

The Formulation of Disjunctivism about ϕ -ing for a Reason*

J.J. Cunningham

Jesus College, University of Oxford

Agents can act and hold attitudes (in general: ϕ) in response to facts they treat as normative reasons to do so. When they do, a rationalising explanation of why they ϕ can be provided using an instance of the following schema, where ‘p’ stands-in for the relevant fact:

$(\neg\psi)$ S ϕ s because p

Agents can also ϕ under the impression that they are responding to a fact, even though they are not. When that’s so, we cannot provide a rationalising explanation of why they ϕ using an instance of the $(\neg\psi)$ schema, but only by using an instance of the following schema:

(ψ) S ϕ s because S believes that p

According to the *Disjunctive View*, there is an important sense in which the explanation providable using an instance of $(\neg\psi)$ is different from the explanation providable using a corresponding instance of (ψ) : an explanation of the former variety is in some important sense available at all only if the relevant fact is on the scene.¹ According to the *Common Kind View*, the two explanations are the same, in the relevant sense.²

*Ancestors of this paper were presented to audiences at the Universities of Warwick and Cardiff. I would like to thank those present for helpful feedback, objections, and suggestions for improvement. Particular thanks go to Matthew Soteriou, Johannes Roessler, Barnaby Walker, Guy Longworth, Maria Alvarez, and my anonymous referee at *The Philosophical Quarterly* whose very detailed feedback helped improve the paper immeasurably.

¹For defences of the Disjunctive View, see Hornsby (2007a, 2007b, 2008), McDowell (2013), and Roessler (2014). Although they are not explicit about it, I suspect that Hyman (1999, 2011, 2015) and Marcus (2012) are also committed to the view.

²Philosophers who accept the Common Kind View include Davidson (1963) and his many followers, as well as Dancy (2000, 2004, 2008, 2011, 2014).

I do not think it has been made adequately clear in what the sense the two explanations are said to be *the same* by the Common Kind Theorist but denied to be so by the Disjunctivist. My first aim in this paper is to remedy that. After laying out some background in §1, in §2 I argue that light can be shed on the clarificatory issue by distinguishing between particular explanations and kinds of explanation. I argue that given that distinction, the relevant sameness thesis can be read in two ways, delivering us two distinct versions of the Common Kind View.

With this clarificatory issue settled, however, it becomes apparent that the only direct arguments for the Disjunctive View available in the literature fail: if sound, they refute only one version of the Common Kind View. This is what I try to show in §3. In §4, however, I provide a fresh argument for the Disjunctive View which avoids that mistake.

1 Preliminaries

I begin with some background material. §1.1 presents the basic distinction between good cases and bad cases in terms of which the debate between Disjunctivists and Common Kind Theorists in this area is to be framed. §1.2 makes explicit my commitment to two epistemic claims about the good case which will colour the discussion. §1.3 extends our understanding of good cases and bad cases, in a way that will prove useful, by developing it in a metaphysical direction. §1.4 makes explicit two ways in which the present discussion has limited scope.

1.1 Good Cases & Bad Cases

Some facts count in favour of our performing certain actions or adopting certain attitudes: they constitute *normative reasons* for us to ϕ . That there is a by-election coming up in my constituency counts in favour of me intending to vote, for example.³ In the *good case* there is a fact on the scene of which the agent is suitably aware, they treat that fact as a normative reason for them to ϕ , and they ϕ in a way that manifests that awareness and manifests their treating the fact as a normative reason to ϕ . That there is a by-election coming up can be a fact on the basis of which I hold an intention to vote, and if I do so, I am in the good case with respect to that intention.

When the agent is in the good case, we can provide an explanation of why they ϕ using a sentence which has the following form:

$(\neg\psi)$ S ϕ s because p

³Normative reasons are *objective* reasons in the sense to be contrasted with the *subjective* reasons of interest to, for example, Schroeder (2007, 2009) and Vogelstein (2012).

The ‘because’ at issue here is a rationalising ‘because’: a ‘because’ which is correctly applicable only when the agent’s ϕ -ing manifests their treating p as a normative reason for their engagement in the act or maintenance of the attitude at issue, and hence contrasts with the merely causal ‘because’ which appears, for example, in: ‘the man feels sick because he recently ate undercooked chicken’.

As well as good cases, we should acknowledge the existence of *bad cases*. These are cases in which the agent believes that p, treats what they believe as a normative reason for them to ϕ and ϕ s in a way that manifests those attitudes, but in which they fail to count as ϕ -ing in light of a fact nevertheless. This might happen simply because what they believe is false. Alternatively, it might happen because, although p is a fact, the agent fails to be appropriately connected to it.

When the agent is in the bad case, the relevant instance of the $(\neg\psi)$ schema is false. But a corresponding instance of the following psychologicistic schema is true where, again, the ‘because’ is a rationalising ‘because’:

(ψ) S ϕ s because S believes that p

I’ve said that in the bad case, the agent ϕ s in a way that manifests their belief that p and their treating p as a normative reason for them to ϕ , so that a (ψ) statement is true of them. This also happens in the good case. It’s just that in the good case something more is going on: the agent also counts as responding to the fact that p, so that, in addition, a $(\neg\psi)$ statement is true of them. So the correct way to state the contrast here is as follows. We have a neutral condition: S rationally ϕ -ing in a way that manifests their belief that p, so that a (ψ) statement is true of them. In the good case, a further condition obtains: the agent counts as responding to the fact that p, so that a $(\neg\psi)$ statement is true of them too. In the bad case, *all that holds* is the neutral condition.⁴

⁴Complete rationalising explanations of why an agent ϕ s in fact involve a complex cluster of factors, not limited to the particular cognitive states (ψ) and $(\neg\psi)$ explanations concern. Building on the work of Lewis (1986) and Swanson (2010), Fogal (2018) develops a plausible contextualism about rationalising explanation according to which when providing rationalising explanations, we ought to pick out that element of the explanatory cluster which, given the conversational purposes at hand, serves as a good representative of the whole cluster (see also Fogal (2017)). The content of this paper is intended to be consistent with that framework. For reasons of space I will restrict myself to two comments explaining how this is so. First, (ψ) and $(\neg\psi)$ statements won’t always be the appropriate ones to assert in response to a request for a rationalising explanation of why the agent ϕ s. Second, insofar as I talk of (ψ) and $(\neg\psi)$ explanations, this can be thought of as short-hand for talk of the cluster which underpins the agent’s ϕ -ing, *but with attention restricted just to the cognitive elements of that cluster associated with (ψ) and $(\neg\psi)$ statements*.

1.2 An Epistemic Development

Next, I want to make explicit my acceptance of two epistemic theses about the good case which will colour the discussion throughout. First, I am going to accept from the outset Hornsby's (2007a; 2007b; 2008) *Reasons-Knowledge Thesis*:

(RKT) Necessarily, if S ϕ s because p then S knows that p.

