



Minimal disturbance: in defence of pragmatic reasons of the right kind

Lisa Bastian¹ 

© The Author(s) 2019

Abstract This paper draws attention to an important methodological shortcoming in debates about what counts as a reason for belief. An extremely influential distinction in this literature is between reasons of the ‘right kind’ and the ‘wrong kind’. However, as I will demonstrate, arguments making use of this distinction often rely on a specific (and not explicitly stated) conception of epistemic rationality. Shifting focus to a reasonable alternative, namely a coherentist conception, can lead to surprising consequences—in particular, pragmatic reasons can, against orthodoxy, indeed be reasons of the right kind for belief.

Keywords Epistemic rationality · Right-kind reasons · Wrong-kind reasons · Coherence · Pragmatic reasons

1 Introduction

This paper draws attention to an important methodological shortcoming in debates about what counts as a reason for belief. An extremely influential distinction in this literature is between reasons of the ‘right kind’ and the ‘wrong kind’. However, as I will demonstrate, arguments making use of this distinction often rely on a specific (and not explicitly stated) conception of epistemic rationality. Shifting focus to a reasonable alternative, namely a coherentist conception, can lead to surprising consequences—in particular, pragmatic reasons can, against orthodoxy, indeed be reasons of the right kind for belief.

✉ Lisa Bastian
lb225@st-andrews.ac.uk

¹ Arché Research Centre, University of St Andrews, 17-19 College Street, St Andrews KY16 9AL, UK

Here is the plan of the paper: I first spell out the distinction between reasons of the right and wrong kind for belief and relate it to an evidentialist conception of epistemic rationality, which seems to naturally suggest itself (Sect. 2). I then show that it is not obvious that this is the only plausible conception—the distinction leaves room for others. My diagnosis of why this conception has become the default is that people generally and implicitly assume that the rationality and the correctness of a belief converge. The coherentist conception of epistemic rationality, which I explore and motivate, challenges precisely this convergence and also provides us with an example of a pragmatic reason of the right kind (Sect. 3). I finish by considering some objections (Sect. 4) and by noting the importance of specifying standards of correctness and rationality conceptions for debates about rationality in general (Sect. 5).

2 Reasons: the right and wrong kind

My aim in this paper is to highlight the importance of specifying the conception of epistemic rationality underlying philosophical discussions before drawing substantial conclusions. Whilst this might sound like nit-picky methodology, there is a real danger attached to not doing so. In particular, one runs the risk of begging the question against subscribers of competing conceptions, thereby weakening the strength of the overall argument. If whatever conclusions are drawn non-trivially depends on which conception epistemic rationality was adopted, these conclusions can hardly be seen as universal. Critics can easily reject them by merely pointing to competing conceptions, instead of engaging with the merits of the argument. Next to disadvantages in the philosophical debating arena, there is a more theoretical danger of drawing misleading, if not outright false conclusions. We might, for example, come up with a false characterisation of a given phenomenon.

I aim to illustrate these points by focussing on a particular debate that has been plagued by this methodological shortcoming, namely the wrong kind of reason debate. The terms *reasons of the right kind* (RKR) and *reasons of the wrong kind* (WKR)¹ have been used to draw a distinction amongst the *normative reasons* for belief in a proposition p , namely reasons that count in favour of the belief that p .² This distinction is motivated by the intuition that even though many considerations count in favour of belief in a proposition, only some provide the right reasons for

¹ These terms have their origin in discussions of fitting-attitudes accounts of value and, more recently, in buck-passing accounts of value (e.g. Brentano 1889; Scanlon 1998). Rabinowicz and Rønnow-Rasmussen (2004) name it the ‘WKR problem’, which takes a slightly different form here: There can be reasons that count in favour of adopting a pro-attitude towards e.g. a particular action, even though this action does not actually seem to be valuable [e.g. Crisp’s (2000) famous example of preferring a saucer of mud if an evil demon would otherwise inflict pain on you]. So there can be a mismatch between what is valuable and what we have reasons to adopt a pro-attitude towards. This creates a problem for fitting-attitudes analyses unless they can rule out such mismatches by identifying these reasons as reasons of the wrong kind. But providing a general rationale for telling apart right from wrong kinds of reasons in this context has proven difficult. This part of the problem is similar to the one that I am concerned with here.

² This follows Scanlon’s (1998) characterisation.

believing it. Take an example: I set myself the goal of running a marathon and started to go on regular long-distance runs. That I have completed this training counts in favour of the belief that I can run a marathon. Now compare the fact that I have trained appropriately to the fact that confidence in my physical ability increases the chances of me actually finishing a marathon. This also counts in favour of my belief that I can run a marathon: if I have this belief, I am more likely to achieve my goal, all else being equal. So both considerations, that I have trained appropriately and that confidence increases chances of success, count in favour of believing that I can run a marathon. But only the first one seems to do so *in the right way*.³

Several attempts have been made to characterise reasons of the right kind.⁴ A popular strategy that I focus on here is the appeal to earmarks. This strategy consists in two steps. First, a list of earmarks is developed that is supposed to capture features of WKR and RKR respectively. Second, token reasons are categorised as WKR or RKR depending on whether they display these earmarks. Schroeder (2012) provides an example of such an earmark strategy. He proposes three⁵ ‘earmarks’ of RKR:

- (1) Right-kind reasons motivate⁶ belief on their basis.
- (2) Right-kind reasons bear on the rationality of a belief.
- (3) Right-kind reasons bear on the correctness of a belief.⁷

These earmarks are often used to exclude pragmatic reasons from the realm of RKR reasons by showing that they do not display some or all of the earmarks and to motivate the view that only *evidential* reasons⁸ can be RKR for belief. This position is commonly known as *Evidentialism*. Before discussing this view in the next section, note that the second earmark explicitly mentions the rationality of a belief. It does so without further specifying the underlying conception of (epistemic) rationality. This is an instance of the practice that I am cautioning against. If we decide that a token reason (or indeed a whole sub-class of reasons) are WKR

³ At this point, one might wonder whether the fact that confidence increases my chances of success really counts in favour of believing that I can run a marathon *in the wrong way*. This might be the case if what we are concerned with is adopting true beliefs or beliefs that are good candidates for knowledge; but maybe not wrong if we are concerned with actually running a marathon. This ‘wrong for what?’-line of thought is compatible with the methodological shortcoming in the debate that this paper draws attention to.

⁴ For overviews, see Gertken and Kiesewetter (2017) and Sylvan (2016).

⁵ Schroeder proposes a fourth earmark: ‘pragmatic reasons for belief have a recognizable ‘flavor’ that makes them feel intuitively like reasons for other attitudes [...]’. He himself thinks that this is ‘arguably not a proper earmark in its own right’ and does not spell it out further (Schroeder 2012: 460).

⁶ This is connected to the debate about doxastic voluntarism, which I bracket. For arguments for and against, see Ginet (2001) and Williams (1973). What’s more, the earmark assumes a strong link between evidential reasons and our ability to form beliefs on their basis. This link has been questioned by e.g. McCormick (2014) and Reiser (2009).

