

Knowledge, Explanation, and Motivating Reasons

Dustin Locke

Claremont McKenna College

Penultimate Draft. To be published in *American Philosophical Quarterly*. Please cite published version.

Abstract

According to a number of recent philosophers, knowledge has an intimate relationship with rationality. Some philosophers hold, in particular, that rational agents do things for good motivating reasons, and that p can be one's motivating reason for ϕ -ing (acting/believing/fearing/etc.) only if one *knows* that p . This paper argues against this view and in favor of the view that p cannot be one's motivating reason for ϕ -ing—in the relevant sense—unless there is an appropriate *explanatory connection* between the fact that p and one's ϕ -ing. I argue that this view offers a better account of the cases alleged to support the knowledge view.

Keywords. Epistemology, Knowledge, Explanation, Reasons, Rationality

1. Introduction

In (1975), Peter Unger argued that knowledge is a necessary condition on having something as one's motivating reason. We'll call this

The Knowledge View. S's motivating reason is that p only if S knows that p .¹

More recent authors have also endorsed the knowledge view, provided that the phrase 'motivating reason' is properly interpreted. John Hyman (2011) suggests that the phrase 'motivating reason' can either mean that the agent *is guided by a certain fact* or that the agent *relies on a certain premise in her reasoning*. Call the former the 'factual guidance interpretation'. Provided that the factual guidance interpretation is at issue, both Hyman (1999, 2011) and Jennifer Hornsby (2007a, 2007b) also endorse the knowledge view.

The knowledge view has plausible connections to broader issues in the theory of rationality. Unger (1975), for example, has used the knowledge view in service of an argument that global skepticism—the view that we know nothing—leads to global irrationalism—the view that we never act/believe/assert/etc rationally. But it is plausible that the knowledge view of motivating reasons has implications for the relevance of knowledge to rationality even if global skepticism is false. Suppose, for example, that the only good motivating reason S could have to ϕ would be that p , but that S does not know that p .² Given that S is rational in ϕ -ing only if S has good motivating reasons for ϕ -ing, it follows from the knowledge view that S cannot ϕ rationally unless S comes to know that p .

Aside from some brief remarks in section four, this paper leaves aside questions of exactly what the knowledge view implies for the role of knowledge in rationality.³ Instead, the paper is primarily concerned with arguing against the knowledge view of motivating reasons and in favor of an alternative. In short, the alternative view is that to have p as one's motivating reason (in the relevant sense), one must stand in the right kind of *explanatory relation* to the fact that p .

The heart of the argument against the knowledge view must wait until section five. Before then, we need to do a fair bit of groundwork. In section two I distinguish *motivating reasons* from *normative reasons* and *explanatory reasons*. In section three I consider the standing dispute over whether attributions of motivating reasons are factive—that is, whether 'S motivating reason is that p ' entails p . Here I suggest that there are two distinct notions of having p as one's motivating reason, one of which is factive while the other is not. This kind of

¹ The knowledge view of motivating reasons should not be confused with the related principles defended by John Hawthorne and Jason Stanley (2008) and Jeremy Fantl and Matthew McGrath (2009). See fn. 9.

² Here and henceforth, ' ϕ -ing' ranges over anything that an agent can do for a motivating reason: acting, believing, fearing, regretting, etc. Similarly, when I speak of the agent 'doing something' or 'deciding (to do something)', I mean to use these phrases in this very broad sense.

³ Prominent skeptics of the idea that knowledge has any important connection to rationality include Mark Kaplan (1985, 2003) and Jonathan Kvanvig (2003).

resolution to the debate over the alleged factivity of reasons attributions can feel a little unsatisfying: surely, one might think, there is *something* of substance at issue in the debate over factivity. I agree, and in sections three and four I try to explain just what this substantive issue is.

In section five I take direct aim at the knowledge view. As we will see, Unger, Hyman, and Hornsby's arguments for the knowledge view all rest on consideration of a certain class of Gettier cases. Consideration of another class of Gettier cases, however, reveals that their diagnosis is mistaken: in these other cases, the agents do not have knowledge of the relevant fact that *p*, and yet they *do* have *p* as their reason (in the relevant sense). I then show that Unger, Hyman, and Hornsby's Gettier cases can be explained without appeal to the knowledge view—they can be explained by appeal to what I call 'the explanatory view'. Sections six and seven offer further developments to the explanatory view. Section seven also considers some potential objections.

2. Explanatory vs. Normative vs. Motivating Reasons

Philosophers typically draw a two-way distinction between normative and motivating reasons. Some philosophers also distinguish both of these from what I will call 'explanatory reasons'.⁴ When one gives an explanation of the behavior of some inanimate object, one gives an explanatory reason. In many cases, to give an explanatory reason is merely to state a cause. For example, (1) offers an explanatory reason why the billiard ball moved.

1. The reason the ball moved was that it was struck by another ball.

We can also give explanatory reasons with respect to agents. For example, (2) offers an explanatory reason why Harry blinks so frequently.

2. The reason that Harry blinks so frequently is that his contacts irritate his eyes.

Motivating and normative reasons, however, are reserved exclusively for agents. These are the kinds of reasons we give when we say things like (3) and (4), respectively.

3. Bridget's reason for going to the store is that she is out of milk.

⁴ Some philosophers use the phrase 'explanatory reason' interchangeably with, or instead of, the phrase 'motivating reason'. I do not follow these practices here.

4. There is a reason for Bridget to go to store—namely, that she is out of milk.

The notion at play in (4) is that of a normative reason. Roughly, to say that p is a normative reason for S to ϕ is to say that p speaks in favor of S 's ϕ -ing. The meaning of (4) is roughly that Bridget's being out of milk speaks in favor of Bridget's going to the store.

Unfortunately, it is not so easy to characterize the notion of a motivating reason, which is the notion at play in (3). Here is Michael Smith.

The distinctive feature of a motivating reason to ϕ is that, in virtue of having such a reason, an agent is in a state that is *explanatory* of her ϕ -ing, at least other things being equal—other things must be equal because an agent may have a motivating reason to ϕ without that reason's being overriding. (1994: 96)

We can simplify this a bit by focusing exclusively on cases where the agent not only has a motivating reason *to* ϕ , but indeed has a motivating reason *for* ϕ -ing. The latter entails that the agent in fact ϕ -s, and does so for her motivating reason. Focusing on such cases allows us to drop the 'all things being equal' qualification and simply say that

The distinctive feature of a motivating reason for ϕ -ing is that, in virtue of having such a reason, an agent is in a state that is explanatory of her ϕ -ing.

Unfortunately, Smith's characterization of a motivating reason does not distinguish motivating reasons from what I above called 'explanatory reasons'. These too have what Smith calls the 'distinctive feature' of motivating reasons. Hence, Smith's characterization does not allow us to capture the different senses of 'reason' at play in sentences like (2) and (3).

