

How many normative notions of rationality? A Critical Study of Wedgwood's *The Value of Rationality*

Giacomo Melis

1. *The importance of rationality*

The main goal of Wedgwood's book¹, expected to be the first instalment of a trilogy, is to defend the claim that the concept of rationality is normative. Among other things, on Wedgwood's understanding, this is supposed to entail that 'we always ought to be as we are rationally required to be' (33).² Since Wedgwood argues that a mentalist variety of internalism is true of rationality, in his picture the demands of rationality may be characterized as the demand that the agent is broadly coherent—that one's way of thinking fit with the mental states and events present in one's mind at the relevant times (4).

Yet, internal coherence can't prevent one from being led astray on occasion, as the possibility of acquiring misleading evidence and the predicament of victims of sceptical scenarios illustrate. If so, as Wedgwood acknowledges, 'we may well doubt whether it must *always* be true that there is something good simply in being coherent' (38, emphasis original). Addressing this concern—that is, explaining what is good about coherence, and why we should always avoid being incoherent—takes up the main part of the book (chapters 4-9).

Wedgwood's answer begins with five careful chapters of stage-setting and preliminary considerations. These chapters develop many philosophical insights and arguments that any philosopher would find worth engaging with independently of the goal that they serve in the book, and include the following broad topics. Chapter 4 offers a discussion of why it is a mistake to think that the most fundamental normative notion is that of reasons; chapter 5 discusses the distinction between different senses of 'ought' and isolates the sense that is taken to be linked to rationality; chapter 6 proposes that the most fundamental normative concepts are evaluative ones and that the concept of rationality is one of them; chapter 7 defends mentalist internalism about rationality, and chapter 8 critiques attempts to explain why rationality matters within frameworks that accept Wedgwood's commitments to internalism and decisiveness of the requirements of rationality.

The specific answer to the question of why rationality matters -- why we should always avoid being incoherent -- arrives in chapter 9, and it is rich in technical details. Simplifying, Wedgwood suggests 'that every type of mental state or mental event that can be rational or irrational has an *aim*, and thinking rationally is a *means* to achieving that aim' (5). Wedgwood takes the relevant aim to be an *external* one: to get things right. The aim is external because the norms that regulate it do not evaluate mental states on the basis of what goes on in the thinker's mind, but on the basis of those mental states' relations to the external world (6, 147, *passim*). In a nutshell, the reason why we should care about the coherence of our mental states is the following:

The general connection between rationality and correctness is this: if your way of thinking is irrational, that is *bad news* (according to what these mental states are 'telling you' about the world) about how correct this way of thinking is; and the more irrational this way of thinking is, the *worse* the news is about this way of thinking's degree of correctness. (213)

¹ *The Value of Rationality*. Oxford University Press, 2017. x + 267 pp.

² Unless otherwise specified, all page references are to *The Value of Rationality*.

The rough thought then seems to be something like: the higher the degree of coherence of your mental states, the higher the chances of getting things right. Thus, we might say that we should care about rationality because it typically leads to correctness.

The account advanced by Wedgwood is based on the assumption that rationality and irrationality come in degrees (so that when we say that a way of thinking is rational *simpliciter*, we really mean that it is no more irrational than any available alternative, 213), and it unpacks the metaphor of ‘irrationality as bad news about correctness’ in terms of a probability measure on the space of relevant possible worlds:

For every thinker and every time, the mental states and events present in the thinker’s mind at that time determine a relevant space of worlds and a probability distribution P over that space of worlds. The degree of rationality exemplified by the way in which the thinker is actually thinking at this time is determined by how well this way of thinking compares with the available alternatives in terms of their expected degree of incorrectness, according to this probability distribution P. The better this way of thinking compares to the available alternatives, in terms of their expected degree of incorrectness according to P, the more rational the way of thinking is. So, in particular, if this way of thinking’s expected degree of incorrectness is *minimal* -- that is, no higher than the expected degree of incorrectness of any available alternative -- then it compares with every available alternative as well as it possibly could. In that case, the way in which the thinker is thinking at this time will be ideally or optimally rational. In a slogan, to be *rational* is to do *as well as possible at minimizing expected incorrectness*. (217)

I cannot do full justice to the richness of Wedgwood’s proposal here but, before I start discussing it, let me at least note two features that may bring us closer a comprehensive account of rationality, and for this reason I find especially praiseworthy. They are the ambition to harmonize internal and external demands on thinking, and the individuation of a common structure in rational belief and rational choice (see, e.g. 229-30).

