

Factivism Defended: A Reply to Howard

J.J. Cunningham

University of Nottingham

In his recent paper *Primary Reasons as Normative Reasons* Nathan Robert Howard argues that normative reasons are not purely facts (*that it's raining; that I promised...*), but also include goals (*to avoid getting wet; to keep my promise...*). Howard supplies three arguments for this thesis. This reply aims to show that those arguments are unsound. It also aims to show that Howard's view is itself subject to a dilemma. §1 introduces the debate; §2 responds to Howard's arguments; §3 presents the dilemma.

First, a preliminary. Howard's focus is on *objective* normative reasons: Reasons *there are* for one to ϕ , as opposed to *subjective* normative reasons: Reasons one is suitably aware of, so that they can bear on what one does. Howard also restricts his attention to normative reasons *for action*. There's more to be said about these restrictions. Due to space, I follow Howard in making them.

1 Factivism vs. the Dual-Aspect View

Normative reasons are items which favour (or disfavour) one acting in some way. That the cat is hungry is a normative reason to feed her; the exposed wire being live is a normative reason not to touch it; their commitment to equality is a normative reason to vote Green... and so on. Normative reasons are typically contrasted with *motivating reasons*: Considerations which one treats as normative reasons for one to ϕ and which thereby motivate one to ϕ .¹

The orthodox view is that normative reasons are *facts*, typically about the world and not our minds. It's *the fact* that the cat is hungry, as opposed to the *event* of her hunger or my *believing* that she's hungry, which favours me feeding her. Likewise, it's *the fact* that the wire is live which is the normative reason not to touch it – not, for example, the *property* of liveness.

¹This may strike one as a controversial thing to say about motivating reasons. But since this use of 'motivating reasons' is also Howard's (2021: §2) usage, I'm dialectically permitted to operate with it.

Amongst proponents of factivism there is a debate about what facts are and hence what reasons are. One view is that they are true Fregean propositions.² Another is that they are states of affairs which function to make propositions true.³ What these versions of factivism have in common is that they identify normative reasons with *propositionally-structured* things: *that* the cat is hungry, *that* the wire is live, *that* the policy promotes equality...and so on. What differentiates them is *which* propositionally-structured entities count. We'll remain neutral on this intramural debate. What will matter to us is whether we should be in the business of identifying normative reasons exclusively with propositionally-structured entities. 'Factivism' is used to denote any theory according to which we should.

No plausible version of factivism is going to identify normative reasons with *single facts*. Rather, factivists should identify normative reasons with *clusters of facts*. To show this, I invoke a double-counting phenomenon to which attention has already been drawn by Raz (2011) and Fogal (2016). Consider:

- (1) The fact that the cat is hungry is a reason to feed her.
- (1*) The fact that the cat is hungry is a reason to feed her and the fact that it's her dinnertime is a reason to feed her. So there are two reasons to feed the cat.

Intuitively, (1) is fine. But (1*) is unnatural: It seems like double counting to say that there are *two* reasons to feed the cat, each constituted by one of the facts specified. Instead, the more natural thing to say is that the two facts are compose a *single* reason. Other examples abound. The fact that the wire is live is a reason not to touch it; the fact that the wire will cause one pain when touched is also a reason not to touch it; but we don't have *two* reasons here – we only have a *single* reason of which the two facts are components. Similarly, the fact that Greens will enact equality-promoting policies is a reason to vote for them; so is the fact that were they elected, the wealth-gap would narrow; but we don't have *two* reasons here – we just have a *single* reason of which the two facts are components. Generalising, the factivist should be understood as saying that normative reasons are complex clusters of facts, not single facts.⁴

The claim that the factivist should identify normative reasons with clusters is controversial. Most factivists in the literature simply go ahead and identify normative reasons with single facts. And, although I've provided an argument for the cluster view, I don't

²Alvarez (2010); Lord (2018).

³Dancy (2000); Mantel (2018).