2. The reason that Harry blinks so frequently is that his contacts irritate his eyes.
3. Bridget's reason for going to the store is that she is out of milk.

Jonathan Dancy (2000: 1) offers us a little more help. Dancy characterizes a motivating reason as a 'consideration in light of which' an agent does what she does. This characterization is substantive enough to distinguish the notions at play in (2) and (3), respectively—billiard balls do not consider things, and so they do not do things in light of considerations. Provide that we

do not read too much into the notion of a ‘consideration’, Dancy’s characterization is also thin enough to not beg the question against live views to be discussed in this paper.⁵

Our topic in this paper is first and foremost *motivating reasons*. Hence from this point forward, when I use the term ‘reason’ without qualification, I will always intend ‘motivating reason’.

3. Factivity

Some philosophers (e.g., Unger 1975) maintain

Factivity. S’s reason is that *p* only if *p*. (E.g., Bridget’s reason for going to the store is that she is out of milk only if she is out of milk.)

Other philosophers (e.g., Dancy 2000) deny it.⁶ One possibility to be taken seriously here is that, even after we’ve restricted our attention to motivating reasons, the phrase ‘S’s reason’ admits of multiple interpretations, one of which is factive and another is not. Dancy (2011) has recently and explicitly, though briefly, considered this suggestion. Here is the entirety of what he has to say.

I mention, only to put aside, a view expressed in an unpublished paper by Dustin Locke. Locke allows that some reasons-explanations are not factive, so that there is a non-factive sense of “S’s reason is that *p*.” But he wants to admit the pressure on the other side, and so he allows that there is a distinct, factive sense of “S’s reason is that *p*.” Those who say that reasons-explanations fail if the supposed reason is not the case are using the factive sense; those who deny it are using the non-factive sense. Therefore in a way everyone is right, and everyone is wrong. My only comment on this is that one cannot resolve philosophical puzzlement in this way by multiplication of senses. (351)

⁵ By not reading too much into the notion of a consideration, I mean in particular that we must not simply assume from the start that a consideration is something that an agent *knows*.

⁶ Following John Turri (2012), we can divide non-factivists into *statists* and *abstractionists*. Statists maintain that motivating reasons are always *mental states* of the agent in question. Different statists offer different views of what the relevant mental state is. For some, the relevant mental state is a complex consisting both of a *belief* (e.g., a belief that one is out of milk) together with a suitably related desire (e.g., the desire for milk). For others, the mental state is a complex mental state that consists of both a non-normative belief (e.g., a belief that one is out of milk) and a suitably related normative belief (e.g., the belief that one *ought* to have milk). Because Bridget can be in the relevant mental state even if the content of that mental state isn’t true—e.g., she can believe that she is out of milk even if she isn’t—statists thus have an explanation of the non-factivity of reasons attributions. In contrast to statists, *abstractionists* maintain that motivating reasons are always propositions. Abstractionists explain the non-factivity of reasons attributions by noting that even when *p* is false, the (false) proposition that *p* nonetheless exists, and so that proposition can nonetheless be the agent’s reason.

Other philosophers have had a different attitude towards Locke's proposal. In Hyman's (2011) reply to the just-quoted paper by Dancy, he agrees with Locke's contention that the phrase 'S' reason' admits of both factive and non-factive interpretations. Hyman maintains that in one sense of the phrase 'S's reason', to say that S's reason for ϕ -ing is that p is to say that p figures as a premise on the basis of which S ϕ -s. But, Hyman insists,

[W]e sometimes use the term "reason" to express a different idea: the idea of a person's being guided by a fact. This is a metaphor, of course, but it is a perfectly familiar one to English-speakers, and it is not difficult to explain. If someone is said to have been guided by a certain fact, this means that he took it into consideration, when he modified his thought or behavior in some way, or decided what to think or what to do. (2011: 360 – 361)

So Hyman claims that in *another* sense of the phrase 'S's reason', to say that S's reason for ϕ -ing is that p is to say that in ϕ -ing, S is guided by the fact that p (2011: p. 360). Call these two interpretations the 'premissing interpretation' and the 'factual guidance interpretation', respectively. Hyman maintains that when he endorses Factivity, he is doing so under the factual guidance interpretation of 'S's reason'.

What, then, of Dancy's claim that 'one cannot resolve philosophical puzzlement in this way by multiplication of senses' (2011: 351). Neither Locke nor Hyman suggests that we can *resolve* philosophical puzzlement by *multiplication of senses*. Rather, Locke and Hyman claim that senses are *already* multiple—the phrase 'S's reason' already has two distinct senses—and that recognizing these distinct senses *dissolves* what might have appeared to be a substantive dispute over Factivity. I take it that no one would deny the second claim here: if there really are two distinct senses of 'S's reason', one of which is factive and the other is not, then what appeared to be a substantive dispute over Factivity is actually no such thing. So Dancy must mean to deny the first claim: Dancy must mean to deny

The Ambiguity Thesis. The phrase 'S's reason is that p ' has (at least) two distinct meanings, one of which entails that p , the other of which does not.

How might we settle this issue? I suppose the question is largely empirical, and that the most relevant kind of empirical study would simply ask subjects whether they hear a contradiction in

sentences like ‘Bridget’s reason for going to the store was that she was out of milk, but Bridget was not out of milk.’ I know of no such studies. The closest thing we have is philosophers who report divergent intuitions about such sentences. Dancy (2001: 132 – 3) says that he simply does not ‘hear’ a contradiction in the sentence ‘His reason for doing it was that it would increase his pension, but in fact he was quite wrong about that’. Unger (1975: 208), on the other hand, maintains that “it is inconsistent to say ‘His *reason* was that the store was going to close, but it *wasn’t* going to close’.” Although Unger concludes from this (alleged) inconsistency that “in this respect, our locution, ‘S’s reason is that p’, is like other important locutions with ‘reason’,” Unger gives no *argument* for the inconsistency. He thus seems to be reporting how things sound to him: the sentence in question just *sounds* like a contradiction. Of course, it is entirely possible that these philosophers are simply in the grips of their own theories about what it is to have a reason. Given this possibility, the fact that they report distinct intuition is a far cry from a conclusive case for the ambiguity thesis.

In lieu of convincing empirical evidence either for or against the ambiguity thesis, I suggest we try a different tact. Let us drop the question of whether the phrase ‘S’s reason’ is really ambiguous in the way that Locke and Hyman suggest. Let us instead ask, first, *why* are philosophers interested in the notion of a reason—that is, a motivating reason—and second, given our answer to the first question, ought philosophers to thus be interested in both notions that Hyman proposes—that is, both the notion of *p figuring as a premise in one’s decision* and the notion of *being guided by the fact that p*. If so, then regardless of whether these notions both deserve to stand as meanings of the phrase ‘S’s reason’, surely they both deserve our attention.

