The Version of Record of this manuscript has been published and is available

in PHILOSOPHICAL EXPLORATIONS 30 Mar 2021

https://www.tandfonline.com/doi/full/10.1080/13869795.2021.1908581

RESEARCH ARTICLE

From the agent's point of view: The case against disjunctivism about rationalisation

Edgar Phillips*

Institut Jean Nicod, Ecole Normale Supérieure, PSL University, Paris, France

A number of authors have recently advanced a 'disjunctivist' view of the rationalising explanation of action, on which rationalisations of the form 'S A'd because p' are explanations of a fundamentally different kind from rationalisations of the form 'S A'd because she believed that p'. Less attempt has been made to explicitly articulate the case against this view. This paper seeks to remedy that situation. I develop a detailed version of what I take to be the basic argument against disjunctivism, drawing on a framework of explanatory proportionality. The disjunctivist cannot reject this framework, I argue, because they need it to respond to another challenge, from psychological individualism. As I explain, however, the proportionality-based challenge is not in principle insurmountable, and I outline a number of ways in which a case for disjunctivism might be developed in response to it. The paper thus clarifies the dialectic around disjunctivism about the rationalisation of action and, specifically, what advocates of the view must do in order to make a compelling case for it.

* Email: edgar.phillips@ens.psl.eu

Keywords: motivating reasons; reasons for action; disjunctivism; rationalising explanation; rationalisation

1. Introduction

We commonly explain people's actions simply by stating facts that are, or were, or that seemed to the agent to be, reasons for taking the action in question. I am staring into a tree. You ask me why. I tell you that there is a golden oriole in the tree. My answer, if true, apparently provides a perfectly good account of why I am doing what I am doing, and apparently does so by giving my reason for doing it. I am looking into the tree because there is a golden oriole up there. In light of the fact that there is a golden oriole up there, it makes sense for me to look into the tree: by doing so I might get a good look at a striking and beautiful bird. It seems that in cases like these, a fact about the world *rationalises* an agent's action, in roughly the sense made standard by Davidson (1963): the fact explains the action 'from the agent's point of view', by giving the agent's reason for doing what they did. I'll call a 'worldly' fact that explains in this kind of way a *worldly reason*,[1] and an explanation of an action or attitude in terms of such a fact a *worldly rationalisation*.

Of course, we sometimes make mistakes. Suppose there is no bird in the tree: I mistakenly think there is a golden oriole up there because a loudspeaker hung in the branches is playing that bird's unmistakable fluting song. You might still advert to a 'worldly' fact – say, the fact that the song of the golden oriole is emanating from the tree – to explain my looking. In so far as the sound explains my looking, though, it's plausible to say that it does so only indirectly,

in virtue of explaining why I *think* there's an oriole up there. If I lacked this belief, I might enjoy the song, but I'd have no reason to look in the tree (I like birds, not audio equipment). If, in this error case, *you* happen to know the truth, that there is no bird, then you won't be able to understand my action simply by looking at how things are in my environment. To understand my looking, you need to consider my action in light of *how I take things to be* – in light, that is, of what I believe. A suitable explanation of my action, in this 'bad case', would be something like: They're looking into the tree because they think there is a golden oriole there. Such an explanation – what I will call a *perspectival rationalisation* – seems to make sense of an action in a similar kind of way to the way that a worldly rationalisation does in the 'good case'. It shows what point the agent sees in doing what they're doing. Yet the two explanations cite very different kinds of facts: ordinary facts about the world in one case, psychological facts about the agent in the other. Nonetheless, they are closely connected, in that the content of the belief that explains the action in the one case corresponds to the fact that explains the action in the other.

Even more strikingly, it seems like the one explanation entails the other. If I'm looking in the tree because there is an oriole – and this is a rationalising explanation – it seems to follow that I'm looking in the tree because I believe there is an oriole. Of course, if you know that there really is an oriole in the tree, citing my belief alone might not be the most informative explanation you could give to a third party. Hence it can sound odd or misleading to give only a perspectival rationalisation when its worldly counterpart is available. This is plausibly explained as a matter of conversational pragmatics: if you yourself weren't sure that there was actually an oriole in the tree, you could still explain

my looking in terms of my belief. If you subsequently found out that I was right, this wouldn't show that your prior explanation was false, just that it didn't tell the whole story. The reverse entailment, however, does not hold, as is obvious from error cases. I can have a false belief that rationalises my action perfectly well even though there is no corresponding worldly fact on the scene. This asymmetrical dependence of one kind of rationalisation on the other might seem to suggest that in some sense what *really* rationalises action is belief, or that what really matters for understanding my action is not how things actually are, but how I take things to be. On this view, what is fundamental for the purposes of rationalisation is not how things actually are 'out there', or whether the agent has the actual facts in view, so to speak, but rather how things *appear* from the agent's perspective. The point of view from which we rationalise an action is the agent's one, narrowly or non-factively construed. I'll call this view *narrow perspectivalism*, or *perspectivalism* for short.

Narrow perspectivalism seems to promise a simple and unified account of the two forms of rationalisation, explaining the asymmetrical dependence of the worldly on the perspectival form, and promises to give a naturalistic account of how a fact about things 'out in the world' could bear on a person's behaviour. As we will see, how exactly it does explain these things in detail is not necessarily entirely straightforward. In particular, we might doubt whether perspectivalism can provide an adequate account of worldly rationalisation. Might it not be that, sometimes at least, we understand an action most basically as a response to the facts of the situation, such that there are irreducibly worldly rationalisations?

Elsewhere, where there is an asymmetrical dependence of a 'success' condition present only in the 'good case' on a neutral condition also present in the 'bad case', some authors advocate *disjunctivist* (or *non-conjunctivist*) views, claiming that there is something special about the success condition that cannot be captured without loss in terms of the neutral condition.[2] Thus knowledge-first theorists argue that what it is to know that p cannot be explained in terms of believing that p plus further conditions being met; disjunctivists about perceptual knowledge claim that when one sees that p, one has grounds for believing that p which entail that p, going beyond those one has when it merely seems to one that p; and naïve realists hold that the phenomenal character of veridical perception is constituted by the mind-independent objects perceived, so that a genuine perceptual experience is an event of a fundamentally different kind from a hallucination.[3]

While the relationship between worldly and perspectival rationalisation has received considerably less attention than these issues in epistemology and the philosophy of perception, there have recently been a few noteworthy attempts to support a disjunctivist view of the rationalisation of action.[4] Even less work has been done, however, to explicitly and directly motivate the opposite, perspectivalist view.[5] Narrow perspectivalism is treated as something like the default view, but most of the arguments to be found in the existing literature are either very brief or concern claims whose relation to the perspectivalism–disjunctivism issue is at best indirect (such as arguments concerning 'causalism' about action, or whether reasons for action are mental states). This paper aims to fill this dialectical gap. The goal is not to establish perspectivalism conclusively, but rather to make it clearer what exactly

proponents of disjunctivism are up against, and hence to indicate what form a compelling case for disjunctivism might have to take.