⁴Raz (2011: Ch.1) construes factivism in this way, and argues for it using the double-counting phenomenon above. For further dialectical uses of that phenomenon see Fogal (2016), Fogal and Worsnip (2021), Maguire and Snedegar (2021), and Fogal and Risberg (Forthcoming).

have space for a full defence of it here.⁵ To the extent that the arguments below are premised on the ascription to the factivist of the cluster view, then, those arguments should be taken as conditional: They are sound *assuming, as seems plausible* that the factivist should endorse the cluster view.

Howard denies factivism in favour of the *dual-aspect view*. The dual-aspect view identifies normative reasons with *fact/goal pairs*. A fact/goal pair is a (cluster of) fact(s), some goal, and a certain relation between them: The goal is such that, given the fact(s), the action favoured promotes the goal. On this view, hat speaks in favour of me feeding the cat is not just the fact that the cat is hungry but that fact *plus* the goal to keep her healthy. What speaks in favour of me not touching the wire is not just the fact that it'd cause me pain if I do so, but that fact *plus* the goal to avoid pain. . . and so on.

Which of these views in the ontology of normative reasons should we prefer? The remaining sections defend factivism.

2 Howard's Arguments Rejected

This section surveys and criticises Howard's three arguments for the dual-aspect view.

2.1 The Argument from the Good Basis View

Howard's first argument takes as its starting point the good basis view of normative reasons.⁶ According to this view, for something to be a normative reason for S to ϕ is for it to be such that it'd be a *good basis* for S to ϕ . But good bases for ϕ -ing are a subset of S's motivating reasons for ϕ -ing: Motivating reasons on which it is good for S to ϕ . So, whatever normative reasons are, they must be identifiable with good motivating reasons.

It follows that if we're to individuate motivating reasons in terms of goals, then normative reasons are to be individuated in terms of goals, too, just as the dual-aspect view says. Howard aims to prove that motivating reasons *are* individuated partly by goals. He asks us to consider:

Cubs Case. The Cubs fan and the Cubs hater are both heading to Wrigley to see the Cubs play. Both go on the basis of the belief that the Cubs are

⁵For example, one challenge would be to explain why the the double-counting phenomenon can't be accommodated by Maguire & Snedegar's (2021) distinction between load-bearing and derivative reasons. The most promising reply would involve demonstrating that even load-bearing reasons are subject to the double-counting phenomenon.

⁶Setiya (2007), Way (2017), and Gregory (2016).

playing. But the goal that each has is different: the fan goes to see the Cubs win; the hater goes to see them lose.

About this case, Howard says:

...if motivating reasons are simply facts, then we're forced to concede, counterintuitively, that the hater and the fan go to Wrigley for the same reason because they act on the basis of the same fact. But that seems clearly wrong – after all, one goes to Wrigley to see the Cubs win and the other goes to see the Cubs lose. (Howard, 2021: 100)

In other words: Cubs-fan and Cubs-hater have different *motivating reasons*. But since they act on the basis of just the same belief it can't be that we're to individuate their motivating reasons just by what they believe. Since they act for different *goals* the most plausible thing to say is that their motivating reasons are partly individuated by different goals. But then Howard has his result.

It doesn't help the factivist here to note that, for them, just as normative reasons are *clusters* of facts and not single facts, motivating reasons are also to be identified with *clusters* of beliefs as opposed to single beliefs. After all, Howard is simply free to write it into his case that it's the very same *cluster* of beliefs which moves each agent to act.⁷

Nevertheless, I think the factivist has a plausible reply to Howard's first argument. Let us start with a distinction. Cubs-fan and Cubs-hater are motivated to go to the game by the very same motivating reason: A cluster of beliefs which includes *that the Cubs will probably lose; that I'll find it entertaining; that I'm vested in how the Cubs do tonight; that I have a season ticket...* and so on. But each agent treats each cluster as constituting a *different normative reason* to go to the game. After all, from Cubs-fan's point of view, the reason they have to go to the game is also a reason to celebrate if the Cubs win. But, from Cubs-hater's point of view, the reason they have to go to the game is a reason to celebrate if the Cubs *lose*. So although each agent is moved to action by the very same set of beliefs, each agent treats that single self-same cluster as constituting a different normative reason for them to go to the game.