Coming to the philosophical merits of Wedgwood’s account, my first concern is the following. I suspect that while the line of argument summarized above may underwrite an account of why agents like us -- inhabitants of a generally cooperative environment -- should care about rationality, it would struggle to account for the notion of rationality that applies to victims of radical sceptical scenarios.

2. *Evil demons*

Consider a subject S_1 who inhabits the actual world w_1 , and her counterpart S_2 who is victim of a Cartesian evil demon in world w_2 . S_1 and S_2 have the same subjective experiences and apparent memories, and go through exactly the same processes of reasoning, forming, maintaining, and revising exactly the same beliefs in exactly the same ways. Following Cohen (1984), Wedgwood maintains that, despite the fact that most of S_1 ’s beliefs are true and most of S_2 ’s beliefs are false, their beliefs are equally rational. Indeed, Wedgwood endorses internalism precisely because he is moved by an argument with this very shape: the so-called ‘argument from hallucination’ (161ff).

However, the more S_2 ’s mental states exhibit the internal coherence in terms of which Wedgwood characterizes rationality, the less they are likely to be correct. To use Wedgwood’s terminology, satisfying the requirement of rationality will lead S_2 to do *less* well with respect to her beliefs’ degree of incorrectness than (some ways of) violating the requirement (215). It would thus seem that the proposal has to face a version of the new evil demon problem. This is typically a problem for externalist accounts of epistemic justification, and it may affect Wedgwood’s internalist account of rationality in virtue of the subordination of the value of rationality to the external aim of correctness.

The issue has a crucial dialectical relevance for Wedgwood, given his commitment to internalism and his endorsement of the argument from hallucination outlined above. But the problem is more than dialectical; it's substantial. Remember that the broader question that Wedgwood is trying to answer is that of explaining why 'it must *always* be true that there is something good simply in being coherent' (38, cf. 33, 202, 207-8, *passim*). The proposed answer is, roughly, that by being coherent one minimizes the risk of incorrectness. The concern is that this answer may apply to agents like us, but it doesn't apply to victims of sceptical scenarios. Hence, it's still not clear why it must *always* be true that there is something good in being coherent.

Someone who sympathizes with the so-called Russellian retreat (cf. Wright 1991: 88) might think that accounting for the significance of the rationality or justification of *our* attitudes is good enough—never mind the attitudes of victims of sceptical scenarios. The thought would be that just like one might be tempted to concede that one doesn't *know* any proposition about the external world and yet insist that one is *justified or rational* in proportioning one's beliefs to one's sensory experiences, so one might concede that it's not clear why *any agent whatsoever* – i.e. including victims of sceptical scenarios – ought to be coherent, and yet insist that *at least for agents like us* – supposedly inhabitants of a cooperative environment -- it is always good to be coherent.

Wedgwood himself appears to come close to taking this line when, in summarising why his proposal allows us to make sense of the claim that rationality requires that our beliefs fit with our sensory experiences (233), he notes that 'the explanation of why rationality imposes this requirement will depend on the nature of the rational probability distribution over the relevant space of possible worlds', and he suggests that the (epistemically) possible worlds (in such space) where one is in a radical sceptical scenario 'only take up a *small* proportion of the whole space of worlds, while the worlds where experiences are veridical take up a much *larger* proportion of the space' (*ibid.*). Since this putative fact is supposed to be reflected in the relevant probability distribution, Wedgwood concludes that

at least so long as your experiences and background beliefs contain no unusual defeating evidence, it is conditionally much more probable, given the assumption that it looks to you as though you are facing a red ball, that you really are facing a red ball than that you are a victim of an evil demon or the like. (*ibid.*)

Yet, for the inquiring agent to substantiate this conclusion, she would need an explanation as to why the possible worlds in which one is victim of a sceptical scenario would take up only a small proportion of the space of possible worlds. After all, the agent's sensory experiences -- the very thing that determines the rationality of the corresponding beliefs -- are indistinguishable in the two cases. So, how can one rationally *claim* that the possibility of being the victim of a sceptical scenario is a tiny one? Things may very well be so, and a third-party observer may well be in the position to know that things are so, but how can one -- in the face of the indistinguishability of misleading and veridical sensory experiences -- take a step back and advance that claim from one's own perspective? (Cf. Wright 1991, 2004).³