The paper proceeds as follows. First, in §2, I outline a more precise formulation of the perspectivalist and the disjunctivist positions, drawing on a recent paper by J.J. Cunningham. In §§3–4 I then look at two arguments against disjunctivism. The argument in §3 is based on a strong *a priori* thesis about psychological explanation, limiting psychologically explanatory factors to intrinsic properties of the agent. This rules out not only worldly reasons as genuinely psychologically explanatory, but arguably even beliefs, in so far as these are individuated in externalist or world-involving ways. Thus it is not clear to what extent even perspectivalism would survive this argument. Thankfully, there is a standard response to this line of argument based on considerations of explanatory proportionality: in essence, some world-involving conditions explain an agent's acting as they do better than any intrinsic conditions of the agent because the worldly conditions are better proportioned to the relevant outcome. Proportionality thus in principle opens a route for disjunctivists to argue that the worldly conditions that they favour are also sometimes genuinely explanatory. In §4, however, I show how considerations of explanatory proportionality can, in light of the asymmetrical dependence between the two forms of rationalisation, be turned against the disjunctivist. This asymmetry, I argue, suggests that narrowly perspectival conditions are better proportioned to the actions they explain than are the corresponding worldly conditions, and this suggests that the latter are not explanatorily relevant to the relevant outcomes. Finally, in §5, I explain how this argument, while powerful, could potentially be resisted without rejecting

the proportionality principles on which it is based, outlining a number of ways in which the disjunctivist might proceed. Whether any of these lines of response might be successful remains a matter for further investigation.

2. What is at issue?

Our topic is the relationship between two kinds of explanation. It is important not to confuse this with any of a number of other, perhaps more familiar, debates about reasons for action. One such debate, for instance, concerns the nature of and relations between 'normative' and 'motivating' reasons for action. Some claim that while normative reasons – the reasons that 'favour' and justify actions – are worldly facts, motivating reasons – those that actually move agents to act – are something quite different, such as mental states.[6] Another related question concerns the relationship between reasons and rationality. Some hold that whenever someone acts rationally they act for a reason, arguing on this basis that in error cases there must therefore be reasons that are not facts. Others argue that it is better to say that someone who acts on a false belief can act rationally even though they do not act for a reason.[7] Both issues are orthogonal to the present question. To demonstrate this, and in order to engage as closely as possible with the concerns of disjunctivists, I will assume a robustly 'worldly' framework on which reasons for action, both normative and motivating, are facts, typically facts about the agent's situation. The *explanans* in a perspectival rationalisation, on this picture, is not (at least not normally) a normative reason for the agent to act as they did, nor is it their motivating reason for acting as they did. Rather, when a belief rationalises an action in the perspectival way, we can think of the *content* of the belief as an

‘apparent reason’: something that seemed to the agent to be a reason for doing what they did (see Alvarez [2010, 2018]; Scanlon [1998]; Sylvan [2015]). Where the belief in question is true, its content may in fact be (or correspond to) a genuine reason to do what the agent did, but the perspectival rationalisation doesn't tell us whether this is the case.

Such a ‘worldly’ conception of reasons is quite consistent with perspectivalism. Put in such terms, perspectivalism simply says that the actual presence or absence of a reason for doing what the agent did is not relevant for the purposes of rationalisation. Rationalisation is about understanding actions in light of the apparent reasons on which agents act. Whether these apparent reasons are, or correspond to, genuine normative reasons (for instance) is a further question that we may very well be interested in but which is extraneous to our understanding of the action as rationally intelligible. The disjunctivist, on the other hand, can claim that the presence of a genuine reason can make a meaningful difference to our understanding of the action itself.

Perspectivalism and disjunctivism have thus far been characterised only in rather vague terms. It will be helpful to bring some precision to the dispute. To that end, I will employ a framework recently developed by Cunningham (2019b), which I briefly outline in the remainder of this section.[8] First, Cunningham draws a distinction between *particular explanations* and *kinds of explanation*. A particular explanation is, on Cunningham's account, a proposition ‘concerning some aspect of the world, individuated at the level of Fregean sense, provided paradigmatically using a sentence of the form “p because q”’ (Cunningham 2019b, 241). Such a proposition is composed out of three elements: two further Fregean propositions – the *explanandum*,

corresponding to 'p', and the *explanans*, corresponding to 'q' – and, thirdly, 'a Fregean sense correspondent to the "because"' (Cunningham 2019b, 241). This last element, Cunningham explains, is 'a mode of presentation of some explanatorily efficacious relation that is represented to hold between the entities corresponding to *explanandum* and *explanans*' (Cunningham 2019b, 241). *Kinds* of explanation are individuated in terms of this last element: different kinds of explanation correspond to different kinds of explanatorily efficacious relations and hence different senses of 'because'. So, for example, (1) *The window broke because the ball struck it* and (2) *The planet is warming because atmospheric levels of carbon dioxide have been increasing* are obviously different particular explanations, but they are (presumably) both explanations of the same kind, efficient-causal explanation. In Cunningham's framework, this means the 'because' in each of (1) and (2) corresponds to the same explanatorily efficacious relation – here, the relation of efficient causation.[9] By contrast, (3) *The ball is hard because it is made of steel* is both a different particular explanation from (1) and (2) and an explanation of a different kind, namely a constitutive explanation. Roughly, (1) and (2) make their *explananda* intelligible by virtue of 'picking out an entity that stands in an efficient-causal relation to the entity picked out by the *explanandum*', while (3) does so by 'picking out an entity that stands in a constitution relation to the entity picked out by its *explanandum*' (Cunningham 2019b, 240–1).

In both worldly and perspectival rationalisations, we can say, in Cunningham's mode, that the *explanandum* – the agent's performing the relevant action – is made intelligible by the *explanans*'s picking out an entity that stands in a *rational-motivation relation* to the action. When a single action

can be explained by either a worldly or a perspectival rationalisation, the *explanandum*, the action, is the same. The interesting questions turn on either the relation between the *explanans* of each, or between the rational-motivation relations at play in each, or both.

Cunningham argues that the issue of disjunctivism turns on whether or not worldly rationalisations and perspectival rationalisations are explanations of the same *kind*: that is, whether there is a single rational-motivation relation that underlies both, or whether each depends on a distinct form of rational-motivation relation. To see why, it helps to bring in an aspect of worldly rationalisations that I have not yet discussed. We've already noted that whenever a worldly rationalisation 'S A'd because p' is true, there is a corresponding true perspectival rationalisation 'S A'd because S believed that p'. While worldly reasons themselves don't normally concern the agent's perspective, then, the fact that a worldly rationalisation is true of a certain action does tell us something about that perspective: it tells us not only that a certain fact obtained but also that that fact appeared to the agent to obtain. However, worldly rationalisations do not hold of all actions rationally motivated by beliefs that happen to be true. Suppose Liz needs a Jersey cow and they have Jerseys at the Hereford market. Liz also believes that they have Jerseys at the Hereford market and goes to Hereford on this basis. However, Liz only believes this because a friend told her that they have Jerseys in Hertford and Liz later mixed the two towns up in her memory. Here it seems that while Liz goes to Hereford because she thinks they have Jerseys there – and this is the 'because' of perspectival rationalisation – and while it's true that they have Jerseys there, it's not true that Liz goes to Hereford *because they have*