What we have here is a distinction which holds of an agent who ϕ s motivated by a reason. On the one hand, there are *the things* treated as a normative reason by them. On the other hand, there is the *way* those things are treated. Two agents might treat the very same set of considerations as constituting a normative reason for them, but treat them in different ways, so that each agent ends up treating that very same set of considerations as a *different* normative reason. For example, one of them might treat those considerations

⁷Thanks to my reviewers for pointing this out.

as supporting a particular course of action whereas the other treats them as supporting an alternative, inconsistent course of action, so that the very same set of considerations is treated as constituting a *different* normative reason by each agent.

The distinction I've just drawn is independently plausible. It is simply an instance of the more general and innocuous distinction between what a thing is and the way that thing is treated. The co-carers of a particular cat might treat that cat quite differently, for example: One treats it like a human infant, the other treats it more austere. The religious ascetic treats physical pain quite differently from the rest of us: They treat it as a deserving punishment from the almighty, whereas the rest of us might treat it as either a nuisance or something to see the doctor about. How a thing is treated by a particular agent is constituted by a complex cluster of dispositions and abilities the agent has towards the thing in question, which will of course typically be grounded in yet further of their personal-level psychological states. The distinction between what an agent treats as a normative reason and how they treat it is merely another instance of this readily acknowledgeable phenomenon.

With this distinction in tow, let us now return to Howard's first argument against factivism. Howard's crucial premise is that Cubs-fan and Cubs-hater act for *different motivating reasons*. But in light of the distinction just defended, Howard's premise is ambiguous. On the one hand, Howard could be claiming that the *consideration which is treated as a normative reason* to go to the game is different across the two agents. On the other hand, Howard could be claiming that each consideration is treated in a different way by each agent, so that each agent treats it as a *different normative reason*.

On the first reading, the thought that Cub-fan and Cubs-hater act for different motivating reasons boils down to the claim that they treat a different set of considerations as their normative reason for action. But as long as it's assumed that it's the very same cluster of beliefs which motivates each agent, *neither* the factivist *nor* Howard will accept the claim on this reading. And, as we've seen, Howard will have to insist that it *is* the very same cluster of beliefs doing the motivational work, or else he falls foul of the simple factivist-friendly response to his argument according to which motivating reasons are really *clusters* of beliefs and not single beliefs.

On the second reading of Howard's crucial premise, the thought that Cub-fan and Cubs-hater act for different motivating reasons boils down to the claim that they each treat the self-same cluster of considerations in a different way – for instance, as favouring different kinds of responses – so that they each treat a different normative reason as existing. However, the factivist can grant that it's *true* that Cubs-fan and Cubs-hater act for different motivating reasons on *this* reading. After all, factivism only says that the things which are treated as normative reasons by each agent are propositionally-

structured; not that, when the same set of propositionally-structured entities motivate a pair of distinct agents, it follows that each agent treats it as the very same normative reason. Two agents, as we've seen, might have a different set of reasoning and response dispositions towards a single set considerations, such that they each treat that set as constituting a different normative reason. The factivist can help themselves to this claim, suggest that Cubs-fan and Cubs-hater are doing just that, and that *this* is what explains why they act for *different* motivating reasons.

2.2 The Argument from the Moral/Prudence Contrast

Let's now turn to the second of Howard's arguments. The argument takes as its starting point the contrast between egoism and altruism:

Consider two agents: the pure egoist and the pure 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. (Howard, 2021: 103)

We stipulate that the egoist is only ever motivated by prudential normative reasons and the altruist by moral normative reasons. Now we arrive at a problem:

If we add the orthodox supposition that normative reasons are facts, . . . these two agents never perform an action on the basis of the same fact. But that's clearly false. For example, each can be moved to save a child because the child is drowning. This is possible because the altruist's and the egoist's goals differ. The altruist is moved, we may suppose, by her goal of respecting and preserving human life. By contrast, the egoist is moved, we may suppose, by her goal of being esteemed and rewarded for saving the child. Nevertheless, it is possible for the same fact, and only that fact, to move each to perform the same act. (*Ibid.*: 103–4)

The egoist and altruist both go to save the drowning child. The only fact that motivates the egoist is the fact that the child is drowning. Likewise, the only fact that motivates the altruist is the fact that the child is drowning. If factivism is true, then the normative reason in response to which each agent acts is exactly the same. But *ex hypothesi* they act for different normative reasons. So factivism can't be right. Moreover, the solution to the problem is obvious: Each acts in the pursuit of a different goal. So we can account for the difference in normative reasons by saying that each reason is partly individuated by a different goal. And this is just the dual-aspect view.