In an earlier passage, Wedgwood does admit that he has no explanation yet as to why the demon-world should take up only a small proportion of the space of possible worlds:

³ I take it that the significance of the sceptical challenge is not merely to attack the possibility of knowledge as such, but rather to generate a situation in which the inquiring agent struggles to lay claim to the knowledge she might have. Accordingly, a satisfactory answer would individuate a way in which the agent may rationally claim to have the relevant piece of knowledge. For some initial discussion, see Wright (2010: 204-205), and Stroud (1994/2002: 147ff).

Since it seems rational to have a high degree of confidence in the contents of one's sensory experiences (at least in the absence of special defeating factors), there is presumably some explanation that can be given of why it is rational. Whatever this explanation is, it should be capable of being adapted to give an explanation of why this space of possible worlds must have this feature. (226)

Wedgwood goes on to note that this provides an example of how big questions of epistemology 'can be reinterpreted as questions about the rational probability measure on this space of worlds' (ibid.). That may be right, but we're still falling short of explaining why 'it must *always* be true that there is something good simply in being coherent', which was advertised as one of the main goals of the book. In the concluding remarks, Wedgwood concedes that in order to assess the significance of the proposed conception of rationality, we will have to consider what light it sheds on traditional epistemological issues, including radical scepticism—something that will be done in the second volume (247). If so, one of the main goals of *this* volume is hostage to a discussion that will appear in the second volume.⁴

3. Two ways in which rationality may have normative significance?

One might think that the issues that arise in relation to the new evil demon problem reveal that there are two ways in which rationality may have normative significance: in an internalist sense, rationality may require coherence between the subject's mental states, and in an externalist sense it may demand that attitudes are adequately supported by one's reasons. Following Wedgwood (11), let's call the first putative kind 'rationality as coherence: RC', and the second 'rationality as reasons-responsiveness: RR'.⁵

While advocates of such suggestion need to account for the putative common core of the two notions, in virtue of which both deserve to be called 'rationality' or 'justification' and have normative significance,⁶ pluralism about the normativity of rationality seems to be at least an available theoretical option. Yet, appealing to internalism, Wedgwood claims that the difference between RR and RC is a distinction without a difference:

[I]f (as RR implies) rationality consists in having attitudes that are 'justified' or 'supported' by one's 'reasons', one must *have* the relevant reasons at the relevant time -- and whether or not one '*has*' a certain reason at this time depends on the *mental states* and *mental events* that are present in one's mind at that time. (11)

Once we grant that whether one has the relevant reasons is determined by the agent's mental states and events -- this is the appeal to internalism -- the issue of the responsiveness of one's attitudes to reasons becomes a question of one's attitudes fitting with the subset of one's mental states constituted by the

⁴ In correspondence, Wedgwood admitted that he has not yet spelled out the explanation of why the evil-demon worlds have low probability according to the 'rational probability function', but insisted that we should not doubt that there is such an explanation to be found. If one shares Wedgwood's optimism that we can explain why the rational probability function—the probability that should be guiding our thinking—is one of the *very many* probability functions in which the evil-demon worlds have low probability, one may also find that the general explanatory question ('what is good about thinking rationally?') considered in this book has been satisfactorily addressed.

⁵ If we keep our focus on the epistemic case, the suggestion goes back at least to Kornblith (1983), and Alston (1985). For current discussion see, among others, Scanlon (2007), Worsnip (2018), Christensen (2010), Pryor (2018), Hughes (fc).

