs at work on scale A – are not themselves at work on scale B.³⁵ Instead, I wish to argue that scale B is solely occupied with certain *facts about the evidence*. That certain facts about the evidence are RKR of suspension is an intuitively plausible idea, and one that has been recurrently noted in the literature. Roeber (2016) and Schmidt (2023), for instance, have both explained the prohibitive balancing of epistemic normativity by treating the fact that the evidence is tied as such an RKR. Furthermore, the account of suspension presented in the previous section shares this prediction. After all, one seems to be perfectly able to justify a negative stance on ‘Am I in a position to tell whether p?’ by pointing to the fact that one’s evidence concerning p is tied. I propose that one should think about the OGRs in favour of decidedness in an analogue fashion. Just like its antagonist, one can *only* justify a positive stance on ‘Am I in a position to tell whether p?’ by pointing to certain facts about the evidence, rather than to the evidence itself. The most straightforward fact about the evidence that would justify the disjunctive option of decidedness is plausibly that one’s set of evidence is *decisive* (i.e., the fact that there is a large swing on scale A).

However, one might question the restriction of RKRs in favour of decidedness to facts about the evidence. Is the first-order evidence itself not also an RKR of decidedness? If one receives evidence that p , can one not become decided rather than suspend *for that reason*?³⁶ While this observation might seem compelling on first sight, one should also consider the following phenomenon: countervailing evidence does not act as a *rebutting defeater* of the supposed RKR to become decided, but as an *undercutting* one. That is, if one would also receive the same amount of evidence that

³⁵ To put it in the idiom of contrastive reasons, evidence of p is a reason to believe rather than disbelieve p , but it is not (on its own) a reason to believe rather than suspend.

³⁶ I thank one of the reviewers of *Philosophical Studies* for raising this concern.

not-*p*, there would not be *any* reason to become decided anymore. This result looks odd on the view that the first-order evidence itself is a reason to become decided: even if we add countervailing evidence to the picture, the first-order evidence does not cease to exist. Yet, this result fits perfectly fine with my view that only facts about the evidence constitute RKR of decidedness. After all, by adding countervailing evidence, the facts about the evidence (e.g., its decisiveness) change. Thus, even in simple cases in which one just receives some evidence that *p*, I take it that the RKR for which you become decided is a fact about that evidence and not the evidence itself.

If these observations are on the right track, one has a very straightforward approach for accommodating prohibitive balancing. If the evidence on scale A is close to being tied, the OGRs in favour of suspension outweigh the OGRs in favour of decidedness on scale B. Thus, in light of the theory presented in the previous section, both ways of being decided become impermissible. This picture of prohibitive balancing is, I argue, significantly simpler than its competitors and bears some intuitive plausibility. The central component which makes such a neat explanation of prohibitive balancing possible is the somewhat novel idea that evidence in favour of *p* – the paradigmatic RKR of beliefs – is not directly weighed against the RKR of suspension. This shift in conceiving of the weighing of epistemic reasons makes my approach immune to two problems that plague its rivals and require them to opt for increasingly complicated pictures.

First, several authors have expressed doubt regarding the idea that one can allude to any such fact about the evidence as an additional epistemic reason.³⁷ The following is a *prima facie* detrimental objection to the appeal to any such metareason concerning the balance of one's first-order evidence: the balance of evidence can only be a *derived reason*, and thus, it should not make any difference to a weighing of reasons which already makes use of those considerations from which the derived reasons are derived, namely the evidence supporting *p* and evidence supporting non-*p*. Adding facts about the evidence to one's weighing of epistemic reasons amounts to an illegitimate form of *double counting*. However, the usage of facts about the evidence in my two-scale model is completely immune to this challenge. After all, at no point do I weigh the fundamental reasons – evidence concerning *p* – against facts about the evidence – that it is balanced or decisive. Instead, both sets of considerations are separated and each is weighed on its own scale. Thus, by not having evidence itself amount to a reason on scale B, I am able to circumvent this severe problem.

Second, accounts of prohibitive balancing recurrently struggle to accommodate two diverging prohibitive balancing scenarios. After all, the evidence concerning *p* can be tied, either, in cases in which one is in the possession of close to no evidence, or in cases in which much evidence is accumulated on either side. In short, the weight of the reasons in favour of suspension should be *flexible* enough to make plausible predictions concerning both scenarios.³⁸ However, if – as in my account – evidence that *p* is not directly weighed against suspension, one does not face any problem in

³⁷ This scepticism has been defended by Berker (2018, p. 450). Earlier invocations might be at work in Schroeder (2012a, p. 276) and Snedegar (2017, p. 125 Fn 13). Special thanks to Chris Tucker for pointing me to some of these references.

³⁸ For discussions of this problem see Snedegar (2017, p. 124 ff.), Brunero (2022), and Tucker (forth.).

accommodating the different ways in which the evidence might be tied. Regardless of the total amount of evidence at work on scale A, the weighing of reasons on scale B – which explains prohibitive balancing – remains the same. After all, it is only the swing of scale A that matters for scale B, not the total weight of evidence on either side.