⁷ Whilst this is often treated as a truism, there are accounts which question whether truth is the standard of correctness for belief (e.g. Gibbons 2013: 90).

⁸ Occasionally, the expression ‘epistemic reasons’ is used as well. I understand evidential reasons as a specification of epistemic reasons, where ‘epistemic reasons’ just is a placeholder for the category of reasons that are reasons of the right kind for belief.

because they do not bear on the rationality of a belief, without specifying which conception of rationality is at work, this categorisation is conditional on endorsing this particular conception. Failure to state this clearly can result in categorisations that appear to be universal when they are in fact contingent, and as a result, we may prematurely dismiss token reasons (or entire sub-classes of reasons) from the realm of RKR.

Before moving on to my discussion of Evidentialism and the potential of pragmatic reasons being RKR, let me briefly consider some related but different methodological concerns that have been raised in the literature.

Hieronymi (2013) takes issue with excessive reliance on earmarks. She argues that we should not treat first characterisations of phenomena as actually capturing the nature of these phenomena. These first characterisations should be treated as nothing more than initial guides that allow us to build an account, which we can then use to revise the first characterisation—rather than treating this characterisation as set in stone. In other words, we should not treat pre-theory as theory.⁹ Schroeder's use of earmarks does not always steer clear of this danger. Whilst he initially describes the earmarks as 'striking features, which do, pretheoretically, characterize the difference [...] (Schroeder 2012: 458) between the two kinds of reasons, he later seems to rely on them as individually necessary and jointly sufficient conditions for being a RKR: 'if these considerations bear all of the marks of right-kind reasons, they *are* right-kind reasons [...]' (ibid.: 466).

Whilst being sympathetic to Hieronymi's words of caution, the point I want to make is a more general one. Hieronymi's criticisms derive their force from an assumed mismatch between pre-theory and theory. We only have to worry about excessive reliance on pre-theoretical earmarks if there is a danger of pre-theory and theory coming apart. So in an ideal case, where our pre-theoretic conceptions perfectly track the correct theory, we could happily rely on earmarks to demarcate phenomena. My point applies even in the ideal case: whenever we use theoretically laden features to characterise or demarcate certain phenomena, it is crucial to specify which conception of these features is employed in order to avoid begging the question. Schroeder's earmark strategy faces this problem, but not uniquely.

There are a number of criticisms in the literature that share my more general concern about the use of specific conceptions in the investigation of a subject matter, in cases where there is disagreement about the correct conception. In her argument for unity of theoretical and practical norms for belief and against the distinction between WKR and RKR, McCormick (2014, 2018) argues that the underlying conception of 'reason' matters greatly for the outcome of this investigation.¹⁰ For example, if the underlying conception of 'reason' already

⁹ In a similar vein, Rabinowicz and Rønnow-Rasmussen (2012) wonder whether Schroeder intends the earmarks to function as a definition of the RKR/WKR-distinction.

¹⁰ A similar point can be found in Gibbons (2009, 2010). Whilst Gibbons is not directly concerned with the RKR/WKR distinction itself, he also acknowledges that the underlying conception of 'reason' impacts the conclusions drawn from an argument. In his case, he argues against a dualist view of reasons that regards normative and motivating reasons as separate by pointing out how such dualism is motivated by a previously accepted conception of 'reason' as perspective-dependent.

includes a commitment to only evidential reasons being RKR, it is not surprising that pragmatic (or in her terminology, practical) reasons qualify as WKR. McCormick argues that such an underlying conception of ‘reason’ is not the only option available, hence opening up the space of RKR. This is what we agree on. But McCormick is doubtful of the distinction as such and thinks that the already familiar distinction between *good* and *bad* reasons is sufficient. It can do all the work that the WKR/RKR distinction is supposed to do.¹¹ In particular, McCormick’s view can allow for pragmatic reasons to be reasons for belief, albeit bad ones. In contrast, the example of a pragmatic reason of the right kind that I develop in this paper is intended to not only be a reason for belief, but also a potentially *good* reason for belief. What’s more, I do not take my points to question the distinction between RKR and WKR per se. Instead, I only intend to cast doubt on one particular characterisation of it.

Further traces of my point can be found in Reisner (2013). Again, his target is a slightly different one. He is concerned with finding the set of genuine rational requirements, and in the course of this, argues that the enkratic principle is not one of them. But he shares with McCormick and myself the belief that previous commitments substantially impact the investigation of a given subject matter. Whilst McCormick focusses on the impact of previously accepted conceptions of ‘reason’, Reisner and I focus on the underlying conception of ‘rationality’ that theorists bring to the table. What matters most for Reisner’s project is whether our starting point is ‘intuitionism’ or ‘algorithmic systematicity’ about rationality¹² (Reisner 2013: 437). On the algorithmic systematicity-conception, rationality is ‘principally concerned with consistency amongst an agent’s mental states’ (ibid.: 438). He goes on to argue that the enkratic principle cannot be constructed as a genuine rational requirement on this conception of rationality. Reisner and I have different subject matters as targets. But we agree that the underlying conception of rationality is crucial for the resulting conclusion and we both aim to show the importance of specifying these conceptions. Reisner’s paper does not endorse the consistency conception of rationality. It merely puts forward a conditional conclusion: if one accepts the consistency conception, then the enkratic principle is not a genuine rational requirement. This parallels my exploration of a coherence conception of epistemic rationality: if one accepts such a coherence conception, then pragmatic reasons can be reasons of the right kind for belief.

2.1 Which reasons are RKR?

With the methodological points on the table, let’s turn to the WKR/RKR debate to trace the shortcomings I raised and to see how this results in a premature dismissal of pragmatic reasons from the realm of reasons of the right kind for belief.

¹¹ Against this, D’Arms and Jacobson (2014) argue that the WKR/RKR and good reasons/bad reasons distinctions cut across each other.

¹² To be exact, Reisner takes these to be stances towards *rational requirements* rather than *rationality simpliciter*. But since he takes rationality to ‘comprise[s] a set of rational requirements’ (Reisner 2013: 461), I take this shortcut in my presentation to be unproblematic.

A striking difference between the two reasons in our example is that one provides evidence for the truth of the belief whereas the other does not. This motivates Evidentialism, which I will ultimately reject. Evidentialism, in its general version, accepts that there can be reasons of the right and wrong kind for belief (e.g. Hieronymi 2005), but holds that only evidential reasons can be reasons of the right kind for belief. Reisner (2009: 258) provides a more modest characterisation of Evidentialism as the view that ‘there is something especially important or central about evidential reasons for belief’.

It should be noted that there exists a different brand of Evidentialism, sometimes called *Strict Evidentialism* (e.g. Shah 2006), which is committed to WKR-scepticism, i.e. the view that there are no reasons of the wrong kind for belief. On this view, a consideration either is or is not a reason for belief; it cannot count in favour of belief in the wrong way. In this paper, I limit my discussion to the weaker version, or Evidentialism proper, so to speak.

The earmark strategy can be used to provide support for Evidentialism. By showing that evidential reasons display the earmarks of reasons of the right kind, whereas pragmatic reasons do not, it is concluded that only evidential reasons can be RKR.