The short answer to the first question is that philosophers are interested in the notion of a reason because it seems that the overall rationality of an agent’s acting as she does is at least sometimes (or perhaps often or always) in part a matter of her reasons for acting as she does. For purposes of this paper I will take this answer for granted so that we can turn to the second question. Our second question thus becomes this: is the overall rationality of an agent’s acting as she does ever (1) at least sometimes in part a matter of which propositions figure as premises upon which she acts, and (2) at least sometimes in part a matter of which facts guide her in acting as she does? Now I suspect that most or perhaps all philosophers will grant that the overall rationality of agent’s acting as she does is at least sometimes in part a matter which propositions figure as premises upon which she acts. The controversy is going to be over whether the overall rationality of an agent’s acting as she does is at least sometimes in part a matter of which facts guide her in acting. In any case, this question will be our focus here. In the next section I will argue that the answer is ‘yes’. The remainder of the paper will focus on

the question of what it takes to be guided by a fact—or, as I shall then put it, *to treat a fact as a normative reason*.

4. Being Guided by a Fact

Before arguing that the overall rationality of agent's acting as she does is at least sometimes in part a matter of which facts she is guided by, I first want to point out that Hyman is not the only philosopher who has employed the notion of *being guided by a fact*. The notion also appears in Jennifer Hornsby's (2007a, 2007b) work, although not explicitly under that description. Moreover, as we will see in the following section, both Hyman and Hornsby make *knowledge* a necessary condition on being guided by a fact.

Hornsby begins by assuming that facts sometimes speak in favor of a given course of action (2007b: 2). In our terminology, this is the assumption that *there are facts that are normative reasons*. Hornsby, however, eschews the phrase 'normative reason', noting that 'normative questions... don't lapse' when we speak of motivating reasons. Instead, Hornsby uses the phrase '(F)-type reason' to refer to a fact that speaks in favor of an agent's doing something. In Hornsby's terminology, the assumption here is simply that *there are (F)-type reasons*.

But Hornsby does not merely assume that *there are (F)-type reasons*: she also assumes that agents sometimes 'act for' (F)-type reasons. This additional assumption is not highlighted in Hornsby's paper, and if one isn't paying careful attention, it might slip in unnoticed.⁷ What exactly does this assumption amount to? An (F)-type reason is a fact that speaks in favor of something—that is, in our terminology, a fact that is a normative reason. So the assumption that agents sometime act for (F)-type reasons is the assumption that agents sometime act for facts that are normative reasons. Now the idea of 'acting for a fact' is not a familiar one. What might this phrase mean? I know of no better answer than that to act for a fact is to be *guided by a fact in acting as one does*.

Let us now return to the question we left off with at the end of the previous section: is the overall rationality of one's acting as one does ever in part a matter of which facts one is guided by? Let us begin to answer this question by having a closer look at the relevant notion of *being guided by a fact*. Recall Hyman:

⁷ The assumption first appears at the beginning of section 1.2.

[The notion of being guided by a fact] is a metaphor, of course, but it is a perfectly familiar one to English-speakers, and it is not difficult to explain. If someone is said to have been guided by a certain fact, this means that he took it into consideration, when he modified his thought or behavior in some way, or decided what to think or what to do. (2011: 360 – 361)

It would be a mistake to understand Hyman as saying that *merely* taking the fact that p into consideration when modifying one's thought/behavior or deciding what to think/do is sufficient to be guided by the fact that p . This would mean that, on the factual guidance interpretation of 'S reason is that p ', S's merely taking p into consideration when deciding to ϕ is sufficient to make it the case that S's reason for ϕ -ing is that p . That cannot be right: suppose that the fact that p is a reason to *not* ϕ , and that S is aware that the fact that p is a reason to not ϕ . In that case, S might take the fact that p into consideration and yet, thinking that there are overriding reasons to ϕ nonetheless, decide to ϕ . In such a case, S takes the fact that p into consideration when deciding to ϕ , but it would be rather odd to say that there is any sense in which S's reason for ϕ -ing is that p (or even that *one of* S's reasons for ϕ -ing is that p). What we need is a way to distinguish those cases in which S takes the fact p into consideration when ϕ -ing, but is not guided by the fact that p , from those cases where S takes the fact that p into consideration when ϕ -ing, and is guided by the fact that p . I propose the following, not as an *analysis* of the distinction, but as a way of making it: in ϕ -ing, S is guided by the fact that p if and only if in ϕ -ing, S *treats the fact that p as a reason—that is, a normative reason—to ϕ .*

With this understanding of the notion of being guided by a fact, we can return to the question immediately at issue: is the rationality of one's acting as one does ever in part a matter of which facts one is guided by? This question now becomes the following: is the rationality of one's acting as one does ever in part a matter of which facts one treats as normative reasons? Provided that agents *do* at least sometimes treat facts as normative reasons, it is hard to see how the answer to this question could be 'no'. If an agent ϕ -s and in doing so treats the fact that p as a normative reason to ϕ , surely the overall rationality of what she has done depends in part on whether the fact that p is actually a normative reason for her to ϕ . Suppose that the fact that p is not a normative reason for her to ϕ . In that case, she is rationally criticizable (barring a good excuse) at least insofar as she treats something that is not a normative reason to ϕ as a normative reason to ϕ .

The above argument rests on the assumption that agents do sometimes treat facts as normative reasons. Anyone who, like Hornsby, maintains that facts sometimes *are* normative

reasons (or, as Hornsby would say, (F)-type reasons) should find this assumption plausible. It would be very strange indeed to maintain that facts sometimes are normative reasons, but at the same time insist that agents never treat them as such. That said, on some views, facts are never normative reasons, because normative reasons are a different sort of beast—they are propositions, states of the agent, or what have you. But on many views facts sometimes are normative reasons. Indeed, on many views facts are the only things that are normative reasons.⁸ This includes Dancy's (2011) view.

Now I will take it that we all agree that good reasons, reasons for doing this rather than that, are facts. But this does little to persuade us that the reason for which one acts must always be a fact... That this is so does not follow from that on which we all agree, namely that reasons for doing one thing rather than another are all facts. (348)

Dancy's point here is that although reasons for doing one thing rather than another—that is, normative reasons—are all facts, it does not follow from this that the reasons for which agents act—that is, motivating reasons—are all facts. We need not dispute this. Neither Locke nor Hyman maintains that motivating reasons are always facts. The question at issue is whether agents sometimes treat facts as normative reasons. And provided that at least some facts *are* reasons, it is highly plausible that agents sometimes treat them as such. Moreover, it is highly plausible that when an agent treats a given fact as a normative reason to ϕ , her overall rationality in ϕ -ing will be (barring a good excuse) in part a matter of whether that fact is actually a normative reason to ϕ .