According to (RKT), part of what is distinctive of good cases is that in them, the agent rationally ϕ s in a way that manifests their *knowledge* of the fact which they treat as a reason for them to ϕ , and not merely a belief they have in that fact. Although this claim has the status of an unguarded assumption here, I think it is defensible.⁵

Indeed, although the contrast I will focus on is the contrast between (ψ) and ($\neg\psi$) explanations, I will implicitly be taking the latter to be identical to explanations which, in the good case, can be provided with:

($\neg\psi_k$) S ϕ s because S knows that p.

This is the second epistemic commitment I wish to make explicit. Why is it plausible? Well I take it that, given the truth of (RKT), ($\neg\psi$) explanations are equivalent to corresponding ($\neg\psi_k$) explanations: necessarily, 'He intends to vote because there is a by-election coming up', for example, is true if, and only if, 'He intends to vote because he knows that there is a by-election coming up' is also true. But I take it also that it would be infelicitous to say: 'he intends to vote because there is a by-election coming up, and also he intends to vote because he knows that there is a by-election coming up', where this is to be interpreted as an assertion of the conjunction of *two* explanations. The reason why it would be infelicitous to say such things, it seems to me, is at least partly that what we have here is *one* explanation being provided by two different statements. The two explanations are not merely equivalent then, but identical.⁶

1.3 A Metaphysical Development

It will help to deepen our picture of what is going on at the metaphysical level in good cases and bad. In the good case, the agent ϕ s in a way that manifests their knowledge that p. We can think of this as consisting in the instantiation of a certain (synchronic) relation that holds between their ϕ -ing and their state of knowledge: a *rational-motivation*

⁵See my Cunningham (Unpublished) for a defence of (RKT).

⁶In terms of the distinction to be drawn in §2.1, I intend this to read as the thought that each *particular* ($\neg\psi$) explanation is identical to a corresponding *particular* ($\neg\psi_k$) explanation.

relation, as I will call it. Across both cases, the agent ϕ s in a way that manifests their belief that p. Again, we can think of this as a certain (synchronic) relation that holds between their ϕ -ing and their belief that p: a rational-motivation relation which binds their ϕ -ing with their state of belief. I label the relation the *rational*-motivation relation because the instantiation of the relation constitutively involves the exercise of the agent's capacity to treat what they believe as a normative reason, and I take it that this capacity is essential to them *qua* rational agent.

Care must be taken to avoid a certain misunderstanding of the minimal metaphysical picture just offered. When I say that there is a rational-motivation relation that binds the agent's ϕ -ing with some psychological state that they are in (knowledge in the good case, belief across both cases), this is not to be read as the implausible psychologistic thesis that *what motivates* the agent to ϕ is some fact about their own psychology. On the contrary, I think that what motivates the agent to ϕ in the good case is the fact known. And across both cases we can correctly say that the agent is motivated by what they believe. Rather, what I claim here is that ϕ -ing in a way that is motivated by either the fact that p or the apparent fact that p *consists in* the holding of a certain relation between one's ϕ -ing and some cognitive mental state. Consistently with that, the relevant state of mind needn't be identified with what motivates the agent.⁷

This metaphysical development allows us to frame the following question: should we think of the rational-motivation which holds in the good case as a different kind of relation to that which holds neutrally, because the former requires for its instantiation the obtaining of the factors distinctive to the good case? This question will occupy centre-stage later.

1.4 Two Limitations

Finally, I want to note two limitations on the scope of this paper. In addition to the distinction drawn in §1.1, one might wish to contrast cases in which the agent ϕ s because p – what I am here labelling 'good cases' – with cases in which the agent ϕ s for the *good* reason that p. In those sorts of cases – which we might label 'good⁺ cases' – the agent does not just ϕ in recognition of a fact which is *treated* as a normative reason for them to ϕ , but they ϕ in recognition of the relevant fact having that normative status. Good⁺ cases differ from good cases not least because an agent might ϕ because p even though p is not really a normative reason for them to ϕ ; consider, for example, agents who act out of racial prejudice. The first limitation is that I will be bracketing off good⁺ cases.

Agents can ϕ in a way that manifests a belief, even though their ϕ -ing does not

⁷Compare Davis (2003, 2005), Setiya (2011), Hieronymi (2011), and Olson and Svensson (2005).

manifest their treating what they believe as a normative reason. Actions performed by young children or certain non-human animals might fit that description, as might certain cases of akratic agency, the sorts of Freudian cases discussed by Velleman (2000), and the sorts of cases of wilful irrationality discussed by Stocker (1979) and Setiya (2007). The second restriction is that I will bracket off that broad class of cases in this paper. My focus is on the contrasting class of cases in which the agent ϕ s in a way that manifests their treating what they believe as a normative reason and, even more narrowly, on how to understand the relation between two sub-sets of that contrasting class: good cases and bad cases of rationally ϕ -ing.

2 Two Conceptions of Rationalising Explanation

To a first approximation, the Common Kind View says that (ψ) explanations are, in some relevant sense, *the same* as ($\neg\psi$) explanations. To a first approximation, the Disjunctive View denies this. To make the distinction between the two views more precise, we need to make it explicit in what sense the Common Kind Theorist asserts, but the Disjunctivist denies, that the two sets of explanations are the same. In the remainder of this section I'll attempt to resolve this interpretive issue. In §2.1, I draw the distinction between particular explanations and kinds of explanation. In §2.2 I apply the distinction to the case of rationalising explanation. In §2.3 I suggest that the Common Kind Theorist's sameness thesis is ambiguous in light of that distinction, and hence that the Common Kind View is correspondingly ambiguous. In §2.4 I identify the Disjunctive View with the denial of both versions of the Common Kind View and unpack what that denial amounts to.

2.1 Particular Explanations & Kinds of Explanation

I want to begin by bringing into focus the basic distinction between a kind of explanation and the particular explanations which exemplify it. Consider the following efficient-causal explanations:

- (i) The tree collapsed because it was hit by a strong wind.
- (ii) The left side of the man's face is numb because the dentist administered local anaesthetic to it.

There is a clear sense in which we have two *different explanations* here: each has different *explanantia* and *explananda*, after all. But there is also an intuitive sense in which

they are of the *same kind*. They both belong to the kind *efficient-causal explanation*. To a first approximation, how the relevant *explanans* makes the relevant *explanandum* intelligible is, in each case, by dint of picking out an entity which stands in an efficient-causal relation to the entity picked out by the *explanandum*.

Now consider the following pair of constitutive explanations:

(iii) The wall is hard because it is made out of bricks.

(iv) The statue weighs half a ton because the clay from which it is made weighs half a ton.

Again, there is a clear sense in which each of these explanations are *different explanations*. But there is also a clear sense in which they are of the *same kind*: they are both *constitutive explanations*. To a first approximation, how the *explanans* of each of (iii)-(iv) makes its corresponding *explanandum* intelligible is by picking out an entity which stands in a constitution relation to the entity picked out by its *explanandum*.

How should we cash-out this distinction between particular explanations and kinds of explanation? This is of course a large question in the theory of explanation which I cannot hope to tackle completely here. Instead, I provide a sketch of an answer which contains enough detail to enable me to proceed.