Jerseys there. Her belief isn't sensitive in the right way to the actual obtaining of that fact for the fact to explain her action. Based on such cases, several authors have argued that worldly rationalisation depends on the agent's *knowing* the relevant worldly fact.[10] Worldly rationalisation, the thought goes, is unavailable in just the kinds of cases in which the agent fails to know the fact in question: where the belief is mistaken, unreasonable, 'Gettiered', and so on. The best explanation for this is that acting for a worldly reason requires knowledge. If this is right, then 'S A'd because p' implies that S knew that p. Actions of which worldly rationalisations hold are actions that manifest knowledge. This seems to establish a close connection between worldly rationalisations and what we might call *knowledge rationalisations*: S A's because p (in the rationalising sense) only if S A's because S knows that p.[11]

Acknowledging such a role for knowledge gives us a helpful alternative way of considering the relation between worldly and perspectival rationalisations. We can now say that each kind of rationalisation makes the action it explains intelligible by identifying a cognitive mental state that stood in a rational-motivation relation to the action. Since the *relata* are of a common kind, the idea that there might be a single relation at play becomes much clearer. It also means we can consider the disjunctivism–perspectivalism question in terms of the relative explanatory roles of knowledge and belief. This will become especially significant in §§4–5. For now, though, it also gives us a useful way to formulate the disjunctivist's basic claim: they claim, and the perspectivalist denies, that when a worldly rationalisation 'S A's because p' holds of an action, this is so in virtue of the obtaining of a rational-motivation relation R between the agent's state of knowing that p and their action of A-ing,

where R is a relation that can *only* hold between an action (or other rational response) and a state of knowledge. Knowledge, for the disjunctivist, explains *qua* factive state.

The perspectivalist denies this. It is important to recognise, though, that this does not necessarily mean denying that the agent's state of knowledge underlies a rationalising explanation in these cases; what it does mean is denying that it explains *qua* knowledge, *qua* factive state. Specifically, the perspectivalist claims that the rational-motivation relation on which worldly rationalisations depend is just the same as that on which perspectival rationalisations depend: it is of a kind which can relate cognitive states and actions whether or not those cognitive states are factive. At most, knowledge explains action, for the perspectivalist, *qua* appearance, rather than *qua* knowledge.[12] My action might be rationally motivated by a state that is in fact one of knowledge, but that it counts as knowledge doesn't matter for the purposes of my action's being rationally motivated, nor, therefore, for the purposes of making my action intelligible. The basic perspectivalist claim, then, is that worldly and perspectival rationalisations are not explanations of different kinds: in both, the sense of the 'because' is a mode of presentation of a rational-motivation relation that can hold between actions and non-factive cognitive states. One form of perspectivalism might insist, further, that any given worldly rationalisation is in fact the same *particular* explanation as its corresponding perspectival rationalisation. They would have to claim that when we say something of the form 'S A'd because p' (where that p is a worldly fact), in so far as this gives a rationalising explanation of S's A-ing, it does so simply by telling us that S A'd because S believed that p. The apparent

difference between the two statements – in particular the fact that 'S A'd because p' on its face asserts an explanatory connection between the fact that p and S's A-ing – would have to be explained or explained away, and it is a nice question whether this can be done in a satisfactory way.[13] A less restrictive form of perspectivalism, though, allows that 'S A'd because p' is a different particular explanation from 'S A'd because S believed that p', whilst nonetheless insisting that the *ways* that the former and the latter make sense of the action are the same, that the relation that each posits between the agent's cognitive perspective and their action is the same.

There is much more to be said about how exactly the perspectivalist might account for worldly rationalisations, and it is arguably a condition on the acceptability of the perspectivalist view that it provide a sufficiently plausible account of how worldly rationalisations work – in particular, one that does not commit us to an error theory about our ordinary rationalising practices.[14] Whether such an account is available is in part what is at issue: the disjunctivist claims that there is something about worldly rationalisation that *cannot* be captured in narrowly perspectival terms. We will return to this issue in §5. For now, I will set this issue aside and assume for the sake of argument that the perspectivalist can at least give some account of how worldly rationalisations explain actions without positing an essentially worldly rational-motivation relation.

3. The argument from individualism

Disjunctivism is a thesis about a certain kind of psychological explanation.[15] One way to challenge the disjunctivist thesis is to argue that worldly facts are

simply not the right kind of facts to give psychological explanations of behaviour. Rationalisations, it is widely assumed, are causal explanations (Davidson 1963), and we might be suspicious of the idea of a 'worldly' condition's causally explaining a person's behaviour. This suspicion can be cashed out in terms of *psychological individualism*, the view that the psychologically explanatory facts about an agent are fixed by how that agent is intrinsically or internally (Crane 1991; Fodor 1991; Jackson 2009). Since the worldly facts do not in general depend on the agent's intrinsic condition, individualism implies that worldly rationalisations are genuinely psychologically explanatory only in so far as they carry information about facts that *do* depend solely on the agent's intrinsic condition. This would seem clearly to rule out the disjunctivist's idea of a necessarily 'worldly' rational-motivation relation, leaving at most the neutral, narrowly perspectival relation. In this section I will briefly present a common line of argument against the kind of restriction that individualism places on psychological explanation. The main point of doing so is to establish the significance of principles of *explanatory proportionality* that we will, in the next section, turn against disjunctivism.

Individualism of the sort that would rule out disjunctivism might derive from something like the following line of thought. Psychological explanations of a person's actions are causal explanations of that person's actions. As such, a psychological explanation imputes causal efficacy to the conditions mentioned in its *explanans*. The immediate or proximal causes of a person's actions are intrinsic conditions of that person. Since a person's actions are immediately caused by intrinsic states of that person, only differences in the agent's intrinsic condition can make a causally explanatory difference to what

they do. Aspects of the world outside the agent can affect what an agent does only mediately, through somehow affecting their intrinsic condition.

In Fodor's (1991) version of this argument, consideration of certain thought experiments reveals that any putative difference in the behaviour of two intrinsically qualitatively identical agents will conceptually depend upon a corresponding difference in the relationally-individuated content of their psychological states. This means that the applicability of certain relational descriptions to an agent's behaviour depends on how things are in the agent's environment. However, Fodor argues, this isn't a genuine causal difference, since at a more basic level of description the behaviour will be the same. For example: Jerry is here on Earth; his intrinsic duplicate 'twin', Gerry, is on Twin Earth. The only difference between Earth and Twin Earth is that on the latter, instead of water, they have twater, a substance qualitatively indiscriminable from water at the 'macro' level but with a different molecular structure. This style of example derives from arguments for semantic externalism, the idea that meanings are not 'in the head': when Jerry says 'Water is wet,' he refers to water, whereas when Gerry says the same thing, he refers to twater. Since this difference in meaning corresponds to no difference in the twins' intrinsic qualities, the meanings of their utterances cannot be fixed by their intrinsic qualities (Putnam 1975). The same kind of argument is extended to support externalism about psychological content, which makes a corresponding claim about what Jerry and Gerry believe (judge, suppose, and so on), rather than just the meanings of the words they utter (Burge 1979, 1986; McGinn 1977).