The mistake made by this argument is that the factivist can perfectly well allow that there are really *two* normative reasons here. The key to noticing this is to recall that for the factivist, normative reasons are best thought of as clusters of facts, not single facts. Given that, all the factivist needs is to find a fact which is a member of the altruist's moral reason cluster which isn't a member of the egoist's prudential reason cluster or *vice-versa*. Such a task is not difficult to achieve. The fact that it will make one look good if one saves the child is a plausible member of the prudential reason cluster but is not a member of the moral reason cluster, for example.

Howard might respond to this as follows: Fact-wise, the egoist and the altruist have exactly the same *motivating reason* – the very same cluster of considerations move them to action. But now suppose that factivism is true and suppose that the *normative reason* in response to which the egoist acts differs from the normative reason in response to which the altruist acts, as per my current objection to Howard. In that case, we couldn't *identify* each agent's normative reason with their motivating reason – after all, their normative reasons differ but their motivating reasons are the same. But when an agent acts in response to a normative reason, their motivating reason *is* identical to that normative reason. So my objection must be wrong.

This response trades on an identity thesis which is too strong. If we accept the plausible cluster views of normative and motivating reasons, there are two identity theses worth contrasting:

Identity-Strong. Necessarily, if an agent acts in response to a normative reason, then:

The fact F is a member of the set which constitutes the agent's normative reason if, and only if, it is identical to one of the considerations which forms a part of the agent's motivating reason.

Identity-Weak. Necessarily, if an agent acts in response to a normative reason, then:

Each member of the set of the agent's motivating reason is identical to a member of the set of facts which constitutes the normative reason.

Identity-Strong is manifestly implausible. Normative reasons are complex clusters of facts. To act in response to a normative reason obviously doesn't require that be moved by *all* the facts which constitute that reason. The fact that London has some superb art galleries is a reason to visit. The fact that London has approximately 1500 art exhibitions is also a reason to visit. We don't have *two* reasons to visit, here. Rather, we have two components of a single reason. But someone need only know one of those facts to count as visiting in response to that reason. Indeed, the second fact, which specifies the approximate number of exhibitions, is probably known by only very few tourists who visit the city for its art.

The weak thesis doesn't run into this problem. That's because it only requires each component of the agent's motivating reason to be identical to a corresponding member of the relevant normative reason and *not vice-versa*. Perfect overlap between motivating a normative reason isn't required by Identity-weak: Only partial overlap.

Of course, Identity-Weak raises questions. For example: According to Identity-Weak, what degree of overlap is required between the agent's motivating reason and the corresponding normative reason, if the agent is to count as responding to the latter *via* being moved by the former? There is a viable response to that challenge to Identity-Weak available, which invokes a know-how-based theory of reasons-responsiveness I've defended at length elsewhere.⁸ But for my current dialectical purposes all I need are the points that *if* there is a plausible identity thesis linking motivating and normative reasons, *then* it is Identity-Weak and *not* Identity-Strong. That point should be acknowledged even there are theoretical challenges which remain open to Identity-Weak. After all, Identity-Strong is subject to straight-up counterexample, as I've shown above.

I conclude that if there is any identity thesis in this area which survives the point that both normative and motivating reasons are clusters, it is Identity-Weak and not Identity-Strong. However, the present response to my objection requires the truth of Identity-Strong. That's because the present response requires that the altruist and the egoist *cannot* have the same motivating reasons whilst responding to different normative reasons. But to say that the two agents *can* act motivated by the same motivating reasons even though their normative reasons differ would be to accept the following possibility: The two agents' normative reasons overlap but don't entirely coincide and the two agents act motivated only by members of this overlapping set. The proponent of the response will therefore have to commit themselves to the impossibility of this latter state of affairs. The only way to argue that overlapping but distinct normative reasons can't motivate two agents who have identical motivating reasons would be to require that a normative reason can motivate only if *all* its components motivate the agent. But this is just to insist on the implausible Identity-Strong.