Let us begin with particular explanations. A particular explanation is best understood as a Fregean proposition: 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'. The proposition in question involves three components. The first component is another Fregean proposition with which we identify the *explanandum* of the explanation. It corresponds to 'p' in the linguistic schema just cited, and is a mode of presentation of the entity, or entities, that is the target of the explanation. The second component is a further Fregean proposition with which we identify the *explanans* of the explanation. It corresponds to 'q' in the linguistic schema just cited, and is a mode of presentation of the entity in terms of which we want to provide our explanation. The third element is a Fregean sense correspondent to the 'because' in the linguistic schema just cited. It is a mode of presentation of some explanatorily efficacious relation which is represented to hold between the entities corresponding to *explanandum* and *explanans*, be it, for example, an efficient-causal relation or a constitutive relation. Two particular explanations are distinct just in case the propositions which constitute their *explananda* or *explanantia* differ, or else because the *because* concepts which partly composes each differ.⁸

⁸Compare Strawson (1985).

Now let us move on to *kinds* of explanation. Kinds of explanations are individuated by different kinds of explanatorily efficacious relations such as the efficient-causal and constitution relations. Thus, two particular explanations are different in kind just in case the *because* which party composes one corresponds to a different kind of explanatorily efficacious relation to the *because* which party composes the other.

We can extend our understanding of what a kind of explanation is in the following way. For any given particular explanation, *E*, we can ask the following how-question:

(H_q) How does *Es explanans* function to explain *Es explanandum*?

At least part of the correct answer to (H_q) will involve a specification of the set of explanatorily active or efficacious properties exemplified by the entity to which *Es explanans* corresponds. In the case of (i), this set of properties would include the momentum of the air molecules which constitute the gust of wind, and their direction of travel. In the case of (iii), the set of properties would include the properties relating to the molecular structure of bricks, such that they ground the instantiation of the property of hardness in entities that are constituted by collections of bricks.

The kind to which *E* belongs partly determines what the correct answer is to (H_q), by restricting which set of properties of the entity to which *Es explanans* corresponds count as explanatorily efficacious. If *E* is efficient-causal, for example, then the properties in question will be restricted to those properties which ground the disposition the entity in question has to causally affect the entity corresponding to *Es explanandum* in the relevant way. If, by contrast, *E* is constitutive, then the explanatorily efficacious properties will be restricted to properties of the entity to which *Es explanans* corresponds which get to ground the instantiation of the relevant properties of the entity picked out by *Es explanandum*, given a constitution relation that holds between the two entities.

Every particular explanation falls under a certain kind. But two particular explanations, *E*₁ and *E*₂, might be distinct even though they are the same in kind. The pairs of explanations (i)-(ii) and (iii)-(iv) suffice to demonstrate this. In terms of the framework just sketched this would amount to the following. Since *E*₁ and *E*₂ are different particular explanations, they differ with respect to the identity of their *explananda*, *explanantia*, or the *because* which composes them. But since *E*₁ and *E*₂ are the same in kind, the *because* which composes each corresponds to the same kind of explanatorily efficacious relation. The upshot of this is that the same restriction on which class of properties count as explanatorily efficacious when it comes to successfully answering (H_q), asked of each explanation, applies to both.

2.2 The Distinction Applied to Rationalising Explanation

With the distinction between particular explanations and kinds of explanations on the table, I now want to map that distinction on to the sort of explanation which interests me here: rationalising explanation.

I'll begin by focusing my attention on (ψ) explanations. A particular (ψ) explanation, RE_ψ , consists in a Fregean proposition the *explanans* of which ascribes a state of believing that p to the agent, the *explanandum* of which ascribes a ϕ -ing to the agent, and the *because* concept of which is a mode of presentation of the kind of rational-motivation relation which binds the agent's ϕ -ing with their state of believing just in case a (ψ) explanation holds of them: a kind of relation which is instantiated neutrally, across both good cases and bad.

The kind to which RE_ψ belongs is individuated by the rational-motivation relation to which its *because* concept corresponds. When we ask (H_q) of RE_ψ , the successful answer to that question will cite the properties of the agent's belief that p which are explanatorily operative with respect to the agent's ϕ -ing. The kind of rational-motivation relation which binds the agent's state of belief with their ϕ -ing just in case a (ψ) explanation holds will place a restriction on which properties of the agent's belief are operative. How we think of the content of the restriction will depend on how we conceive of that sort of relation, for example: whether we think of it as efficient-causal.

But since the kind of rational-motivation relation associated with (ψ) explanations holds across both good cases and bad, we will surely have to say that part of the answer to (H_q) , asked of RE_ψ , is that the agent's belief explains *qua* constituting the *appearance* of the fact that p to agent, where p is what is treated as a normative reason. Moreover, even if the belief is also an article of knowledge possessed by the agent, we will have to say that it does not explain why the agent ϕ s *qua* knowledge. Its status as a state of *awareness* of the fact that p , if it does have that status, is not part of how the agent's state of belief that p gets to explain why the agent ϕ s, when a (ψ) explanation holds; only its status as an *appearance* of the fact that p gets to play that explanatory role.

Let me now turn my attention to $(\neg\psi)$ explanations. The framework I have developed here leaves us with the following three options. According to the first option, although (ψ) and $(\neg\psi)$ *statements* differ, the *particular explanations* they can be used to provide do not. According to the second option, particular $(\neg\psi)$ explanations are the same *in kind* as corresponding (ψ) explanations, but are not identical to them, for their *explanantia* differ. According to the third option, $(\neg\psi)$ explanations are indeed different in kind from (ψ) explanations: the *because* of each corresponds to a different kind of rational-motivation relation. I will say much more concerning these three ways of conceiving

$(\neg\psi)$ explanations momentarily, for I will be suggesting that the Common Kind View should be identified with either the first or the second option, and that the Disjunctive View should be identified with the third.

2.3 The Common Kind View

When the proponent of the Common Kind View says that $(\neg\psi)$ explanations are *the same* as corresponding (ψ) explanations it is possible to read this in either of two ways. On the one hand, what they have in mind could be that *particular* $(\neg\psi)$ explanations are identical to corresponding *particular* (ψ) explanations: the first option listed at the end of §2.2. Alternatively, it could be that particular $(\neg\psi)$ explanations are the same *in kind* as corresponding particular (ψ) explanations, even though the former are not identical to the latter: the second option listed at the end of §2.2.

The first view is stronger than the second because, like the second view, it is committed to the claim that (ψ) and $(\neg\psi)$ explanations are the same in kind, but it is also committed to something more: that the relevant particular explanations are identical. Let us codify these two versions of the Common Kind View:

Common Kind View_S For any particular $(\neg\psi)$ explanation, that explanation is identical to the corresponding particular (ψ) explanation.

Common Kind View_W For any particular $(\neg\psi)$ explanation, that explanation is the same in kind as the corresponding particular (ψ) explanation, but is not identical to it.