Fodor grants these externalist conclusions; however, he claims that the differences in meaning and content are not relevant for the purposes of

psychological explanation. True, when Jerry wants water and thinks there is water in the kitchen, he goes to the kitchen to get water, not twater, and when Gerry wants twater and thinks there is twater in the kitchen, he goes to the kitchen to get twater, not water, so the difference in mental content corresponds to a difference in behaviour. But this difference, Fodor argues, is not causal or psychologically explanatory, but merely conceptual: the content of an intention in action is conceptually fixed by the contents of the mental states that produce it. If we switched the twins' places, Jerry would still intend to get water and Gerry would still intend to get twater, but in a real sense what each would actually *do* would be exactly what the other would have done in his place. Since the difference is merely conceptual, it is non-contingent, hence not causally explanatory. Or so Fodor argues.

For individualism to support perspectivalism (rather than a more extreme error theory about rationalisation in general) we would have to assume that the beliefs to which perspectival rationalisations refer can play their explanatory role understood as purely intrinsic states of the agent. The point of Fodor's argument is to say that even if beliefs have externalist content, this content does not play a psychologically explanatory role: for psychologically explanatory purposes, we can, as it were, strip off any putatively world-involving aspect of belief ascriptions. What does the real work is whatever purely intrinsic state underlies the relevant belief. However, it's not clear that examples like Fodor's really do establish this claim, and there is reason to think that the world-involving aspects of externalist beliefs might be explanatorily significant after all. To explain why, though, we need to introduce the notion of proportionality.

The idea here is that whether a certain condition is explanatorily relevant to a certain outcome depends on how well-proportioned the condition is to the outcome. Consider, for example, the following two explanations of a person's injuries: (4) her leg was broken because she was hit by a car; (5) her leg was broken because she was hit by a 2007 Peugeot 206. While (5) is more specific, there is an intuitive sense in which (4) is a more proportionate explanation, in large part because the extra specificity in (5) is largely irrelevant to the outcome that's being explained. The 2007 Peugeot 206 does not, presumably, have any special leg-breaking power that other models lack. This rough idea is nicely sharpened and formalised by Yablo in articulating his proportionality theory of causal relevance (Yablo 2003). The theory is inspired by counterfactual theories of causal relevance, but aims to resolve certain problems that such theories face concerning the counterfactual relevance of properties that, when we consider causal relevance, seem unnecessarily weak or (like being a 2007 Peugeot 206 in our example) unnecessarily strong. Yablo (2003, 324) rules out such cases as follows:

- A property P of x is egregiously weak (relative to effect y) iff some more natural stronger property of x is better proportioned to y than P is.
- A property P of x is egregiously strong (relative to effect y) iff some as natural weaker property of x is better proportioned to y than P is.

'Natural', here, is used in the sense common in the metaphysics of properties, made standard by Lewis (1983). Roughly, the idea is that more natural properties are more fundamental. Perfectly natural properties are absolutely fundamental. Less natural properties are more 'arbitrary' or

'gerrymandered'. So, for instance, the property of being a resplendent quetzal is (presumably) more natural than that of being the national bird of Guatemala; the property of being circular is more natural than that of being either circular or pointy; and so on.[16]

Proportionality is assessed in terms of counterfactuals:

- Q– is better proportioned to *y* than Q+ iff *y* would still have occurred, had *x* possessed Q– but not Q+.
- Q– is worse proportioned to *y* than Q+ iff *y* would not have occurred, had *x* possessed Q– but not Q+.

Q– and Q+ being, respectively, weaker and stronger properties of *x*.

We can now state the proportionality theory of causal relevance:

- A property P of *x* is causally relevant to effect *y* iff
  a) had *x* lacked P, *y* would not have occurred
  b) P is not egregiously weak or strong. (Yablo 2003, 324, apparent error corrected)

Supposing rationalisation is a species of causal explanation, considerations of causal relevance will presumably be significant for assessing the explanatory relevance of different conditions for the purposes of rationalisation. The proportionality-based response to the individualist argument is, in brief, that relational conditions are sometimes better proportioned to the outcomes we are explaining than is any intrinsic condition of the agent. If this is right, then the individualist's claim that relational conditions cannot be genuinely explanatory is mistaken.

Take the ordinary explanation 'Liz went to Hereford because she believed there were Jerseys there'. 'Liz believes that there are Jerseys in Hereford' is, if

we accept content externalism, a relational or world-involving ascription: to have a belief with this content, Liz has to stand in a suitable relation to Jersey cows and to Hereford (whatever relation it is that, on our preferred externalist account, puts Liz in a position to think about Jerseys and about Hereford respectively). Call the belief that there are Jerseys in Hereford *B*. Call Liz's (also relational) action of going to Hereford *A*. Plausibly, B could be realised by any of a diverse range of intrinsic states, I_1, … I_n such that 'the' state I that realises B would be, to borrow a phrase from an earlier Fodor, 'wildly disjunctive' (Fodor 1974, 103). Much as the same intrinsic state underlies Jerry's and Gerry's respective beliefs about water and twater, two people in different circumstances who have acquired the belief that there are Jerseys in Hereford in different ways might have quite different intrinsic states realising that shared belief. If so, it is plausible that B is better proportioned to A than is the intrinsic state I_x that in fact realises B: A would still have occurred had Liz possessed B but not I_x, so B is not egregiously weak (even if less natural than I_x); and A would not have occurred had Liz lacked B. Hence B is causally relevant according to the proportionality framework.

The argument as presented is, admittedly, extremely brief, and these issues merit further discussion. They are discussed in much greater detail elsewhere.[17] The point here is not, however, to refute individualism, but instead, first, to suggest that the mere fact that worldly rationalisations appeal to environmental conditions does not immediately show that disjunctivism must be false and, second, to show how considerations of proportionality can be brought to bear in assessing the relative explanatory roles of different conditions.

4. Perspectival rationalisations as proportionate explanations

The basic argument for perspectivalism rests on the already noted asymmetric dependence of worldly rationalisation on perspectival rationalisation. As Davidson puts it: 'Your stepping on my toes neither explains nor justifies my stepping on your toes unless I believe you stepped on my toes, but the belief alone, true or false, explains my action' (Davidson 1980, 8). Because of this, perspectival rationalisation seems to capture relevant similarities between cases – similarities in the actions of agents and in the psychological factors that are apt to explain those actions – that worldly rationalisation seems to miss.

Take the following three cases:

(6) You stepped on my toes and I know that you did.

(7) You stepped on my toes but my belief is 'Gettiered' (say because, although I didn't feel you stepping on my toes, a small child, whom I didn't notice because they were below my line of sight, stepped on the toes of my other foot at the same time).

(8) You didn't step on my toes and I merely think that you did (perhaps because of the child again).