As an example, take the fact that I have completed my training. This is an evidential reason that also displays the earmarks: having completed a training plan motivates my belief that I can run a marathon. It also bears on the rationality of this belief because it is conducive to a conventional epistemically important property, namely truth. And since evidential reasons bear on the truth of a belief, it is almost a truism that they also bear on the correctness of belief.

With this in place, reasons of the wrong kind would include all non-evidential reasons. Pragmatic reasons belong to this category.¹³ Pragmatic reasons are similar to evidential reasons in that both count in favour of belief in some way. But instead of being conducive to believing truly, pragmatic reasons are conducive to various other interests, such as maximising utility, producing desirable consequences or prudence. Since the consequence-based reason in the example is non-evidential, we should expect it to lack the earmarks: it does not motivate belief as easily;¹⁴ it is not necessarily conducive to believing truly and hence does not bear on the belief’s rationality¹⁵ and neither on its correctness. And so, following Schroeder’s strategy

¹³ I am inclined to think that pragmatic reasons exhaust the category of non-evidential reasons but I will not argue for this here. It does not matter for my purposes whether there are other non-evidential reasons that are not also pragmatic reasons.

¹⁴ Strict Evidentialists like Shah (2006) will consider a lack of the first earmark decisive. For Shah, it is a conceptual truth of belief that it is transparent, i.e. that deliberating *whether to believe p* immediately gives way to the question *whether p*. He adds a deliberative constraint which holds that in order for something to be a reason for belief, it must be able to motivate belief. The combination of transparency and this constraint leads him to conclude that only evidential reasons can be reasons for belief. On this account, if pragmatic reasons cannot motivate belief, they could not be reasons for belief at all – not even of the wrong kind.

¹⁵ Kelly (2002) and Shah (2006) share the view that pragmatic reasons do not bear on belief’s rationality. Both argue that it is part of the nature of belief that only evidential reasons can make belief rational.

of identifying RKR with reasons that display the earmarks, our non-evidential reason is a reason of the wrong kind for belief because it lacks the earmarks.

Note at this point that this line of argument simply uses ‘rationality’ as it features in the second earmark to draw a conclusion about the realm of RKR without specifying which conception of rationality is used here. Argumentative moves like this display precisely the methodological shortcoming that this paper is concerned with.

3 Broadening right-kind reasons: pragmatic reasons as RKR?

Since I want to motivate the possibility of pragmatic reasons being reasons of the right kind, I am generally sympathetic to Schroeder’s argument,¹⁶ as he also construes the category of RKR more broadly to include some pragmatic reasons, albeit for withholding belief. Unlike Schroeder, I think that there can be room for pragmatic RKR for belief, not just for its absence. This is where we disagree: Schroeder thinks that there is ‘a distinctive dimension of rational assessment of beliefs—sometimes called *epistemic* rationality—that is affected by the epistemic reasons of which the subject is aware but not affected by the pragmatic reasons of which the subject is aware’ (Schroeder 2012: 459). But I do not think that it is necessarily the case that there is such a distinct domain of epistemic rationality on which pragmatic reasons have no bearing. Quite the opposite: I want to show that there *is* a plausible and venerable conception of the rationality of belief which allows for pragmatic reasons. If this is the case, then I have shown that pragmatic reasons can display one of the earmarks, namely the second one. What’s more, this highlights the importance of specifying the conception of epistemic rationality at play, because it leads to different substantial conclusions. So let’s turn to the second earmark and explore an alternative conception of epistemic rationality.

3.1 A coherentist conception of rationality

The conception of rationality that I will motivate is broadly coherentist. I’ll now provide a short characterisation of *Coherentism* in general. I’ll then introduce two popular ways of cashing out ‘rationality as coherence’ in order to lend additional support to the coherentist approach. In the next section, I’ll look closer at the origins of these views in order to draw some parallels that will help me formulate my example of a RKR pragmatic reason for belief.

Coherentism can apply to a number of theories. Prominent varieties include coherentism about justification, truth and, what’s most important for this paper, rationality. Their basic claim can be summarised as ‘A justified/true/rational belief is a coherent belief’. Spelling out what ‘coherence’ amounts to is a research field in

¹⁶ More precisely, Schroeder argues against the idea that the distinction between state- and object-given reasons tracks the distinction between RKR and WKR. His argument relies on showing that there are state-given reasons that have the earmarks of RKR. The state-/object-given distinction is not my concern here.

its own right. But Bonjour's (1985) five principles of coherence provide a nice overview, as they combine various different approaches:

- (1) A system of beliefs is coherent only if it is logically consistent,
- (2) A system of beliefs is coherent in proportion to its degree of probabilistic consistency,
- (3) The coherence of a system of beliefs is increased by the presence of inferential connections between its component beliefs and increased in proportion to the number and strength of such connections,
- (4) The coherence of a system of beliefs is diminished to the extent to which it is divided into subsystems of beliefs which are relatively unconnected to each other by inferential connections,
- (5) The coherence of a system of beliefs is decreased in proportion to the presence of unexplained anomalies in the believed content of the system (Bonjour 1985: 95, 98).

(1) and (4) relate to the view that some coherence relations are basically deductive relations. In contrast, (2) and (3) are connected to the idea that some coherence relations are inductive relations; (5) is closer to the thought that coherence is in part about explanatory relations.¹⁷

The conception of rationality as coherence amongst attitudes corresponds to one of the two main views within the debate about the nature of rationality (Hinchman (2013); the other one being rationality as responsiveness to reasons). Within this conception, Reisner (2013) draws a further distinction between understanding rationality as 'strict consistency' on the one side and rationality as 'unity' on the other.

On the strict consistency understanding, rationality is concerned with constraining the 'relations amongst the contents of an agent's mental states' (ibid.: 437), where the relevant constraints are provided by rigid formal systems like logic or probability calculus.¹⁸ Hence, this conception reflects Bonjour's first and second principle. In the realm of theoretical rationality, closure requirements are a good example for the kinds of demands strict consistency conceptions of rationality make:

You are rationally required that (if you believe x and you believe if x then y , then you believe y).¹⁹

All this requirement does is to demand that the contents of an agent's belief conform to modus ponens. The content of the requirement is provided by a rule of inference in classical logic. In the practical realm, an example of a strict consistency requirement is

¹⁷ For a helpful overview that cuts the debate up in this way, see Murphy (2019).

¹⁸ Most notably, Bayesianism can be seen as a coherence conception of rationality that relies on probability calculus. It considers beliefs (or rather credences) rational when they are coherent, where coherence is understood as compliance with the axioms of probability theory in combination with a conditionalization constraint along the lines of Bayes' Theorem.

¹⁹ Reisner (2013: 441).