Let's take stock. I began the previous section by rehearsing a standing dispute over whether 'S's reason is that p ' can be true even when p is false. I then suggested that the relevant parties to the dispute have been relying on distinct readings of the phrase 'S's reason' and noted that this suggestion has been endorsed by at least one party to the dispute, namely, Hyman. More specifically, I noted that Hyman has stated that when he endorses Factivity, he is interpreting the phrase 'S's reason' to mean 'S is guided by the fact that p '. Nonetheless, I suggested that there is a substantive dispute in the vicinity of Factivity. The substantive dispute is over whether the notion of *being guided by a fact* is a philosophically important notion. I argued that it is. My argument proceeded in two steps. The first step maintained that the notion of *being guided by a fact* is equivalent to (because it is a metaphor for) the notion of *treating a fact as a normative reason*. The second step argued that any theorist who maintains that facts

⁸ See, e.g., Maria Alvarez (2010) and Joseph Raz (2011).

sometimes *are* normative reasons should find it plausible that agents sometimes treat them as such, and that, when they do, their overall rationality will typically be in part a matter of whether the facts that they treat as normative reasons to ϕ really are reasons to ϕ . Hence, regardless of whether the notion of *being guided by a fact* deserves to serve as one reading of the phrase ‘S’s reason’, the notion deserves our attention.

The remainder of this paper will be dedicated to giving it that attention, and our question will be this: what are (some of) the necessary conditions of being guided by a fact? Equivalently, what are (some of) the necessary conditions for *treating a fact as a normative reason*? So that we don’t need to keep asking this one question in two different ways, let’s make a terminological decision. In this section I have suggested that the following phrases, under their intended interpretations, are equivalent to one another.

S is guided by the fact that p (in ϕ -ing).

S ϕ -s for the fact that p .

S treats the fact that p as a normative reason (in ϕ -ing, to ϕ).

For the sake of consistency, I want to settle on one of these phrases and stick to it throughout the remainder of this paper. I suggest we stick to the last phrase: ‘S treats the fact that p as a normative reason’. With respect to the phrase ‘S is guided by the fact that p ’, the phrase ‘S treats the fact that p as a normative reason’ has the advantage that it can be taken *literally*. And with respect to the phrase ‘S ϕ -s for the fact that p ’, the phrase ‘S treats the fact that p as a normative reason’ has the advantage of being rather more familiar from ordinary English. That said, the reader should feel free to substitute in his/her favorite phrase throughout the remainder of the paper.

5. The Knowledge View Considered and Rejected

What does it take to treat a fact as a normative reason? One natural thought is that *treating the fact that p as a reason* is just what happens when one relies on p as a premise and, in fact, p is true.

The Naïve View. In ϕ -ing, S treats the fact that p as a normative reason if and only if p figures as a premise in S’s decision to ϕ and p is true.

Unfortunately, the naïve view won't do. Suppose that a fortune teller tells Jack that his great uncle has just passed away and left him millions of dollars in inheritance. Jack, being gullible, believes what the fortune teller says. Relying on the premise that he has just inherited millions of dollars, Jack buys a new car. Now suppose that, by some highly unlikely coincidence and unbeknownst to Jack, one of Jack's relatives has just passed away and Jack has just inherited millions of dollars. Has Jack then treated *the fact* that he has just inherited millions of dollars as a normative reason to buy a new car? It seems not.

What more is required for treating a fact as a reason? Cases like Jack's suggest an answer: what Jack doesn't have is *knowledge* that he has just inherited millions of dollars. This, anyway, is the lesson that Unger (1975), Hyman (1999), and Hornsby (2007a, 2007b) draw from such cases. Let's call this

The Knowledge View. In ϕ -ing, S treats the fact that p as a normative reason only if S knows that p .⁹

In Unger's terminology: S's reason for ϕ -ing is that p only if S knows that p .

In Hyman's terminology: In ϕ -ing, S is guided by the fact that p only if S knows that p .

In Hornsby's terminology: S acts for the fact that p only if S knows that p .

Why think that one cannot treat a fact as a normative reason if one does not have knowledge of that fact? After all, in Jack's case, Jack presumably lacks not only *knowledge* but *justification* for his belief that he has just inherited millions of dollars. Why, then, think that the lack of *knowledge* and not the lack of *justified belief* explains why Jack does not treat the relevant fact as a normative reason?

Strikingly, Unger, Hyman, and Hornsby have *independently* offered nearly identical arguments for the Knowledge View: each appeals to a certain type of Gettier case. Here are their respective cases, slightly modified for ease of exposition.

⁹ It is important not to confuse the knowledge view with superficially similar principles. For example, John Hawthorne and Jason Stanley (2008) defend the principle that it is *appropriate*—that is, rationally permissible—for S to treat p as a reason only if S knows that p . Similarly, Jeremy Fantl and Matthew McGrath (2009) argue that if S knows that p , then p is warranted enough to justify S in ϕ -ing, for any ϕ . Unlike the knowledge view, these principles are distinctly *normative*—they are principles concerning when it is *appropriate* to treat p as a reason, or when p is warranted enough to *justify* someone in doing something. In this essay we are exclusively concerned with what it takes to have p as one's reason, regardless of whether it is *appropriate* to have p as one's reason or whether p *justifies* one in doing what one does. I consider the other principles elsewhere [omitted].

Quick Evaporation (Unger 1975: 209 - 210). Mary has just walked outside. Everything looks as if it has been raining, and indeed it *has* just been raining. However, unbeknownst to Mary, after it rained, the temperature shot up to 130 F and evaporated all of the rainwater. Shortly thereafter, some spray and sprinkle trucks came by and covered the area with water again. It is this water that Mary sees, and it is because she sees this water that Mary believes that it was raining. And because she believes that it was raining, Mary believes that the crops will grow. **Unger's intuition:** Mary does not treat *the fact* that it rained as a normative reason.

Re-run (Hyman 1999: 447 – 448). Henry is watching television on the day of the men's Wimbledon final. The television shows McEnroe beating Connor, and indeed McEnroe *has* just beaten Connor for this year's championship. However, unbeknownst to Henry, what he is watching is a re-run of last year's Wimbledon final. It is on the basis of seeing that match that Henry believes that McEnroe is this year's Wimbledon champion. And because he believes this he decides to place a bet on McEnroe winning the upcoming US Open. **Hyman's intuition:** Henry does not treat *the fact* that McEnroe is this year's Wimbledon champion as a normative reason.

Thin Ice (Hornsby 2007a: 94). Edmund has just asked his mother whether the ice on the pond is thick enough to skate. His mother tells him that it is too thin, and indeed it *is* too thin. However, unbeknownst to Edmund, his mother believes that the ice is thick enough to skate, but is trying to trick Edmund into thinking that the ice is too thin. On the basis of his mother's testimony, Edmund believes that the ice is too thin. And because he believes this he stays off the ice. **Hornsby's intuition:** Edmund does not treat *the fact* that the ice is too thin as a normative reason.

Unger, Hyman, and Hornsby take these cases to support the knowledge view. The idea is that the most natural explanation of why, in each case, the agent does not treat the relevant fact as a reason is that in each case the agent does not *know* the fact to obtain.

