What the Cluster View Can Do for You[1]
Daniel Fogal & Alex Worsnip
Forthcoming in *Oxford Studies in Metaethics* – please cite published version if possible

1. Prelude: a not-so-strange story

You need to leave the house, but it's raining outside. And you don't want to get your clothes wet, since you have an important meeting to attend. Your umbrella will prevent your clothes from getting wet. So you should take your umbrella with you!

Of late, philosophers of normativity like to talk a lot about reasons. What is the (normative) *reason* for you to take your umbrella, in this case? The case is hardly complicated, by the standards of philosophical examples. And yet, there is a bewildering variety of candidates. Some were explicitly mentioned:

1. You need to leave the house.
2. It's raining outside.
3. You don't want to get your clothes wet.
4. You have an important meeting to attend.
5. Your umbrella will prevent your clothes from getting wet.

At least one was left implicit:

6. If you show up to the important meeting in wet clothes, something bad will (or might) happen.

There are still other candidates we could list, including ones inferable from those above. For example:

7. By taking an umbrella, you will avoid showing up to your important meeting wet.
8. By taking an umbrella, you will prevent something bad from happening.

No doubt we could go on.

So: which of these facts is *the reason* for you to take the umbrella with you? If that sounds like a silly question to you: good. We're on your side.

You might initially think it sounds like a silly question because these are *all* good reasons to take your umbrella with you. But that's not quite right. It isn't that these are all *separate*

reasons to take your umbrella with you. Contrast another candidate reason to take your umbrella:

9. Your umbrella is super stylish.

The relationship between facts (1)-(6) seems importantly different from the relationship between any of (1)-(6) and (9). (1)-(6) form a *single* coherent rationale for carrying an umbrella, whereas (9) seems like it's introducing a *new*, distinct rationale for doing so. As we'll say, (1)-(6) provide a single "strain" of normative support for carrying an umbrella, where (9) plus other facts provide another.

2. The atomic view and the cluster view

What's the point of the foregoing observations? It's to put pressure on a conjunction of theses that—despite the myriad controversies about reasons—are frequently taken for granted:

(i) Reasons are sources of normative support for actions, attitudes, etc.

(ii) Reasons, at least paradigmatically and in simple cases, consist in atomic facts.

To see how this conjunction is taken for granted, consider the sorts of things that typically get cited as illustrative examples of reasons. It's said, for example, that the fact that you are howling is a reason for me to believe that you are in pain. And that the fact that you are in pain is a reason for me to help you. And that the fact that it is raining is a reason for you to take your umbrella. These facts, notice, are atomic—that is, they are not conjunctions or disjunctions of facts. And, it's thought, in saying that these facts are reasons for attitudes and actions, we are saying that they count in favor of—provide normative support for—these actions and attitudes. Call the conjunction of (i) and (ii) the **atomic view**.

The atomic view is simple and, in some respects, attractive. It's not just an artifact of the philosophical literature that atomic facts get cited as reasons: this is also the case in ordinary thought and talk. If I ask what reason there is for me to take an umbrella, you might just reply 'well, it's raining outside', and leave it at that. You are unlikely to tiresomely recite all the facts from (1)-(6).

And yet the umbrella case puts some *prima facie* pressure on the atomic view. For it seems that none of (1)-(6) *on their own* provide much, if any, support for taking an umbrella. The fact that you have an important meeting to attend, for example, doesn't provide any support for taking your umbrella if it's not raining outside. And if you don't need to leave the house, the fact that it's raining won't support carrying an umbrella as you potter around the house.

This suggests that it's really the cluster of facts (1)-(6)—and perhaps more—that *jointly* provide normative support for taking an umbrella with you.

In response, the proponent of the atomic view could concede that a whole cluster of facts enters into the explanation of why there's normative support for your taking an umbrella. They could maintain, however, that only one of these facts—perhaps, the fact that it's raining outside—plays the role of constituting *the reason* for you to take an umbrella with you. The others, by contrast, merely serve as "background conditions": things that have to be in place *in order* for the fact that it's raining outside to be a reason for you to take an umbrella with you, or that help to *explain why* the fact that it's raining outside is a reason for you to take an umbrella with you.

Yet, we submit, this picture is in full generality a somewhat incredible one. It returns us to the idea that just one of the facts from (1)-(8) plays the privileged role of being *the reason* to take an umbrella with you, with the others relegated to the status of background conditions. But it seems very hard to find any principled way of identifying which of the facts from (1)-(8) has this privileged role. After all, for each fact there are contexts where it's felicitous to cite it as a reason to take an umbrella. For example, if it's already common ground that it's raining, and I nevertheless ask you what reason there is for you to carry an umbrella, you might cite the fact that you have an important meeting to attend. If (somehow) I don't know that umbrellas prevent one's clothes from getting wet, you might cite that fact. And so on.

Reflections on cases like this motivate rejecting the atomic view and replacing it with what we'll call the **cluster view**. On this view, even in the simplest cases, and even restricting ourselves to a single "strain" of normative support for an action or attitude, the normative support for an action or an attitude is typically provided by a whole cluster of facts. Moreover, many of these facts are *on an explanatory par*—they each play the same kind of role in explaining why there is normative support for the action or attitude in question, as opposed to one having privileged status as *the reason* and the others serving as background conditions. This isn't necessarily to say that there are *never* any background conditions. Rather, it's to say that the number of facts playing the non-background-condition role—i.e., that are front and center in providing a single strain of normative support—will rarely if ever be just one; rather, a cluster of facts will play this front-and-center role.[2]

A few additional clarificatory remarks are in order. First, recall that the atomic view consists of two claims: (i) that reasons are sources of normative support for actions, attitudes, etc; and (ii) that reasons, at least paradigmatically and in simple cases, consist of atomic facts.

---

[2] A forerunner of the cluster view is Fogal (2016), though the focus is on linguistic considerations. We also briefly introduced the view in Fogal & Worsnip (2021). Other views that overlap aspects of the cluster view (without fully committing to it) include Wedgwood (2015) and Titelbaum (2019).

The cluster view can *either* be understood as rejecting claim (i) or claim (ii), depending on how we use the count-noun 'reason(s)'. If we stipulatively use the term to refer to sources of normative support, then the cluster view rejects (ii): reasons then consist in clusters of facts, not atomic facts, even in simple cases. On the other hand, if we use 'reason(s)' to refer to the (typically atomic) facts that are *cited* as reasons in ordinary language, then the cluster view rejects (i): reasons do not, on their own,[3] constitute sources of normative support. This is a largely terminological issue: either way, what's important is that *sources of normative support* are clusters of facts rather than atomic ones. In what follows, we'll usually talk in this more neutral way.