You are rationally required that (if you intend to ϕ and believe that ϕ only if ψ , then you intend to ψ),

which Reisner (2013: 442) presents as the practical analogue to the previous theoretical requirement. But as Reisner notes, rationality is not only concerned with strict consistency demands, such as can be derived from rigid formal systems like logic. Often, we are concerned with ‘a certain kind of psychological consistency’ that cannot be captured by strict consistency alone (ibid.: 440). This requires a broader understanding of rationality as coherence, which he captures as ‘unity’: ‘Consistency between one’s normative beliefs, which are a form of normative self-ascription, and one’s related attitudes constitutes a rationally successful unification of an agent’s psychology’ (ibid.: 445). This is closer to the coherentist conceptions of rationality that are advocated by e.g. John Broome (2013) and Smith (2004). Reisner (2013: 444–445) proposes so-called ‘matching attitude requirements’ as examples for rational coherence requirements along the lines of the unity understanding. They follow a general template:

You are rationally required that: if you believe that you ought to have attitude *A* towards contents *C*, then you have attitude *A* towards contents *C*.

Depending on what we fill in for attitude *A*, these requirements are theoretical or practical. For the purposes of this paper, it does not matter whether we subscribe to a strict consistency or a broader unity understanding of rationality. Both can be captured as coherentist conceptions of rationality. What they have in common is that they require or prohibit certain patterns of attitudes: e.g. the belief that *p*, the belief that if *p* then *q*, and the absence of the belief that *q*. So if we want to assess the rationality of, say, the belief that *q*, looking at that belief alone is not enough. We also have to take into account its relation to other beliefs—in this case its deductive relations—to be able to assess whether the prohibited pattern has been instantiated. Similarly, evaluating the rationality of an individual attitude *A* is not possible. In order to decide on its rationality, we have to consider whether it would instantiate the prohibited pattern of believing that you ought to have attitude *A*, but not having it. In other words, we have to factor in how attitude *A* is related to other attitudes.

3.2 Quinean inspirations

These features of coherentist conceptions of rationality possess striking parallels to some claims about the nature of scientific theories as put forward by W.V.O. Quine, whose holism has been one of the most influential accounts of coherentist thought, albeit primarily within the Philosophy of Science. Given my focus on the rationality of belief, in what follows, I draw parallels between two features of scientific theory-building and belief-formation.

The first parallel concerns the impossibility of individual evaluation. Quine (1951) holds that since individual scientific sentences are underdetermined by evidence, it is impossible to evaluate a single sentence individually for its truth. If I want to evaluate the sentence ‘The pond freezes at temperatures below 0 °C’, I also have to evaluate previous data points and the reliability of my equipment and my

reasoning. In the more recent Quine and Ullian (1978: 16), this idea is explicitly applied to belief: ‘Often in assessing beliefs we do best to assess several in combination. [...] It is in light of our full body of beliefs that candidates gain acceptance or rejection; any independent merits of a candidate [belief] tend to be less decisive’. For example, if I want to evaluate the rationality of my belief that my roommate cooked dinner for me, I have to take into account many other related beliefs: beliefs about my roommate’s habits but also about the reliability of my sensory experience. This extends to deductive relations. With rationality requirements like deductive closure requirements, evaluating a single belief is equally impossible—we have to consider other beliefs as well as the (deductive) relations between them.

This brings us to the second parallel. Structurally, the impossibility of individual evaluation naturally suggests that an individual belief is embedded in a network of other beliefs, which are themselves connected. The parallel to theory-building would be to recognise that the sentences within a theory are members of a complex system whose individual elements bear manifold connections to the others, as opposed to a clear-cut tier system in which a few sentences provide the foundation for others. Thinking of scientific theories and belief systems as structured in this way recommends taking every element—a belief or a scientific sentence—as potentially revisable. If I were to encounter a conflict between my theory and the evidence, I could not identify an individual element of the theory as the exact source of the conflict. Hence, any of the sentences in the system that makes up a theory can be revised to make theory and evidence fit. As Quine and Ullian (1978: 16, 22) put it, ‘[w]hen a set of beliefs is inconsistent, at least one of the beliefs must be rejected as false; but a question may remain open as to which to reject. Evidence must then be assessed, with a view to rejecting the least firmly supported of the conflicting beliefs. [...] When an observation shows that a system of beliefs must be overhauled, it leaves us to choose which of those interlocking beliefs to revise.’

For example, if my thermometer indicates an outside temperature of -5°C but I also see that the water in the pond is not frozen at all, I could revise any of the sentences in my system to avoid this conflict. I might give up the belief in the correct functioning of my thermometer, I might start to doubt my eyesight or I might postulate a different freezing point of water. According to Quine, even the ‘profoundest laws of atomic physics or even of pure mathematics and logic’ could potentially be revised to preserve fit with experience. They are no exception to ‘[t]he totality of our so-called knowledge or beliefs [which is] a man-made fabric which impinges on experience only along the edges’ (Quine 1951: 116).

These parallels between belief-formation and scientific theory-building naturally suggest themselves on a coherentist approach to rationality that relies on requirements like deductive closure and means-end coherence. Given this, I think we can go a step further and use some of the criteria that guide the process of theory-formation to also guide the process of belief-formation and -revision.

3.3 Minimal disturbance: a pragmatic right-kind reason

If all sentences in a scientific theory have to be treated as part of systems rather than as single elements, and no single one is immune to revision, which part of the system should we change on encountering a conflict with evidence? Which sentences to drop, alter or add? Here, Quine endorses a ‘thorough pragmatism’ (Quine 1951: 126). The process of adapting our system of scientific sentences is guided by *simplicity* and *conservatism*. The system is to remain as simple as possible and the less invasive way of accommodating the evidence is preferred to alternatives which would require a more drastic revision. We have a ‘natural tendency to disturb the total system as little as possible [...]’ (ibid.: 120). In his later work, Quine explicitly states conservatism as a virtue of hypotheses: ‘In order to explain the happenings that we are inventing it to explain, the hypothesis may have to conflict with some of our previous beliefs; but the fewer the better. [...] The less rejection of prior beliefs required, the more plausible the hypothesis—other things being equal’ (Quine and Ullian 1978: 66–67). Given the parallels, I’ll suggest to adopt the same strategy for belief-change.

But there is additional reason to think that this is a sensible strategy in the belief case. Beliefs about the world and about what we should do or believe play a crucial role for agency. We rely on them all the time, for example when we make plans,²⁰ judgments or decisions. Without such beliefs, we would have a much harder time navigating the world and going about our lives. So there is a need for a fairly stable and continuous belief set. What’s more, our computational capacities are limited. To constantly question and revise our beliefs about the world would create a burden for us as agents that our limited computational capacities would not allow us to shoulder. Given this, and assuming that most of the time our belief-formation method is not entirely and systematically flawed, it makes sense to adopt a conservative approach to revising our belief system. Otherwise, we risk unnecessarily overthrowing large parts of our belief set. Assuming that it takes time to form new beliefs²¹ and to restore order in our belief set, this would create a state of abeyance that, in extreme cases, might even amount to a period in which we find ourselves being temporarily being stripped of our agency.