In any of (6)–(8), we could explain my stepping on your toes by saying that I thought you stepped on my toes, but only in (6) is it true that I stepped on your toes because you stepped on my toes. It thus seems, first, that if the fact that p gives a worldly rationalisation of an agent's A-ing, then the agent's A-ing in the same case is also explained by their believing that p, and second, that their believing that p would explain their A-ing even if it were not the case that p. Moreover, both explanations rationalise the action: both explain the action by

showing what point the agent saw in it, or how it made sense from the agent's point of view to take that action. When we can give a worldly rationalisation, the agent's having a corresponding belief is, but the fact's obtaining is not, necessary for the action to be rationalised. This, the perspectivalist argues, suggests that the fact's obtaining is not a necessary part of the rationalisation of the action, even when the rationalisation is given in the worldly form.

This line of thought can be sharpened through the same framework of proportionality outlined above. Since Yablo's tests for relevance are expressed in terms of properties, using them to compare the causal relevance of beliefs and worldly reasons will require us, somewhat awkwardly but not, I think, problematically, to represent the obtaining of a worldly reason as a property, such as *being in a world in which p*. How does this property compare with the agent's believing that p with respect to causal relevance, in a case in which the agent's A-ing can be given a worldly rationalisation, 'S A'd because p'? Where S A's because p, the agent's being in a p-world does seem causally relevant. The worldly rationalisation suggests that, had it not been the case that p (had S not been in a p-world), S would not have A'd. Whether S's being in a p-world is egregiously weak or strong will depend on the case and what we fill in for 'p', but there is no general reason to expect that S's being in a p-world will be egregiously strong relative to S's believing that p. In general, neither property will be weaker or stronger than the other, since neither, as a rule, entails the other. It seems, therefore, that when a worldly rationalisation can be given, the agent's believing that p and her being in a world in which p will in general both be causally relevant.

This might seem to support the disjunctivist. Unfortunately, that would be too quick. The real point of the perspectivalist argument is not that worldly reasons are never causally relevant to actions. Rather, it is that a worldly rationalisation explains an action only because it tells us something relevant about how things seemed to the agent. The argument, then, is not best put by comparing the relevance of the *obtaining* of a worldly reason with that of the agent's having a corresponding belief, but rather by comparing the relevance of the agent's having the belief in question with that of the agent's being such that she is in a position to act for the worldly reason in question. In other words (maintaining the assumption we made earlier), we need to compare the relevance of the agent's believing that p with that of their knowing that p.

In normal cases, S's knowing that p will likely meet the first condition for causal relevance: it will often be true that had the agent not known that p, they would not have A'd. After all, it will often be the case that had the agent not known that p, they wouldn't have believed that p. However, knowledge, the perspectivalist claims, fails the second condition: it is egregiously strong, because believing that p is better proportioned to S's A-ing, in that S would still have A'd had she believed that p but not known that p (because her belief was false, Gettiered, or whatever) and believing that p is, we may suppose, at least as natural a property as knowing that p (see Yablo [2003]). If that is right, then knowledge fails our test for causal relevance.[18]

We need to remember, of course, that a rationalisation of an action does more than just identify conditions under which a certain kind of event occurs. A rationalisation is not *just* a causal explanation: it explains the action *from the agent's point of view*, showing what point the agent saw in taking it, and thus

making it ‘rationally intelligible’. However, the perspectivalist will argue, the same observations still tell against disjunctivism even when we turn our attention to the rational intelligibility of the action. The agent's beliefs, after all, make the action rationally intelligible. Indeed, they constitute the ‘point of view’ from which we understand action. Where the agent's knowing something makes their action rationally intelligible, the action would still be rationally intelligible had they merely believed and not known. Again, the stronger condition looks egregiously strong. The weaker belief condition is better proportioned to the action's being rationally intelligible than is the stronger condition, knowledge.

5. Resisting the argument

How might the disjunctivist respond to this line of argument? One way would be to reject Yablo's tests for relevance. The challenge here would be to find a well-motivated way of assessing explanatory relevance that allowed us both to maintain the proportionality-based rejection of individualist restrictions without at the same time supporting the proportionality-based argument for perspectivalism. It is unclear what such a test might look like or what might motivate it. I will set this line of response aside, anyway, since there are other means by which the disjunctivist might resist the perspectivalist conclusion and which, I believe, connect more directly with the things that actually motivate disjunctivists. In this section I offer a brief overview of potential argumentative routes the disjunctivist might take, along with obstacles that arise on those routes. The goal is not to say conclusively whether any one of these routes will in fact lead to a compelling case for disjunctivism, but merely

to offer a sketch of the dialectical landscape as shaped by the proportionality-based argument for perspectivalism.

A first option is to argue that knowledge is in fact more natural than belief, rendering belief egregiously weak in the good case. Notably, the perspectivalist argument presented in the previous section does not depend on the view that knowledge is *less* natural than belief. It is hence consistent with the idea that knowledge cannot be analysed in terms of belief and that knowledge is a distinctive kind of mental state in its own right (Williamson 2000). However, the argument does depend on the idea that belief is at least as natural as knowledge. The most obvious way to motivate the idea that knowledge is more natural than belief would be to analyse belief in terms of knowledge. There have been some proposals along these lines. Williamson, for instance, suggests that believing that p is 'roughly, treating p as if one knew p' (Williamson 2000, 47). Somewhat similar accounts are suggested by Nagel (2017) and Hyman (2017), and a somewhat more sophisticated version of the idea is developed by Wimmer (2019). The challenge for the disjunctivist wishing to pursue this line will be, simply, to show that something like this is correct, which is no mean feat.[19]

A different approach might appeal not to a theory of belief but to a metaphysical account of something else that underlies rationalising explanations, such as the rational capacities that the agent exercises in responding to reasons. We might, for instance, be attracted to a conception of rationality as the capacity to respond to reasons – where 'reasons' here are normative reasons and these are worldly facts (Raz 2005). On this view, we might say that the agent's rational capacities are only properly or 'non-

defectively' exercised when the agent acts in a way that is subject to worldly rationalisation, and that when an action can only be rationalised in the narrowly perspectival way, this shows that their rational capacities were exercised only 'defectively' (compare McDowell 2013; Roessler 2014). If this is right, then a certain view of the metaphysics of capacities – in particular, a view on which a capacity is 'ontologically dependent' on (hence, perhaps, less natural than) its non-defective exercise and on which the defective exercise of the capacity is in turn ontologically dependent on the capacity – might take us to a disjunctivist view (Kalderon [2018] applies such a view of capacities to perception).