5.2 Live error possibilities as reasons to suspend

In this subsection, I wish to draw the reader's attention to a different sort of fact about the evidence – namely the fact that it is able or unable to exclude some *salient error possibility*. I use the term '*live error possibilities*' (LEPs) to denote an error possibility which, on the one hand, is made salient by some contextual factors, and, on the other hand, is not ruled out by one's evidence. Just like being balanced and being decisive, the existence of an LEP is a fact about one's evidence. Furthermore, there is good reason to believe that this different sort of fact about the evidence is also at work on scale B. After all, in doxastic deliberation, pointing towards an LEP seems to have the envisaged effect: it rationalises the third doxastic stance of suspension over belief and disbelief. I regard this to be an independently plausible datum. Furthermore, it sits well and thus supports the idea that suspension is governed by a question along the lines of 'Am I in a position to tell whether p?'. After all, an LEP can justify a negative judgement on that question.³⁹

Suspension's general sensitivity to LEPs entails some interesting upshots. First, it explains the odd behaviour of bare statistical evidence. In Blue Bus, learning about the distribution of buses in the area makes one error possibility salient which is in principle not ruled out by the statistical evidence: perhaps the bus involved in the accident was one of the 10% of vehicles that are not part of the Blue Bus Company.⁴⁰ Therefore, there is an OGR in favour of suspension which can explain why the evidence in question is not sufficient for rationalising a belief. After all, if the LEP in question makes suspension win out on scale B, believing ceases to be a rational option. It is important to note that this behaviour of statistical evidence does not preclude it from *sometimes* rationalising beliefs. This is in consonance with much of the literature on statistical evidence, where scholars have attempted to explain why some statistics appear to be ill foundations for rational beliefs while others appear fine.⁴¹ Provided that statistical evidence can be more or less capable of making possibilities of error salient, one can explain this fact in terms of the existence of LEPs – and thus, by pointing towards the OGRs in favour of suspension.

³⁹ I consider there to be different ways to spell out the resulting picture, which are, in general, compatible with my account. For instance, as one reviewer of *Philosophical Studies* suggested, a single LEP should already be sufficient for rationalising suspension. However, one might also opt for a more lenient approach, according to which believing in the face of some LEPs can still be allowed.

⁴⁰ An explanation of this feature of statistical evidence in terms of salient error possibilities is proposed by Jackson (2020). According to her, however, it is belief itself that is sensitive to these open possibilities.

⁴¹ Consider, for instance, Smith's (2016) normic support view. In cases in which statistical evidence satisfies the normality constrain, it suffices to rationalise belief. For a diverging criterion of what demarcates rationalising statistical evidence see Silva (2023).

Second, I maintain that intuitions regarding cases of future and easily available evidence are also best explained in terms of LEPs: learning about the availability of good evidence on a question *only* constitutes an RKR against believing *insofar* as it makes certain possibilities of error salient which are not yet excluded by one's evidence.⁴² Consider the example of detective D from Sect. 3, who learns about forthcoming DNA results. This information should make the possibility of the DNA results *defeating* D's case against the gardener G salient; and by extension, some error possibilities of how G did not commit the murder should also become salient, such as the possibility that the butler framed G. A similar story can be told about easily available evidence. The fact that I know that getting out of bed and looking out the window can easily generate a defeater for my evidence that it will rain today makes the error possibility that the weather forecast was wrong salient.⁴³

Third, suspension's sensitivity to LEPs can also yield an explanation of pragmatic and moral encroachment. The idea that the stakes involved in paradigmatic cases of encroachment broaden the range of salient error possibilities is one standard approach to this supposed aspect of epistemic normativity.⁴⁴ In the pair of cases from Sect. 3, Hannah does not behave rationally in High Stakes because she has not yet ruled out some salient error possibility. The possibility that the bank changed its opening hours since she last checked is made salient due to the practical aspects of the context. In Low Stakes, Hannah is not required to exclude such a distant possibility. Therefore, in High Stakes, an LEP generates an additional OGR in favour of suspension which, in turn, explains why she counts as irrational in forming her belief.

5.3 Higher-order defeat revisited

The idea that an appeal to suspension and its OGRs can help us to understand the nature of higher-order defeat is already proposed by Lord and Sylvan (2021). According to them, treating higher-order defeaters as reasons to suspend judgement is a plausible position in the debate about higher-order evidence and is in line with their analysis of what it is to suspend judgement. However, I propose a slightly different connection between higher-order defeat and the OGRs of suspension, which is based on the following two observations.⁴⁵

⁴² This approach to forthcoming evidence has the additional perk of explaining at least some of the diverging intuitions that people have towards the rationality of the subject's belief. If learning about the evidence does not make the corresponding error possibility salient to you, it is only natural to regard the subjects in question as being rational in maintaining their beliefs.