Let me now explain what each of these views amounts to, starting with Common Kind View_S. According to Common Kind View_S, $(\neg\psi)$ statements provide exactly the same particular explanations as their (ψ) counterparts. Thus, $(\neg\psi)$ explanations have the character I attributed to (ψ) explanations in §2.2. On Common Kind View_S, there will surely remain a difference between (ψ) and $(\neg\psi)$ *statements*. The proponent of Common Kind View_S is free to cash out this difference in whatever way they prefer, so long as it remains consistent with the idea that in so far as each statement provides a rationalising explanation, it is the very same (particular) explanation provided. They might, for example, say that the difference between the two sorts of statement is entirely pragmatic. They might, on the other hand, say that there is a semantic difference between the two: the truth of a $(\neg\psi)$ statement requires p to be the case, and for the agent to know that p . But if they opt for the latter view, they will have to say that *in so far as the two statements provide rationalising explanations*, there is no difference between them with respect to under what conditions they are true.

Now let me move on to Common Kind View_W. On this view, the *explanans* of a given ($\neg\psi$) explanation, RE $\neg\psi$, differs from the *explanans* of a corresponding (ψ) explanation. The former includes the fact that the agent knows that p, whereas the latter includes only the fact that the agent believes that p. ($\neg\psi$) statements provide different particular explanations to corresponding (ψ) explanations.

Nevertheless, on Common Kind View_W, the two explanations are the same *in kind*: the *because* which partly composes each corresponds to the very same kind of rational-motivation relation, namely: the kind of rational-motivation relation which holds neutrally, across both good cases and bad. A consequence of this is that the proponent of Common Kind View_W will have to think that in answering (H_q), asked of RE $\neg\psi$, the very same restriction applies as when answering (H_q) of the corresponding (ψ) explanation: the state of mind picked out by the *explanans* of the explanation functions to explain why the agent ϕ s *not* by dint of constituting the agent's *awareness* of a fact treated as a normative reason, but merely by dint of constituting the *appearance* of such a fact. Thus, according to Common Kind View_W, the agent's knowledge that p doesn't function to explain why the agent ϕ s *qua* knowledge of the fact that p, only *qua* constituting the appearance of that fact. But that the agent knows *is* included in the *explanans* of the explanation nevertheless.

Notice that both Common Kind Views agree about what is going on at the metaphysical level. They both say that there is only a single kind of rational-motivation relation, and it is the kind which is instantiated neutrally, binding the agent's ϕ -ing with a state of mind that's required only to constitute the *appearance* of a fact to them. Across both good cases and bad, the views agree that the agent's ϕ -ing is linked by the rational-motivation relation to their state of believing that p: that is the state which plays the role of constituting the appearance of the fact that p to them at the neutral level. And about the good case, the views agree that an instance of that very kind relation binds the agent's ϕ -ing with a state of knowledge that p: the state which constitutes the appearance of the fact that p to them, as well as being a state of believing, is a state of knowledge of the fact treated as a normative reason in the good case. The disagreement arises concerning what each view says is going on at the explanatory level: Common Kind View_W accepts, whereas Common Kind View_S denies, that that the agent knows that p is part of the *explanans* of the ($\neg\psi$) explanation. Common Kind View_W goes on to insist, however, that the agent's knowledge that p in the good case does not function to explain why the agent ϕ s *qua* knowing, only *qua* believing.

Finally, notice also that both views ascribe a *composite structure* to the good case. According to both, we get to the good case by starting out with a motivational structure that's neutral: a rational-motivation relation which binds the agent's ϕ -ing with their

state of believing that p , and which is of a kind which is instantiated across both good cases and bad. We treat this motivational structure as basic, and then account for the good case by appending to that neutral motivational structure certain worldly and epistemic facts: we add that the state of belief in question is a state of knowing, and hence that what's believed is a fact. ϕ -ing in the light of the fact that p is a matter of standing in the same sort of motivational relation to p as one does to one's apparent reason in the bad case, but where there are, in addition to that relation holding, certain worldly and epistemic conditions which hold. Those conditions are *additional* to the motivational structure in the sense that the motivational structure does not require the obtaining of those further worldly and epistemic conditions, and nor could it be grounded in the obtaining of some other motivational structure which does require them to obtain.

2.4 The Disjunctive View

The Disjunctivist denies that (ψ) and $(\neg\psi)$ explanations are the same, in whatever way the Common Kind View wishes to insist that that's so, so they will have to deny both Common Kind Views. To do this, they will have to deny not just that corresponding instances of (ψ) and $(\neg\psi)$ provide the same particular explanation, but that they provide the same kind of explanation: they will have to accept the third view presented at the end of §2.2.

But the Disjunctivist doesn't simply deny the sameness of kind claim. They also have something positive to say about the kind of explanation available in the good case. I will identify Disjunctivism with the following thesis:

The Disjunctive View In the good case, the rationalising explanation to which S 's ϕ -ing is subject is of a kind which requires p to be fact, known by S .

According to the Disjunctive View, $RE_{\neg\psi}$ has a different *explanans* to its (ψ) counterpart: the former includes the fact that the agent knows that p , whereas the latter includes only that the agent believes that p . But unlike on Common Kind View_W, the Disjunctivist says that $RE_{\neg\psi}$ is different *in kind* from its (ψ) counterpart: the *because* which partly composes $RE_{\neg\psi}$ picks out a different kind of rational-motivation relation to the relation picked out by the *because* which partly composes its (ψ) counterpart. Let me now unpack this.

According to the Disjunctivist, we should be *pluralists* about the rational-motivation relation. First, we should say that there is a kind of rational-motivation which binds the agent's ϕ -ing with their state of believing across both good cases and bad. That kind of rational-motivation relation holds just in case a (ψ) explanation holds. But second, we

should also say that there is a kind of rational-motivation relation which binds the agent's ϕ -ing with their belief that p in the good case alone. This second rational-motivation relation is distinct in kind from the first because necessarily, it holds between the agent's ϕ -ing and their state of believing *if, and only if*, the agent is in the good case. Hence, the second kind of rational-motivation *requires* for its instantiation that what the agent believes is true, and that they know it.

On the Disjunctive View, it is this second kind of rational-motivation relation, with its essentially factive and knowledge-involving character, which holds just in case a $(\neg\psi)$ explanation holds, and the *because* which partly composes the relevant $(\neg\psi)$ explanation is a mode of presentation of that relation. The second kind of rational-motivation relation individuates for us the kind to which $RE_{\neg\psi}$ belongs. The Disjunctivist wishes to say that the correct answer to (H_q) , asked of $RE_{\neg\psi}$, will in part be that the agent's state of knowing that p explains *qua* state of awareness of the fact that p , and not merely by dint of constituting the appearance of the relevant fact: it explains *qua* knowledge, and not merely *qua* belief. This thought separates the Disjunctive View from Common Kind View_W.

Earlier, I attributed to the Common Kind View the thought that the good case exemplifies a composite structure: to account for the good case, we take a motivational structure that's neutral, and then add extra worldly and epistemic features to it. The Disjunctive View precisely rejects this composite conception. According to the Disjunctive View, there is a motivational structure particular to the good case: a kind of relation that holds between one's act or attitude and the psychological state which constitutes the appearance of a fact, which can hold *only if* the state of mind is not merely the *appearance* of a fact (namely: a belief), but an *awareness* of it (namely: a state of knowing).