Second, although the cluster view is itself a metaphysical view, it is naturally paired with a kind of contextualist view about ordinary uses of the count noun 'reason(s)'.[4] On this view, when a cluster of facts normatively supports Φ-ing, which of these facts it's apt to *cite* as a reason will depend on the conversational and informational context. In particular, it will depend on how informative different facts belonging to the cluster are in helping one grasp why there's normative support for Φ-ing. Most obviously, if someone is *already aware* of certain facts within the cluster, but not others, it'll typically be more informative to cite the facts that they're not aware of. Crucially, though, the cluster view counsels against mistaking these pragmatic considerations about what's most informative, and therefore most apt, to cite as a reason for *metaphysical* asymmetries in the role that the different facts play in explaining why there's normative support for Φ-ing.

Relatedly, it's often apt to cite a fact inferable from more than one of the facts in the cluster as a reason, as a kind of "summary" or "representative" of the cluster as a whole. This is what happens in our opening example if (7) or (8) is cited as a reason. These facts are inferable from facts in the original cluster (1)-(6), but aren't part of the cluster itself—they don't *combine* with (1)-(6) to generate the relevant support.[5]

Finally, though the cluster view itself concerns what *normatively supports* responses, we hold a similar view about what *motivates* responses.[6] Just as philosophers often talk about normative reasons in a way that presupposes an atomic view, so too they often talk about motivating reasons in a way that presupposes an atomic view. On this view, there's often a single atomic fact (or consideration) that on its own motivates a response. Again, a rival

---

[3] The cluster view can still accept that reasons in this second sense are *partial* sources of normative support, in the sense that they combine with other reasons to provide normative support.
[4] See Fogal (2016).
[5] See also Johnson King (2019), who rightly observes that you can felicitously cite things from different metaphysical "levels" as 'the reason'. The cluster view goes further, claiming that *within a given level* it's typically clusters of facts that provide normative support, and that any member of the cluster can be felicitously cited as 'the reason', given the right context.
[6] Cf. Fogal (2018).

picture is one whereby responses are typically motivated by a cluster of considerations (even within a single "strain" of motivation), all of which are on an explanatory par even though some will be more apt to cite for conversational or communicative purposes.

We think that once the dispute between the atomic view and cluster view is laid bare, the cluster view is extremely plausible. It's just not very plausible that there'll typically be a unique fact that plays the privileged role of being *the* reason for Φ-ing (even within a single strain of normative support), with everything else relegated to background condition status. In fact, perhaps you're now finding the cluster view so plausible that you doubt anyone would deny it. And yet, as we'll show, the atomic view is frequently presupposed in debates about reasons. We'll aim to show that this presupposition generates unnecessary (psuedo-)puzzles that can be largely (dis)solved with the cluster view in place. The cluster view's ability to do this constitutes a further argument in its favor, beyond the intuitive considerations we've adduced so far. We'll then consider whether there are any compelling reasons to reject the cluster view. In particular, we'll examine Mark Schroeder's arguments against what he calls the "no background conditions" view and see whether they have any force against the cluster view.

3. What the cluster view can do for you

3.1 *Substantive rationality and epistemic constraints*[7]

It's widely held that (substantive[8]) rationality consists in correctly responding to one's reasons.[9] Yet it's also widely held that one cannot be rationally required to respond to reasons that are entirely outside one's epistemic ken.[10] To take a notorious example due to Broome (2007), the fact that the fish contains salmonella is a reason for you to refrain from eating it; yet, if you don't know—and couldn't possibly have known—that the fish contains salmonella, it seems absurd to say that you are *irrational* for eating it. Judgments of *rationality*—perhaps in contrast to judgments about what one ought or has reason to do—are necessarily constrained by the epistemic situation of the agent to whom they apply.

To deal with examples like this, proponents of the view that (substantive) rationality consists in correctly responding to reasons typically say that it consists in correctly responding to a particular *kind* or *subclass* of reasons, variously called one's "evidence-relative", "possessed", "available", or "apparent" reasons. The details differ between accounts, but the broad idea is that rationality only requires one to respond to reasons that meet an epistemic constraint:

---

[7] This subsection overlaps Fogal & Worsnip (2021: §6).
[8] As contrasted with structural rationality, which is concerned, broadly speaking, with coherence. Cf., e.g., Fogal & Worsnip (2021).
[9] E.g., Kiesewetter (2017), Lord (2014, 2018), Worsnip (2021: ch. 2).
[10] E.g., Broome (2007), Sylvan (2015), Kiesewetter (2017), Lord (2018), Worsnip (2021: §2.1).

roughly, those reasons that one's epistemic situation makes available to one, or puts one in a position to know. Since the fact that the fish contains salmonella is outside your ken, you aren't rationally required to respond to it. That's why it's not irrational for you to eat the fish.

However, some (Lord 2014, 2018; Sylvan 2015) charge that this epistemic constraint is not restrictive enough to account for the full range of intuitive data. Consider:

> **Lois's Fish.** Lois just ordered fish from her favorite seafood restaurant. Right before she digs in, the waiter comes out to inform her that the fish contains salmonella. Lois has the unfortunate belief that salmonella is one of the many bacteria found in food that is harmless to humans. And, indeed, this belief is rational. A renowned food scientist told her so. So she goes ahead and forms an intention to eat the fish and eats the fish. (Lord 2018: 98)

Lois is aware that

10. The fish contains salmonella.

So, it seems that this reason meets the relevant epistemic constraint for bearing on (substantive) rationality. And yet Lois is still not irrational for eating the fish.

Lord's solution to this (supposed) problem is to add a "practical" condition on one's "possessing" a reason (and, thereby, on that reason's bearing on rationality).[11] Not only must one be in a position to know the fact R that constitutes the reason; one must also be "in a position to manifest knowledge about how to use R as a reason to Φ" (Lord 2018: 121). Only then will R bear on verdicts of substantive rationality. Since Lois doesn't meet this practical condition with respect to her reason for eating the fish, she's off the hook for being rationally required to respond to it.