⁴³ One recurrent reply to Schroeder's cases is that one must distinguish the rationality of one's belief from the rationality of deliberating about a certain matter (Hieronymi, 2013; Shah & Silverstein, 2013). Thus, one might maintain that all that the above cases show is not that believing is irrational, but only that one should *wait* or *postpone* one's doxastic decision making (McGrath, 2021a, b; Crawford, 2022). While being sympathetic to a reply along these lines, my account comes with a distinct advantage: it offers a straightforward way to accommodate subjects who have already made up their mind before learning about the forthcoming evidence.

⁴⁴ For a recent explanation of the encroachment mechanism along these lines see, for instance, Mueller (2021).

⁴⁵ I should again emphasize that this position on higher-order defeat is not essential to my position and is only meant as a reply to the thesis that higher-order defeat has a bearing on the RKR of belief and

First, while higher-order defeat leaves the evidential support relation intact, this does not mean that it does not work like an undercutting defeater in other regards. Consider the case of the laced drink from Sect. 3. My evidence on the logic puzzle is still very good and the way in which this reason bears on the question ‘Does p entail q ?’ is not altered. Furthermore, my evidence is quite decisive. *However*, the degree to which my evidence on whether p entails q is decisive does not seem to speak in favour of an affirmative answer to ‘Am I in a position to tell whether p ?’. Since I received a higher-order defeater in the form of information about my laced drink, the decisiveness of my evidence no longer appears to be a reason in favour of decidedness.⁴⁶ Thus, concerning the weighing on scale A, higher-order defeat does not have any effect. Nevertheless, on scale B it acts as a *disabler* or *attenuator* in the sense of Dancy (2004, Chap. 3): the information about my drink is not itself a reason against being decided, but it undermines or modulates the support relations in which other reasons stand to this option.⁴⁷ Therefore, just as in the case of prohibitive balancing, an individual who is subject to higher-order defeat might lack any positive reason in favour of being decided on an issue.

What remains to be demonstrated to accommodate the objection, is that there are also some OGRs in favour of suspension in these cases. This brings me to my second observation, which is that pointing to LEPs seems to be an appropriate way to justify suspension in cases of higher-order defeat. The possibility that there is an instance of p being true without q being true is clearly ruled out by my evidence on whether p entails q . Yet, given the higher-order defeater at work in these cases, this error possibility seems to be live nonetheless. The information about the laced drink strikes me as re-enabling those once-excluded possibilities of error. Thus, next to being a disabler or attenuator of the OGRs in favour of decidedness, there is a good case to be made that higher-order defeat also functions as an *enabler* or *intensifier* of some OGRs in favour of suspension. Therefore, in cases of higher-order defeat, there should be a special preponderance of OGRs in favour of suspension and, given RKOG, an additional RKR against believing.

6 Permissive balancing and suspending intentions

The previously specified account is able to defend RKOG against the objections that stem from epistemic normativity by *extending* the set of RKRs against believing. Yet, the following pressing question remains: how does RKOG fare with regard to other attitudes? In particular, one might worry that RKOG *overgenerates* RKRs against *intentions*. After spelling out this objection in detail (6.1), I point towards a reasonable and independently motivated rejoinder on the side of RKOG (6.2). I argue that

disbelief.

⁴⁶ Even though it might be less easy to see, I suppose that the same holds true for facts about the balance of evidence. If I am in a scenario of tied evidence and receive information about the drug in my drink, then the balance of evidence does not seem to speak in favour of suspension in the same way it did before.

⁴⁷ By contrast, an undercutting defeater works as a disabler or attenuator on the reasons weighed on scale A (Schroeder, 2011). Thus, my position is still able to demarcate an important difference between these two forms of defeat.

the desired difference between doxastic and practical deliberation can be explained by a difference in the kind of attitudes at work.

6.1 A remaining worry

As I mentioned in Sect. 3, prohibitive balancing is supposed to be a feature that is peculiar to epistemic normativity. By contrast, practical normativity is said to balance in a *permissive* manner in that two options that are equally supported by one's reasons both emerge as permissible. Now, one might consider the previous account of suspending judgement to raise an issue when applied to *conative* states or *intentions*. According to some authors, just as there is a third, neutral doxastic state, there is also a third, neutral conative state, namely the *suspension of intentions*.⁴⁸ Just as one can remain undecided about how the world is, one can also remain undecided with regard to what to do. Thus, *prima facie*, the suspension of intentions should receive the very same treatment as the suspension of beliefs. However, if this point holds up to scrutiny, there is a case to be made that practical normativity exemplifies the very same kind of prohibitive balancing as epistemic normativity.