Since conservatism allows us to maintain agency and continuity in our belief-system, I will now explore how it could be put to work in the context of changing a belief system, which includes adding, dropping, or altering beliefs.²² Conservatism demands that the system in its entirety be changed as little as possible in order to accommodate incompatible evidence. The consequences of changing one belief on others is something that bears on the rationality of belief. To see how this criterion

²⁰ See Bratman (1987) on the importance of plans: they ‘facilitate coordination both socially and within our lives, [and] help enable prior deliberation to shape later conduct’. Plans are needed because ‘we are not frictionless deliberators’ (ibid.: 28) and plausibly require a stable belief set, too.

²¹ See for example Schroeder (2009).

²² I do not intend to develop a full theory of belief-formation here. For one, the notions of ‘importance’ and ‘centrality’ of beliefs are a lot more complex than I present them. I take the points I make to provide just the necessary background for my example of a pragmatic RKR.

could be transformed into a decision rule, consider the following toy-example of a belief system (Fig. 1):

The boxes b_1 – b_5 denote individual beliefs. The whole net of belief, including their connections, is the belief system S . Note that the beliefs differ with regard to their connections. Some of them are more interconnected than others. An outgoing arrow describes a support-connection,²³ for example b_1 supports b_2 and b_3 and is supported by b_2 and b_4 . b_4 has more outgoing connections than b_1 and hence plays a more important role in S . Similarly, b_1 can be considered more important than b_5 . Now assume that an agent with belief system S encounters a new piece of information that is incompatible with S . Now, she is required to make changes in S in order to accommodate this new piece of information.

On our coherentist picture, every single one of her beliefs is a potential candidate for revision. So which belief should the agent alter, drop or add? This is where conservatism steps in as a guiding principle for her decision. It demands that S be changed as little as possible. This suggests something like the following decision rule:

(MD) Minimise disturbance of S when resolving the conflict.

This rule favours minimal disturbance by recommending the change that leads to the fewest changes in S and solves the conflict. It provides the example of a pragmatic RKR: that a particular change would minimise disturbance of S is a reason for this change. But in order to apply this decision rule, we need a way of measuring and comparing the overall changes caused by changing an element of S . An intuitive indicator for this is the number of outgoing support-connections of beliefs. Call this C . Then,

$$\begin{aligned}Cb_1 &= Cb_2 = Cb_3 = 2; \\Cb_4 &= 3 \\Cb_5 &= 0.\end{aligned}$$

This allows us to specify (MD) a bit more:

(MD*) Change S such that (i) C is minimised for all b involved and (ii) the conflict is resolved.

Now, (MD*) can provide specific reasons for the change which minimises disturbance and resolves the conflict.

(MD*) places the number of outgoing support-connections at the centre of the deliberation about how to change the belief system in order to accommodate the

²³ Note that support-connections are a weaker notion than entailment. They also should not be understood as evidential relations (see Sect. 4. For why even an evidential understanding of support connections does not threaten my argument). Support-connections can take a number of forms, similar to the counting-in-favour-relations featured in the initial example about running a marathon: the belief that I have stuck to a training plan supports the belief that I can run a marathon; so does believing that confidence in my physical ability increases my chances of success.




Fig. 1 Belief system

conflict. Whether an agent has a reason to add or drop a belief (or multiple ones) will depend on which of these changes minimises the overall C -score.²⁴

For illustrative purposes, I assume for now that dropping one belief would already resolve the conflict. So we do not have to consider the options of changing multiple beliefs or belief-addition²⁵ for the moment. The question then becomes: which belief to drop? With regard to S and given these assumptions, (MD*) provides a reason for dropping b_5 because Cb_5 is the lowest. Dropping b_5 does not require any additional change in S , as it has no outgoing support-connection. Compare this to, say, altering b_4 . This could potentially require three other belief alterations which in a more complex S could set off a chain reaction and require even more changes of beliefs. Not so for dropping b_5 : this would not remove any support for other beliefs. But dropping b_4 instead would make three other beliefs less supported.

It should be noted that, unlike in my toy-example, (MD*) might recommend changing a number of beliefs over changing a single one. These will be cases where changing a single belief either does not suffice to resolve the conflict or does not result in overall C -minimisation. According to (MD*), the agent should opt for the change that resolves the conflict *and* minimises the C -score of the beliefs involved, even if that requires multiple belief changes.

²⁴ For present purposes, it will suffice to think of the overall C -score as the simple addition of the individual beliefs' C -scores. This does not exclude the possibility of a more sophisticated mode of aggregation in a more detailed version of the toy-example.

²⁵ I discuss an example featuring belief addition at the end of this section.

One might want to object at this point that (MD*) does not actually minimise disturbance. Would it not disturb the system less if we first identified the belief whose change would resolve the conflict and then directly moved towards changing that one? There are two things I want to say about this putative alternative to (MD*). One, on the coherentist picture, we cannot establish a one-to-one correspondence between a single belief and a piece of incompatible information, because beliefs are embedded in a system and impossible to evaluate individually. There might not be a single belief in the system that is responsible, so to say, for the conflict and even if there was, we would have no principled way of finding out which one it is. And so the alternative to (MD*) outlined above would fail at its first stage. Two, it is too quick to think that changing multiple beliefs causes more disturbance than changing a single one. When measuring disturbance, we have to take into account the belief's support connections. It might be less disturbing to change a couple of beliefs with low C , as that would require only few additional changes down the line, than changing only one belief with a high C , as this would potentially require many changes down the line. For these reasons, I do not consider the above suggestion a plausible alternative to (MD*).

My picture has the further advantage that C is also an indicator for the importance of a belief in the system, in addition to measuring the consequences of a particular belief alteration: the higher the number of outgoing support-connections, the more important a role it plays. This has the desirable consequence of accommodating an important objection against coherentist approaches in general. Because these approaches hold that all elements in a system are potentially subject to revision, laws of nature or logical truths do not enjoy any kind of special standing. If required, they can be changed just like any other element. Introducing a decision rule like (MD*) does not eliminate this possibility, but makes it the case that beliefs which take logical truths or laws of nature as their content are in fact much less likely to be revised or given up. That is because their C will be one of the highest in the whole system. A belief b_g about the law of gravitation for example will support other beliefs about the interplay between physical forces, all the way to the belief that the apple will fall on the floor if I drop it. Hence, it is very unlikely that b_g will ever be revised because there will be many other beliefs with a lower C that will be changed first. And if the encountered conflict with experience is so profound that revising all the other beliefs with lower C did not resolve it and b_g would be the next on the list, then the conflict must be of such fundamental nature that it deserves this kind of radical revision. But again, conflicts of this kind are rather unlikely to occur.

To see how (MD*) could provide a reason for adding instead of dropping a belief, let's consider a historical example: Dr. Ignaz Semmelweis' famous observations about the connections between childbed fever and handwashing (Ginnivan 2014). In the 1840s, maternal mortality rates after giving birth in one of Vienna's two public hospitals were a lot higher than maternal mortality rates after having given birth outside of these hospitals, such as at home, in the streets or on the fields. There also was a significant difference between the two hospitals, with approximately 1 in 25 women dying in the public hospital that trained midwives, as opposed to approximately 1 in 10 women dying in the public hospital that trained

medical students. Faced with this evidence, chief resident Semmelweis famously postulated a connection between having contact with cadavers and risk of infections. This would explain why mortality rates were a lot higher in the hospital that trained medical students, who had regular contact with cadavers during autopsies. As a result, he decided to make chlorine handwashing mandatory before assisting with a birth, resulting in drastically decreasing maternal mortality rates.