Note, however, that this discussion rests on an unarticulated assumption—namely, that the fact that the fish contains salmonella enjoys a privileged status as *the reason* for Lois to refrain from eating the fish, a status that the fact that

11. Salmonella is harmful to humans.

doesn't share. Presumably, (11) is taken to be merely a background condition.

But why think this? We submit that there's no good reason for thinking so. (10) and (11)—and perhaps other facts—plausibly *come together* to provide a coherent rationale for refraining from eating the fish; without either, the rationale seems crucially incomplete. Moreover, we

---

[11] Cf. also Sylvan (2015).

don't see any reason to hold that the two facts play importantly different roles in explaining the normative support for refraining from eating the fish. What's misleading here is that for most of us (though evidently not Lois), the fact that salmonella is harmful to humans is an extremely familiar fact, one that couldn't easily fail to obtain. This tends to put it in the background of our *thinking* about Louis and the fish, as something we simply take for granted. But as suggested earlier, this isn't a good reason to think that *metaphysically* speaking, it's merely in the background. Consequently, the proponent of the cluster view should say that (10) and (11) *together* provide support for refraining from eating the fish.

But now we get a much simpler explanation of why Lois is not irrational for eating the fish: Lois is unaware that salmonella is harmful to humans, and this fact is a crucial member of the cluster of facts that provides normative support for refraining from eating the fish. Lois thus fails to meet the epistemic condition for this cluster's bearing on her substantive rationality. Being aware that the fish contains salmonella simply doesn't suffice for meeting the *epistemic* condition, once we acknowledge that this fact *on its own* fails to provide normative support. Thus, there's no need to introduce any kind of practical constraint on one's possessing a reason, or on its bearing on (substantive) rationality. We simply need to reformulate the epistemic condition so that it doesn't presuppose the atomic view. Roughly, the reformulated epistemic condition says that one just needs to be in a position to know the *full* cluster of the facts that provide the relevant strain of normative support (or, at least, a representative subset thereof).[12]

This explanation has at least three advantages over Lord's. First, it's simpler: it doesn't introduce a new complex "practical condition" to deal with the modified fish case; instead, it just adverts to the original epistemic condition (formulated to cast off the atomic view). Second, it's truer to the intuitive source of Lois's problem. Intuitively, Lois's problem is epistemic, not practical: she doesn't know that the fish contains salmonella. Indeed, her problem seems symmetric to someone in the opposite position—that is, someone who knows that salmonella is harmful to humans but isn't aware that the fish contains salmonella. Saying that Lois fails to meet Lord's practical condition seems a roundabout way to explain her problem. Third, our explanation avoids various other problems that the practical condition introduces. We have explained these problems elsewhere (Fogal & Worsnip 2021: §6.1); briefly, the main problem is that such a condition makes it *too easy* for agents to avoid irrationality in a range of other cases where they culpably fail to acknowledge the normative

---

[12] We include the parenthetical qualification to allow that sometimes one might be in a position to see that a cluster of facts supports some response without being aware of *every* single fact in that cluster. But this is obviously not the case with respect to the case at hand: when unaware that salmonella is harmful to humans, Lois lacks something that's *crucial* to seeing how the cluster as a whole supports refraining from eating the fish.

upshot of some set of considerations, and thereby fail to meet the practical condition. As such, this is an example of how switching to the cluster view brings a variety of advantages.

3.2 *Proper basing* I: *acting with moral worth*

We turn now to the debate about when an action has moral worth. Moral worth is a status that not all morally right actions attain—one might perform the right action by accident, for example, in which case one is not morally praiseworthy or creditworthy *for having done the right thing*. It's this evaluation, roughly speaking, that is expressed by saying that one's action lacks moral worth.

A natural proposal for explaining the relevant kind of non-accidentality is to say that a morally right action has moral worth when it is done *for the right reasons*—that is, when the right reasons serve as one's *motivating* reasons for performing the action. But what are the right reasons? There are two competing views. On one, to act for the right reasons (and hence for one's action to have moral worth) is to be motivated by the features of the action that *make* it morally right.[13] On another, it's to be motivated by the fact that the action in question is morally right.[14]

In a recent article, Keshav Singh (2020) contends that neither of these views—which he labels the 'Right-Making Features View' and the 'Rightness Itself View', respectively—succeeds in securing the relevant kind of non-accidentality. He argues by counterexample. Let's begin with his counterexamples to the Right-Making Features View.[15]

Singh's first case is as follows:

> **Venom.** Jack, a surgeon, is hiking when he sees a stranger get bitten by a venomous snake and faint. He immediately makes an incision near the bite so that the venom will drain out. Making the incision is the right thing to do, and Jack's reason for doing it (that it will allow the venom to drain out) is part of what makes it right. But Jack doesn't have any particular concern for doing the right thing in this case, nor does he conceive of his reason as one that makes his action right. He is simply intrinsically interested in draining venom out of wounds. (Singh 2020: 162)

Though Jack does the morally right thing, his action doesn't have moral worth. In particular, there's a kind of accidentality involved: the right thing to do just happens to line up with what

---

[13] Cf. Arpaly (2002) and Markovits (2010).
[14] Cf. Sliwa (2016) and Johnson King (2020). The view is commonly attributed to Kant.
[15] Cf. also Lord (2017: 438-9), Way (2017: 210-11), and Howard (forthcoming: 2-3).

Jack is "intrinsically interested in"—namely, draining venom out of wounds. Yet, Singh claims, Jack's (motivating) reason for making an incision near the bite is

12. Making an incision will allow the venom to drain out.

And this is a genuine right-making feature of his action. Thus, the case initially seems to be a counterexample to the Right-Making Features View.[16]

But, as Singh acknowledges, (12) is only *part* of what makes the action right. Another crucial part is that allowing the venom to drain out will improve the stranger's condition: if draining the venom won't help (or if it'll make things worse), then making the incision *won't* be the right thing to do. Singh doesn't actually say what the benefit of allowing the venom to drain out is. So to fix ideas, let's stipulate that:

13. Allowing the venom to drain out will save the person's life.

We now have at least two facts that help make the act of making an incision right—(12) and (13). The question is then whether either enjoys some privileged status over the other as *the reason* for making an incision. Unsurprisingly, we think the answer is: they do not. Instead, the two facts (perhaps with others) *jointly* provide support making an incision.