2.3 The Argument from Weighting Explanations

Howard's third argument takes the form of a dilemma for the factivist. First, we'll follow Howard in making the general assumption that if an overall ought or permission of a certain kind (moral, prudential, epistemic, aesthetic...) holds of a certain agent, this is always to be thought of as grounded in facts about the weight of reasons of the same kind which are present. For example, if an agent prudentially ought to ϕ , then this will be

⁸[REDACTED]

because there are prudential reasons for them to ϕ which *jointly outweigh* the competing prudential reasons taken together. Likewise, if an agent morally *may* ϕ , then this is because there are moral reasons on the scene which are *jointly as strong as* the moral reasons against ϕ -ing taken together.

Next, Howard asks us to consider:

Rich Parents. Suppose that two children, qualitatively identical in all morally relevant respects, are drowning and only one can be saved. Consequently, you may save either, but you must save one. Moreover, one child is from a rich family; the other child is not. The person who saves the first child will be richly rewarded. Not so for the second. (Howard, 2021: 106)

We are to consider what should be done from the moral point of view here. The two children are the same in all morally salient respects but not in all prudentially salient respects – one child has rich parents, the other doesn't. We'll grant Howard's plausible verdict that, morally speaking, one *may* save either and hence one *may* save the rich child.

Howard makes an assumption about what the factivist is committed to saying about Rich Parents:

Supposing that reasons are facts, the fact that the rich child is drowning is a moral reason to save them. It's also a prudential reason to save them. (*Ibid.*: 106–7)

The assumption is that, for the factivist the fact that the rich child is drowning is a *single* normative reason of *two types*: A moral reason to save that child *and* a prudential reason to save them.

Now a dilemma supposedly arises for the factivist:

This creates a dilemma. Either [Option A:] we include... [the fact that the rich child is drowning] when determining the balance of moral reasons, or [Option B:] we do not. If we do include it, we'll have allowed a prudential reason to distort the balance of moral reasons... However, if we exclude the fact, then we'll have ignored a morally relevant reason, thereby distorting the balance of moral reasons... (*Ibid.*: 107)

Option A is supposed to be problematic because if we include the fact that the rich child is drowning in our determination about what morally to do, we'll have included something which doesn't just add *moral* weight to the situation but also *prudential*

weight – the factivist, remember, is supposedly committed to the claim that the fact at issue is *both* a normative reason *and* a prudential reason. But then the combined weight of the moral reasons to save the rich child is *greater* than the combined weight of the moral reasons to save the other child, thereby delivering us the incorrect verdict that we morally ought to save the former.

But Option B is also problematic: If we don't include the fact that the rich child is drowning in our determination about what morally to do then we'll have ended up leaving out a highly relevant moral consideration. We'll thereby end up having to conclude, again incorrectly, that there isn't sufficient moral reason to save the rich child.

So if factivism is true, we get the wrong overall verdicts about Rich Parents whether or not the fact that the rich child is drowning is part of what determines the overall moral verdict. The dual-aspect view doesn't run into this dilemma because it precisely denies that the fact that the rich child is drowning is *itself* a single normative reason of both a moral and prudential type.

I don't think this argument is successful. There are two defensible responses to it. The first is by now familiar: The factivist is *not* committed to saying that *single* facts are to be identified with normative reasons, but *clusters* of facts. Thus, the assumption that the factivist treats the fact that the rich child is drowning as itself a single normative reason of two types, moral and prudential, is mistaken. What they really say is that there is a cluster of facts with which we identify the prudential reason and a second cluster with which we identify the relevant moral reason. As long as these two clusters are non-identical, the dilemma is avoided. And these clusters aren't identical: The prudential cluster, but not the moral cluster, will include the fact that the child has rich parents.