The rationale for this conclusion is as follows: intentions, just like beliefs, are judgement-sensitive attitudes.⁴⁹ Instead of being sensitive to considerations that bear on the question 'Is p the case?', OGRs of intentions justify answers to the question 'Is it right to φ ?'.⁵⁰ For the sake of simplicity, I only consider two-option cases; therefore, intending to φ and intending not to φ are the only relevant attitudes at play. Given the assumption above that the doxastic and conative attitudes work in parallel, intending to φ and intending not to φ are also *ways of being decided* with regard to a particular choice. Qua manifesting decidedness, intentions also answer a secondary deliberative question – namely 'Am I in a position to tell whether it is right to φ ?'. The other relevant attitude, which is sensitive to a negative judgement on this question, is the suspension of intentions. Thus, practical deliberation can also be characterised by a two-scale weighing model: on *scale A*, intending to φ competes with intending not to φ , and the relevant OGRs are considerations about the rightness of φ -ing. On *scale B*, however, the relevant OGRs should again be *facts about one's evidence*, such as the degree to which one's evidence that φ -ing is right is decisive, or LEPs concerning the rightness of φ -ing.

Now, one might argue that this model of practical deliberation is able to provide the correct predictions in the case of Buridan's ass. After all, more than one option can be a right thing to do. Thus, the donkey is in a position to determine whether eating one bale of hay rather than the other is a right thing to do. Therefore, it has good

⁴⁸ An explicit discussion of such a state can be found in Snedegar (2017, Chap. 6). Even though Schroeder (2012b) does not discuss the third doxastic option, his parallel treatment of belief and intention seems to suggest that, if there is such a form of suspension in the former case, there should also be a similar form of suspension in the latter case. Finally, while Forcehimes (2021) talks about *desire* rather than intentions, he acknowledges a third, flat-out desiring state – the withholding of desire.

⁴⁹ For the purposes of this paper, I ignore other takes on the nature of intentions. See Setiya (2022) for an overview.

⁵⁰ Now, the exact analysis of the deliberative question guiding intentions is a controversial matter. For different proposals see Hieronymi (2005), Shah (2008), or McHugh and Way (2022, Chap. 3) among others.

reasons to be decided on the matter and, thus, suspension does not interfere with permissive balancing.⁵¹ However, consider the following scenario:

Miners

After an earthquake, a group of miners is trapped in a mine, which water is now entering. The miners are either in shaft A or shaft B. If the team of rescuers on the surface does not flood one of the shafts in due time, the miners are not going to survive. However, the rescuers do not know which shaft the miners are in.

In this case, only one of the two options appears to be the *objectively right* thing to do. However, the rescuers are not in a position to tell whether flooding shaft A, for example, is the objectively right option. Their evidence is balanced, and just as in the doxastic case, suspension should win out on scale B. Therefore, the rescuers ought to suspend their intentions, and both intending to flood shaft A and intending to flood shaft B are irrational. To conclude, if the preceding line of thought is on the right track, my defence of RKOg with regard to epistemic normativity appears to commit me to a problematic position on balancing cases of practical normativity.

6.2 Solution: suspending intentions differs from suspending beliefs

A defender of my take on doxastic deliberation could employ various strategies to avoid such a problematic outcome. After all, several assumptions in the above rationale can be subjected to doubt. For instance, one might question whether intentions aim at objectively rather than *subjectively right* actions. In that case, the miners case would merit the same treatment as Buridan's ass.⁵² Alternatively, one could deny that there exists something like the suspension of intentions, or that suspending intentions does not amount to a judgement-sensitive attitude at all.⁵³ However, my preferred response to the challenge is to keep intentions and beliefs aligned and allow for the existence of a judgement-sensitive attitude of suspending intentions, but to deny that suspending intentions and being decided on what to do answer to the same kind of question as suspending judgement and being doxastically decided. As I demonstrate below, this position is not ad hoc and bears some independent motivation.

I provide a rough draft of how I view the suspension of intentions: in contrast to suspending judgement, suspending intentions is not connected to a judgement answering the question 'Am I in a position to tell whether ϕ -ing is a right thing to do?'. Instead, suspending intentions is linked to the following much more *practical*

⁵¹ Of course, this manner of handling Buridan's ass is completely contingent on my assumption about the deliberative question governing intentions. If one opts for some less permissive variant, such as 'Ought I to ϕ ?', then one might already face the issue that I derive below.

⁵² Such a reply might be questioned by objectivists about obligations (Graham, 2010). At the very least, it makes a uniform treatment of objectivist and subjectivist deontic notions across epistemology and ethics impossible. For a comprehensive depiction of the merits of such a unified framework see Whiting (2022).

⁵³ Again, responding along these lines comes with significant costs. After all, as intended by Schroeder, one should expect that most of the problematic cases of RKR against believing can be reproduced in the case of intentions. Therefore, a proponent of RKOg would need to appeal to the OGRs of a third conative state anyway.

question: should I commit to a plan or form an intention regarding some choice of course of action? The same holds true for conative decidedness. Thus, the OGRs on our practical scale B are *practical considerations* about whether one should reach a decision on some issue. Such a shift in the normative profile of suspension and decidedness yields a straightforward account of Miners. The balance of evidence concerning which action is the right one is not necessarily a reason in favour of suspending one's intentions. Quite the opposite, the practical need to commit to some course of action given the time constraints of the situation might be weighty OGRs in favour of decidedness. Thus, practical normativity balances in a permissive fashion.