To be clear, the Disjunctivist does not deny that the neutral kind of rational-motivation relation holds between the agent's ϕ -ing and their belief that p , so that a (ψ) explanation holds of them, in the good case. Rather, they will say that also, a second kind of rational-motivation relation holds between the agent's ϕ -ing and their belief that p in the good case, so that a $(\neg\psi)$ explanation holds too: one which is essentially factive and knowledge-involving in character. Moreover, I take it that it would be implausible to suggest that the instantiation of both of these relations in the good case is a coincidence, from the Disjunctivist's point of view. Rather, they will say that the obtaining of the essentially factive, knowledge-involving relation grounds the obtaining of the neutral rational-motivation relation. How such grounding works will fall out of whatever positive account the Disjunctivist offers us of the former relation, and the associated account they offer of the latter relation. I leave the question of what these accounts should look like for another occasion.

3 Two Direct Arguments for the Disjunctive View Rejected

Having arrived at the correct formulation of the distinction between the Common Kind and Disjunctive Views, the question naturally arises: what hangs on the debate between them? There is much that can be said in answer to this question. But one thing to say immediately is this. Plausibly, ϕ -ing *rationally* partly consists in what I have been calling the rational-motivation holding between one's ϕ -ing and a state of mind which constitutes the appearance of a fact one treats as a reason to ϕ . If the Disjunctive View is correct, there is a kind of rational-motivation relation which can be instantiated at all only if the state of mind at issue constitutes awareness of the relevant fact. This puts the Disjunctivist in a position to argue that in the good case, ϕ -ing rationally constitutively involves awareness of a fact, typically about the external world: a significant and interesting result about the nature of rationality.

Given that it would be an interesting result were the Disjunctive View true, this raises the question of what can be said in its favour. In the contemporary literature, arguments for the Disjunctive View are few in number. This section outlines what I regard as the only arguments explicitly advertised by their proponents as arguments for the Disjunctive View available in the literature: an argument found in the work of Hornsby (2007a, 2007b, 2008), endorsed by McDowell (2013), which I discuss in §3.1, and an argument from Roessler (2014), which I discuss in §3.2. I suggest in each case that the argument only succeeds in refuting Common Kind View_S, and leaves open the truth of Common Kind View_W. Hornsby and Roessler might respond by suggesting that the latter is not a coherent position. In §3.3 I attempt to deal with that response.

3.1 Hornsby's Argument

According to Hornsby, we should accept (RKT): knowing that p is necessary for ϕ -ing because p . It is then inferred that the explanation present in the good case is different from the explanation present in the bad case, for the former requires that the agent knows that p but the latter doesn't, and hence that the Disjunctive View is true.

I have already committed myself to the truth of Hornsby's (RKT). What interests me is not that epistemic claim, but Hornsby's inference from it to the truth of the Disjunctive View. Is this inference valid? I suggest not, in light of the distinction between the two ways of understanding the Common Kind View outlined in §2.3. What the proponent of Common Kind View_W could say in response is that knowing that p is indeed required for the truth of the rationalising explanation to which S 's ϕ -ing is subject in the good case:

they would suggest that the particular explanation present in the good case has included in its *explanans* the fact that the agent knows that p, so that that very explanation is not present in the bad case. But, they would continue, this isn't a requirement of the *kind* of explanation present in the good case: the rational-motivation relation which binds the agent's ϕ -ing with their state of knowing is a relation that only requires the *appearance* of a fact.

3.2 Roessler's Argument

According to Roessler (2014: 4-7), we are pre-philosophically committed to the claim that when we ϕ in response to the fact that p, the fact in question is part of what explains why we ϕ . Roessler's argument for this is that when one engages in an episode of theoretical reasoning or practical deliberation which results in one's ϕ -ing, the aim of what one is doing, from one's own point of view, is to get the normatively significant facts about the situation right, and then to ϕ in a way that is determined or controlled by one's fulfilling that aim. From the reasoner's own point of view, then, the resultant ϕ -ing will count as being the upshot of getting what normative reasons there are to ϕ right, and thus as a ϕ -ing which is (at least in part) explained by the existence of facts treated as normative reasons by the agent. Since the proponent of the Common Kind View cannot say that a fact treated as a normative reason is part of the *explanans* in the bad case, and they are committed to saying that the explanation in the bad case is the same as that in the good case, they are, thinks Roessler, committed to denying that the *explanans* in the good case could ever include the fact the agent treats as a normative reason too. It is supposed to follow that the Disjunctive View fits best with what agents are pre-philosophically committed to.

I don't want to quibble with Roessler's argument for the claim that we are pre-philosophically committed to the thought that facts treated as normative reasons can form part of what explains why we ϕ . The problem is that this doesn't add up to a problem for the Common Kind View. That's because the proponent of Common Kind View_W could agree that the fact that p is (part of) the *explanans* in the good case, so that the explanation present in the good case is not identical to the explanation present in the bad case, whilst denying that this is part of the *kind* of explanation in question. All that's required, they will say, by the rational-motivation relation which holds in the good case is that there be an apparent fact on the scene, but, consistently with this, a *bone fide* fact can correctly be thought of as forming a part of what explains why the agent ϕ s.⁹

⁹§3 of Roessler's paper might be read as a second and separate argument for the Disjunctive View, and

3.3 Is Common Kind View_W Coherent?

Central to this objection to Hornsby and Roessler is the availability of Common Kind View_W. They might respond by pressing the worry that Common Kind View_W is not a coherent position. And they might argue for this by appeal to the following line of thought. Recall that if Common Kind View_W is true, then how it is that the agent's belief that p in the bad case gets to explain why the agent ϕ s is the same as how the agent's knowledge that p gets to do so in the good case: they both do so by dint of constituting the *appearance* of the fact that p to the agent. But if that's right, then Common Kind View_W implies that the fact that p , and the agent's knowledge of it, are not *explanatorily efficacious* factors with respect to the agent's ϕ -ing in the good case: its *appearing to the agent as if p* is always what does the explanatory work with respect to their ϕ -ing in such cases. But then it looks as if there is no sense in which the worldly and epistemic conditions present in the good case could form part of the *explanans* of the relevant ($\neg\psi$) explanation in the first place. If this is right, then the distinction between our two Common Kind Views collapses.