Again the challenge is to explain what motivates this metaphysical picture. Even granting the initial conception of rationality as the capacity to respond to (worldly) reasons, it is not clear that we need to treat its non-defective exercise as basic. That idea seems, potentially at least, to be threatened by a natural move made by, for instance, Mantel (2018, Chap. 8), who distinguishes the epistemic and motivational aspects of the capacity to respond to worldly reasons. On Mantel's account, that capacity is a complex competence consisting in more basic sub-competences, respectively to form true beliefs, to be motivated to do what the reasons represented in one's beliefs favour, and to execute those motivations. The distinction seems well-motivated: forming a belief (which may or may not constitute knowledge) and being rationally motivated by that belief (or knowledge) are, it seems natural to say, two distinct mental 'acts', employing distinct capacities. Even if we accept the priority of the non-defective exercise of a capacity over its defective exercise, what is relevant for the obtaining of a rational-motivation relation between a cognitive

state and an action is surely the motivational sub-competence specifically. But if this is distinct from the competence to form true beliefs, whose proper exercise results in knowledge, we might assume that the former can be exercised non-defectively even where the latter is exercised defectively. One can be rationally motivated to do what one's beliefs favour even when those beliefs are not knowledgeable.

This comes to the heart of what motivates some disjunctivists (most clearly McDowell 2013; Roessler 2014): they doubt that we can understand the capacity to be rationally motivated by a belief independently of the idea of being rationally motivated by a worldly reason. The motivational aspect of the capacity to respond to reasons is, in Mantel's account, understood as the competence to be motivated by one's beliefs 'to do what is favoured by the represented reasons' (Mantel 2018, 43). But what exactly does this mean? If normative reasons are worldly facts, then it is not immediately clear in what sense, if any, beliefs, in particular false beliefs, might 'favour' taking a particular course of action. The fact that it's raining, for example, counts in favour of my putting on my raincoat. Does my belief that it's raining count in favour of my putting on my raincoat? If it's not actually raining, then we might think not. As we put it earlier, I have an *apparent* reason to put on my raincoat, but here it is *merely* apparent, not a genuine reason.

There is, of course, a very familiar sense in which beliefs can be said to 'favour' actions, whether those beliefs are true or false: our beliefs can make certain courses of actions *rational* to take. But what does it mean for a belief to make an action rational? The idea that beliefs provide apparent reasons fits neatly with a 'transparency' view of rationality as, roughly, apparent

normativity: a view on which what it is for our beliefs to make a certain course of action rational is, more or less, for it to seem to us, in having those beliefs, that there is good reason for us to act in that way (see Kolodny [2005] for a detailed account; compare also Sylvan [2015]). If something like this is correct, then it seems natural to hold that our understanding of rational motivation depends on a prior conception of normative, worldly, reasons. When we are rationally motivated by our beliefs, we act as we do because, roughly speaking, there seems to us to be some fact that favours our so acting. So, when we understand an action as rationally motivated by a belief, we understand it as, roughly speaking, something that there would be a certain kind of reason to do were the belief in question true (see Roessler and Perner [2013]). As Roessler, who presses an argument along these lines, puts it: 'The ability to make sense of actions in terms of non-factive rationalizing attitudes depends on the more basic ability to find actions intelligible as something the agent has a normative reason to do (and is doing for that reason)' (Roessler 2014, 353). This seems to suggest that the latter property – being something the agent has a normative reason to do and is doing for that reason – is more the natural one.[20]

Many perspectivalists might want to resist this line of thought early on. They might reject the idea that normative reasons are worldly facts. Or they might insist that rationality has its own normativity. Be that as it may, there is, I think, a gap in the argument even if we grant these premises. We need to distinguish two claims: first, that we can only understand the way that beliefs rationalise actions in relation to the way that normative reasons favour actions; second, that we can only understand the way that beliefs rationalise actions in relation to the way that worldly reasons (or more specifically normative

reasons) rationalise actions. The disjunctivist needs to support the second claim, but the transparency view of rationality seems immediately to support only the first.

The perspectivalist might flesh out the first claim in something like the following way. Rational agents take actions that, as it seems to them, they have good reason to take. Understanding an action as rational therefore requires a sufficient degree of agreement between us and the agent about what sorts of facts count as reasons for what sorts of responses. This is an agreement about something objective, something we see as independent of the agent's perspective. But it is inherently general: it is not primarily about what to do in the particular case, but about how to respond to certain *kinds* of facts. In the particular case, we recognise that even if, as it happens, the agent is motivated by what is in fact a genuine reason to do what they are doing, nonetheless the same grasp of normative reasons on the agent's part could be manifested were they acting on a 'mere' belief.

Roessler's case, then, depends crucially on the clause in his parenthesis: the key disjunctivist idea is that our ability to make sense of actions as motivated by belief depends on a more basic ability to find actions intelligible as something the agent has a normative reason to do *and is doing for that reason*. What I'm suggesting is that this claim does not flow immediately from a 'transparency' account of rationality. What the disjunctivist needs to explain, then, is just what is missing from our understanding of rationality if we do not treat this idea – not just the idea of a fact's being a reason to act in a certain way, but the idea of someone's acting in that way in response to that fact – as (relatively) fundamental. The idea that our understanding of rationality

depends on a prior understanding of normative reasons doesn't quite take us that far.[21]

6. Conclusion

The more or less unchallenged orthodox status of the view I've called perspectivalism goes a long way to explaining the lack of worked-out arguments in its favour. Recent arguments for the opposing disjunctivist view motivate a more careful examination of possible grounds for accepting the orthodoxy. I've argued that considerations of explanatory proportionality plausibly establish perspectivalism as the default view. However, this argument is not necessarily decisive, and I have indicated some ways in which the disjunctivist might seek to resist it. The key question is whether a compelling enough positive case can be made for the relative naturalness or fundamentality of worldly rationalisation. I have indicated some possible routes to that idea and argued that these merit further attention because while they are incomplete as they stand, they need not, if worked out adequately, be undermined by the basic argument for perspectivalism.

Acknowledgements

This paper is derived in part from a section of my PhD thesis, which itself drew from work done for my MPhil Stud thesis. I would like to thank Ulrike Heuer, Fabrice Teroni and an anonymous reviewer for another journal for helpful feedback on earlier versions, and especially Mike Martin for his many incisive questions and helpful suggestions. I am very grateful to two anonymous referees for this journal whose generous comments substantially improved the paper, and I gratefully acknowledge the financial support of the AHRC and the London Arts and Humanities Partnership, which funded my PhD at UCL, and the Swiss

National Science Foundation, which funded my postdoctoral position on the *Modes and Contents* project at the University of Fribourg and in the Thumos research group in Geneva.