While I agree with Unger, Hyman, and Hornsby's intuitions about these cases, I think their diagnosis is misguided. The cases can instead be explained by

The Explanatory View. In ϕ -ing, S treats the fact that p as a normative reason only if (part of) what explains S's ϕ -ing is that p .

In *Quick Evaporation*, the fact that it was raining plays no part in the explanation of why Mary believes that the crops will grow—what does play a part in the explanation is the fact that the spray and sprinkle trucks covered the area with water. In *Re-run*, the fact that McEnroe is this year’s Wimbledon champion plays no part in the explanation of why Henry bets on McEnroe to win the upcoming US Open—what does play a part in the explanation is the fact that McEnroe won last year’s Wimbledon. Finally, in *Thin Ice*, the fact that the ice is too thin plays no part in the explanation of why Edmund stays off the ice—what does play a part in the explanation is the fact that Edmund’s mother wants Edmund to believe that the ice is too thin. The explanatory view alone is thus enough to explain why, in these cases, the agents do not treat the relevant facts as normative reasons.

In the sense that there is no explanatory connection between the relevant fact and the agent’s belief, the above Gettier cases are all akin to Gettier’s original cases,¹⁰ and in this respect they stand in contrast to a class of Gettier cases now known collectively as ‘fake-barn cases’. Carl Ginet’s original fake-barn case was created as a counterexample to Alvin Goldman’s (1967) view that to know that *p* is to have a belief that *p* that stands in the appropriate causal relation to the fact that *p*.¹¹ In Ginet’s fake-barn case, there is an appropriate causal relation between the fact that *p* and the subject’s belief that *p*, and yet most theorists—Goldman (1976) included—agree that the agent still does not know that *p*. For this reason, fake-barn style cases can serve as tests for whether knowledge is really required to treat a fact as a normative reason: if it is, then agents in fake barn cases ought *not* to be able to treat the relevant facts as reasons, since agents in those cases do not have knowledge of those facts. Let’s run this test by constructing “fake-barn” variants of Unger, Hyman, and Hornsby’s cases.

Fake-rain Country. Mary has just walked outside. Everything looks as if it has been raining, and indeed it *has* just been raining. However, unbeknownst to Mary, she is in fake-rain country. On most days in fake rain country, it does not rain, but the citizens use spray and sprinkle trucks to make it *look* as though it has rained. Just by chance, it is one of the few days when it really has rained in fake-rain country, and Mary is seeing the actual rainwater. Because she sees the rainwater, Mary believes that it was raining. And because she believes that it was raining, Mary believes that the crops will grow.

Intuitively, Mary does treat the fact that it rained as a normative reason.

¹⁰ Edmund Gettier (1963).

¹¹ Fake barn cases were invented by Ginet but originally published by Alvin Goldman (1976).

Re-run Hotel. Henry is watching television on the day of the men's Wimbledon final. The television shows McEnroe beating Connor, and indeed McEnroe *has* just beaten Connor for this year's championship. However, unbeknownst to Henry, he is in re-run hotel. Nearly all of the rooms in re-run hotel are showing re-runs of last year's Wimbledon final. Just by chance, Henry has been assigned to one of the few rooms in re-run hotel that is playing this year's match. It is on the basis of seeing this match that Henry believes that McEnroe is this year's Wimbledon champion. And because he believes this he decides to place a bet on McEnroe winning the upcoming US Open. **Intuitively, Henry does treat the fact that McEnroe is this year's Wimbledon champion as a normative reason.**

Fake-mother Invasion. Edmund has just asked his mother whether the ice on the pond is thick enough to skate. His mother tells him that it is too thin, and indeed it *is* too thin. However, unbeknownst to Edmund, his home has just been invaded by aliens disguised as perfect duplicates of his mother. If Edmund had spoken with any one of these aliens, he would have been told that the ice is perfectly safe. Just by chance, Edmund is talking to his actual mother, who tells him that the ice is too thin, and she tells him this because she inspected it this morning and found it to be so. On the basis of his actual mother's testimony, Edmund believes that the ice is too thin. And because he believes this he stays off the ice. **Intuitively, Edmund does treat the fact that the ice is too thin as a normative reason.**

In each of these cases, the agent is guided by the identified fact and yet does not know the fact to obtain. Hence, each is a counterexample to the knowledge view.

In discussion of these cases, I have occasionally heard philosophers offer a certain kind of defensive move on behalf of the knowledge view. The idea is to insist that in both the original cases and the "fake-barn" variants, the facts that the agents treat as reasons are facts that are *directly* known to them. With respect to the fake-rain cases, for example, the suggestion might be that in both Unger's original case and the fake-barn variant, Mary is guided by the fact that it *appears to have rained* (or, perhaps, the fact that *the grass is now wet*). It is important to note that this sort of move is no defense at all unless it is coupled with the claim that, even in the fake-barn variants, the agent does not *also* treat the indirectly known fact as a normative reason—in the fake-rain case, the claim would have to be that Mary does not *also* treat the fact

that it rained as a normative reason for believing that the crops will grow. Somewhat tellingly, I have never heard an actual proponent of the knowledge view make this kind of defensive move. There is, I think, good reason for this. Proponents of the knowledge view want to insist that in non-Gettier variants of cases like Unger, Hyman, and Hornsby's, the agent *does* treat the relevant fact as a reason, despite that fact being only indirectly known to the agent. That way, they can insist that in the Unger, Hyman, and Hornsby's cases it is specifically the lack of *knowledge* that prevents the agent from treating the fact as a reason. If, however, it were generally true that agents treat only directly known facts as reasons, then there would be no reason to think that a lack of knowledge was crucially at play in Unger, Hyman, and Hornsby's Gettier cases, and so there would be no argument from such cases to the knowledge view.

Several philosophers have reported to me that while they share the stated intuition in each of the fake-barn variants of Unger, Hyman, and Hornsby's cases, the strength of the intuition is not constant across them: it is strongest in *Fake-rain Country* and weakest in *Fake-mother Invasion*. The explanatory view predicts as much. As we move through the cases, the proximity of the fact to the ϕ -ing decreases—that is, the length of the explanatory chain from the fact to the ϕ -ing increases. If treating a fact as a normative reason is in part essentially a matter of having one's ϕ -ing explained by that fact, then we should expect that our intuitions about whether an agent treats a fact as a normative reason to fade as we consider cases in which the fact becomes more and more remote from the ϕ -ing.