This example illustrates how (MD*) could provide a reason for adding a belief. Semmelweis had to reconcile a number of his beliefs with incompatible information, for example the belief that midwives are better trained to deal with birth-related complications than medical students, that medical students are better trained to deal with birth-related complications than persons without medical training and that receiving medical attention is safer than not receiving it. The incompatible information is constituted by the drastically higher mortality rates connected to births within rather than outside hospitals.

Dropping or changing any of these beliefs would likely have demanded many more changes to Semmelweis' belief system. Presumably, he had good reason to believe that receiving medical attention is better than not receiving it, such as sick people being more likely to survive in hospitals than on the streets. Changing his belief about the competences of midwives and medical students, respectively, also would have required many more revisions, for example beliefs about how the content of medical training impacts skill.

But adding the belief that contact with cadavers increases the risk of infection does not require any of those revisions. It accommodates the incompatible information whilst maintaining the other beliefs. This is an example of how (MD*) could favour adding a belief over altering or dropping one, thereby giving us a *pro tanto* reason *for* belief.²⁶

3.4 (MD*) and AGM

At this point, the reader might feel reminded of the AGM framework for belief revision. And indeed, there exist a number of parallels between the view I sketched and the framework. But they also differ in important ways. I will explore some of these similarities and differences in this section.²⁷ The most important point to take away from this discussion however is that the similarities to an established and well-worked out framework like AGM only lend more support to my methodological claim. Even if one was not convinced by the example of a pragmatic RKR that I

²⁶ My remarks here should not be understood as saying that adding a new belief *always* is the best way of accommodating the conflict. For example, a newly added belief might be incompatible with existing beliefs and cause more disturbance in that way. In many cases, changing or dropping a belief would be the option favoured by (MD*), as illustrated by the toy example I gave earlier. Whether (MD*) recommends the dropping or adding of a particular belief (or multiple ones) will depend on its final specification, which in turn will depend on how we specify the toy example to obtain a fully fledged model of a coherentist belief system. As I mentioned earlier, this is a task that I cannot take up in this paper.

²⁷ My presentation of the AGM framework relies mostly on Kyburg and Teng (2001), which largely follows Gärdenfors (1988). The main tenets of the AGM framework have their origins in Alchourrón and Makinson (1982), Alchourrón et al. (1985), and Gärdenfors and Makinson (1988).

sketched via (MD*), there exist alternative examples, similar in spirit, which can also give rise to pragmatic RKR by exploiting the same methodological lacuna in adopting a broadly coherentist conception of rationality.

The AGM framework represents belief systems as logical theories, in which each individual belief is represented as a sentence. This feature entails one of the two main similarities between AGM and my picture: both share a preoccupation with coherence. More precisely, AGM requires that the belief system remain logically consistent. In this regard, it can be seen as a specification of what Reiser called ‘strict consistency’ conceptions of rationality, which is not equivalent to, but compatible with broader coherentist conceptions, such as the (MD*) picture.

Like in my toy-example, belief systems, or theories, can be changed in three different ways. We can either add a belief (or sentence), we can drop a belief or we can change a theory. But these three options take on a slightly different form in the AGM framework. To see this, consider first another feature that sets AGM apart from my picture. Since AGM represents belief systems as logical theories, these systems are deductively closed, i.e. they also contain all the logical consequences of the individual beliefs. The belief system on my picture only contains the beliefs the agent actually holds and not, for example, beliefs in all disjunctive statements that follow from the other beliefs. Now, if we gain new information and need to add a belief to our belief system, AGM has us add not only this belief but also all its consequences to the system, in order to maintain deductive closure. On my picture, we would only need to add the new belief itself. Whenever a belief is dropped, AGM requires that we remove not only this particular belief but also other beliefs from which the dropped belief could be derived. And in the case of changing the system, AGM assumes that a specific belief in the system has been contradicted. In order to accommodate this and to maintain consistency, we have to remove the contradicted belief and add a belief containing its negation. Further changes might be required to ensure that there are no inconsistencies between the consequences of the newly added belief in the negation and its consequences or the consequences of the existing beliefs.

This leads us to the second main similarity between my picture and AGM. They share a concern for conservatism. On my picture, this took on the form of the decision rule (MD*). On AGM, this concern becomes relevant when we drop beliefs or change theories. In both cases, there might be multiple options available that would yield the desired change and preserve consistency and deductive closure. AGM demands that we choose the option which yields minimal change, where this is understood as maximising closeness to the original belief system. It is recognised that this choice cannot be made with the logical structure of the system alone but needs to appeal to ‘extralogical structure’ instead (Kyburg and Teng 2001: 167). This is further spelled out as ‘epistemic entrenchment’, which tracks the reliability of the existing beliefs. The higher the epistemic entrenchment of a belief, the more reliable it is. When choosing amongst the options, less epistemically entrenched beliefs have to be preferred.

(MD*) also demands conservatism, but it does so in terms of the outgoing support connections of beliefs. It is not concerned with their reliability or epistemic entrenchedness. However, it does care about how a particular belief is connected to

others. The more connected, the less likely it is to be recommended for change by (MD*). So interconnectedness and epistemic entrenchedness perform similar functions. In contrast, the number of outgoing support connections is a meaningless measure in AGM because every belief has infinite outgoing support connections, given that it entails infinite further (disjunctive) beliefs.²⁸ So whilst both pictures share a concern for conservatism, they do so in very different ways.

In general, AGM and my picture are similar in spirit but AGM is obviously the more restrictive framework. This is mainly because it represents belief systems as deductively closed logical systems and understands the connections between beliefs as governed by logical laws. The (MD*) picture, despite already being highly idealised, is closer to an actual representation of a belief system, which need not be deductively closed (we rarely believe all the consequences of our beliefs) and in which beliefs can be connected in all kinds of ways (our reasoning usually does not strictly follow the laws of logic). I take it that AGM in turn is best understood as a model of belief revision for fully rational agents. Given the similarities, it might be possible to achieve similar precision for my view by capturing (MD*) in the AGM framework. Alternatively, I would suggest to embrace the imprecision of my picture. After all, our subject matter here are reasons, which are notoriously messy. The support that reasons provide for belief comes in many different ways and is not exhausted by the the logical connections that exist between the believed propositions. This is part of their appeal, which might be lost if we could apply the notion of ‘reason’ only to very neat but restrictive frameworks like AGM.

