

PRAGMATIC BELIEVING AND ITS EXPLANATION

WARD E. JONES

Department of Philosophy
Rhodes University
w.jones@ru.ac.za

SUMMARY: Most explanations of beliefs are *epistemically* or *pragmatically* rationalizing. The distinction between these two types involves the explainer's differing *expectations* of how the believer will behave in the face of counter-evidence. This feature suggests that rationalizing explanations portray beliefs as either (i) a consequence of the believer's *following a norm*, or (ii) part of a *sub-intentional goal-oriented* system. Which properly characterizes pragmatic believing? If there were pragmatic norms for believing, I argue, they would not be consciously *followable*. Yet an unfollowable norm is not a norm at all, and so I conclude that there are no such norms and that pragmatic believing is a sub-intentional, and not a norm-driven, process.

KEY WORDS: rationalizing explanation, causal explanation, norm, goal

RESUMEN: La mayoría de las explicaciones de las creencias racionalizan de forma *epistémica* o *pragmática*. La distinción entre estas dos involucra las diferentes *expectativas* del que explica acerca de cómo se comportará el creyente frente a evidencia contraria. Este rasgo sugiere que las explicaciones racionalizadoras toman las creencias o bien (i) como una consecuencia de que el creyente *siga una norma*, o bien (ii) como parte de su sistema *subintencional orientado a fines*. ¿Cuál caracteriza de manera apropiada el creer pragmático? Aquí argumento que si hubiera normas pragmáticas para creer, no se *podrían seguir* conscientemente; sin embargo, una norma no seguible no es una norma en ningún sentido, por lo que concluyo que no existen tales normas y que el creer pragmático es un proceso subintencional y no un proceso guiado por normas.

PALABRAS CLAVE: explicación racionalizadora, explicación causal, norma, fines

Introduction

One potentially fruitful method for discovering the nature of believing is to investigate how we explain beliefs. If folk psychological doxastic explanations essentially present beliefs as having a certain property, then this provides a *prima facie* reason to

think that the process of believing is itself so characterized. In this paper, I address two questions, one general and the second more specific, concerning the nature of beliefs and their explanations.

The more general concern is with the nature of *rationalizing* doxastic explanations, those that, as I explain in Section 1, work by showing that it is *appropriate* for the believer to believe what she does. In Section 2, I argue that giving a rationalizing doxastic explanation of someone's belief involves having certain expectations of how she (or her belief) will behave in the face of counter-evidence. If we offer an *epistemic* explanation of someone's belief, then we expect the believer to be evidence-sensitive. When we offer a *pragmatic* explanation of someone's belief, presenting it as, say, a wishful thought or *idée fixe*, then we will have different expectations of her doxastic behavior.

This feature suggests that doxastic explanations are one of three types of explanations, discussed in Section 3. If in explaining someone's belief, we have expectations about the future behavior of the belief, then doxastic explanations present the belief either (i) as *occurrently dependent* upon something else, (ii) as a consequence of the believer's *following a norm*, or (iii) as part of a *sub-intentional goal-oriented* system. The first option can be quickly dismissed; occurrent dependence does not in itself involve anything that could function as the rationalizing element found in most doxastic explanations. A rationalizing doxastic explanation, I conclude, portrays a belief as being the result of the believer's following of a norm or as dependent upon a sub-intentional teleological process.

In the fourth and final section of the paper, I focus more closely on the second concern of this paper, the nature of pragmatic believing. The nature of norms dictates that they must be consciously *followable*, even if not consciously *followed* on any one occasion. However, pragmatic believing cannot be carried out in full awareness. It follows that the notion of a pragmatic norm for believing does not make any sense. I conclude that pragmatic believing is a sub-intentional, and not a norm-driven, process.

1. *Rationalizing Explanations of Beliefs*

A doxastic explanation gives an account of why an agent holds a particular belief at a particular time. It is an answer to the question, ‘Why does Jones believe that p ?’

Why and when do we want doxastic explanations? Perhaps the most significant purpose for giving or seeking a doxastic explanation is that evaluative properties —e.g., *rational* or *justified*— are conferred upon beliefs largely in virtue of the explanation given for them. We need to know why someone believes something before we can know whether his believing it is justified. There are exceptions to this rule. We often need not know why someone holds an *a priori* justified belief or a belief about her own mental states in order to determine whether her holding such beliefs is justified or not. Such beliefs can simply be presumed to be justified, and prior explanation is not needed. However, other (and I suspect most) types of beliefs are considered to be rationally held only if they are held in virtue of appropriate determinants. Doxastic explanation is, by and large, a condition for doxastic evaluation.