Importantly, however, Jack *isn't* motivated by (12) and (13). A mismatch thus results between the cluster of facts that provide normative (specifically, moral) support for Jack's action and those that motivate him, since (13) is a crucial part of the former but is missing from the latter. Indeed, though Singh says that Jack's motivating reason for making an incision near the bite is (12), this too is questionable. Rather, what really motivates Jack is (12) *along with* something like:

14. Draining venom out of wounds is interesting (to me/Jack).

And (12) and (14) *don't* provide moral support for making an incision. Although (14) may not be something that explicitly occurs to Jack in deliberation, it's still an essential part of the full motivating story—i.e., the distinctive person-level, rationalizing explanation that

---

[16] Singh acknowledges that the Right-Making Features View could be revised or formulated to avoid this first counterexample. His suggested revision holds that only *intrinsically* morally significant considerations—by which he seems to mean 'considerations that are necessarily morally significant—count as right-makers in the relevant sense. But it's a bit of a mystery why it's only morally worthy to be motivated by "intrinsically" morally significant considerations: if one has a grasp of why some set of contingently morally significant considerations is morally significant, and this grasp plays an important role in motivating one, this seems to suffice for acting with moral worth. We thus think that our own response, drawing on the Cluster View, is preferable.

motivating reasons characteristically provide—when it comes to Jack's action.[17] Jack makes the incision, after all, because it allows him to do something he's intrinsically interested in doing—namely, drain venom out of a wound. So while it's natural to say that Jack made the incision because it will allow the venom to drain out, it's also natural to say that he made the incision because he's intrinsically interested in draining venom out of wounds. Fortunately, given the cluster view, we don't face the question of which of the descriptions picks out the *real* reason Jack made the incision—they are complementary rather than competing claims, with each highlighting a different element of the cluster that provides the full motivating explanation of Jack's action.

The upshot is that once we switch from a tacitly atomic framing to one that appropriately takes account of the cluster view, the Right-Making Features View can accommodate **Venom**: Jack's action fails to be morally worthy because there's a mismatch between what motivates his action—(12) and (14)—and what makes his action right—(12) and (13).

However, Singh offers another, somewhat more challenging counterexample to the Right-Making Features View:

> **Venom*.** Jack's fundamental motivating reason for making the incision is that it will save the stranger's life [...] But this is not because he sees anything particularly morally significant about saving lives. He has no concern for doing the right thing. Rather, he is simply intrinsically interested in saving lives in the same way he was intrinsically interested in draining venom in **Venom**." (Singh 2020: 162)

Here, Jack is motivated by both (12) and (13). Yet his action still seems to lack moral worth.

With the cluster view in hand, there are two possible responses to this revised case. The first begins by noting that, just as we should say that Jack is moved by (12) and (14) in **Venom**, so we should really say in **Venom*** that Jack is motivated by (12), (13) *and*

15. Saving lives is interesting (to me/Jack).

---

[17] It might be objected that Jack need not be motivated by the *fact* or *proposition* (16), but instead might be motivated by *his interest in* draining venom out of wounds (which is more like a desire). We agree. (Much more generally, the fact that the full story about an agent's motivation typically involves both (beliefs about) facts *and* desire-like states, whereas (*pace* the view considered in sec. 3.4 below) normative reasons do not, constitutes a much more general complication for the view that motivating and normative reasons "mirror" each other, and for accounts of statuses like moral worth in terms of a "match" between the two.) Nevertheless, it's still plausible that moral worth requires one's action not to be crucially motivated by desires that don't in some sense correspond to a feature of normative reality (e.g., to what is of value), and Jack's desire doesn't.

But now, again, there's something crucial to Jack's motivation that doesn't reflect anything that morally supports making the incision. If his motivation were to mirror moral reality, (12) and (13) *on their own* would suffice for motivating him.[18] The fact that (15) is also necessary shows that what motivates him fails to mirror moral reality—the facts about what makes his action right—and hence that his action lacks moral worth.

A second reply to **Venom*** is also possible. Perhaps the moral support for making the incision is provided by (12) and (13) *together* with a general moral fact. For simplicity, suppose it's the following:

  16. Saving lives is, other things being equal, the right thing to do.

If what really makes the incision the right thing to do is (12), (13) and (16), then Jack's action lacks moral worth in **Venom*** for reasons parallel to those for which it lacked moral worth in **Venom**. What motivates him is (12), (13), and (15), when what supports making the incision is (12), (13), and (16).[19]

Admittedly, we enter controversial territory here. Some, including leading proponents of the Right-Making Features View, think that facts about what is morally right do not themselves constitute moral reasons: rather, moral reasons consist only in the (non-moral) facts that *make it the case* that an action is morally right.[20] However, Zoë Johnson King (2019) persuasively argues against this contention. Moreover, note that while the fact that making the incision is morally right can't help make it right to make the incision—a fact can't make itself the case—the more *general* fact that saving lives is pro tanto morally right—viz., (16)—arguably *is* part of what makes making the incision the right thing to do.[21] If so, the Right-Making Features View could be revised to allow both moral *and* non-moral facts that are part of the explanation of why Φ-ing is morally right to be part of what motivates an agent who Φ's with moral worth.

This is a sort of concession to the Rightness Itself View, since it allows moral facts to be part of what motivates those who act with moral worth. But it strikes us as an attractive compromise. Indeed, given the cluster view, we can now see the kernel of truth in the

---

[18] Again, this claim is complicated by the fact that we should expect *some* kind of desire-like state to play a role in the full motivating or rationalizing explanation of an action. See n. 17 above for discussion.

[19] For simplicity we're following much of the literature in treating moral worth as all-or-nothing—actions either have it or don't. However, we're sympathetic with the idea that moral worth is a graded property, coming in degrees. Being motivated by the full cluster is at best required for *full* moral worth, with partial credit still available to those motivated by a not-too-incomplete subset. Jack may get partial credit in **Venom***.

[20] Cf., e.g., Dancy (2000: 165-167); Zimmerman (2007: appx. 2).

[21] Cf. Enoch (2019); Fogal & Risberg (2020).