⁶ See Alston (2005) for an attempt to articulate pluralism about epistemic justification.

possessed reasons. But that is just a version of the broad notion of coherence articulated by Wedgwood; hence it would seem that RR, modulo internalism, is entailed by RC.⁷

The appeal to internalism is central in the way Wedgwood defuses another challenge to the suggestion that there is a unique sense in which rationality is normative: the one arising from cases where the agent forms a rational false belief that she ought to Φ , together with the assumption that normative reasons consist in facts that do not concern the agent's mental states. In such cases, RC appears to demand that the agent goes on to Φ ⁸, and yet since the belief that she ought to Φ is false, it cannot presumably provide a genuine normative reason to Φ ;⁹ thus, it would seem that RR demands that the agent does not Φ . Invoking internalism allows to rebuke this challenge because it leads to the rejection of the assumption that normative reasons are provided by facts external to the agent's mind: the relevant type of internalism has it that normative reasons supervene on one's non-factive mental states. For an internalist there are thus no obvious obstacles to allowing that rational false beliefs may provide genuine normative reasons,¹⁰ and the puzzle we are considering dissipates. One way out is to say that the agent in the case outlined ought to follow RC -- the real requirement of rationality -- and Φ , even if in the specific case Φ -ing is not the right thing to do.¹¹ After all, fallibility is a feature of human cognition and deliberation.

We have thus seen that internalism plays a crucial role in Wedgwood's reduction of RR to RC. One might thus now expect a discussion of the arguments provided in its support in chapter 7, but I prefer to continue the discussion on whether there might be more than one sense in which rationality is normative. Let's try to take seriously the option that normative reasons consist in facts that may lie outside of the agent's mind, and see how we might deal with the cases outlined.

Once we avail ourselves of the distinction between normative reasons (reasons that speak in favour of making some move in reasoning or deliberation) and motivating reasons (reasons on whose basis the agent makes the relevant move in reasoning -- possibly because she takes them to be normative reasons), we can provide the following explanation about the two problematic cases considered above. We can say that while in the unproblematic cases where the evidence is *not* misleading one's motivating reasons match one's normative reasons, in cases of misleading evidence about what one ought to do and in sceptical scenarios, motivating reasons do not match the relevant normative reasons. To account for the sense in which the move in reasoning of the two agents in the problematic cases elicits a positive reaction in the observer, we might wish to add that while they are not justified/rational in Φ -ing, they are excused for doing so (see, for example, Littlejohn *fc*, and Williamson *fc*).

On this approach, normative reasons are cashed out in agent-independent terms, and guide the agent's reasoning by becoming motivating reasons. Thus, the agent's response to normative reasons is mediated by motivating reasons: in the good cases, normative reasons also play the role of motivating reasons.¹² The thing to emphasize here is that since normative reasons are entirely agent-independent,

⁷ A similar move is made by old-fashioned coherentism about justification such as Bonjour (1985), and subjective Bayesianism of the sort outlined by Titelbaum (2010). For the assimilation of RC to RR—the guiding idea being that genuine reasons can never support attitudes that are incoherent—see Kolodny (2005), and Lord (2014).

⁸ If the rationality that the agent's belief that she ought to Φ is beyond question, the only way to maintain overall attitudinal coherence would be to Φ (or to form the intention to Φ).

⁹ On the assumption, shared by Wedgwood (see. p. 99), that only facts or true proposition may provide genuine normative reasons.

¹⁰ Note that this is compatible with the claim that only facts or true propositions can provide normative reasons; it's just that the relevant facts or true propositions pertain to the agent's mental states.

¹¹ For Wedgwood's own (much more articulated!) discussion of the case at hand, see especially §§ 2.4, and 4.2.

¹² Some philosophers who describe good motivating reasons as a subset of normative reasons (those normative reasons had by the subject) are Lord (2017) and Kiesewetter (2018). See also Ginsborg (2006: § 1), Sylvan (2016: § 3.2), Worsnip (2016: 347-8), and Schroeder (2011). The position I am considering here is a declination of the view that in chapter 4 Wedgwood calls 'Reasons First', and attributes to Raz (2011), Dancy (2000), Scanlon

RR needn't be reduced to RC, but may stand as an independent requirement. Hence, the possibility of two distinct requirements of rationality is on the table.

4. Normative and motivating reasons

Wedgwood does not think that normative reasons may be reduced to facts external to agents' mental states, which in the good case may play the role of motivating reasons too (see §§ 4.4-4.6).¹³ One of the arguments advanced suggests that reducing motivating reasons to normative reasons (characterized as facts about the agent's environment) leads to a picture of rational deliberation that overestimates the centrality of outright belief (§ 4.5). While I am not sure that Wedgwood's argument works, I think that focusing on doxastic attitudes other than beliefs (more specifically, suspension of judgment) may be a rather insightful move.