The converse is true as well. We not only explain beliefs in order to justify them, but we also show beliefs to be justified in order to explain them. Most, although perhaps not all, doxastic explanations, that is, explain by showing that the believer is, in some sense, *right* to believe what she does. Such explanations I call *rationalizing doxastic explanations*.

A rationalizing doxastic explanation is a statement that purports to give information about the determination of a subject’s belief by maintaining that the determinants for the belief, in some sense or other, make the belief an *appropriate* one for the believer to possess.

The notion of appropriateness is the key to understanding rationalizing explanations.¹ In offering a rationalizing explanation

¹ The notion of ‘appropriateness’ is one that Philip Pettit uses repeatedly in his very general characterization of (rationalizing) *action* explanations. See the first two pages of his 1979. The work of Donald Davidson has been most

of Jones's belief, one presents Jones as believing because there is some sense in which it is *suitable* or *fitting* for him to do so. It is this element—that of appropriateness or suitability—that makes an explanation rationalizing. Rationalizing explanations make a descriptive claim, namely an explanation of why the believer believes what he does, by way of making an evaluative claim. The agent is said to believe because *in a minimal sense* he *should* believe what he does. The qualification 'in a minimal sense' must be added because an explainer who gives a rationalizing explanation of Jones's belief that *p* need not *endorse* Jones's belief that *p*. She need not agree that Jones ought to believe that *p*. However, the explainer must make some sort of positive evaluation of the belief. Rationalizing explanations work precisely by showing that the belief in question is one that it is *proper* for the believer to hold.

It may be that not all doxastic explanations are rationalizing. Some everyday doxastic explanations that cite emotions or habitual behavior, for example, may not involve reasons for the explained belief.

Jones believes that some noise was a snake because he is so afraid of them.

Jones believes that his boss's expression is scornful because he resents her so much.²

Jones believes that God exists because he has belonged to a church, and engaged in its rituals, his entire life.

These explanations do not, as written, *deny* that the respective believers should believe what they do. Indeed, it may be that in any one case, when more details are filled in, the explanations turn out to be rationalizing; or it may be clear from the context that the rationalizing details are implicit, that these explanations are intended to be rationalizing. Nevertheless, it is easy

influential in bringing rationalizing explanations to the forefront of recent discussion; see, for example, his 1982.

²These first two examples are reminiscent of Rosalind Hursthouse's examples of non-rationalizing *action* explanations. See her 1991.

to imagine the above explanations turning out not to claim that the believer holds the belief *because* the believer is right to do so, and instead describe a case of brute causation: that some emotion, or repeated behavior, simply determined the belief in question. However, most doxastic explanations —and the ones with which I will be concerned in the remainder of this paper— are to be contrasted with the above explanations in that they assert that the believer holds the belief she does because it is proper for her to do so.

Rationalizing doxastic explanations divide quite naturally into two types, those that are *epistemically* rationalizing and those that are *pragmatically* rationalizing; consequently, beliefs themselves divide into two types —epistemic and pragmatic— depending upon how they are to be explained. Epistemically rationalizing explanations appeal to the epistemic appropriateness of the belief: its status as likely to be true or knowledge. Pragmatically rationalizing explanations, on the other hand, account for the explanandum belief by describing how the belief is non-epistemically appropriate: believing what she does benefits the believer in some way that is not derivative of its status as true or epistemically justified.

At first glance, one might think that the division between these two types of beliefs and their explanations falls straightforwardly in line with the *determinants* or *causes* of beliefs. According to this thought, epistemically rationalizing doxastic explanations are epitomized by statements like the following, in which a belief is said to have been brought about by ‘epistemic causes’, usually evidence-taking or perception:

Jones believes that he has the measles because he has discovered red spots on his stomach.

Jones believes that I am eating garlic because he can see (or smell or hear) that I am doing so.

In the former case, it is claimed that Jones believes that he has the measles because Jones believes that some other state of affairs —namely that there are red spots on his stomach—

holds, a state of affairs that Jones thinks is correlated with the measles.³ In the latter case, Jones believes that I am eating garlic because he is having a perceptual experience as of my eating garlic. Pragmatically rationalizing doxastic explanations, on the other hand, appeal to ‘pragmatic causes’, usually the believer’s desires, needs, or anxieties:

Jones believes that he will get the job because he desperately needs more income.