Howard comes close to recognising the availability of this response. He distinguishes between:

- (1) The child is drowning
- (2) The child is drowning and I'll be rewarded if I save them

And he says this of the distinction:

...it's natural to think that while (1) gives a moral reason to save them, (2) doesn't; it's a purely prudential reason... However, observations about (2) do not show what must be shown to rebut the argument. That argument concerns (1), and its problematic dual significance, not (2). So unless it can be shown through additional argument that facts about (2) imply that (1) is not problematic, discussion of (2) rather than (1) is merely a distraction. (Howard, 2021: 108)

To the extent that this can be interpreted as a rejoinder to the cluster-based response I've just defended it seems to misunderstand that response. The response is this: There's independent reason for the factivist to deny the claim that facts like (1) are normative reasons at all. Instead, they should say that it's clusters of facts – like (2), albeit much more complex – which are normative reasons. Howard therefore cannot help himself to the claim that (1) is of dual significance.

There is also a second, theoretically richer, response to Howard's dilemma available. Suppose we grant that if factivism is true, then the single fact about the rich child is both a reason of prudence and a moral reason to save them. According to the first horn, the factivist is therefore precluded from including that fact in the determination of what one is morally to do. Howard reasons thus:

- (P1) If the factivist allows the fact that the rich child is drowning to be part of what determines what one morally ought to do, then they'll have to say that it contributes its weight *as a whole* to determining what one ought to do.
- (P2) If that fact contributes its whole weight to the grounding of what one morally ought to do, then overall one morally ought to save the rich child.
- (P3) It's not the case that overall one morally ought to save the rich child.

(C) It's not the case that the factivist can allow the fact that the rich child is drowning to be a part of what determines what one morally ought to do.

But the factivist can plausibly deny (P1). To do so, the factivist would have to say: We can carve the weight of a reason of dual significance in two – there's its moral weight, and there's its prudential weight. We then say that when it comes to the grounding of overall *moral* verdicts only the moral weight of the reason is to be taken into account. And we say that when it comes to the grounding of overall *prudential* verdicts, only the prudential weight of the reason is to be taken into account. Generalising: What grounds an overall verdict of a certain kind (moral, prudential, epistemic...) is all and only reasons of that very same kind *qua* the weight that they have *of that kind*. On this picture, (P1) is false: We are simply to deny that weighting explanations work in the way (P1) codifies.

On Howard's alternative picture, we're not to think that reasons can have *moral* weight and, in addition, *prudential* weight; instead, they have a degree of weight *simpliciter*. Reasons of *dual* significance, if there are such things, will then be understood

as items which contribute their weight *simpliciter* to the grounding of *both* prudential *and* moral overall verdicts. This view implies (P1).

The availability of the first picture of the weight of reasons – on which (P1) is false – might be occluded by the metaphor of *weight* itself. The weight of an object, after all, doesn't come in a series of distinct modes or flavours. Instead, there's just the weight of the object *simpliciter*. If we use the metaphor of weight to think about the feature of normative reasons we're interested in, then, we'd naturally end up led to the second, (P1)-friendly, position. But as Fogal and Risberg (Forthcoming: §4.3) point out, *weight* isn't the only metaphor in terms of which we can conceive the normative pressure exerted by a reason. We can also fruitfully think of it in terms of *force*. And there's no problem with understanding how a single entity could exert multiple different kinds of forces at once.

Whether we accept (P1) turns on which of these pictures of the weight or force of reasons we accept. Which should we prefer? Well, one might reasonably think that this is an issue which our favoured theory of reasons is going to have to decide. Indeed, Howard himself doesn't offer an argument for his preferred picture. The factivist would be within their rights to say, then, that whether we should accept (P1) turns partly on whether we should be factivists in the theory of reasons in the first place.

This is all well and good, one might think: Howard's third argument against factivism is not itself convincing. Nevertheless, there remains an *explanatory challenge* in the vicinity for the factivist to answer. The challenge is that they need to provide a compelling answer to the question of what makes a reason prudential as opposed to moral. Even if we agree that, for the factivist, normative reasons are clusters of facts, we still don't have a good answer to that question. After all, presumably it's possible for a single cluster to be both a moral and a prudential reason. Howard's dual-aspect view, however, provides a ready response to the question.