Wedgwood notes that for a fact about the agent's environment to play the motivational role, the agent must have access to that fact. Having access to the relevant fact involves representing it by presentational attitudes like beliefs, experiences or episodic memories (100). But, Wedgwood objects, this suggestion ignores the role of the agent's *partial* beliefs in rational deliberation. Consider cases where one ought rationally to suspend judgment on p . Since p and not- p cannot both be true, at most one of them may be a genuine normative reason; since by assumption the rational response is one of suspension, the relevant fact (be it p or not- p) is not accessible to the subject and cannot play the motivational role. What plays the motivational role then? Wedgwood suggests that, in forming the attitude of suspension, you may be guided 'by the fact that you have a 0.5 degree of belief in p and a 0.5 degree of belief in "not- p "' (101). He adds:

Being guided by such degrees of belief cannot be interpreted as responding appropriately to any outright beliefs in the true propositions that explain how you are rationally required to think. (ibid.)

This strikes me as debatable. According to some views, rational suspension of judgment demands outright belief in (higher-order) propositions such as <the evidence does not support p and it does not support not- p >.¹⁴ If so, rational suspension of judgment may be accounted for in terms of responding appropriately to reason-constituting facts in which agents have outright belief. If we like Wedgwood's emphasis on the role of partial belief, we may characterize such facts in terms of the true proposition that the evidence supports belief in p and belief in not- p to the same degree: 0.5. Alternatively, we can characterize the relevant true proposition as <the evidence does not support p , and the evidence does not support not- p >.¹⁵

(2014) and Schroeder (2007): all of them think that the very same item may play both the role of normative reasons and that of motivating reasons.

¹³ Things are much more subtle than I can represent here. Wedgwood actually talks of items that 'play the normative-explanation role', contrasts them with items that play the 'ideal-motivation role', and argues they are two kinds of normative reasons. Talking of the difference between normative and motivating reasons simpliciter is a necessary simplification given the constraints on space, but see § 4.2 of *The Value of Rationality* for Wedgwood's own introduction of the two notions of reasons under discussion.

¹⁴ See, for example, Bergmann (2005), Crawford 2004, Friedman (2013: §7), Raleigh (fc).

¹⁵ In a different vein, Wedgwood argues that rationality cannot be accounted for fully in terms of responding to reason-constituting facts in which agents have outright beliefs by considering that 'within a Bayesian framework, so long as two agents had different "prior probabilities", their now having exactly the same outright beliefs... is quite compatible with their also having different probability assignments or partial beliefs' (101, emphasis mine). There is no room to discuss this, but it seems to me that, in such cases, the compatibility of different rational responses despite the same outright beliefs would be explained by the different prior probabilities, which determine different epistemic predicaments. Is it surprising that, given different starting points, taking in new information may deliver different results?

Yet, that is a belief whose content is a normative fact concerning what the evidence supports. As such, the belief in question is not independent of what one's general account of normativity is. If the relevant general account of normativity is one that, like Wedgwood's, allows that the factors that play the role of normative and motivating reasons may have an independent source of normativity, then showing that the normative fact in question may play the motivational role won't help Wedgwood's opponents. On the other hand, if the general account of normativity is one that finds the primary source of normativity in what may be properly described as normative reasons -- in turn characterized as agent-independent facts -- then one wishing to advocate the view of rational suspension of judgment outlined above would have to explain how normative reasons that involve explicitly normative facts may arise from normative reasons that do not do so (i.e. agent-independent facts).¹⁶

Be that as it may, I take the aforementioned view on suspension of judgment to offer support for the claim that good motivating reasons cannot be reduced to (normative reasons characterized as) facts about the external world. That's because the account of rational suspension that I am invoking has it that justified suspension on p requires that the agent *judges* that the evidence does not support neither p nor not- p . That judgment is what provides the motivating reason for suspension, and while its content may correspond to a state of affairs obtaining outside of the agent's mind -- that of a body of evidence not clearly pointing towards p or not- p -- the judgment itself is what allows the agent to find a stable doxastic attitude while being pulled towards the two opposite poles of inquiry (belief in p ; belief in not- p) by equally strong evidence. Hence, the judgment is part and parcel of the good motivating reason for suspension.¹⁷

We cannot engage with the debate on the nature of rational suspension of judgment now. The point here is that, even if the considerations advanced by Wedgwood turned out to be unpersuasive, he may be right that investigating the nature of rational suspension of judgment could reveal that the notion of normative reasons characterized as facts that do not concern the agent's mental states cannot play all the roles that reasons need to play in rational deliberation.