It seems to me that there is plenty the proponent of Common Kind View_W can say in reply here. First, we do not in general find explanations which appeal to factors that are not explanatorily efficacious problematic. To take some of Jackson and Pettit's (1990) examples, we can explain why the glass broke by appeal to its fragility, even though the efficacious feature is not its fragility, but its particular molecular structure. Likewise, the temperature of the water explains why the glass containing it cracked, but the efficacious feature here is the momentum of the relevant set of the water's molecules. Moreover, we have a model of how an explanation which appeals to non-efficacious factors can function to explain why something is the case: the *program explanations* which Jackson and Pettit (1990) explore. Program explanations are explanations which operate by appealing to properties that necessitate – program for – the existence of further properties which are explanatorily efficacious. The former get to be part of the *explanans* precisely because they program for properties that are efficacious. The examples just given are plausibly construed as cases of program explanation. The proponent of Common Kind View_W is free to say that ($\neg\psi$) explanations are themselves a variety of program explanation: they explain why the agent ϕ s by appeal to fact-involving conditions that necessitate – program for – some explanatorily efficacious property, in this case: that it appears to the agent as if p . Of course, even if the notion of a program explanation

indeed one that targets Common Kind View_W. Whatever the correct interpretation, the argument presupposes Kolodny (2005) and Raz's (2011) controversial claim that there are no subjective oughts. Whilst I am sympathetic to that idea, it is likely that the proponent of the Common Kind View will want to reject it, so that presupposing it as part of an argument for the Disjunctive View is not dialectically effective.

turns out to be problematic, still: we do in general acknowledge explanations which do not appeal to efficacious factors, and it's not obvious why the proponent of Common Kind View_W couldn't simply help themselves to whatever model we decide to give of that explanatory phenomenon.

Second, providing rationalising explanations in terms of beliefs is arguably not to provide an explanation which appeals to an explanatorily efficacious feature: arguably, there is some neural state the agent is in, on which their belief supervenes, which is the explanatorily efficacious feature here. But since Hornsby and Roessler would allow, as indeed is highly plausible, that beliefs can be the *explanantia* of rationalising explanations, they will have to acknowledge that non-efficacious factors can nevertheless do explanatory work, in which case the objection is undermined.

At this point, Hornsby and Roessler might protest further that, intelligible though Common Kind View_W might be, the view is nevertheless false. In particular, they might argue that no explanation of the character Common Kind View_W ascribes to ($\neg\psi$) explanations should be admitted to be possible anyway, because there is no context in which citing one of them could count as providing a *better* explanation of why the agent ϕ s than a corresponding (ψ) explanation.

Again, I think there is a reply available here. Take the explanation: *he intends to vote because there is a by-election coming up*. Contrast this with the corresponding neutral explanation: *he intends to vote because he believes there is a by-election coming up*. It seems to me that citing the former, at least in many ordinary contexts, is more informative than citing the latter, in a number of ways, and that gives us a sense in which it is a better explanation, in those contexts. First of all, citing the former reveals that there is a fact on the scene, which the agent knows. This goes some way to demonstrating that the agent is warranted in ϕ -ing.¹⁰ Moreover, citing the former reveals that there is a dependency of some kind of the agent's intention on a certain fact: there is a non-accidental fit between the agent's practical attitude and a fact which is taken to favour maintaining it. No such dependence of ϕ -ing on the facts treated as reasons is made explicit insofar as a belief of the agent's is cited, only a dependence on a certain narrowly supervening state of mind. These points demonstrate that in ordinary contexts, ($\neg\psi$) explanations are better because more informative than (ψ) explanations, and the proponent of Common Kind View_W is perfectly well in a position to agree with all this.

¹⁰Compare Davidson (1963: 8) and Fogal (2018: 20).

4 A New Argument for the Disjunctive View

What's needed, if we're to motivate the adoption of the Disjunctive View, is an argument that's effective against Common Kind View_S and Common Kind View_W. This is what I want to provide in this closing section. My strategy is to attack directly the claim that we should ascribe what I earlier called a *composite structure* to the good case; an idea which is common to both Common Kind Views. To this end, I want to present a case in which the agent knows that p , they ϕ because they believe that p , but in which they do not ϕ because p . In order to explain why the agent fails to be in the good case, I will argue that we need to say that the Disjunctivist's essentially factive and knowledge-involving rational-motivation relation is missing.

My verdict about the relevant case relies on certain principles concerning conscious reasoning and knowledge. Those principles are specified, and supplied with some defence, in §4.1. In §4.2 I present the supposedly problematic case, and argue in §4.3 that it's possibility lends support to the Disjunctive View in the way advertised.

4.1 Knowledge, Good Cases, and Conscious Reasoning

There are two principles on which I will rely are these:

- (i) Knowing that p is consistent with there being a modally close-by circumstance in which the agent truly believes that p but where the method by which they arrived at their belief has changed such that they no longer count as knowing that p .
- (ii) ϕ -ing because p requires the possession of the ability to consciously reason from the known fact that p , to the conclusion to ϕ .

Let me examine each claim in turn, starting with (i). The plausibility of this thesis can be elicited by appeal to examples. I might know that the ice in the middle of the pond is too thin for skating because I've just seen it crack under the weight of another skater. But there is a close-by circumstance in which I do not see this happen and instead ask my trusted friend, who I know has already skated on the pond, whether the ice in the middle is too thin. My friend has no idea how thin the ice in the middle is, because they never dared venture away from the edge. But, because they want me to stick with them at the edge, they tell me that the ice in the middle is indeed too thin, and I come to truly believe that it is on this basis. In this close-by circumstance in which my doxastic method has changed I do not know that the ice in the middle is too thin for, on this occasion, my friend is untrustworthy. Still, this is compatible with me knowing that the ice is too thin at the actual world, at which I have seen its thinness manifest itself.

Let's now move on to (ii), and let me immediately make a clarificatory point concerning it. Consider two seasoned athletes specialising in the 100m sprint. Let us suppose that the second is suffering from temporary severe numbness in their right leg, bought about as a result of their having slept on it at an awkward angle. There is a sense in which both athletes *can* sprint 100m, and also a sense in which the second *cannot*. We can capture what is going on here by following many in the literature on abilities in drawing the distinction between having the *general ability* to ϕ and *having the option* to ϕ .¹¹ Both athletes have the general ability to sprint 100m: they both have the “motor skill, competence, or know-how” (Clarke, 2015: 893) necessary to sprint 100m. But the second athlete doesn't have the option to do so: they cannot do so *here-and-now*. (ii) should be read as claiming that ϕ -ing because p requires the *general ability* to consciously reason from the known fact that p, to the conclusion to ϕ , *not* as the claim that ϕ -ing because p requires one to have the option to engage in such reasoning. So understood, (ii) is entirely consistent with the possibility of one ϕ -ing because p and yet finding oneself in a situation in which one is suffering from some temporary psychological malady which disables one from carrying out the relevant inference, such as a severe migraine or concussion.

The plausibility of (ii) can be best brought out by beginning with the following intuitively plausible claim: ϕ -ing because p requires that one can be held responsible for one's ϕ -ing – that one can be blamed if one's ϕ -ing is not as it should be, and credited if it is. Plausibly, ϕ -ing responsibly requires that one has the general ability to engage in action which constitutes exercising control over whether one ϕ s. Equally plausibly, this control is exercised in the form of conscious reasoning about whether to ϕ . So: ϕ -ing because p requires a general ability to consciously reason about whether to ϕ . However, if one ϕ s *because p*, this will presumably have an effect on the character of the conscious reasoning one has the general ability to carry out: one will be disposed, during that reasoning, to draw an inference from p to the conclusion to ϕ . And, given (RKT), it's plausible that this reasoning will have the character of reasoning from one's knowledge of the fact that p, to the conclusion to ϕ . This gets us to thesis (ii).