Jones believes that God exists because he wants to live forever.

According to the causal account of the division between epistemic and pragmatic believing, they are distinguished by their characteristic causes. Epistemic beliefs are determined by the possession of evidence or a perceptual experience, as well as perhaps other epistemic causes; pragmatic beliefs are determined by the believer’s pro-attitudes or some other pragmatic cause.

While I endorse the distinction between epistemically and pragmatically rationalizing doxastic explanations, and the correlative distinction between epistemic and pragmatic believing, I deny that these distinctions can be drawn solely in terms of the causes or determinants of beliefs. As we will see in the next two sections, rationalizing doxastic explanations cannot be ‘mere’ causal explanations. Epistemic and pragmatic doxastic explanations —and with them epistemic and pragmatic beliefs— are distinguished not (only) by the determinants to which they appeal, but (also) by the kinds of *expectations* the explainer has of the believer, expectations that are intimately related to the *appropriateness* that (as we have seen) the rationalizing explainer attributes to the belief. As I will later put it, the rationalizing explainer sees the explanandum belief either as a state of an *agent*, or as a state of a *system* belonging to an agent.

³ For this account of evidence-taking (as well as the particular example I use), see Longino 1979.

2. *Expectation and Rationalizing Doxastic Explanations*

Imagine that Jones has seen a number of posters around town advertising a concert, to take place at some time in the near future. Because of these posters, Jones believes that the concert will take place, and he consequently makes plans to attend. Given only this information, we would have no reason to hesitate in explaining his belief epistemically: Jones believes that there will be a concert at such-and-such a time because he has seen, and believed, posters advertising it.

Now imagine that Jones is challenged with respect to this belief. On the day of the concert a friend of Jones's, Rachael, tells him that she has just read in the local newspaper that the signs advertising this particular concert are hoaxes. Taking Rachael and the local papers to be fairly trustworthy, Jones takes the proposition that there will not be a concert tonight to be fairly well-supported. As a result of hearing her claim, his belief that there will be a concert tonight is severely weakened, and he decides to phone the venue of the supposed concert to confirm for himself. Rachael has presented a challenge to Jones's belief, and he, in turn, has appreciated her challenge. His belief has responded accordingly, by weakening. It has been negatively affected by Rachael's counter-evidence.

Most of the time, our beliefs behave like Jones's. When we appreciate and accept evidence for positions incompatible with our own, our beliefs are, usually, affected. We believe them less strongly, or we no longer believe them at all. We search out ways of discovering whether our original beliefs, or the newly suggested alternatives, are right. It is hard to imagine what our doxastic lives would be like if we did not in general respond to what we saw as cogent epistemic challenges. Seeing other people as susceptible to counter-testimony thoroughly shapes how we interact with them; *trust*, *testimony*, and *argumentation* all depend in different ways upon our seeing each other's beliefs as being susceptible to counter-evidence. If we did not think the same of ourselves, nothing like *investigation*—setting out to discover what the world is like— would be possible. So, in

his response to Rachael's claims, Jones behaves like a typical believer and his belief behaves like a typical belief.

We originally explained Jones's belief epistemically, and the fact that he responds as he does to Rachael does not affect the correctness of that original explanation. We would still agree that Jones's belief that there is a concert on soon is an epistemic belief. However, the tendency for our beliefs to weaken in the face of counter-evidence (as Jones's does in the face of Rachael's challenge) is *only* a tendency. Like many other folk psychological 'laws', this is no more than a rough generalization, and it leaves open the possibility that some of our beliefs will not be responsive to arguments against them, even if we accept, or appreciate, such challenges. And when this possibility occurs, when we see that a belief does not respond to counter-evidence, we withdraw our epistemic explanation of it. Evidence-sensitivity is a condition on epistemically rationalizing explanation.

To see this, we can imagine an alternative scenario in which Jones appreciates Rachael's challenge to his belief (he takes her to be trustworthy, etc.), but his belief is *not* affected by this appreciation. His belief is in this case what we might call an *idée fixe*. An *idée fixe* is a belief that is not affected by the subject's own recognition of evidence against it; it is held even when the subject comes to believe considerations against what the *idée fixe* says is true. Jones's *idée fixe* is a belief Jones holds in spite of his being faced with, and appreciating, a challenge to p . His own epistemic assessment of the belief (that is, his appreciation of the counter-evidence against it) does not affect his commitment to the belief itself. As an *idée fixe*, Jones's belief is, we might say, *independent* of Jones's own view of its (lack of) groundedness.