4 Objections and replies

One might object that Minimal Disturbance does not provide pragmatic reasons. In the case of altering or dropping a belief, the objection would have it that the fact that changing b_5 causes only minimal disturbance to the belief system really is just an evidential reason. This might be a tempting line of thought since Minimal Disturbance relies on the support-connection between beliefs. If these are interpreted as evidential relations, it seems plausible to consider the reasons provided by Minimal Disturbance to be evidential since they pertain to evidential relations. But on closer inspection, this does not entail that Minimal Disturbance provides evidential reasons. On an intuitive understanding of ‘evidential reasons’ for belief, these reasons bear on the truth of belief. But the reasons provided by Minimal Disturbance do not. To see this, note that b_5 is well-supported in the system. If we assume that the belief-system is well-adjusted to the truth, then it would be surprising, to say the least, if evidential reasons recommended changing or altering a belief that is evidentially well-supported. Of course, belief systems are not always well-adjusted to the truth. But even on a somewhat distorted belief system, Minimal Disturbance’s reasons cannot be interpreted as evidential reasons. That is

²⁸ This assumes that outgoing support-connections could be represented as entailments on AGM, which is a possible but very restrictive reading of ‘support-connection’, as I mentioned earlier.

because there is no direct correspondence between the reason it provides (namely that changing b_5 would minimally disturb the system) and the truth of b_5 . Instead, this reason is entirely pragmatic—it is grounded in the consequences of changing b_5 .

The objection would run similarly for adding the belief that contact with cadavers increases the risk of infection. Again, it would be argued that Minimal Disturbance actually provides an evidential reason. This is a very tempting conclusion to draw in this case, because the progress of science has provided additional support for Semmelweis' postulation. We now know for a fact that cadavers carry bacteria that can cause infection. But this is just a coincidence in which a pragmatic reason produced a true belief and cannot be taken as proof that the reason was evidential in the first place. Minimal Disturbance did not count in favour of adding the belief about risk of infection because of its truth, likely or not, but only because this addition would lead to minimal revisions in the belief system. It would have equally recommended a false belief, if this had the consequences of minimally disturbing the belief system.

Nevertheless, it is likely that following Minimal Disturbance will generally produce true beliefs. If we accept my earlier stipulation that our belief formation method is not systematically flawed, then it is reasonable to assume that most of the beliefs in our belief system will be true, or at least evidentially supported. Following (MD*) has the effect of preserving the coherence of our belief set, which presumably contains many true beliefs. And so the changes that it recommends will be conducive to the belief system's coherence, and, as a likely side effect, lead to more true beliefs as well. But importantly, this is only a side-effect. (MD*)'s reasons are ultimately pragmatic and coherence-based.

As it stands, Minimal Disturbance does not provide evidential reasons. It does not pertain to the truth of the belief changes it counts in favour of. Instead, it only takes into account the consequences of the potential change and recommends the minimally invasive change. I do not want to deny that there could be conceptions which construct the very notion of 'evidence' in a coherentist way. I take it that this would involve understanding evidence as bearing on the coherence of a belief instead of its truth. On coherence theories of truth,²⁹ which understand truth as coherence with other propositions, these two might even coincide. But this is a somewhat radical position that has to be motivated in its own right if it is supposed to function as an objection to the pragmatic status of Minimal Disturbance's reasons. Still, my claims are compatible with the commitments of a proponent of a coherence theory of truth or evidence. In a sense, they inhabit the area at the far end of the theoretical space that I have carved out. Here, reasons that we traditionally consider pragmatic are subsumed under evidential reasons. This is in harmony with my intention to broaden the category of RKR for belief to reasons that are not strictly truth-related. Whether we call them pragmatic or evidential (on a coherence theory of evidence) or epistemic does not matter. What's more, this possibility only lends additional motivation to the second claim I argued for—namely the importance of specifying the conceptions of epistemic rationality at play. Without

²⁹ For an overview, see Young (2018).

such specifications, it is possible that we actually agree on appropriate candidates for RKR (for example Minimal Disturbance and truth-related reasons), whilst being stuck in a superficial disagreement caused by the labels 'pragmatic' and 'evidential' or 'epistemic'. Theorising about and using distinctions like RKR/WKR then runs the risk of being confusing and pointless.

Now, with the pragmatic status of Minimal Disturbance's reasons established, one might object that precisely this result casts doubt on the plausibility of the coherentist conception. Any conception of rationality which holds that pragmatic reasons could be RKR should immediately be dismissed, or so the objection goes.

In reply to this, I can only refer back to the methodological shortcoming of this debate that I have alluded to at various points in the paper: the underspecification of the conception of rationality at issue. To illustrate, consider the structure of the debate so far. We started out with a list of characteristics of reasons of the right kind which we dubbed 'earmarks'. We then used this list to organise the domain of reasons into right-kind and wrong-kind reasons for belief. The objection departs from this structure. It wants to use the resulting categorisation of reasons as an argument for or against a conception of rationality.

But this argument can only work if we can treat the set of RKR as a criterion for a plausible conception of rationality. A conception of rationality is only plausible if it results in this specific set of reasons as RKR, so the thought goes. There are two problems with this. One, prior to investigation, we simply do not have the set of RKR. The whole point of the WKR/RKR debate and the use of earmarks is to determine which reasons belong to the set of RKR. So we cannot use something that we do not yet have to judge conceptions of rationality.

Two, even if we have agreed on a set of RKR, we cannot use this set as an arbiter between plausible and implausible conceptions of rationality. That is because the concept that we are trying to evaluate (rationality) is contained in the criterion we use to evaluate. We only arrive at a set of RKR through consideration of the earmarks, one of which contains rationality. This is unhelpful and potentially circular. Compare this to a situation in which the explanandum is already contained in the explanans. It would not be very enlightening to answer the question 'Why is it raining?' with 'Because there is rain'. It would be similarly pointless to try to shed light on what plausible conceptions of rationality there are by stipulating that they have to characterise the reasons that bear on rationality as RKR.

This points to yet another difficulty with the objection's strategy: the second earmark is neutral with regards to the conception of rationality at play. The objection exploits this neutrality by assuming a more evidentialist conception of epistemic rationality for the second earmark. This then leads to a different set of RKR, which obviously does not contain pragmatic reasons. But this should not be surprising. Of course, the coherentist conception looks unattractive if we evaluate it with a set of RKR that has been provided by an evidentialist conception. Given that in the dialectic of the objection, we are in the business of finding plausible conceptions of rationality, we cannot assume that the best interpretation of the second earmark is an evidentialist conception. This would be deciding on an answer before we settled the question. Given these three problems, I contend that the objection does not pose a problem for my argument.

5 Concluding remarks

Let me reconnect the preceding points to the aim of this paper. I wanted to highlight a methodological shortcoming in the RKR/WKR debate by showing how exploiting this shortcoming opens up the possibility of pragmatic reasons bearing on belief's rationality. In order to do so, I explored and motivated a coherentist conception of rationality, which paved the way for a decision rule (MD*) that provides pragmatic, non-evidential reason to alter, add or drop a certain belief because this would minimise overall change in one's belief system and hence is associated with less theoretic cost. So (MD*) gives us a pragmatic reason that has an earmark of RKR.

This is not to say that all pragmatic reasons bear on belief's rationality, understood as being appropriately connected to and supported by the other elements in the belief system and being the result of rational system-adaptation. I also do not want to say that pragmatic reasons cannot be outweighed by other (evidential) reasons. But even if that were the case, they would still bear on belief's rationality. So unlike Schroeder and the Evidentialists, I think there is at least one plausible conception of epistemic rationality on which pragmatic reasons bear on the rationality of belief.