4.2 A Case of Rational Incapacitation

I am now in a position to present the case which I think causes a problem for those wishing to ascribe a composite structure to the good case. Consider first of all the following standard good case:

Wimbledon Fanatic I. Thelma is a huge fan of the Wimbledon tennis cham-

¹¹My terminology here is borrowed from Maier (2015).

pionship. This year, the tournament is set to take place on the 2nd–15th of July, and Thelma knows this because she read it on the official Wimbledon website. That Wimbledon will take place on the 2nd–15th of July is a fact Thelma treats as a reason to postpone visiting her parents until after the 15th of July. Indeed, Thelma holds an intention to do just that in light of the fact in question.

With respect to her intention to postpone visiting her parents until after the 15th of July, Thelma is in the good case.

Now consider the following scenario, which takes place at a point in time after Thelma has formed the relevant intention:

Wimbledon Fanatic II. Unbeknownst to Thelma, she underwent a piece of brain surgery whilst she slept yesterday evening which has the following upshot: whenever she engages in an episode of conscious reasoning directed towards answering the question: *When should I visit my parents this summer?*, this triggers an apparent memory of a news reader telling her that Wimbledon will be on the 2nd–15th of July this year and causes her to forget reading it on the official website. This memory happens to be correct: Wimbledon will in fact take place on those dates this year, and so she would continue to believe truly were she to engage in such reasoning. But the memory is unreliable: even if Wimbledon were to take place on different dates, that apparent memory would still be triggered when she engages in that deliberation. Still, having already settled the matter of when to visit her parents, Thelma never consciously raises the question again, and so her new psycho-neural disposition never gets triggered.

In *Wimbledon Fanatic II*, Thelma continues to hold her intention because she *believes* that Wimbledon will take place on the 2nd–15th of July: her intention is subject to a (ψ) explanation. That, I think, is uncontroversial. But there are two further claims I want to defend about the case: that Thelma continues to know that Wimbledon will take place on the 2nd–15th of July and that Thelma's intention is no longer subject to a ($\neg\psi$) explanation – she ceases to be in the good case, with respect to the intention at issue.

First, there's the claim that Thelma continues to know that Wimbledon will take place on the the 2nd–15th of July. This could be denied, it seems to me, only if principle (i) is rejected. With (i) rejected, one could argue that Thelma doesn't know that Wimbledon will take place on the 2nd–15th of July by appeal to the thought that there is a modally close-by circumstance, one in which the neural disposition is triggered,

in which she ceases to know but continues to believe relative to an unreliable method, namely the unreliable memory. But I have already argued for (i).

Next, there's the claim that Thelma no longer counts as intending to postpone her visit *because* Wimbledon will take place on the 2nd–15th of July. This follows from (ii), defended in §4.1, in conjunction with the thought that Thelma has lost her general ability to consciously reason from the known fact that Wimbledon will take place on the 2nd–15th of July to the conclusion to postpone her visit. Having already argued for the former, I will now argue for the latter.

Let us pause to reflect, first, on what it takes to possess a general ability to ϕ , and let us take as our example, again, the general ability to sprint 100m. Plausibly, whether an agent counts as possessing such a general ability depends partly on their neuro-physical constitution. Suppose, for example, that the athlete with the severely numb right leg has sustained a brain injury so that they suffer from the severe numbness permanently. We would not want to say, in that case, that the athlete continues to possess the general ability to sprint 100m, but merely fails to have the option to do so. We would not want to say this any more than we would want to say that a patient suffering from right optic nerve damage causing blindness in their right eye continues to possess the general ability to see with their right eye, but merely fails to have the option to do so. Rather, we would want to say that the athlete ceases to so much as possess the general ability to sprint 100m.

In *Wimbledon Fanatic II*, there is an aspect of Thelma's neuro-physical constitution that precludes her from reasoning from her knowledge that Wimbledon will take place on the 2nd–15th to the conclusion to postpone her visit: given her neuro-physical constitution, were she to consciously raise the question of when to visit her parents, she would cease to *know* that Wimbledon will take place on the 2nd–15th, after all. It seems to me that Thelma's situation in *Wimbledon Fanatic II* is saliently analogous to the case of the athlete suffering from brain injury causing permanent severe numbness. As a result, we should say that Thelma fails to possess the general ability to reason from her knowledge to the conclusion to postpone her trip, not merely that she fails to have the option to do so.

It might be objected that my verdict about Thelma presupposes an implausibly fine-grained conception of general abilities. It is implausible, so this objection goes, to say that the general ability I have characterised as an ability to consciously reason from knowledge that Wimbledon will take place on the 2nd–15th of July to the conclusion to postpone visiting one's parents is really distinct from an ability characterised in more general terms, such as, for example, the ability to reason about the timing of tennis matches, or, even more generally, the ability to consciously reason about the dates on

which future entertaining events will occur. But, whatever more general ability we opt for, presumably Thelma continues to possess *that* ability.

Suppose we agree to a coarse individuation of the general ability to consciously reason from knowledge that p to the conclusion to ϕ . My reply to the objection is to point out that that would only mean that we should no longer read (ii) as the claim that ϕ -ing because p requires the general ability to consciously reason from the known fact that p to the conclusion to ϕ *simpliciter*. Rather, we should read it as the claim that ϕ -ing because p requires possession of that general ability *and* possession of a more specific second-order general ability to realise the first general ability in the specific way at issue: with the particular values for p and ϕ filled-in appropriately. Since I have already effectively argued that Thelma fails to satisfy that conjunctive condition, interpreting (ii) in this way does not undermine the argument just presented for the claim that Thelma is not in the good case. Moreover, interpreting (ii) in this way would not undermine the soundness of the argument I presented for it in §4.2.

4.3 Against the Composite Conception of the Good Case

What all of this gets us is the result that Thelma, in *Wimbledon Fanatic II*, does not intend to postpone visiting her parents *because* Wimbledon will take place from 2nd–15th of July. This is so, even though she knows that Wimbledon will take place from 2nd–15th of July and she holds her intention because she *believes* that Wimbledon will take place from 2nd–15th of July. Now, proponents of the Common Kind Views ascribe a composite structure to the good case: they say that we should conceive of good cases by starting out with a motivational structure common to good cases and bad – a ϕ -ing, a belief that p , and a neutral kind of rational-motivation relation binding the two – and then add independent worldly and epistemic factors to that structure – the fact that p , and the agent’s knowledge of it. The problem is that if this is the correct way to understand the good case, then Thelma appears to have everything she needs to be in it.

The Disjunctivist eschews ascribing a composite structure to the good case in favour of thinking that there is a kind of rational-motivation relation which can hold between the agent’s ϕ -ing and their belief that p just in case they ϕ because p . They are thereby in a ready position to explain why Thelma fails to ϕ because p : they can simply say that instantiating the essentially factive and knowledge-involving kind of rational-motivation relation requires that one possesses the general ability to consciously reason from the known fact p , to the conclusion to ϕ .