Rightness Itself View. According to the simplest version of such a view, Jack acts with moral worth iff Jack is motivated by the following:

17. Making the incision is the morally right to do.

But notice that (17) is *inferable* from (12), (13) and (16), assuming that no other morally relevant features are at play. Thus, if Jack is motivated by (12), (13) and (16), it'll typically also be felicitous to cite (17) as his motivating reason. However, the reverse doesn't hold: just because Jack is motivated by a fact like (17), it doesn't follow that he can correctly be described as being motivated by (12), (13) and (16). For he might not be tracking what *makes* the incision morally right. Thus, it doesn't follow that he acts with moral worth. This explains why the simplest version of the Rightness Itself View is subject to counterexamples, such as one that Singh offers:

> **Moving.** Simon's friend is in a tough spot and needs last minute help moving. Simon helps him, which is the right thing to do. And he's motivated to do so by the consideration that it's the right thing to do. According to Simon's conception of morality, what's right is what benefits one's friends and harms one's enemies. (Singh 2020: 166)

Singh takes Simon's motivating reason to be:

18. Helping Simon is the right thing to do.

And that's not an infelicitous thing to say. But it would be equally correct to describe Simon's motivation as consisting of the following *two* facts:

19. Helping Simon is an instance of benefitting my friend.
20. It's morally right to benefit one's friends and not one's enemies.

And (19) and (20) do not in fact support helping Simon—since (20) is false.

Summarizing, while *some* agents who can be felicitously described as being motivated to Φ by the thought that Φ-ing is morally right can also be felicitously described as being motivated by the underlying considerations that make Φ-ing morally right, others, like Simon, cannot. On our view, this is what explains why the former agents act with moral worth, but Simon doesn't. The cluster view thus gives us a version of the Right-Making Features View that identifies the kernel of truth in the Rightness Itself View, while also being immune to the counterexamples Singh offers to both.

Finally, it's worth briefly considering Singh's own positive view. The lesson Singh draws from his counterexamples is that merely performing an action for the right reasons is insufficient for moral worth. Instead, he holds that an action has moral worth iff it is performed for the

right reasons—viz., in his view, the right-making features of the action—*as such*: that is, "under the guise of their being sufficient to make [one's] action right" (171). Like ours, Singh's view is a sort of compromise between the Right-Making Features View and the Rightness Itself View. However, we find it to be either overintellectualized or mysterious. Non-philosophers rarely form beliefs about whether some consideration is sufficient to make their action right. It could be replied that they don't have to *believe* that the consideration is a right-making feature; they merely have to *tacitly represent* it as such. But this is murky: what does it take to tacitly represent a consideration as a right-making feature? All things equal, it is preferable to avoid this murkiness, and the cluster view allows us to do so.

3.3 Proper basing II: *believing with doxastic justification*

A structurally analogous issue in epistemology arises about what makes for "doxastic justification." Doxastic justification is a status that token beliefs (and other doxastic states), as held by particular people, can have or lack. Whereas a belief's being *propositionally* justified is roughly a matter of its being *a justified attitude to have*, its being doxastically justified is roughly a matter of *one's having it in a justified way*. The two statuses are thus broadly analogous to morally right action and morally worthy action, respectively.

These rough characterizations, of course, need sharpening. Here's a simple and attractive view: a belief that *p* is propositionally justified when one has sufficient evidence for *p*; and it is doxastically justified when it is *based* on the evidence that propositionally justifies it.[22] This is structurally analogous to a simple account of the moral worth, whereby the right action is that supported by the moral reasons, and a morally worthy action is one performed *for* those reasons. But, as in the moral case, epistemologists have objected to the simple account, holding that basing one's belief on the evidence that propositionally justifies it does *not* suffice for its being doxastically justified. For example, Carter & McKenna offer the following putative counterexample:

> **Francophile Cartographer.** Rae (irrationally) believes that French cartographers are the *only* reliable sources of cartographical information. Rae's Francophilia so strongly influences her assessment of the reliability of maps that she distrusts all information written by non-French authors. Rae stumbles upon several pieces of evidence (*E1*, *E2*, and *E3*) for believing cartographical claim X. But, simply because *E1*, *E2*, and *E3* were written by Italian cartographer Giacomo Gastaldi, Rae disregards this evidence and so does not come to believe claim X on the basis of it. Later that day, Rae encounters the *very same* pieces of evidence (*E1*, *E2*, and *E3*) for cartographical claim X, but this time they were

---

[22] Note that this simple view allows for basing to be more complex than just a simple causal relation: what it takes for an attitude or action A to be based on grounds G is a complex and separate question, and the answer to it can be plugged into the simple account.

written by French cartographer Pierre Desceliers. Because—and only because—Desceliers is French, she accepts this evidence, and comes to believe X on the basis of it. (Carter & McKenna 2020: 709-10)[23]

According to Carter & McKenna, Rae is not doxastically justified in believing X even though she bases her belief on the very same evidence that propositionally justifies it—namely, *E1*, *E2*, and *E3*. If this is so, we have a counterexample to the simple account.

But is that right? We think not. First, consider what Rae bases her belief on. The cluster view suggests that the basis of Rae's belief in X is broader than Carter & McKenna acknowledge. Rather, the full cluster is (something like):

21. Desceliers writes that *E1*, *E2* and *E3*.
22. Desceliers is a French cartographer.
23. *All* and *only* French cartographers are reliable sources of cartographical information.

But the cluster (21)-(23) does *not* adequately support believing X, because (23) is false (and, by Carter & McKenna's own stipulation, irrational for Rae to believe).

By contrast, a cluster that *does* support believing X would, in place of (22), need to include a fact about some other property that Desceliers has, such as being a *qualified* or *expert* cartographer—notably, properties that Gastaldi shares—that, unlike his being French, actually establishes his title to reliability. And in place of (23), it would need to include a fact to the effect that cartographers with these properties are reliable sources of cartographical information.

Carter & McKenna's counterexample, then, rests on the claim that *E1*, *E2* and *E3* are the only propositions that constitute the basis of Rae's belief. They take it that the propositions that Desceliers wrote these claims, that he is French, and that French cartographers are the only reliable sources of cartographical information, are *not* part of Rae's basis, but rather only explain *why* Rae forms her belief (solely) on the basis of *E1*, *E2*, and *E3*. But, absent the atomic view, we don't see what justifies this. The upshot: once we accept the cluster view, we can explain why Rae isn't doxastically justified *without* abandoning the simple account of doxastic justification, according to which to be doxastically justified is to believe on the basis of the evidence that propositionally justifies one in believing.

---

[23] See also Way (2017: 210), though Way's examples involve additional complications that we cannot address here due to space.