If it is right that good motivating reasons cannot always be reduced to normative reasons understood as facts beyond the subject's mind, then the externalist's general account of the relation between reasons and rationality sketched in § 3 above is in trouble. If so, presumably, the way out of the problematic cases inspired by the externalist view (see § 3 above again) would also be revealed to be a non-viable option. Does that make Wedgwood's suggestion that the requirements of rationality-as-reasons-responsiveness (RR) may be reduced to those of rationality-as-coherence (RC) more palatable?

5. *Internalist dilemmas*

We saw that Wedgwood rejects the externalist account proposed in the previous section by invoking a multiplicity of factors that play the roles reasons need to play in rational deliberation. One might wonder whether this pluralism about (normative) reasons is in tension with monism about the normativity of rationality. It may thus be helpful to say something about why it is not.

Wedgwood argues that while linguistic evidence suggests that there is a plurality of 'oughts' (and to each of them corresponds a normative reason),¹⁸ identifying the specific nature of the 'ought' in

¹⁶ The real issue for Wedgwood's opponent here is that it is hard to see how explicitly normative facts (corresponding to propositions that include an evaluative or normative term) may be characterized in agent-independent terms. Doesn't normativity require an agent that is subject to its requirements? On the face of it, the evidence supports p (or not- p) *for agent A* (in relation to some epistemic goal, etc); similarly, q is a reason to Φ *for subject S* (in circumstances C, etc).

¹⁷ By contrast, no analogue judgment need to be part of good motivating reasons for belief and disbelief.

¹⁸ See § 5.2 for an initial list of the various 'oughts' that we come across in ordinary discourse, and the discussion at §§ 2.4 and 4.2 for the association between 'ought' and normative reasons with the focus on the difference

question requires an assessment of the parameters that constrain the semantics of ‘ought’.¹⁹ Only when the parameters include the *distinctively agential* value of rationality (see § 6.4), the relevant ‘ought’ will be the distinctively subjective ‘ought’ of rationality.

While in cases like the ones we have discussed there might be conflicts between different ‘oughts’, in Wedgwood’s picture these are not conflicts between requirements of rationality. Rather they are conflicts between a subjective ‘ought’ that is constrained by the information available to the agent and expresses the requirement of rationality, and an objective ‘ought’ that is not constrained by the information available to the agent (and may be said to express the requirement of correctness).²⁰ And for Wedgwood there is no question about what the rational agent would do in such cases: to the extent that she is rational, she will be guided by the subjective ‘ought’ that coincides with the requirements of rationality. Indeed, the agent is not even in the position to see the conflict: the objective ‘ought’ is out of her view, and she can only see the subjective one. In virtue of this, the dominance of the subjective ‘ought’ in the problematic cases is a feature of internalism about rationality. Unfortunately, it’s not clear that sticking to internalism is enough to rule out the possibility of rational dilemmas.

There are at least two types of cases that seem to generate genuine rational dilemmas, and yet do not rely on externalism about responsiveness to reasons. The first of such cases concerns situations where the agent cannot tell, in principle, what the relevant evidence exactly is:

Suppose I have good evidence that supports having credence 0.8 in p . I’m also confident that p is either logically equivalent to q or logically equivalent to $\text{not-}q$; but I’m not sure which. So I’m not sure whether my credence in q should be 0.8 or 0.2. Aware of my own fallibility, I split the difference and have credence 0.5 in q . In that case, I’d have credences that are probabilistically incoherent. But can’t I be reasonable in doing so? (Pryor 2018: 132)

Here the agent ends up taking a credence towards q when she has very strong reason to believe that this is not the correct credence. So, she ends up having to accept a conjunction of the form ‘I have credence n in q , and the evidence suggests that n is not the correct credence to have in q ’. This situation sounds very much like a violation of the internal coherence championed by Wedgwood. To avoid it, the agent could perhaps suspend on both p and q , but that would seem to come at the price of violating RR—the requirement to follow the evidence or reasons. After all, the evidence clearly suggest that p is true (and that it is logically equivalent to either q or $\text{not-}q$). A similar scenario emerges from Williamson’s clock belief cases. These are cases such that whatever the agent’s evidence exactly is, the agent has a strong reason to think that the evidence will be different from what she takes it to be.²¹