Why can’t the Common Kind Theorist respond to this argument by saying that possession of the relevant ability is a further independent factor, to be added to the mo-

tivational structure common to both good cases and bad, in accounting for what ϕ -ing because p consists in, along with p 's being a fact and the agent knowing that p ? They can then acknowledge that Thelma isn't in the good case by simply pointing out that she fails to possess the relevant general ability.

My reply to this is that the Common Kind Theorist is not in fact in a position to deny that Thelma fails to possess the relevant ability in the first place. Given that the Common Kind Theorist wants to say that ϕ -ing because p factors into ϕ -ing because one believes that p , plus p 's being true and one knowing that p , it would be *ad hoc* for them to resist making an analogous move with respect to the general ability to reason from one's knowledge that p to the conclusion to ϕ . It would be *ad hoc*, that is, for them to deny that possessing that ability consists in possessing the ability to reason from one's *belief* that p to the conclusion to ϕ , plus p 's being true and one's knowing that p . But if that's so, then they are not in a position to so much as deny that Thelma fails to possess the general ability to consciously reason from one's knowledge that p to the conclusion to ϕ in the first place. My argument for the Disjunctive View therefore survives this objection.

References

- Clarke, Randolph (2015). "Abilities to Act". *Philosophy Compass*, 12 (10): 893–904.
- Cunningham, J.J. (Unpublished). "Knowledgeably Responding to Reasons". Unpublished Manuscript. URL: https://drive.google.com/file/d/0B7if-X8oAj_vczdEU1pUQ08tQkE/view.
- Dancy, Jonathan (2000). *Practical Reality*. Oxford: Oxford University Press.
- (2004). "Two Ways of Explaining Actions". In: *Agency and Action*. Ed. by John Hyman & Helen Steward. Cambridge: Cambridge University Press: 25–42.
- (2008). "On How to Act – Disjunctively". In: *Disjunctivism: Perception, Action, Knowledge*. Ed. by Adrian Haddock & Fiona Macpherson. Oxford: Oxford University Press: 262–279.
- (2011). "Acting in Ignorance". *Frontiers of Philosophy in China*, 6 (3): 345–357.
- (2014). "On Knowing One's Reasons". In: *Epistemic Norms*. Ed. by JJJohn Turri & Clayton Littlejohn. Oxford: Oxford University Press: 81–96.
- Davidson, Donald (1963). "Actions, Reasons, and Causes". *Journal of Philosophy*, 60 (23): 3–19. Reprinted in Davidson (2001): 3–20. All page references are to the reprint.
- (2001). *Essays on Actions and Events*. Oxford: Oxford University Press.

- Davis, Wayne A. (2003). "Psychologism and Humeanism". *Philosophy and Phenomenological Research*, 67 (2): 452–459.
- (2005). "Reasons and Psychological Causes". *Philosophical Studies*, 122 (1): 51–101.
- Fogal, Daniel (2017). "Reasons, Reason, and Context". In: *Weighing Reasons*. Ed. by Errol Lord and Barry Maguire. Oxford: Oxford University Press.
- (2018). "Deflationary Pluralism about Motivating Reasons". In: *The Factive Turn in Epistemology*. Ed. by Veli Mitova. Cambridge: Cambridge University Press.
- Hieronymi, Pamela (2011). "Reasons for Action". *Proceedings of the Aristotelian Society*, 111 (3): 407–427.
- Hornsby, Jennifer (2007a). "Knowledge, Belief and Reasons for Acting". In: *Explaining the Mental*. Ed. by Carlo Penco, Michael Beany and Massimiliano Vignolo. Newcastle: Cambridge Scholars Publishing: 88–105.
- (2007b). "Knowledge in Action". In: *Action in Context*. Ed. by Anton Leist. Berlin: Walter de Gruyter: 285–302.
- (2008). "A Disjunctive Conception of Acting for a Reason". In: *Disjunctivism: Perception, Action, Knowledge*. Ed. by Adrian Haddock and Fiona Macpherson. Oxford: Oxford University Press: 262–279.
- Hyman, John (1999). "How Knowledge Works". *Philosophical Quarterly*, 50 (197): 433–451.
- (2011). "Acting for Reasons: Reply to Dancy". *Frontiers of Philosophy in China*, 6 (3): 358–368.
- (2015). *Action, Knowledge, and Will*. Oxford: Oxford University Press.
- Jackson, Frank and Pettit, Philip (1990). "Program Explanation: A General Perspective". *Analysis*, 50 (2): 107–117.
- Kolodny, Niko (2005). "Why be Rational?" *Mind*, 114 (455): 509–563.
- Lewis, David (1986). "Causal Explanation". In: *Philosophical Papers*. Vol. II. Oxford: Oxford University Press: 214–240.
- Maier, John (2015). "The Agentive Modalities". *Philosophy and Phenomenological Research*, 90 (1): 113–134.
- Marcus, Eric (2012). *Rational Causation*. Camb. Mass.: Harvard University Press.
- McDowell, John (2013). "Acting in the Light of a Fact". In: *Thinking About Reasons*. Ed. by David Bankhurst, Brad Hooker and Margaret Olivia Little. Oxford: Oxford University Press: 13–28.
- Olson, Jonas and Svensson, Frans (2005). "Regimenting Reasons". *Theoria*, 61 (3): 203–214.

- Raz, Joseph (2011). *From Normativity to Responsibility*. Oxford: Oxford University Press.
- Roessler, Johannes (2014). “Reason Explanation and the Second-Person Perspective”. *Philosophical Explorations*, 3: 1–12.
- Schroeder, Mark (2007). *Slaves of the Passions*. Oxford: Oxford University Press.
- (2009). “Means-end Coherence, Stringency, and Subjective Reasons”. *Philosophical Studies*, 2 (143): 223–248.
- Setiya, Keiran (2011). “Reasons and Causes”. *European Journal of Philosophy*, 19 (1): 129–157.
- Setiya, Kieran (2007). *Reasons Without Rationalism*. Cambridge, Mass.: Princeton University Press.
- Stocker, Michael (1979). “Desiring the Bad: An Essay in Moral Psychology”. *Journal of Philosophy*, 12 (76): 738–758.
- Strawson, P.F. (1985). “Causation and Explanation”. In: *Essays on Davidson: Actions and Events*. Ed. by Bruce Vermazen and Merrill B. Hintikka. Oxford: Oxford University Press: 115–135. Reprinted in Strawson (1992):109-131. All page references are to the reprint.
- (1992). *Analysis and Metaphysics*. Oxford: Oxford University Press.
- Swanson, Eric (2010). “Lessons from the Context Sensitivity of Causal Talk”. *Journal of Philosophy*, 5 (107).
- Velleman, J. David (2000). *The Possibility of Practical Reason*. The University of Michigan Library.
- Vogelstein, Eric (2012). “Subjective Reasons”. *Ethical Theory and Moral Practice*, 2 (15): 239–257.