3.4 The ontology of normative reasons I: "dual aspect" views

Nathan Howard (2021a) argues against the "orthodox" view of normative reasons as facts and in favor of a "dual aspect" view of normative reasons, according to which reasons are fact-goal pairs. In arguing against the orthodox view, Howard presents a series of cases where a single fact seems to correspond with multiple reasons.[24] For example, he asks us to consider:

> **The Egoist and the Altruist.** The pure egoist cares only for herself; she is motivated *only* by prudential reasons. The pure altruist cares only for others; she is motivated *only* by moral reasons. As a result, these two agents are motivated in wholly different ways [and so] never act the same way for the same normative reason. (Howard 2021a: 7)

Howard notes that given the orthodox assumption that normative reasons are facts, the pure egoist and pure altruist are never motivated by the same fact. But, he thinks, that implication is false. For example, he thinks, each can be motivated by to save a child because the child is drowning. On Howard's view, what differs in the egoist and the altruist's motivations is only their *goals*: of respecting and preserving human life, and being esteemed and rewarded for saving the child, respectively. There is no difference in the *facts* that motivate them (7-8).

Call this case *Drowning Child*. Howard's argument can be reconstructed as follows:

P1 If normative reasons are facts, then the pure altruist and the pure egoist are never motivated by the same fact.

P2 In *Drowning Child*, the pure altruist and the pure egoist are motivated by the same fact—namely, the child is drowning.

C Therefore, normative reasons are not facts.

We think defenders of orthodoxy should be unmoved by Howard's argument. For it crucially relies on the assumption that the normatively relevant fact is *atomic*—namely:

24. The child is drowning.

Given the cluster view, however, we should expect additional facts to be involved in providing the relevant normative support for saving the child. For example, rather than being restricted to just (24), the facts motivating Howard's pure altruist plausibly include something like (25) as well:

---

[24] In his (2021a), Howard presses this point with respect to moral and prudential reasons. In his (2021b), he presses much the same point with respect to "right-kind" and "wrong-kind" reasons. For brevity, we focus on the former , but our responses generalize.

25. Saving a child from drowning is a way of respecting and preserving human life.

This contrasts with the cluster of facts motivating Howard's pure egoist, which plausibly includes something like (26) in addition to (24):

26. Saving a child from drowning is a way of being esteemed and rewarded.

The cluster view thus provides principled grounds for rejecting P2 above—the claim that in *Drowning Child*, the pure altruist and the pure egoist are motivated by the same fact. For although there is *overlap* between the facts that motivate the altruist and the egoist—both are motivated, in part, by (24)—there is also divergence—the altruist is motivated by (24) *together with* (25), whereas the egoist is motivated by (24) *together with* (26).

The orthodox view of reasons (*qua* sources of normative support) as facts, when paired with the cluster view, can therefore explain cases like *Drowning Child*. Moreover, the orthodox-friendly explanation is preferable to Howard's for a number of reasons. First, it's more parsimonious: reasons are identified with (clusters of) *facts*, rather than something newfangled like fact-goal pairs. Second, it involves a more thoroughgoing rejection of the atomic view: although Howard rejects the letter of the atomic view, he retains its spirit by identifying reasons with pairs of *atomic facts* and goals. Many of the reasons canvassed at the outset for rejecting the atomic view are thus reasons to reject Howard's view. And third, it's more general. Howard' often talks about 'normative reasons' without qualification, but his discussion focuses almost entirely on reasons for actions, rather than reasons for attitudes. And the dual-aspect view is quite implausible with respect to reasons for attitudes (especially for belief).[25] Given the natural (but not incontrovertible) assumption that reasons for actions belong to the same ontological category as reasons for attitudes, this poses a problem for Howard's view.

3.5 *The ontology of normative reasons II: contrastivism*[26]

A final example of presupposition of the atomic view leading to an overcomplicated ontology of reasons is provided by one of Justin Snedegar's (2017) core arguments for his "contrastivist" view of reasons. According to contrastivism, nothing is a reason for Φ-ing *simpliciter*; rather, facts are reasons for Φ-ing *rather than* performing some alternative act Ψ. Snedegar offers the following case, inspired by Ross (2006):

---

[25] This is partly because many doubt that goals are involved in reasons for belief at all (Kelly 2003). Moreover, even if they are, they may not play the same role they play in paradigmatic desire-based reasons for action.

[26] This subsection overlaps Worsnip (2019).

Dinner Invitation. "Suppose I have three dinner invitations, and that I must select exactly one: Invitation A is for Armenian with Ara. Invitation B is for burgers with Burt. Invitation C is for Chinese with Charlie. I love spending time with Ara, really like spending time with Burt, but can barely tolerate spending time with Charlie. On the other hand, I love Chinese, really like burgers, but can barely stomach Armenian." (Snedegar 2017: 52)

Snedegar's challenge for the non-contrastivist is as follows. Plausibly, I ought to accept invitation B. But if the reason-relation is non-contrastive, it's unclear what consideration constitutes a reason for accepting invitation B. The fact that invitation B is for a meal with Burt doesn't seem to favor accepting invitation B *simpliciter*, because it doesn't favor accepting invitation B over accepting invitation A. And the fact that invitation B is for burgers doesn't seem to favor accepting invitation B *simpliciter*, because it doesn't favor accepting invitation B over accepting invitation C. If that's so, there seems to be no reason to accept invitation B. But that's hard to square with the verdict that I *ought* to accept invitation B.

Conversely, the contrastivist can handle the case nicely. There's a strong reason for me to accept invitation B *rather than* invitation A, given that I greatly prefer burgers to Armenian. The reason is that invitation B is for burgers (perhaps conjoined with the fact that invitation A is for Armenian). Similarly, there's a strong reason for me to accept invitation B *rather than* invitation C, given that I greatly prefer Burt's company to Charlie's. The reason is that invitation B is with Burt (perhaps together with the fact that invitation C is with Charlie). Moreover, Snedegar endorses the principle that one ought to Φ, out of some set of relevant alternatives, iff for every relevant alternative Ψ one has most reason to Φ rather than Ψ. Thus, I ought to accept invitation B.

However, Snedegar implicitly assumes that the non-contrastivist has to find a *single* atomic fact (or perhaps a conjunction of just two facts, since he seems to allow his own view to do this) that favors invitation B *simpliciter*. Once we switch to the cluster view, the problem disappears. For we can then say that something like the following cluster of facts provides support for accepting invitation B, simpliciter:

27. Invitation A is for Armenian with Ara.
28. Invitation B is for burgers with Burt.
29. Invitation C is for Chinese with Charlie.
30. I greatly prefer burgers to Armenian.
31. I only slightly prefer Chinese to burgers.
32. I greatly prefer Burt's company to Charlie's.
33. I only slightly prefer Ara's company to Burt's.