Notes

1. What characterises a worldly reason, on the present usage, is that it is a fact about the agent's situation to which the agent responds or by which the agent is rationally motivated (in a broad sense of 'rationally'). It is a further question whether the fact is a *good* (normative) reason for doing what the agent does.
2. For the distinction between disjunctivism and non-conjunctivism, see Williamson (2000). For simplicity I will henceforth refer simply to 'disjunctivism'. For present purposes this should be taken to include non-conjunctivism.
3. See, respectively Williamson (2000); McDowell (1998); Martin (2004, 2006). See also Hinton (1973); Snowdon (1980).
4. In particular Cunningham (2019a, 2019b); Hornsby (2008); McDowell (2013); Roessler (2014).
5. A notable exception is Dancy (2000), who argues for a version of perspectivalism. While my thinking about these issues owes a great deal to Dancy's work, I won't discuss his arguments in detail. Dancy is primarily concerned not with the 'S A'd because p' form of worldly rationalisation but with the construction 'S A'd for the reason that p', and he, perhaps idiosyncratically, claims that the latter construction is not factive for 'p'. This leads to an idiosyncratic way of framing the issue. It also means that Dancy's arguments for perspectivalism struggle to engage with those who take more seriously the 'worldliness' of worldly rationalisation; see for instance the exchange between Dancy (2011) and Hyman (2011).
6. See for example Smith (1994); also Davidson (1963). Dancy (2000) argues that normative and motivating reasons cannot be fundamentally distinct in this way. See Mantel (2014, 2018) for recent sceptical discussion of Dancy's argument.
7. For the former, see for instance Comesaña and McGrath (2014); for the latter, Alvarez (2018).
8. Cunningham uses the labels 'the Disjunctive View' for what I am calling disjunctivism and 'the Common Kind View' for what I am calling (narrow) perspectivalism.
9. An interesting question Cunningham doesn't raise is whether there might be different kinds of efficient-causal explanation (different efficient-causation relations) and if so how those different kinds relate to one another. The possibility of subdividing kinds of explanation might merit further discussion, especially for the disjunctivist, who wants to say that there are two different kinds of rationalisation. A helpful model here might be the discussion of different species–genus relations in Ford (2011).
10. Cunningham (2020); Hornsby (2008); Hyman (1999, 2011, 2015) and McDowell (2013) argue that only facts one knows can rationalise one's actions. Hughes (2014) and Locke (2015) argue that a weaker condition is sufficient. For present purposes, I will assume that the knowledge account is correct.
11. Cunningham (2019b, 238) argues for the stronger claim that a worldly rationalisation and the corresponding knowledge rationalisation are the same particular explanation.
12. I take this '*qua*' formulation from Cunningham's discussion. The formulation might seem obscure if we take explanation to relate facts rather than particulars (as urged, for instance, by Strawson [1992]). Cunningham seeks to give an account of how particulars (or at least 'entities') and relations between particulars under explanations. Cunningham takes it that when S knows that p, S's state of knowing that p is identical to S's state of believing that p, and that this state is an 'entity' of a sort that can underly (rationalising) explanations by standing in a rational-motivation relation to actions. Perspectivalist and disjunctivist can agree that knowledge-states thus relate to actions, grounding worldly rationalisations. The perspectivalist denies that the rational-motivation relation here depends on the knowledge-state's actually being a state of knowledge: the explanation holds in virtue of the knowledge-state but not in virtue of its *being* a knowledge-state.

'Knowledge explains but not *qua* knowledge' is shorthand for this idea. Thanks to an anonymous reviewer for this journal for encouraging me to address this issue.

13. See for instance Dancy (2011) and the response in Hyman (2011).
14. Thanks to an anonymous reviewer for this journal for raising this concern.
15. I assume here a broadly Davidsonian conception of psychology as the sphere of rationalising explanation. See in particular Davidson (1980, Chap. 11–13).
16. For a very helpful overview, see Dorr (2019).
17. In particular, versions of the argument are given by Peacocke (1993); Stalnaker (1989, 1990); Williamson (2000) and Yablo (1992, 2003). Compare also Davidson (1980, Chap. 11).
18. At least, its being a state of knowledge is not causally relevant; it is not causally relevant *qua* knowledge. If the agent's knowledge that p is identical with their belief that p, then it might be causally relevant qua belief). See again Cunningham (2019b).
19. For some challenges to such 'knowledge-first' approaches to belief see for instance McGlynn (2014, Chap. 2); Hawthorne et al. (2016); Rothschild (2020).
20. Perhaps 'naturalness' is no longer the right way to think of things when we enter the normative realm. But it does seem fitting to talk about fundamentality here, and this could presumably play a similar role in the argument.
21. Cunningham (2019b) gives an argument that is significantly different from Roessler's, but which occupies a broadly similar dialectical space. Both arguments, I think, merit further attention.

Notes on contributor

Edgar Phillips is a postdoctoral researcher in philosophy, currently based at Institut Jean Nicod in Paris, working on a project supported by the Leverhulme Trust. He works in the philosophy of mind, philosophy of action, moral psychology and sometimes epistemology and philosophy of psychology, mostly on issues relating to interpersonal understanding.

References

Alvarez, Maria. 2010. *Kinds of Reasons*. Oxford: Oxford University Press.

Alvarez, Maria. 2018. "Reasons for Action, Acting for Reasons, and Rationality." *Synthese* 195: 3293–3310. doi:10.1007/s11229-015-1005-9.

Burge, Tyler. 1979. "Individualism and the Mental." *Midwest Studies in Philosophy* 4 (1): 73–121. doi:10.1111/j.1475-4975.1979.tb00374.x.

Burge, Tyler. 1986. "Individualism and Psychology." *The Philosophical Review* 95 (1): 3–45. doi:10.2307/2185131.

Comesaña, Juan, and Matthew McGrath. 2014. "Having False Reasons." In *Epistemic Norms: New Essays on Action, Belief, and Assertion*, edited by Clayton Littlejohn and John Turri, 59–79. Oxford: Oxford University Press.

Crane, Tim. 1991. "All the Difference in the World." *The Philosophical Quarterly* 41 (162): 1–25. doi:10.2307/2219783.

Cunningham, J.J. 2020. "Knowledgeably Responding to Reasons." *Erkenntnis* 85: 673–692. doi:10.1007/s10670-018-0043-3.

Cunningham, J.J. 2019a. "Is Believing for a Normative Reason a Composite Condition?" *Synthese* 196: 3889–3910. doi:10.1007/s11229-017-1630-6.

Cunningham, J.J. 2019b. "The Formulation of Disjunctivism About ϕ-ing for a Reason." *The Philosophical Quarterly* 69 (275): 235–257. doi:10.1093/pq/pqy019.

Dancy, Jonathan. 2000. *Practical Reality*. Oxford: Oxford University Press.

Dancy, Jonathan. 2011. "Acting in Ignorance." *Frontiers of Philosophy in China* 6 (3): 345–357. doi: 10.1007/s11466-011-0143-5.

Davidson, Donald. 1963. "Actions, Reasons, and Causes." *The Journal of Philosophy* 60 (23): 685–700. doi:10.2307/2023177.

Davidson, Donald. 1980. *Essays on Actions and Events*. Oxford: Oxford University Press.

Dorr, Cian. 2019. "Natural Properties." In *The Stanford Encyclopedia of Philosophy* (Fall 2019 Edition), edited by Edward N. Zalta. https://plato.stanford.edu/archives/fall2019/entries/natural-properties/.

Fodor, Jerry A. 1974. "Special Sciences (or: The Disunity of Science as a Working Hypothesis)." *Synthese* 28: 97–115. doi:10.1007/BF00485230.

Fodor, Jerry A. 1991. "A Modal Argument for Narrow Content." *The Journal of Philosophy* 88 (1): 5–26. doi:10.2307/2027084.