One objection that might be raised against such a view of suspending intentions is already raised by Schroeder (2012b, p. 476): given that we make suspension sensitive to practical considerations, how can we avoid the emergence of a *WKR's problem*. Following Kavka's (1983) toxin puzzle, the standard position on the WKRs of intentions is that they are practical considerations that favour the *intentions* to act rather than the *act* itself. After all, a monetary incentive to *intend* to drink a toxin does not make the *act* of drinking the toxin any more or less right. However, given my depiction of the question governing practical decidedness and suspension, such monetary incentives appear to make a difference: a monetary incentive *to be decided* on whether one should drink a toxin falls under the OGRs of decidedness and, given RKOG, should be an RKR to be decided.

My response to this concern consists of two parts. First, it is important to note that an OGR on scale B is *not* an OGR on scale A. Thus, allowing for incentives to be decided by manifesting *either* the intention to φ *or* the intention not to φ as RKRs in favour of becoming decided does not entail that a monetary incentive to intend φ is a RKR in favour of intending φ . Therefore, my account does not open the floodgates to WKRs of intentions. The only Kavkaesque scenarios that it allows for look quite different from the original toxin puzzle. Here, a billionaire might offer you a monetary incentive to *decide whether* you plan to drink the toxin tomorrow. It is perfectly possible to take the billionaire's money and make the rational decision not to drink the toxin.

Second, and in stark contrast to Schroeder's assessment, I take this result to be a feature of my theory rather than a bug: counting such monetary incentives as RKRs to come to a decision is the correct prediction. After all, paradigmatic RKRs to become decided concerning a choice or to suspend one's intentions look quite similar to these monetary incentives. Above, I alluded to the *time constraints* of a given situation: given that postponing one's decision comes with some practical downsides, one ought to decide sooner rather than later. These practical downsides might pertain to the action under investigation, such as the pressure to act in Miners, but they might also be related to the success of some other action of yours, or even actions of other people in cases of coordination. In fact, I do not consider any stark difference to exist between these kinds of practical considerations and the monetary incentive of the billionaire envisaged above. Therefore, instead of constituting an objection to my

account, reflection on the RKR of practical decidedness brings some independent motivation of my position to light.⁵⁴

7 Conclusion

I consider the contribution of this paper to be threefold. First, it introduces a new, innovative approach to weighing epistemic reasons and yields a natural picture of how suspension and (dis-)belief rival each other as doxastic options. The twofold weighing of reasons that concern different questions strikes me as a more adequate model of our deliberation than any of the simple weighing pictures in the literature, which run into all sorts of issues by cramping suspension, belief, and disbelief onto the same scale. Second, by making use of this picture of suspension, one is able to accommodate all of the supposed features of epistemic normativity, which are said to cause trouble for RKOG. The OGRs of suspension and decidedness can explain prohibitive balancing, the behaviour of statistical evidence, pragmatic and moral encroachment, the significance of forthcoming evidence, and higher-order defeat. Finally, as the previous section has demonstrated, the two-scale weighing model can plausibly be extended to the realm of practical deliberation without generating a form of prohibitive balancing.

In closing, I wish to flag some routes for further research, which the foregoing discussion has left open. First, while this work has highlighted the potential behind the additional resources of the two-scale model, many of the details, especially concerning the reasons on scale B had to remain vague. Thus, it remains an avenue for future research to investigate whether there are any additional RKRs to be found on this scale, as well as how those RKRs interact.⁵⁵ Second, the final result of this paper is also a clear gateway to a further analysis of practical deliberation. There might be many puzzling features of practical normativity which can also be resolved by pointing to the OGRs of suspension and decidedness. Therefore, there might be further theoretical payoffs down the line. Third, the inquiry of suspension of other forms of attitudes is still very much an open field. Hence, it would be interesting to investigate whether one can extend the present picture of suspension to these attitudes as well.⁵⁶ Furthermore, given that I argue that suspending intentions differs from doxastic suspension in a crucial way, it remains to be seen whether the suspension of any other attitudes is more in line with the former or the latter.

⁵⁴ Note that by making suspending intentions sensitive to practical considerations one is also able to handle the extension of some of the five puzzles in Sect. 3 to intentions. For instance, as Schroeder (2012b) argues, forthcoming information should constitute an RKR against intending. An account like mine can accommodate this challenge by pointing to the practical benefits of having a better-informed decision.

⁵⁵ Concerning the former question, one might envisage cases in which one receives testimony that one is or is not in a position to tell whether *p* (Turri, 2012). Such testimony, one might argue, should be weighed on scale B as well. Zinke (2021) points towards vagueness and objective chance as two further RKRs of suspension. The issue of practical RKRs to suspend, which have been envisaged by Snedegar (2017) and Lord (2020), is discussed by me in Vollmer (Unpublished Manuscript).