It's clear that this full cluster of facts provides a single, unified rationale *for accepting invitation B.* There's simply no need to subdivide the cluster and say that one part of it constitutes a reason for accepting B rather than A and another part constitutes a reason for accepting B rather than C.

Of course, contrastive claims about reasons are intelligible too, and the cluster view can make sense of them. It's plausible, for example, that the cluster of facts consisting of (27), (28), and (30) provide significant support for accepting B, while failing to provide such support for accepting A. This cluster thus constitutes a reason—i.e. a source of normative support—for accepting B rather than A. To the extent that our ordinary talk about reasons is often elliptical for such claims, it may (often) be implicitly contrastive. But that falls short of establishing that the reason (or support) relation is necessarily, fundamentally, metaphysically contrastive—that it *only* makes sense to talk of facts as supporting responses *rather than* other responses, and that talk of facts supporting a response simpliciter is necessarily incomplete. On the contrary, given the cluster view, we can explain contrastive claims about reasons in fundamentally non-contrastive terms.

4. Reasons to reject the cluster view?

So why don't philosophers accept the cluster view? One reason is that ordinary thought and talk seems to presuppose the atomic, rather than the cluster, view. But, as noted at the outset, this is far from probative—ordinary thought and talk about reasons is at best an imperfect guide to the theoretical roles that reasons are standardly supposed to play. The only other source of resistance we can locate is Schroeder's (2007) critical discussion of what he calls the "no background conditions" (NBC) view. NBC denies that there there are any background conditions, i.e. facts that help explain why other facts provide normative support. Any would-be background condition is actually part of the cluster of facts playing the front-and-center role of providing normative support.

NBC is an extreme version of cluster view. The cluster view says that sources of normative support (typically) consist of clusters of facts rather than individual atomic facts, and posits that elements of these clusters are often mistaken for background conditions. NBC adds, more strongly, they there are no background conditions *at all.* Thus, while NBC has some plausibility, the cluster view isn't committed to it, and we remain neutral on whether to accept it, viewing this as an intramural debate between cluster theorists. Nonetheless, it's worth looking at Schroeder's arguments against NBC, to see whether they apply to the cluster view more generally.

One concern Schroeder voices is that NBC "denies a distinction with respect to reasons that is a perfectly good distinction in many other domains"—namely, the distinction "between the facts necessary in order to explain why something is an F, and what the thing is, which is the

F" (24). For reasons we don't have space to get into in detail here, we're not convinced this is a good objection to NBC.[27] But in any case, the objection clearly doesn't apply to the (weaker) cluster view, since the latter is compatible with allowing background conditions and thus with Schroeder's distinction.

A second worry raised by Schroeder concerns an alleged tension between NBC and what he calls the "Deliberative Constraint" (DC):

> **DC**:   One's reasons are the kinds of thing that one ought to pay attention to in deliberation. (33)

To illustrate the tension, he considers the case of Ronnie, who likes to dance, and the reason he has to go to a dance party:

> "The fact that Ronnie desires to dance has to be mentioned in the complete explanation of why it is the case that there is a reason for Ronnie to go to the party. Ronnie is acting in a manner that is objectionably self-regarding if [in deliberation] the thought that there will be dancing at the party doesn't move him to go there, until he remembers that he desires to dance. If Ronnie genuinely desires to dance, then *all it should take* for him to be moved to go to the party is the thought that there will be dancing there." (27)

We can reconstruct the alleged tension as follows:

1. The fact that Ronnie likes to dance is part of the reason he has to go to the party. (From the description of the case together with NBC)
2. In deliberating about what to do, Ronnie ought to pay attention to the fact that he likes to dance. (From (1) together with DC)
3. In deliberating about what to do, it's not the case that Ronnie ought to pay attention to the fact that he likes to dance.

Schroeder takes (3) to follow from a more general principle:

> **OSR**: Facts about desires are (typically) *not* the kinds of thing that one ought to pay attention to in deliberation. (To do so would be objectionably self-regarding.)

We thus seem to be faced with a choice:

---

[27] Very briefly: if reasons to Φ are the things that *explain* why you ought to Φ or why there's normative support for Φ-ing, then since any would-be background condition is plausibly still part of the explanation of why you ought to Φ, they are also part of the reason or cluster. Thus, at least given a certain view of reasons, the proponent of NBC has a principled explanation of why the distinction doesn't apply here.

A. Reject NBC, and hence (1).
B. Reject DC, and hence (2).
C. Reject OSR, and hence (3).

Although the "standard move", according to Schroeder, is to reject DC, he accepts it (28). He does so because he wants to defend the broadly Humean thesis that (all of) our reasons are desire-dependent while also getting as much of the "intuitive data" correct as possible, which he takes to include OSR together with DC (28). That leaves NBC as the odd man out.

Interestingly, unlike Schroeder's first argument, this argument against NBC generalizes to any version of the cluster view that holds that facts about desires sometimes belong in the clusters that provide normative support for actions (rather than being relegated to the status of background conditions). For example, it applies to any view that holds that the fact that Ronnie desires to dance—together, presumably, with the fact that there will be dancing at the party—is part of the reason [or cluster of facts that generate support] for him to go to the party. That amounts to (1)—the cluster theorist needn't go as far as denying that there are *any* background conditions to be subject to Schroeder's argument.

Our response to Schroeder's argument is to deny *either* DC *or* OSR (or both), depending on how the two are understood. As stated above, DC is somewhat imprecise. On one precisification of DC, it says that if something is a reason (or a part thereof), then one ought to take it into account in deliberation: that is, it would be *impermissible not to take it into account in deliberation*. This claim is too strong. As Schroeder himself affirms (95-96), reasons are ubiquitous: there are (very weak) reasons to do almost anything. It seems far too demanding to say that agents ought to pay attention to *all* these reasons in deliberation. Moreover, even if the claim is restricted to *weighty* reasons, we think that advocates of the cluster view (including advocates of NBC) should reject the claim that correct deliberation requires agents to pay attention to *every* fact in the cluster that provides support for acting. Some can operate in the background of their psychology, if at all.