I will conclude this section by noting that, even prior to considering the fake-barn variants of Unger, Hyman, and Hornsby's cases, we might have expected them to be counterexamples to the knowledge view. One of the original lessons of the fake-barn case was that whether one knows that p depends in part on things that are going on “outside” of the relations between the agent and the fact that p —these included things such as whether there are any fake barns around. According to many philosophers, this is because (1) if you have gotten *lucky* in being right about p , then you do not know that p , and (2) stuff going on external to the relations between the agent and the fact that p matters to whether the agent got lucky in being right about p . Yet there seems to be no analogous anti-luck requirement on treating a fact as a normative reason: whether you treat the fact that p as a normative reason does not depend on whether you have gotten lucky in being right about p . To treat a fact as a normative reason is to *respond* to that fact in a certain way. The agents in the fake-barn variants are responding to the relevant facts (in the relevant way), and this is why they, unlike their counterparts in Unger, Hyman, and Hornsby cases, treat the relevant facts as normative reasons.

6. The Explanatory-Representation View

A proponent of the knowledge view might insist that, apparent counterexamples aside, there are reasons to prefer the knowledge view to the explanatory view. What might these be? One natural suggestion is that the explanatory view alone cannot account for the difference between cases in which an agent treats the fact that p as a normative reason, on the one hand, and cases in which the fact that p is (merely) an explanatory reason that something happens, on the other. Suppose that one billiard ball strikes another and as a result the second is set into motion. Clearly, the second billiard ball does not treat the fact that it was struck by the first as a normative reason to move. Billiard balls do not treat facts as normative reasons. But they do things for explanatory reasons—that is, there are facts that explain what they do. As it stands, the explanatory view offers us no account of why these cases are merely cases of explanatory reasons, and not cases of treating facts as normative reasons.

How can we strengthen the explanatory view to account for this distinction? The knowledge view offers a natural suggestion here: billiard balls cannot treat facts as reasons because they do not have knowledge of facts. But a much weaker thesis works just as well: billiard balls do not treat facts as reasons because they do not treat facts as anything, and they do not treat facts as anything simply because they do not mentally represent facts in *any* way—e.g., they do not have *perceptions of* or *beliefs about* facts. The idea here is that to treat a fact as a normative reason, the fact must not only play a part in the explanatory history of one's ϕ -ing, but the explanatory history must involve some sort of mental representation of that fact. The picture is as follows.

the fact that $p \rightarrow$ S's mental representation of the fact that $p \rightarrow$ S's ϕ -ing

where $x \rightarrow y$ means that x is part of the explanation of y . Let us call this

The Explanatory-Representation View. In ϕ -ing, S treats the fact that p as a normative reason only if (1) S mentally represents the fact that p , (2) the fact that p is part of the explanatory history of S's representation of the fact that p , and (3) S's representation of the fact that p is part of the explanatory history of S's ϕ -ing.

It is important to note that, like the knowledge view, the explanatory-representation view merely places necessary conditions on being guided by a fact. Strengthening these conditions so that

they provide both necessary and sufficient conditions would presumably involve spelling out clauses (1), (2) and (3) in more detail. To spell out (1), we would need a theory of just what kind(s) of mental representation(s) of the fact that p are necessary for treating the fact that p as a reason. Must one have a *belief* that p ? Or is, say, a pre-doxastic perceptual experience that p sufficient? With respect to (2) and (3), presumably not just any way of figuring in the relevant explanatory history is sufficient. What, then, is the right way? It is worth noting that a proponent of the knowledge view faces a similar issue: to strengthen her account so that it provides both necessary *and sufficient* conditions on treating a fact as a reason, she will need a theory of just how the knowledge that p must relate to the agent's ϕ -ing—merely knowing that p while ϕ -ing will not suffice for treating p as a normative reason.

Save some brief comments below, I will not engage in the larger project of transforming the explanatory-representation view into a view that provides both necessary and sufficient conditions on treating a fact as a normative reason. My purpose has been to show that there is an alternative to the knowledge view that better accounts for the cases offered in its favor. That alternative is in general the explanatory view, and in particular the explanatory-representation view. Both the explanatory-representation view and the knowledge view explain why agents in Unger, Hyman, and Hornsby's Gettier cases do not treat the relevant facts as reasons, but the explanatory-representation view does so without facing counterexamples from other kinds of Gettier cases. I will conclude by considering some objections to the explanatory-representation view.

7. Objections and Replies¹²

Objection: What about cases where agents frequently lack the appropriate explanatory connection to the relevant fact, but where occasionally they have it? Don't agents have the same reasons across those cases?

Consider again Unger's *Quick Evaporation* case. Suppose that last year Mary woke up and walked outside to find water puddles on one-hundred different days. Suppose that on ninety-nine of those days she was in Unger's *Quick Evaporation* case, but on one of those days the water she was seeing really was rainwater. It seems plausible that on all one-hundred days, Mary's reason for believing that the crops will grow was the same. But the explanatory-representation account suggests otherwise: the account suggests that on the one morning when the water that Mary saw really was rainwater, Mary treats the fact that it rained as a normative

¹² In this section I consider some potential objections raised by a pair of helpful referees.

reason, while on the other ninety-nine mornings, Mary does not treat the fact that it rained as a normative reason.

Reply: Such agents do have the same reason, in a sense, but not in the relevant sense.

First, it should be stressed that the explanatory-representation view does not entail that on the day Mary sees actual rainwater, she treats the fact that it rained as a normative reason. The explanatory-representation view can have no such implication, because the explanatory-representation view merely places a *necessary* condition on treating facts as normative reasons. Nevertheless, were the explanatory-representation view strengthened to offer both necessary and sufficient conditions, and were it strengthened in a plausible way (see below), I suspect that the view would entail that in the case where Mary sees actual rainwater, she is guided by the fact that it rained (at least given a natural way of filling in the details of the case). So the spirit of the objection remains: doesn't Mary have the same reason regardless of whether she is seeing the actual rainwater?

The answer is that, *in a sense*, yes. Above I noted that Locke and Hyman distinguish two senses of the phrase 'S's reason is that *p*', only one of which is that S treats the fact that *p* as a normative reason. The other sense, according to Hyman, is that *p* figures as a premise in S's reasoning. There is thus a sense of 'Mary's reason' such that, yes, on each day, Mary's reason for thinking that the crops will grow is the same: on each day, the proposition that it rained (on that day) figures as a premise in her reasoning about whether the crops will grow. But this is perfectly compatible with the explanatory-representation view, which is a view about what it takes to *treat a fact as a normative reason*. And the verdicts of that account seem intuitively plausible here: when Mary is seeing the actual rainwater, she is indeed responding to *the fact* that it rained, and so she is indeed treating *that fact* as a reason. But when she isn't seeing rainwater, but she is instead seeing the water from the spray and sprinkle trucks, she is not responding to *the fact* that it rained, and so she is not treating the fact that it rained as a normative reason.¹³

¹³ Moreover, it is not clear that this kind of objection does not apply equally well to the knowledge view. But the issue here is tricky, since it is a tricky matter whether Mary knows, on the one day that she sees actual rainwater, that it was raining. If she does, then, since she doesn't know that it was raining on the other days, the objection here applies as much to the knowledge view as it does to the explanatory-representation view: the latter, just like the former, would seem to imply that on the day that it really rained Mary is guided by the fact that it rained, but on the other days she is guided, if at all, by other facts. However, a defender of the knowledge view might push back, arguing that this is in effect a fake-barn case: since Mary is so often in the *Quick Evaporation* case, she does not know that it rained even on the one day that it really has.