⁵⁶ In his recent paper, Fritz (2022) defends the idea of an *alignment* between doxastic attitudes and affective ones. Thus, given the results of the previous discussion, an affirmative answer to the abovementioned question appears to be very much in his spirit.

Acknowledgements Parts of this work have been presented at the *ICPR Doctoral Workshop* at the University of Innsbruck, the *Contemporary Issues Across Ethics and Epistemology* graduate conference at the University of Pavia, the *Goethe Epistemology Meeting* at the University of Frankfurt, as well as the *Southampton-Humboldt Normativity Conference* at the University of Southampton. I am grateful to all the feedback I received on these occasions. Furthermore, I thank the reviewers of *Philosophical Studies* for their scrutiny, which helped to improve this paper in several ways.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Competing interests There is no potential conflict of interest to disclose.

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 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

- Antill, G. (2019). Evidence and self-fulfilling beliefs. *American Philosophical Quarterly*, 56(4), 319–331. <https://doi.org/10.2307/48563046>.
- Archer, A. (2022). The questioning-attitude account of agnosticism. *Synthese*, 200(498), 1–15. <https://doi.org/10.1007/s11229-022-03971-w>.
- Berker, S. (2018). A combinatorial argument against practical reasons for belief. *Analytic Philosophy*, 59(4), 427–470. <https://doi.org/10.1111/phib.12140>.
- Bolinger, R. (2020). Varieties of moral encroachment. *Philosophical Perspectives*, 34, 5–26. <https://doi.org/10.1111/phpe.12124>.
- Booth, A. R. (2014). Two reasons why epistemic reasons are not object-given reasons. *Philosophy and Phenomenological Research*, 89, 1–14. <https://doi.org/10.1111/j.1933-1592.2012.00631.x>.
- Brunero, J. (2022). Practical reasons, theoretical reasons, and permissive and prohibitive balancing. *Synthese*, 200(2), 1–23. <https://doi.org/10.1007/s11229-022-03628-8>.
- Buchack, L. (2014). Belief, credence, and norms. *Philosophical Studies*, 169(2), 1–27. <https://doi.org/10.1007/s11098-013-0182-y>.
- Christensen, D. (2010). Higher-order evidence. *Philosophy and Phenomenological Research*, 81, 185–215. <https://doi.org/10.1111/j.1933-1592.2010.00366.x>.
- Cohen, S. (2016). Reasons to believe and reasons to act. *Episteme*, 13(4), 427–438. <https://doi.org/10.1017/epi.2016.22>.
- Crawford, S. (2004). A solution for russellians to a puzzle about belief. *Analysis*, 64(3), 223–229. <https://doi.org/10.1111/j.0003-2638.2004.00489.x>.
- Crawford, L. (2022). Suspending judgement is something you do. *Episteme*, 19(4), 561–577. <https://doi.org/10.1017/epi.2022.40>.
- Dancy, J. (2004). *Ethics without principles*. Oxford University Press.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52, 913–929. <https://doi.org/10.2307/2107917>.
- DiPaolo, J. (2018). Higher-order defeat is object-independent. *Pacific Philosophical Quarterly*, 99(2), 248–269. <https://doi.org/10.1111/papq.12155>.
- Ferrari, F., & Incurvati, L. (2022). The varieties of agnosticism. *Philosophical Quarterly*, 72(2), 365–380. <https://doi.org/10.1093/pq/pqab038>.