A different type of cases has been outlined by Worsnip (2018: 24ff), who argues along the following lines. Grant that higher-order evidence does matter for what first-order attitude it is rational to take; yet, arguably higher-order evidence matters *less strongly* for the question of what first-order attitude one should take towards p than it does for the question of what related higher-order attitude one should take. Similarly, one could think that first-order evidence bears *more strongly* on the first-order

between what Wedgwood calls ‘objective’ and ‘subjective’ ‘oughts’ (see especially the discussion of the example of the putative soldier approaching at pp. 92-93).

¹⁹ Specifically, they are: ‘the domain f that maps each world of evaluation onto a domain of available worlds that are up for assessment by means of [the relevant] kind of “ought”, a probability distribution E that captures a kind of epistemic perspective that is built into the truth conditions of sentences involving this “ought”, and a value function V that measures the value of the propositions in a partition of the relevant domain of available worlds’ (136). See chapter 5 for details on the semantics of ‘ought’.

²⁰ See especially §§ 2.4, 4.2, and 5.2.

²¹ See Williamson (2014), and Christensen (2010) for some initial discussion.

question whether p than it does on the higher-order question of what the evidence supports.²² Putting the two things together, we should expect there to be cases at the margins where the first-order evidence outweighs the higher-order evidence at the level of the first-order attitude, while the higher-order evidence outweighs the first-order evidence at the level of the second-order attitude. Such cases would be cases where one may be rational to believe conjunctions such as $\langle p \rangle$, and the evidence does not support $p \rangle$. Again, these would seem to be cases where respecting the ideal of coherence between one's attitudes can only be done at the price of violating the ideal of proportioning one's attitude to the evidence or reasons available.

As far as I can see, the types of cases just mentioned do not rely on externalist assumptions about rationality or reasons. If so, we have some reason to question Wedgwood's reduction of RR to RC, even if we are willing to embrace internalism about rationality. Of course, one cannot be expected to discuss all issues related to the position one is articulating, but Wedgwood's defence of monism about the normativity of rationality would have benefited from a discussion of the cases just mentioned.²³ As things stand, I still need some persuasion that the correct account of rationality is not one that allows for the possibility of conflicts between different requirements of rationality. And if that is a live option, then Wedgwood's characterization of the normativity of requirements of rationality as demanding that one *always* conforms with them is also in question.

I hope that, by focusing only on a small number of the issues discussed in *The Value of Rationality*, I have managed to give a sense of the richness, depth, ambition, and far-reaching implications of a volume that will surely become highly influential in the years to come. There is no question that any philosopher will benefit greatly by engaging with this book. And apart from the merits in enhancing our understanding of the notion of rationality, the volume stands as a model for constructive theory building on a scale that it is rarely seen these days in philosophy. May it be an inspiration to other philosophers!²⁴

The University of Stirling

Stirling

FK9 4LA

giacomo.melis@stir.ac.uk

giacmelis@gmail.com

Funding

This work was supported by the John Templeton Foundation [grant number: 58450], and the Knowledge Beyond Natural Science project at the University of Stirling.

References

Alston, W. P. 1985. Concepts of Epistemic Justification. *The Monist* 68 (1): 57-89.

Alston w. P. 2005. *Beyond Justification: Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.

²² See (Worsnip 2018: 26-8) for details.

²³ At one point (p. 12, footnote 16) Wedgwood anticipates a discussion of Worsnip's 'ingenuous cases' in § 2.4, but the discussion of § 2.4 does not address directly Worsnip's cases. Clock-beliefs type of cases are never mentioned.

²⁴ Thanks to Kent Hurting, Alan Millar, Joshua Thorpe, and Ralph Wedgwood for comments on an earlier draft. I am also indebted to Rowan Cruft, Philip Ebert, Kent Hurtig, Alan Millar, Peter Milne, Joshua Thorpe, and Crispin Wright for discussion.