Ford, Anton. 2011. "Action and Generality." In *Essays on Anscombe's Intention*, edited by Anton Ford, Jennifer Hornsby and Frederick Stoutland, 76–104. Cambridge, MA: Harvard University Press.

Hawthorne, John, Daniel Rothschild, and Levi Spectre. 2016. "Belief is Weak." *Philosophical Studies* 173: 1393–1404. doi:10.1007/s11098-015-0553-7.

Hinton, J.M. 1973. *Experiences: An Inquiry into Some Ambiguities*. Oxford: Clarendon Press. doi:10.1093/acprof:oso/9780198244035.001.0001.

Hornsby, Jennifer. 2008. "A Disjunctive Conception of Acting for Reasons." In *Disjunctivism: Perception, Action, Knowledge*, edited by Adrian Haddock and Fiona Macpherson, 244–261. Oxford: Oxford University Press. doi:10.1093/acprof:oso/9780199231546.001.0001.

Hughes, Nick. 2014. "Is Knowledge the Ability to φ for the Reason That P?" *Episteme* 11 (4): 457–462. doi:10.1017/epi.2014.16.

Hyman, John. 1999. "How Knowledge Works." *Philosophical Quarterly* 49 (197): 433–451. doi:doi.org/10.1111/1467-9213.00152.

Hyman, John. 2011. "Acting for Reasons: Reply to Dancy." *Frontiers of Philosophy in China* 6 (3): 358–368. doi:10.1007/s11466-011-0144-4.

Hyman, John. 2015. *Action, Knowledge, and Will*. Oxford: Oxford University Press.

Hyman, John. 2017. "Knowledge and Belief." *Aristotelian Society Supplementary Volume* 91 (1): 267–288. doi:10.1093/arisup/akx005.

Jackson, Frank. 2009. "Primeness, Internalism, Explanation." In *Williamson on Knowledge*, edited by Patrick Greenough and Duncan Pritchard, 109–121. Oxford: Oxford University Press. doi:10.1093/acprof:oso/9780199287512.001.0001.

Kalderon, Mark Eli. 2018. "Experiential Pluralism and the Power of Perception." In *The Philosophy of Charles Travis: Language, Thought, and Perception*, edited by John Collins and Tamara Dobler, 222–235. Oxford: Oxford University Press. doi:10.1093/oso/9780198783916.001.0001.

Kolodny, Niko. 2005. "Why Be Rational?" *Mind* 114 (455): 509–563. doi:10.1093/mind/fzi509.

Lewis, David. 1983. "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61 (4): 343–377. doi:10.1080/00048408312341131.

Locke, Dustin. 2015. "Knowledge, Explanation, and Motivating Reasons." *American Philosophical Quarterly* 52 (3): 215–232.

Mantel, Susanne. 2014. "No Reason for Identity: On the Relation Between Motivating and Normative Reasons." *Philosophical Explorations* 17 (1): 49–62. doi:10.1080/13869795.2013.815261.

Mantel, Susanne. 2018. *Determined by Reasons: A Competence Account of Acting for a Normative Reason*. New York: Routledge.

Martin, M.G.F. 2004. "The Limits of Self-Awareness." *Philosophical Studies* 120: 37–89. doi:10.1023/B:PHIL.0000033751.66949.97.

Martin, M.G.F. 2006. "On Being Alienated." In *Perceptual Experience*, edited by Tamar Szabo Gendler and John Hawthorne, 354–410. Oxford: Oxford University Press. doi:10.1093/acprof:oso/9780199289769.001.0001.

McDowell, John. 1998. "Criteria, Defeasibility, and Knowledge." Chap. 17 in *Meaning, Knowledge, and Reality*. Cambridge, MA: Harvard University Press.

McDowell, John. 2013. "Acting in the Light of a Fact." In *Thinking About Reasons: Themes from the Philosophy of Jonathan Dancy*, edited by David Bakhurst, Brad Hooker, and Margaret Olivia Little, 13–28. Oxford: Oxford University Press. doi:10.1093/acprof:oso/9780199604678.001.0001.

McGinn, Colin. 1977. "Charity, Interpretation, and Belief." *Journal of Philosophy* 74 (9): 521–535. doi:10.2307/2025795.

McGlynn, Aidan. 2014. *Knowledge First?* Palgrave Macmillan.

Nagel, Jennifer. 2017. "Factive and Nonfactive Mental State Attribution." *Mind & Language* 32 (5): 525–544. doi:10.1111/mila.12157.

Peacocke, Christopher. 1993. "Externalist Explanation." *Proceedings of the Aristotelian Society* 93: 203–230. doi:10.1093/aristotelian/93.1.203.

Putnam, Hillary. 1975. "The Meaning of 'Meaning'." *Minnesota Studies in the Philosophy of Science* 7: 131–193.

Raz, Joseph. 2005. "The Myth of Instrumental Rationality." *Journal of Ethics and Social Philosophy* 1 (1): 1–28. doi:10.26556/jesp.v1i1.1.

Roessler, Johannes. 2014. "Reason Explanation and the Second-Person Perspective." *Philosophical Explorations* 17 (3): 346–357. doi:10.1080/13869795.2014.946954.

Roessler, Johannes, and Josef Perner. 2013. "Teleology: Belief as Perspective." In *Understanding Other Minds: Perspectives from Developmental Social*

Neuroscience, edited by Simon Baron-Cohen, Helen Tager-Flusberg, and Michael V. Lombardo, 35–50. Oxford: Oxford University Press.

Rothschild, Daniel. 2020. "What it Takes to Believe." *Philosophical Studies* 177: 1345–1362. doi:10.1007/s11098-019-01256-6.

Scanlon, T.M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.

Smith, Michael. 1994. *The Moral Problem*. Oxford: Blackwell.

Snowdon, Paul. 1980. "Perception, Vision and Causation." *Proceedings of the Aristotelian Society* 81: 175–192.

Stalnaker, Robert. 1989. "On What's in the Head." *Philosophical Perspectives* 3: 287–316. doi:10.2307/2214271.

Stalnaker, Robert. 1990. "Narrow Content." In *Propositional Attitudes: The Role of Content in Logic, Language, and Mind*, edited by C. Anthony Anderson and Joseph Owens, 131–145. Stanford: CSLI Publications.

Strawson, P.F. 1992. *Analysis and Metaphysics: An Introduction to Philosophy*. Oxford: Oxford University Press. doi:10.1093/acprof:oso/9780198751182.001.0001.

Sylvan, Kurt. 2015. "What Apparent Reasons Appear to Be." *Philosophical Studies* 172: 587–606. doi:10.1007/s11098-014-0320-1.

Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.

Wimmer, Simon. 2019. "Knowledge-first Believing the Unknowable." *Synthese*. Advance online publication. doi:10.1007/s11229-019-02314-6.

Yablo, Stephen. 1992. "Mental Causation." *The Philosophical Review* 101 (2): 245–280. doi:10.2307/2185535.

Yablo, Stephen. 2003. "Causal Relevance." *Philosophical Issues* 13: 316–329. doi:10.1111/1533-6077.00016.