- Forchimes, A. T. (2021). Attitudinal strength as distance to withholding. *Philosophical Studies*, 178, 963–981. <https://doi.org/10.1007/s11098-020-01467-2>.
- Friedman, J. (2013). Suspended judgement. *Philosophical Studies*, 162, 165–181. <https://doi.org/10.1007/s11098-011-9753-y>.
- Friedman, J. (2017). Why suspend judging? *Nous*, 51(2), 302–326. <https://doi.org/10.1111/nous.12137>.
- Friedman, J. (2019). Inquiry and belief. *Nous*, 53(2), 296–315. <https://doi.org/10.1111/nous.12222>.
- Fritz, J. (2020). Moral encroachment and reasons of the wrong kind. *Philosophical Studies*, 177, 3051–3070. <https://doi.org/10.1007/s11098-019-01359-0>.
- Fritz, J. (2022). Encroachment on emotion. *Episteme*, 19(4), 515–533. <https://doi.org/10.1017/epi.2022.38>.
- Gertken, J., & Kieseewetter, B. (2017). The right and the wrong kind of reasons. *Philosophy Compass*, 12, <https://doi.org/10.1111/phc3.12412>.
- Graham, P. (2010). In defense of objectivism about moral obligation. *Ethics*, 121, 88–115. <https://www.jstor.org/stable/10.1086/656328>.
- Harman, G. (2004). Practical aspects of theoretical reasoning. In P. Rawling, & A. Mele (Eds.), *Oxford handbook of rationality* (pp. 45–56). Oxford University Press.
- Hieronymi, P. (2005). The wrong kind of reasons. *Journal of Philosophy*, 102(9), 437–457. <https://doi.org/10.5840/jphil2005102933>.
- Hieronymi, P. (2013). The use of reasons in thought (and the use of earmarks in arguments). *Ethics*, 124, 114–127. <https://doi.org/10.1086/671402>.
- Howard, C. (2019). The fundamentality of fit. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics. Volume 14* (pp. 216–236). Oxford University Press.
- Howard, C. (2023). Fitting attitude theories of value. In E. Zalta & U. Nodelman (Eds.), *Stanford encyclopedia of philosophy (spring 2023 edition)*. <https://plato.stanford.edu/archives/spr2023/entries/fitting-attitude-theories/>.
- Jackson, E. (2020). Belief, credence, and evidence. *Synthese*, 197, 5073–5092. <https://doi.org/10.1007/s11229-018-01965-1>.
- Jackson, E. (2023). A permissivist defence of Pascal’s wager. *Erkenntnis*, 88, 2315–2340. <https://doi.org/10.1007/s10670-021-00454-1>.
- Kavka, G. (1983). The toxin puzzle. *Analysis*, 43, 33–36. <https://doi.org/10.1093/analys/43.1.33>.
- Kieseewetter, B. (2017). *The normativity of rationality*. Oxford University Press.
- Kim, B. (2017). Pragmatic encroachment in epistemology. *Philosophy Compass*, 12, <https://doi.org/10.1111/phc3.12415>.
- Kolodny, N. (2005). Why be rational? *Mind*, 114(455), 509–563. <https://doi.org/10.1093/mind/fzi509>.
- Lasonen-Aarnio, M. (2014). Higher-order evidence and the limits of defeat. *Philosophy and Phenomenological Research*, 88(2), 314–345. <https://doi.org/10.1111/phpr.12090>.
- Leary, S. (2017). In defense of practical reasons for belief. *Australasian Journal of Philosophy*, 95(3), 529–542. <https://doi.org/10.1080/00048402.2016.1237532>.
- Lee, W. (2022). Fittingness first? Reasons to withhold belief. *Philosophical Studies*, 179(12), 3565–3581. <https://doi.org/10.1007/s11098-022-01848-9>.
- Lee, W. (2023). What is structural rationality? *Philosophical Quarterly*. <https://doi.org/10.1093/pq/pqad072>.
- Lilly, W. (2019). Constitutive reasons and the suspension of judgement. *Dialogue*, 58(2), 215–224. <https://doi.org/10.1017/s0012217319000027>.
- Lord, E. (2018). *The importance of being rational*. Oxford University Press.
- Lord, E. (2020). Suspension of judgement, rationality’s competition, and the reach of the epistemic. In S. Schmid, & G. Ernst (Eds.), *The ethics of belief and beyond. Understanding mental normativity* (pp. 126–145). Routledge.
- Lord, E., & Sylvan, K. (2021). Suspension, higher-order evidence, and defeat. In M. Simion, & J. Brown (Eds.), *Reasons, justification, and defeat* (pp. 125–145). Oxford University Press.
- Masny, M. (2020). Friedman on suspended judgement. *Synthese*, 197, 5009–5026. <https://doi.org/10.1007/s11229-018-01957-1>.
- McGrath, M. (2021a). Being neutral: Agnosticism, inquiry, and the suspension of judgement. *Nous*, 55, 463–484. <https://doi.org/10.1111/nous.12323>.
- McGrath, M. (2021b). Epistemic norms on waiting. *Philosophical Topics*, 49(2), 173–201. <https://doi.org/10.5840/philtopics202149221>.
- McHugh, C., & Way, J. (2016). Fittingness first. *Ethics*, 126, 575–606. <https://doi.org/10.1086/684712>.
- McHugh, C., & Way, J. (2022). *Getting things right: Fittingness, reasons, and value*. Oxford University Press.