Finally, let me finish with some considerations about why the predominant view takes pragmatic reasons to be reasons of the wrong kind. A possible explanation of why authors have been so quick to dismiss the possibility of pragmatic RKR might be found in the third earmark—correctness. Here, we deal with a different standard of evaluation. If what is at issue is the correctness of belief, surely evidential reasons are the right kind of reasons. Pragmatic reasons do not bear on a belief's correctness. But evaluating a belief for correctness is not the same as evaluating a belief for rationality. If the standard of evaluation is rationality, truth just is not all that matters—at least not on some plausible conceptions of rationality—and this opens up the domain of RKR for non-evidential, pragmatic ones. By classifying non-evidential reasons as WKR, the standards of evaluation might have been conflated: we are evaluating rationality with the tools of correctness. Once we realise that the second and third earmarks correspond to different standards of evaluation, it is not very surprising that pragmatic reasons can satisfy one standard but not the other.

This casts doubt on characterising right and wrong kinds of reason with respect to different standards of evaluation. It might be more plausible to make the distinction specific to the relevant standard. But exploring this suggestion is beyond the scope of this paper. I conclude that on a coherentist conception of rationality, pragmatic reasons as provided by (MD*) can bear on belief's rationality and be RKR.

Acknowledgements I'd like to thank Miguel Egler, Lewis Ross and Justin Snedegar for their help and comments on the many versions of this paper. Special thanks also to audiences at the University of St Andrews and UNAM and an anonymous reviewer for their suggestions and insights.

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/>.

Funding Funding was provided by Arts and Humanities Research Council (Grant No. 1796583).

References

- Alchourrón, C. E., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2), 510–530.
- Alchourrón, C. E., & Makinson, D. (1982). On the logic of theory change: Contraction functions and their associated revision functions. *Theoria*, 48(1), 14–37.
- BonJour, L. (1985). *The structure of empirical knowledge*. Cambridge: Harvard University Press.
- Bratman, M. (1987). *Intention, plans, and practical reason*. Cambridge: Harvard University Press.
- Brentano, F. (1889). *The origin of our knowledge of right and wrong (1969)*. London: Routledge.
- Broome, J. (2013). *Rationality through reasoning*. Oxford: Oxford University Press.
- Crisp, R. (2000). Value ... and what follows By Joel Kupperman New York: Oxford University Press. *Philosophy*, 75(3), 452–462.
- D'Arms, J., & Jacobson, D. (2014). Wrong kinds of reason and the opacity of normative force. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (pp. 215–244). Oxford: Oxford University Press.
- Gärdenfors, P. (1988). *Knowledge in flux: Modelling the dynamics of epistemic states*. Cambridge: MIT Press.
- Gärdenfors, P., & Makinson, D. (1988). Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the second conference on theoretical aspects of reasoning about knowledge* (pp. 83–96).
- Gerken, J., & Kiesewetter, B. (2017). The right and the wrong kind of reasons. *Philosophy Compass*, 12(5), 1–14.
- Gibbons, J. (2009). You gotta do what you gotta do. *Noûs*, 43(1), 157–177.
- Gibbons, J. (2010). Things that make things reasonable. *Philosophy and Phenomenological Research*, 81(2), 335–361.
- Gibbons, J. (2013). *The norm of belief*. Oxford: Oxford University Press.
- Ginet, C. (2001). Deciding to believe. In M. Steup (Ed.), *Knowledge, truth and duty* (pp. 65–76). Oxford: Oxford University Press.
- Ginnivan, L. (2014). The dirty history of doctors' hands. *Method Quarterly* 2014/1. Retrieved August 29, 2019 from, <http://www.methodquarterly.com/2014/11/handwashing/>.
- Hieronymi, P. (2005). The wrong kind of reason. *The Journal of Philosophy*, 102(9), 437–457.
- Hieronymi, P. (2013). The use of reasons in thought (and the use of earmarks in arguments). *Ethics*, 124(1), 114–127.
- Hinchman, E. (2013). Rational requirements and 'rational' akrasia. *Philosophical Studies*, 166(3), 529–552.
- Kelly, T. (2002). The rationality of belief and other propositional attitudes. *Philosophical Studies*, 110(2), 163–196.
- Kyburg, H. E., & Teng, C. M. (2001). *Uncertain inference*. Cambridge: Cambridge University Press.
- McCormick, M. S. (2014). *Believing against the evidence: agency and the ethics of belief*. London: Routledge.
- McCormick, M. S. (2018). No kind of reason is the wrong kind of reason. In K. McCain (Ed.), *Believing in accordance with the evidence* (pp. 261–276). Berlin: Springer.
- Murphy, P. (2019). Coherentism in epistemology. *The Internet Encyclopaedia of Philosophy*. Retrieved August 30, 2019 from, <https://www.iep.utm.edu/coherent/>.
- Quine, W. V. O. (1951). Two dogmas of empiricism. In R. Bluhm & C. Nitz (Eds.), *From a logical point of view* (pp. 55–127). Stuttgart: Reclams Universal-Bibliothek.
- Quine, W. V. O., & Ullian, J. S. (1978). *The web of belief*. New York: McGraw-Hill, Inc.
- Rabinowicz, W., & Rønnow-Rasmussen, T. (2004). The strike of the demon: On fitting pro-attitudes and value. *Ethics*, 114(2), 391–423.

- Rabinowicz, W., & Rønnow-Rasmussen, T. (2012). *Ethics discussions at PEA Soup: Rabinowicz and Rønnow-Rasmussen on Schroeder*. Retrieved August 29, 2019 from, <http://peasoup.typepad.com/peasoup/2012/06/i.html>.
- Reisner, A. (2009). The possibility of pragmatic reasons for belief and the wrong kind of reasons problem. *Philosophical Studies*, 145(2), 257–272.
- Reisner, A. (2013). Is the enkratic principle a requirement of rationality? *Organon F*, 20(4), 436–462.
- Scanlon, T. (1998). *What we owe to each other*. Cambridge: Belknap Press of Harvard University Press.
- Schroeder, M. (2009). Means-end coherence, stringency, and subjective reasons. *Philosophical Studies*, 143(2), 223–248.
- Schroeder, M. (2012). The ubiquity of state-given reasons. *Ethics*, 122(3), 457–488.
- Shah, N. (2006). A new argument for evidentialism. *Philosophical Quarterly*, 56(225), 481–498.
- Smith, M. (2004). Instrumental desires. *Instrumental Rationality, Aristotelian Society Supplementary*, 78(1), 93–109.
- Sylvan, K. (2016). Epistemic reasons I: Normativity. *Philosophy Compass*, 11(7), 364–376.
- Williams, B. (1973). Deciding to believe. In B. Williams (Ed.), *Problems of the self* (pp. 136–151). Cambridge: Cambridge University Press.
- Young, J. O. (2018). The coherence theory of truth. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2018 Edition). Retrieved August 29, 2019 from, <https://plato.stanford.edu/entries/truth-coherence/#Ver>.

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