Objection: What about facts that play no explanatory role in our arriving at our beliefs in them? Are we not still capable of treating such facts as normative reasons?

According to the explanatory-representation view, facts that play no part in the explanatory history of one's representation of those facts cannot be facts that one treats as normative reasons. This might seem problematic in a couple of closely related cases. First, consider cases where the fact that *p* and the agent's representation of the fact that *p* have some *common cause*. Suppose, for example, that Tom wakes up and sees that his bedside clock reads '8:00 a.m.'. Tom lives in a windowless apartment, so that he cannot see that the sun is up. Still, he believes that the sun is up, and, acting on this belief, he gets out of bed. In this case it might seem that Tom treats the fact that the sun is up as a normative reason, and yet the fact that the sun is up plays no part in the explanatory history of Tom's believing that the sun is up. As one referee notes, this is "an instance of a general kind of case: we come to know a fact based on inductive evidence and then use that fact as a reason and yet that fact doesn't explain our belief or our use of it as a reason."

A similar but importantly different kind of case involves facts about abstracta. Consider, for example, mathematical facts. There are those who deny that mathematical facts stand in explanatory relations to any concrete facts, including facts about what we believe, do, or represent. However, it might also seem that we treat mathematical facts as normative reasons. Suppose you offer me a bet on whether one plus one is two. If the stakes are anything other than extremely lopsided, I will treat the fact that one plus one is two as a reason for betting that it is.¹⁴

Reply: In such cases we should stand by the verdicts of the explanatory-representation view.

With respect to both kinds of cases discussed in the objection, I think we should stand by the verdict of the explanatory-representation view. We should agree that in Tom's case, Tom *relies on the premise that the sun is up* in deciding to get up. But, as I was at pains to stress above, that does not entail that Tom *treats the fact that the sun is up as a normative reason to get up*. To treat a fact as a reason is to *respond* to the fact in a certain way. It seems odd, to say that least, that Tom is *responding* to a fact that he has no contact with—namely, the fact that the sun is now up. The facts that Tom is responding to are more plausibly taken to be facts about what the clock says and facts about past correlations between clock readings and the sun's location in the sky. I am not suggesting that Tom is consciously or even sub-consciously *reasoning* about those facts. Still, it is those facts that he treats as a normative reasons.

¹⁴ Thanks to [omitted] for this example.

That said, it would be nice to have something to say to those who simply don't share the relevant intuitions here—that is, those who, even after taking into account the above considerations, still feel that Tom does treat the fact that the sun is now up as a normative reason. Fortunately, I think we can, if we like, accommodate this position. Doing so requires slight modification of the explanatory-representation view, but that modification leaves the spirit of the view intact. In the context of considering the next objection, I will show how this can be done. Before then, let us briefly address the case of mathematics.

The case of mathematics (and other abstracta) is a bit more complex, if only for the fact that it is far more controversial whether (1) mathematical facts do play a role in explaining our mathematical beliefs and whether (2) we do treat mathematical facts as reasons. The objection rests on the denial of (1) and the affirmation of (2). This is a somewhat uncomfortable position to maintain. To hold (2) is in part to hold that there *are* mathematical facts and that the mathematical facts are (largely) what we take them to be. Hence, one is rational in holding (2) only if one is rational in holding one's mathematical beliefs. However, it is irrational to hold one's mathematical beliefs and simultaneously maintain that we have no "epistemic access" to the mathematical facts. But if, as the objection maintains, (1) is false, then it would seem that we do not have epistemic access to the mathematical facts. The argument here is not a new one (see Paul Benacerraf 1973 and Hartry Field 1989).

There are of course places in this argument where one might push back. Suppose, for example that one were to provide a convincing argument that it is rational to maintain one's mathematical beliefs while at the same time denying that mathematical facts played any role in shaping those beliefs. While I doubt that such an argument could be provided, we need not settle this issue here. Should such an argument be provided, and should, more generally, we become convinced that mathematical facts do not play a role in the explanation of our mathematical beliefs, then it will be plausible to maintain that while we do *rely on mathematical premises in our decision-making*, we are not *responding to* mathematical facts, and so we do not treat mathematical facts as normative reasons.

Objection: What about cases involving "deviant" explanatory chains?

Consider the following case, which was suggested by an anonymous referee. Suppose that after observing that p , a journalist reports that p to her newspaper. The newspaper, however, misprints the report as *not- p* . But, just by chance, a reader misreads the report, taking it to say that p , and comes to believe that p . The reader then acts on her belief that p . In a case like this, the fact that p is part of the explanatory history of the reader's belief that p , and the reader's

belief that p is part of the explanatory history of the reader's action. However, intuitively, it seems that the reader does not treat the fact that p as a normative reason.

Reply: The explanatory-representation view merely offers necessary conditions on treating a fact as normative reason and so is consistent with the case as described. Moreover, the explanatory-representation view can be strengthened so that it is not only consistent with, but indeed accounts for, the case as described.

The newspaper case serves to highlight the fact, noted above, that the explanatory-representation view merely places *necessary* conditions on being guided by a fact: the view does not say that meeting conditions (1), (2), and (3) is *sufficient* for treating a fact as a normative reason. Above I noted that to turn the right-hand-side of the explanatory-representation view into a sufficient condition, we will in part have to spell out just what *kind* of explanatory relation must hold between the representation that p and the agent's representation of the fact that p —that is, we will need to specify the precise explanatory relation at issue in clause (2).

While I do not here wish to commit the explanatory-representation view to any particular theory of the explanatory relation at issue in clause (2), it is worth considering whether we might appropriate the explanatory relation offered by Goldman (1967). Goldman offered his explanatory relation—which I will explain below, but for now simply call 'G'—as a necessary and sufficient condition on *knowledge*: according to Goldman, S knows that p if and only if S's belief that p stands in certain explanatory relation G to the fact that p . As noted above, fake-barn style Gettier cases were originally designed to show that Goldman's relation G is not sufficient for knowledge. Since we have seen that fake-barn style cases also suggest that knowledge is not necessary for treating a fact as a normative reason—equivalently, that treating a fact as a normative reason is not sufficient for knowledge—one cannot help but wonder: is G the precise relation required, as far as clause (2) is concerned, for treating a fact as a normative reason?