- Mueller, A. (2021). *Beings of thought and action. Epistemic and practical rationality*. Cambridge University Press.
- Muñoz, D. (2021). Three paradoxes of supererogation. *Nous*, 55(3), 699–716. <https://doi.org/10.1111/nous.12326>.
- Parfit, D. (2011). *On what matters*. Oxford University Press.
- Raleigh, T. (2017). Another argument against uniqueness. *Philosophical Quarterly*, 67(267), 326–346. <https://doi.org/10.1093/pq/pqw058>.
- Raleigh, T. (2021). Suspension is believing. *Synthese*, 198, 2449–2474. <https://doi.org/10.1007/s11229-019-02223-8>.
- Rinard, S. (2019). Believing for practical reasons. *Nous*, 53(4), 763–784. <https://doi.org/10.1111/nous.12253>.
- Roeber, B. (2016). Reasons not to believe (and reasons to act). *Episteme*, 13(4), 439–438. <https://doi.org/10.1017/epi.2016.23>.
- Rosa, L. (2020). Suspending judgement the correct way. *Inquiry : A Journal of Medical Care Organization, Provision and Financing*. <https://doi.org/10.1080/0020174X.2020.1850344>.
- Schleifer McCormick, M. (2015). *Believing against the evidence. Agency and the ethics of belief*. Routledge.
- Schmidt, E. (2023). Facts about incoherence as non-evidential epistemic reasons. *Asian Journal of Philosophy*. <https://doi.org/10.1007/s44204-023-00075-1>.
- Schroeder, M. (2011). Holism, weight, and undercutting. *Nous*, 45(2), 328–344. <https://doi.org/10.1111/j.1468-0068.2010.00767.x>.
- Schroeder, M. (2012a). Stakes, withholding, and pragmatic encroachment on knowledge. *Philosophical Studies*, 160, 265–285. <https://doi.org/10.1007/s11098-011-9718-1>.
- Schroeder, M. (2012b). The ubiquity of state-given reasons. *Ethics*, 122(3), 457–488. <https://doi.org/10.1086/664753>.
- Schroeder, M. (2015). What makes reasons sufficient? *American Philosophical Quarterly*, 52(2), 159–170. <https://www.jstor.org/stable/24475447>.
- Schroeder, M. (2021). *Reasons first*. Oxford University Press.
- Setiya, K. (2022). Intention. In E. Zalta & U. Nodelman (Eds.), *Stanford encyclopedia of philosophy (winter 2022 edition)*. <https://plato.stanford.edu/archives/win2022/entries/intention/>.
- Shah, N. (2003). How truth governs belief. *Philosophical Review*, 112(4), 447–482. <https://doi.org/10.1215/00318108-112-4-447>.
- Shah, N. (2008). How action governs intention. *Philosophers' Imprint*, 8, 1–19. <http://hdl.handle.net/2027/spo.3521354.0008.005>.
- Shah, N., & Silverstein, M. (2013). Reasoning in stages. *Ethics*, 124, 221–235. <https://doi.org/10.1086/671387>.
- Shah, N., & Velleman, J. D. (2005). Doxastic deliberation. *Philosophical Review*, 114(4), 497–534. <https://doi.org/10.1215/00318108-114-4-497>.
- Sharadin, N. (2016a). Nothing but the evidential considerations? *Australasian Journal of Philosophy*, 94(2), 343–361. <https://doi.org/10.1080/00048402.2015.1068348>.
- Sharadin, N. (2016b). Reasons wrong and right. *Pacific Philosophical Quarterly*, 97(3), 371–399. <https://doi.org/10.1111/papq.12089>.
- Silva, P. (2016). How doxastic justification helps us solve the puzzle of misleading higher-order evidence. *Pacific Philosophical Quarterly*, 98, 308–328. <https://doi.org/10.1111/papq.12173>.
- Silva, P. (2023). Merely statistical evidence: When and why it justifies belief. *Philosophical Studies*, 180(9), 2639–2664. <https://doi.org/10.1007/s11098-023-01983-x>.
- Silva, P., & Tal, E. (2021). Knowledge-first evidentialism and the dilemmas of self-impact. In K. McCain, et al. (Eds.), *Epistemic dilemmas* (pp. 182–199). Routledge.
- Smith, M. (2016). *Between probability and certainty: What justifies belief*. Oxford University Press.
- Snedegar, J. (2017). *Contrastive reasons*. Oxford University Press.
- Staffel, J. (2021). Transitional attitudes and the unmooring view of higher-order evidence. *Nous*, 57, 238–260. <https://doi.org/10.1111/nous.12400>.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford University Press.
- Tucker, C. (Forthcoming). Withholding by default: A difference between epistemic and practical rationality. *Philosophical Studies*.
- Turri, J. (2012). A puzzle about withholding. *Philosophical Quarterly*, 62, 355–364. <https://doi.org/10.1111/j.1467-9213.2011.00043.x>.

- Vollmer, M. (2023). Suspension of judgement: Fittingness, reasons, and permissivism. *Episteme*. <https://doi.org/10.1017/epi.2023.34>.
- Vollmer, M. (Unpublished Manuscript). Against zetetic encroachment.
- Wagner, V. (2022). Agnosticism as settled indecision. *Philosophical Studies*, 179, 671–697. <https://doi.org/10.1007/s11098-021-01676-3>.
- Whiting, D. (2019). Whither higher-order evidence? In M. Skipper, & A. Steglich-Petersen (Eds.), *Higher-order evidence* (pp. 246–264). Oxford University Press.
- Whiting, D. (2022). *The range of reasons in ethics and epistemology*. Oxford University Press.
- Worsnip, A. (2021). *Fitting things together: Coherence and the demands of structural rationality*. Oxford University Press.
- Zinke, A. (2021). Rational suspension. *Theoria*, 87(5), 1050–1066. <https://doi.org/10.1111/theo.12320>.

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