This explanatory challenge can be answered. Properly developed, the factivist is going to want to say that reasons are clusters of facts which stand in the relation of favouring to some agent and some (prospective) response; this relation has a *valency* (pro or con); it has a certain *strength*; and, relatedly, it displays a certain *pattern of normative support*. The pattern of normative support displayed by a reason is the way that reason determines how one ought to respond, across all contexts in which it appears. To completely map the pattern of normative support of a given reason would be to draw-up a complete specification of under what conditions the reason is undercut; its force attenuated and to what degree; is outweighed; is reduced from a decisive to a sufficient reason... and so on.

Here is how the factivist might answer the lingering explanatory challenge: What

differentiates moral reasons from prudential reasons is not that each sort of reason is individuated by a different sort of goal. Rather, what differentiates them is the distinctive pattern of normative support associated with each. *Moral* reasons are clusters of facts which display a distinctively moral kind of pattern of support: They are reasons which are (typically) weightier than prudential reasons; they sometimes silence competing non-moral reasons; they function as side-constraints in the production of value... – with the precise details worked out *via* a first-order theory of morality. *Prudential* reasons are clusters of facts which display a *distinct* pattern of normative support: They are reasons which are typically outweighed or even undercut by moral reasons; their force depends on (some sub-set of) the agent's contingent pro-attitudes such that it can be modified (attenuated/strengthened) by corresponding changes to those pro-attitudes... – again, the precise details being worked out *via* a first-order theory of prudence. This sort of picture is perfectly available to the factivist, and it is the beginnings of an adequate response to this lingering explanatory challenge.

3 A Dilemma for the Dual-Aspect View

I hope to have shown that the dual-aspect theory is unmotivated, for all Howard has said. I now want to show that the view is independently objectionable by confronting it with a dilemma.

The dilemma is premised on the following principle:

Permissible Goals Only. If normative reasons are partly individuated by goals, then a goal partly individuates a normative reason only if that goal is a permissible one to adopt.

Some goals are permissible to adopt, some are not: Some are such that we *may* adopt them, some are such that we *may not* adopt them. It's permissible in general to adopt the goal to help others who are in need; it isn't permissible in general to adopt the goal of inhibiting further those in need, for example. According to Permissible Goals Only, not just any old goals can partly individuate normative reasons – only those goals which one *may* be in the business of pursuing can do that ontological work. If this is right, the dual-aspect theory isn't just committed to the claim that goals partly individuate normative reasons, but that *permissible* goals do.

Permissible Goals Only is a plausible principle. First, it's simply intuitive. Intuitively, the impermissible goal *to inhibit further those in need* cannot be part of what counts in favour of one performing an action which promotes the goal. Why not? Be-

cause it's not a goal one may adopt. And second, about propositionally structured entities we want to say: For any such entity, it will have to be something it is *correct* to believe, if it's to partly individuate a normative reason. For example, the proposition *the wire is live* will have to be *true* if it's to be part of a normative reason not to touch the wire. But it would be very odd to accept a correctness condition on the relevant propositionally structured entities whilst denying an analogous correctness condition on goals.

With Permissible Goals Only established, I can now state the dilemma:

The Dilemma. EITHER (Option A): What makes a goal a permissible one to pursue is that it has some further normative property; OR (Option B): What makes a goal a permissible one to pursue is its having some *non-normative* property.

These options appear jointly exhaustive. I'll now argue that each is problematic for the dual aspect view.

First, take Option A. On this option, we are to offer a theory of the permissibility of a goal which doesn't leave the normative domain. That is, we take a further kind of normative property – value, reasons, fittingness... – and we say that for a goal to be permissible consists in its instantiation of some condition essentially involving that property. Now, my own view is that if such an analysis of the permissibility of a goal is to work it will have to be cast in terms of normative reasons. It'll have to be an analysis of the permissibility of a goal according to which for a goal to be permissible is for one to have sufficient normative reason to adopt it. This is of course controversial, and I do not have space to defend it here. But it's acceptable to make the assumption not least because Howard himself shows sympathy towards it: As we saw in connection with his third argument, above, Howard assumes that we're to analyse the permissibility of an action in terms of there being sufficient normative reason to adopt it.