To spell out G in full detail would take more space than what we have here, but, very roughly, S's belief that p stands in G to the fact that p if and only if S's belief that p was formed via a perception that p , a memory that p , or a "correctly reconstructed" and "appropriate" causal chain linking S's perception/memory to the fact that p . On Goldman's view, "appropriate" causal chains come in two types: either the fact that p was a part of the causal history of S's belief that p or there was some common fact in the causal histories of both S's belief that p and the fact that p . Goldman included the second type of causal chain in attempt to account for knowledge of facts that played no role in the explanatory history of our coming to believe them. This would include, for example, Tom's belief that the sun is now up. Goldman thought it

plausible that in a case like that, Tom does *know* that the sun is now up, and so he designed his account of knowledge to permit an agent to acquire knowledge via correctly reconstructing a pair of causal chains—one leading to the belief and the other leading to the fact believed—that share a common link.

As indicated above, while I do think that Tom acts on his belief that the sun is now up, I don't think it is plausible that Tom treats the fact that the sun is now up as a normative reason. Hence, if we are to appropriate Goldman's account to an account of treating facts as normative reasons, I think we should drop the part that allows for cases of common cause. Nevertheless, I am aware that some would not agree: some would insist that Tom *does* treat the fact that the sun is now up as a normative reason. Such theorists will want to incorporate Goldman's second type of causal chains into the account. This would in a certain sense represent a weakening of the explanatory-representation account, but the spirit is preserved: at bottom, what counts are certain kinds of explanatory connections. But as I said, in my view it would be a mistake to follow this route: we should deny that Tom treats the fact that the sun is now up as a normative reason.

Aside from dropping the clause that allows for cases of common cause, I think we should make two additional tweaks to Goldman's relation before appropriating it. Both involve generalizing the relation just a bit. First, let us allow the relation to hold not just between a *belief* that *p* and the fact that *p*, but between any type of mental representation of the fact that *p* and the fact that *p*. This will allow for cases of treating the fact that *p* as a normative reason without forming a belief that *p*—for example, cases where one has a pre-doxastic perception that *p* and acts on the basis of that perception. Second, let us formulate the relation in terms of *explanatory chains* rather than *causal chains*, leaving it as an open question whether there are any instances of the former that are not also instances of the latter.¹⁵ The result is a relation I will call G*:

Definition of G*. S's representation of the fact that *p* stands in G* to the fact that *p* if and only if S's representation was formed via (or just is) a perception that *p*, a memory that *p*, or a correctly reconstructed causal chain linking a perception/memory of S's to the fact that *p*.

¹⁵ Some potential cases I have in mind here are (1) cases of governance by natural law and (2) supervenience on the microphysical. Both are cases of possible *explanation* without *causation*: the laws (the microphysical) *explain* why things are the way they are, but they do not *cause* them to be the way they are. These are controversial matters that cannot be discussed here. I thus formulate the thesis in terms of the more general notion—*explanation*—so as to remain neutral.

The proposal we are now considering is that clause (2) of the explanatory-representation view should be strengthened to say that ‘S’s representation of the fact that p stands in G^* to the fact that p ’.

I do not have space to consider the merits of this proposal here. But its merits are certainly worth considering, if only because it gives us a nice way of accounting for why in cases of “deviant” causal chains like that in play in the newspaper case, the agent typically does not treat the relevant fact as a normative reason. In the newspaper case, the reader does not treat the fact that p as a normative reason because she presumably has mistaken beliefs about the causal chain linking her perception of the newspaper to the fact that p . When thinking about a case like this, we implicitly make the natural assumption that the reader (non-occurently) believes that the fact that p led to the newspaper printing p and that this led to her believing p . But this is not the way things went down. The reader thus has not correctly reconstructed the causal chain leading from the fact that p to her belief that p , and her belief thus does not stand in G^* to the fact. We can test this account by modifying the case: suppose that the reader believes (on good grounds) both that the newspaper has printed the negation of what the reporter reported *and* that she has misread the newspaper in such a way as to take it to say the negation of what it actually says. Here it seems that the reader very well might treat the fact that p as a normative reason. Again, this suggests that what is going in the original version of the case is that the reader has false (non-occurrent) beliefs about the causal chain between the fact that p and her belief that p .

There remains work to be done here. Nonetheless, it seems to me that the explanatory view generally, the explanatory-representation view specifically, and the G^* version of the explanatory-representation view even more specifically, are plausible and powerful enough to merit further exploration.

Acknowledgements

[Omitted for blind review.]

References

- Alvarez, Maria. 2010. *Kinds of Reasons: An Essay in the Philosophy of Action*. New York: Oxford University Press
- Benacerraf, Paul. 1973. 'Mathematical Truth'. *Journal of Philosophy* 60: pp. 661 – 679
- Dancy, Jonathan. 1993. *Moral Reasons*. Oxford: Blackwell
- Dancy, Jonathan. 2000. *Practical Reality*. Oxford: Oxford University Press
- Dancy, Jonathan. 2011. 'Acting in Ignorance'. *Frontiers of Philosophy in China* 6: 345 – 357
- Fantl, Jeremy and Matthew McGrath. 2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press
- Field, Hartry. 1989. *Realism, Mathematics, and Modality*. Oxford: Blackwell.
- Gettier, Edmund. 1963. 'Is Justified True Belief Knowledge'. *Analysis* 23: 121 - 123
- Goldman, Alvin. 1967. 'A Causal Theory of Knowing'. *The Journal of Philosophy* 64: 357 - 372
- Goldman, Alvin. 1976. 'Discrimination and Perceptual Knowledge'. *The Journal of Philosophy* 73: 771 – 791
- Hawthorne, John and Jason Stanley. 2008. 'Knowledge and Action'. *The Journal of Philosophy* 105: 571 - 590
- Hornsby, Jennifer. 2007a. 'Knowledge, Belief, and Reasons for Acting' in *Explaining the Mental*. Carlo Penco, Michael Beaney, Massimiliano Vignolo (eds). Cambridge Scholars Publishing: 88 – 105

- Hornsby, Jennifer. 2007b. 'A Disjunctive Conception of Acting for Reasons' in *Disjunctivism: Perception, Action, Knowledge*. Adrian Haddock and Fiona MacPherson (eds). London: Birkbeck ePrints. Available at: <http://eprints.bbk.ac.uk/433>
- Hyman, John. 1999. 'How Knowledge Works'. *The Philosophical Quarterly* 49: 433 – 451
- Hyman, John. 2011. 'Acting for Reasons: Reply to Dancy'. *Frontiers of Philosophy in China* 6: 358 – 368
- Kaplan, Mark. 1985. 'It's Not What You Know that Counts'. *The Journal of Philosophy* 82: 350 – 363
- Kaplan, Mark. 2003. 'Who Cares What You Know?' *Philosophical Quarterly* 53: 105-116
- Kvanvig, Jonathan. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Raz, Joseph. 2011. *From Normativity to Responsibility*. Oxford: Oxford University Press
- Smith, Michael. 1994. *The Moral Problem*. Oxford: Blackwell Publishers
- Turri, John. 2012. 'Reasons, Answers, and Goal'. *Journal of Moral Philosophy* 9: pp. 491 – 499
- Unger, Peter. 1975. *Ignorance: A Case for Skepticism*. Oxford: Oxford University Press