- Bergmann, M. 2005. Defeaters and Higher-Level Requirements. *Philosophical Quarterly* 55 (220): 419-36.
- Bonjour, L. 1985. *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Christensen, D. 2010. Rational Reflection. *Philosophical Perspectives* 24 (1): 121-40.
- Cohen, S. 1984. Justification and Truth. *Philosophical Studies* 46: 279-95.
- Crawford, S. 2004. A Solution for Russellians to a Puzzle About Belief. *Analysis* 64 (3): 223-29.
- Dancy, J. 2000. *Practical Reality*. Oxford: Oxford University Press.
- Friedman, J. 2013. Suspended Judgment. *Philosophical Studies* 162 (2): 165-81.
- Ginsborg, H. Reasons for Belief. *Philosophy and Phenomenological Research* 72 (2): 286 – 318.
- Hughes, N. Forthcoming. Dilemmic Epistemology. *Synthese*, online first: <https://link.springer.com/article/10.1007/s11229-017-1639-x>.
- Kiesewetter, B. 2018. How reasons are sensitive to available evidence. In *Normativity: Epistemic and Practical*, eds. C. McHugh, J. Way, and D. Whiting. Oxford: Oxford University Press, 90-114.
- Kolodny, N. 2005. Why be Rational? *Mind*, 114 (455): 509–63.
- Kornblith, H. 1983. Justified Belief and Epistemically Responsible Action. *Philosophical Review* 92 (1): 33-48.
- Littlejohn, C. Forthcoming. A Plea for Epistemic Excuses. In *The New Evil Demon*, eds. J. Dorsch and J. Dutant. Oxford: Oxford University Press.
- Lord, E. 2014. The Coherent and the Rational, *Analytic Philosophy* 55 (2):151-75.
- Lord, E. 2017. What You're Rationally Required to Do and What You Ought to Do (Are the Same Thing!). *Mind* 126 (504): 1109-54.
- Pryor, J. 2018. The Merits of Incoherence. *Analytic Philosophy* 59 (1): 112-41.
- Raleigh, T. Forthcoming. Suspending is Believing. *Synthese* online first: <https://link.springer.com/article/10.1007%2Fs11229-019-02223-8>
- Raz, J. 2011. *From Normativity to Responsibility*. Oxford: Oxford University Press.
- Scanlon, T. 2007. Structural Irrationality. In *Common Minds: Themes From the Philosophy of Philip Pettit*, eds. G. Brennan, R. Goodin, F. Jackson, and M. Smith. Oxford: Clarendon Press.
- Scanlon, T. 2014. *Being Realistic about Reasons*. Oxford: Clarendon Press.
- Schroeder M. 2007. *Slaves of the Passions*, Oxford: Oxford University Press.
- Schroeder, M. 2011. What Does it Take to Have a Reason? In *Reasons for Belief*, eds. Reisner and Steiglich-Petersen. Cambridge: Cambridge University Press, 201-22.

- Stroud B. 1994 (2002). Scepticism, 'Externalism' and the Goal of Epistemology. In *Understanding Human Knowledge: Philosophical Essays*. Oxford: Oxford University Press, 139-154.
- Sylvan, K. 2016. Epistemic Reasons I: Normativity. *Philosophy Compass* 11 (7): 364-37.
- Titelbaum, M. G. 2010. Not Enough There There: Evidence, Reasons, and Language Independence. *Philosophical Perspectives* 24: 477-528.
- Williamson, Timothy. 2014. Very Improbably Knowing. *Erkenntnis* 79: 971-99.
- Williamson, T. Forthcoming. Justification, Excuses and Sceptical Scenarios. In *The New Evil Demon*, eds. J. Dorsch and J. Dutant. Oxford: Oxford University Press.
- Worsnip, A. 2016. Moral Reasons, Epistemic Reasons and Rationality. *Philosophical Quarterly* 66: 341-61.
- Worsnip, A. 2018. The Conflict of Evidence and Coherence. *Philosophy and Phenomenological Research* 96 (1): 3-44.
- Wright, C. 1991. Scepticism and Dreaming: Imploding the Demon. *Mind* 100 (397): 87-116.
- Wright, C. 2004. Warrant for Nothing (and Foundations for Free)? *Aristotelian Society Supplementary Volume* 78 (1): 167-212.
- Wright, Crispin. 2010. Crispin Wright: an Epistemological Self-Profile. In *A Companion to Epistemology: Second Edition*, eds. J. Dancy, E. Sosa Ernst, and M. Steup. Oxford: Blackwell, 204-210.