The upshot is that if the dual aspect theorist takes Option A, they'll have to say that what makes a goal a permissible one is that there is sufficient normative reason for one to adopt it. But now a problem arises. The dual aspect theorist wants to individuate normative reasons partly by goals and I've argued that they'll have to say that the goals in question are permissible ones. But if what makes a goal a permissible one is that there are jointly sufficient normative reasons to adopt it, then for any normative reason that's individuated by a goal, there will have to be some further normative reason in existence which is a sufficient normative reason to adopt the goal which partly individuates our original normative reason. This additional normative reason will itself either be partly individuated by a goal or it won't. If it is, then in the light of what's already been agreed, there will have to be some *further* normative reason on the scene which is a

sufficient reason to adopt that second goal. . . and we are off on a regress. So it must be that this additional normative reason is not itself individuated by a goal. But the dual-aspect theorist precisely wants to say that *all* normative reasons are partly individuated by goals.

Now, take Option B. If the dual aspect theorist endorses this option, they'll end up having to endorse a principle to the effect that the permissibility of a goal is grounded in some particular set of non-normative properties of the goal or its adoption. This property might be *that the adoption of the goal maximises happiness* or *that the goal is consistent with treating others as ends in themselves*, for example. But notice that this would mean that the dual aspect theorist has ended up committing, in the course of attempting to provide a theory merely of the ontology of normative reasons, to the existence and defence of some full-blown reductive account of the deontic property of permission.

Now, it's surely a theoretical *cost* borne by a theory of how to individuate normative reasons if that theory commits us to some full-blown reductive account of the deontic property of permissibility. Perhaps this isn't a decisive objection to the theory of reason individuation in question. But if there is a competing theory which enables us to remain neutral on whether there is a reductive account of permission available, then that competing theory is to be preferred, all-else-equal. And factivism is just such an alternative theory: After all, since goals are not part of the individuation of normative reasons on the factivist picture, the question never arises of how we're to account for the permissibility of such reason-individuating goals. Since the dual aspect theorist who goes for Option B carries this theoretical burden not shared by its opponent, the upshot of Option B is that, if the dual aspect theorist takes it, the onus of proof will be on them.

This concludes my dilemma for Howard's dual aspect view: Since it'll have to have it that it's *permissible* goals which individuate normative reasons, either the view is regressive or else the onus is on its defenders to make a case for it.

References

- Alvarez, Maria (2010). *Kinds of Reasons*. Oxford: Oxford University Press.
- Dancy, Jonathan (2000). *Practical Reality*. Oxford: Oxford University Press.
- Fogal, Daniel (2016). "Reasons, Reason, and Context". In: *Weighing Reasons*. Ed. by Errol Lord and Barry Maguire. Oxford: Oxford University Press.
- Fogal, Daniel and Risberg, Olle (Forthcoming). "Explaining Normative Reasons". *Noûs*.
- Fogal, Daniel and Worsnip, Alex (2021). "Which Reasons? Which Rationality?" *Ergo* (8).

- Gregory, Alex (2016). “Normative Reasons as Good Bases”. *Philosophical Studies*, 173 (9): 2291–2310. DOI: 10.1007/s11098-015-0609-8.
- Howard, Nathan Robert (2021). “Primary Reasons as Normative Reasons”. *Journal of Philosophy*, 118 (2): 97–111.
- Lord, Errol (2018). *The Importance of Being Rational*. Oxford: Oxford University Press.
- Maguire, Barry and Snedegar, Justin (2021). “Normative Metaphysics for Accountants”. *Philosophical Studies*, 178 (2): 363–384. DOI: 10.1007/s11098-020-01435-w.
- Mantel, Susanne (2018). *Determined by Reasons: A Competence Account of Acting for a Normative Reason*. New York, USA: Routledge.
- Raz, Joseph (2011). *From Normativity to Responsibility*. Oxford: Oxford University Press.
- Setiya, Kieran (2007). *Reasons Without Rationalism*. Cambridge, Mass.: Princeton University Press.
- Way, Jonathan (2017). “Reasons as Premises of Good Reasoning”. *Pacific Philosophical Quarterly*, 98 (2). DOI: 10.1111/papq.12135.