However, some of what Schroeder says suggests a weaker version of DC that merely says that if something is a reason, then it's *permissible* to pay attention to it in one's deliberation. We find this weaker reading of DC more plausible. But to preserve the argument against NBC, we must now correspondingly *strengthen* OSR, so that it says not just that facts about desires are (often) such that **it's not the case that we ought** to pay attention to them in deliberation, but rather that facts about desires are (often) such that **we ought not** to pay attention to them in deliberation.

And now, we find OSR too strong. While we agree that agents aren't *required* to think about their desires while deliberating, we think it's often perfectly permissible for them to do so. In the case of Ronnie, for example, we think it's perfectly permissible for him to think, when

considering whether to go to the party, "well, I really want to dance!". Similarly, in our umbrella case discussed at the outset, we think it's fine to think about how you really don't want to get wet, in considering whether to take an umbrella with you. While there's plausibly something objectionably self-regarding about *only* thinking about one's desires, this is compatible with it being permissible for desires to often be *part* of what one thinks about in deliberation.

Thus, Schroeder has not given us a good reason to exclude facts about the agent's desires from the cluster of facts that provides normative support for responses. In Ronnie's case, we think, the right thing to say is that both the fact that there will be dancing and the fact that Ronnie desires to dance are part of a cluster of facts that jointly provide normative support for his going to the party.[28]

Thus, we are unmoved by both of Schroeder's objections. We conclude that if, as we have argued, the cluster view can solve, dissolve, or at least simplify a variety of problems in the theory of normativity, there is a strong case for accepting it.

References

Arpaly, N. 2002. "Moral Worth." *Journal of Philosophy*, 99(5), 223-245.

Behrends, J. 2016. "Normative Source and Extensional Adequacy." *Journal of Ethics and Social Philosophy*, 10(3).

Broome, J. 2007. "Does Rationality Consist in Responding Correctly to Reasons?" *Journal of Moral Philosophy*, 4(3), 349-374.

Carter, J.A. & McKenna, R. 2020. "Skepticism Motivated: On the Skeptical Import of Motivated Reasoning." *Canadian Journal of Philosophy*, 50(6), 702-718.

Chang, R. 2013. "Grounding Practical Normativity: Going Hybrid." *Philosophical Studies* 164(1), 163-187.

---

[28] Or, to qualify this slightly, we think this is the right thing to say conditional on accepting a view according to which Ronnie's desire plays *some* role in explaining why there's normative support for him to go to the party. Some would deny this. The cluster view itself is neutral on the dispute as to whether desires are always, sometimes, or never part of the explanation of normative support. But, for what it's worth, whereas Schroeder's [Humean] view is 'always' and the monistic hedonistic utilitarian's view is 'never', our own view is 'sometimes' (cf. the "hybrid" view developed by Chang (2013) and Behrends (2016)).

Dancy, J. 2000. "Should We Pass the Buck?" *Royal Institute of Philosophy Supplement*, 47, 159-73.

Enoch, D. 2019. "How Principles Ground," *Oxford Studies in Metaethics*, 14, 1-22.

Fogal, D. 2016. "Reasons, Reason, and Context." In E. Lord & B. Maguire (eds.), *Weighing Reasons*. Oxford University Press.

Fogal, D. 2018. "Deflationary Pluralism about Motivating Reasons." In V. Mitova (ed.), *The Factive Turn in Epistemology*. Cambridge University Press.

Fogal, D. & Risberg, O. 2020. "The Metaphysics of Moral Explanations." *Oxford Studies in Metaethics*, 15, 170-194.

Fogal, D. & Worsnip, A. 2021. "Which Reasons? Which Rationality?" *Ergo*, 8(11).

Howard, N. 2021a. "Primary Reasons as Normative Reasons," *Journal of Philosophy*, 118(2), 97-111.

Howard, N. 2021b. "Ambidextrous Reasons (or Why Reasons First's Reasons Aren't Facts)." *Philosophers' Imprint*, 21(30).

Howard, N. forthcoming. "The Goals of Moral Worth." *Oxford Studies in Metaethics*. Page references are to the draft at https://philpapers.org/archive/HOWTGO-7.pdf.

Johnson King, Z. 2019. "We Can Have Our Buck and Pass it Too." *Oxford Studies in Metaethics*, 14, 167-188.

Johnson King, Z. 2020. "Accidentally Doing the Right Thing." *Philosophy and Phenomenological Research*, 100(1), 186-206.

Kelly, T. 2003. "Epistemic Rationality as Instrumental Rationality: A Critique," *Philosophy and Phenomenological Research*, 66(3), 612-640.

Kiesewetter, B. 2017. *The Normativity of Rationality*. Oxford University Press.

Lord, E. 2017. "On the Intellectual Conditions for Responsibility: Acting for the Right Reasons, Conceptualization, and Credit." *Philosophy & Phenomenological Research*, 95(2), 436-464.

Lord, E. 2018. *The Importance of Being Rational*. Oxford University Press.

Markovits, J. 2010. "Acting for the Right Reasons." *Philosophical Review*, 119(2), 201-42.

Ross, J. 2006. *Acceptance and Practical Reason*. PhD Dissertation, Rutgers University.

Schroeder, M. 2007. *Slaves of the Passions*. Oxford University Press.

Singh, K. 2020. "Moral Worth, Credit, and Non-Accidentality." *Oxford Studies in Normative Ethics*, 10: 156-181.

Sliwa, P. 2016. "Moral Worth and Moral Knowledge." *Philosophy and Phenomenological Research*, 93(2), 393-418.

Snedegar, J. 2017. *Contrastive Reasons*. Oxford University Press.

Sylvan, K. 2015. "What Apparent Reasons Appear to Be." *Philosophical Studies*, 172(3), 587-606.

Titelbaum, M. 2019. "Reason without Reasons For." *Oxford Studies in Metaethics*, 14, 189-215.

Way, J. 2017. "Creditworthiness and Matching Principles." *Oxford Studies in Normative Ethics*, 7: 207-228.

Wedgwood, R. 2015. "The Pitfalls of 'Reasons'." *Philosophical Issues*, 25, 123-143.

Worsnip, A. 2019. Review of Justin Snedegar's *Contrastive Reasons*. *Philosophical Review*, 128(3), 367-371.

Worsnip, A. 2021. *Fitting Things Together*. Oxford University Press.

Zimmerman, M. 2007. "The Good and the Right." *Utilitas*, 19, 326-353.