If we are convinced that Jones's belief is an *idée fixe*, then how will we explain it? While I suspect that not all *idées fixes* will be explained in the same way, we can be sure that whatever explanation we offer will not be epistemically rationalizing. The behavior of Jones's belief —its being an *idée fixe*— shows us that Jones does not believe that p in order to gain a true

belief, or knowledge, about the subject matter at hand. For if he were after some such epistemic goal, his own appreciation of Rachael's counter-evidence would affect his belief that p . His belief would be sensitive to the counter-testimony before him, which by his own lights is not insignificant. Jones's belief will not move in the face of counter-evidence, it would seem, because truth-seeking has nothing to do with why he has it. His believing what he does seems not to be an epistemic affair at all. It looks rather to be a case of Jones's believing in order to gain some pragmatic benefit, or —if such a thing is possible— a case of brute causation.

When we discover that Jones's belief is not responsive to counter-evidence, then, we know that the right explanation of it must be either a pragmatically rationalizing or non-rationalizing one. Choosing between these possibilities requires more information about Jones. Were we to discover, for example, that he has been greatly looking forward to seeing this particular concert, then we might declare his belief to be a case of wishful thought, which is a type of pragmatically motivated believing. Alternatively, we might come to think that Jones's belief has been non-rationally determined, perhaps by some emotion or habitual behavior. Whichever explanation we end up endorsing for Jones's *idée fixe*, the important point is that Jones's lack of doxastic responsiveness has affected our explanation of his belief. When we offer an epistemic explanation, we have certain expectations as to how his belief will behave, and when those expectations are unfulfilled, we seek to change the explanation.

This is true in spite of the fact that Jones's belief was, we thought, brought about by an 'epistemic cause', namely his viewing posters advertising the concert. Jones's subsequent doxastic behavior shows us to have been wrong in this, that it was not a simple case of trusting an advertisement in order to find out about the world. In order for Jones's belief that p to deserve an epistemically rationalizing explanation, it is not enough that Jones's belief be 'triggered' by an epistemic reason or a perceptual experience. It is partly a matter of the origin of a belief not necessarily telling us why it is held. Features of the *origin*

of Jones's belief may not by themselves reveal that Jones is *aspiring for*, as it were, a true belief or knowledge. But the point runs deeper than this: it reveals the inadequacy of conceiving of rationalizing explanations of beliefs in terms of mere causality. The fact that expectations are built into rationalizing doxastic explanations reveals that a rationalizing answer to the question 'Why does Jones believe that *p*?' involves more than simply speaking of something that 'pushed' Jones into a state of believing. Rationalizing doxastic explanations give a sense of what Jones is *seeking* or *gaining* in believing what he does. Simply describing how Jones's belief that *p* was originally determined (e.g., by his viewing of posters) is wholly inadequate to capture this information. Saying what 'triggers' a belief falls far short of saying why an epistemic (or, as we will see below, a pragmatic) belief is held. Indeed, any such 'triggering' will be consistent with Jones's holding an epistemic or pragmatic belief.⁴

Viewing Jones's believing that *p* as an epistemic commitment involves seeing Jones's believing as a result or part of his seeking a truth about the subject matter at hand. And built into the very description of Jones as *seeking a truth* is the expectation that he will be responsive to other propositions about the world which he holds to be true. But this means that if Jones's belief is not so responsive, if it is something like an *idée fixe*, then it must not be an epistemic belief. There must be another story, other than an epistemically rationalizing one, about why Jones believes that *p*. In sum: as a doxastic explainer, I expect epistemically explicable beliefs to be evidence-sensitive; a necessary condition on giving an epistemically rationalizing explanation of *S*'s belief that *p* is that I take it to be evidence-sensitive.⁵

Donald Davidson comes near to describing this feature of doxastic explanation when he writes:

⁴ As I will discuss at length below, some of the notions I use in the above paragraph —*aspiring for*, *seeking*, *gaining*— should not be taken to pick out conscious or intentional actions.

⁵ There may be some few propositions —e.g., those of mathematics— for which the believer does not think that there could be such a thing as counter-evidence. For these, the above criteria for pragmatic believing will not apply; their firmness should not be seen as indicative of non-epistemic believing.

The general concept [of irrationality] applies also to unchanges. A person is irrational if he is not open to reason —if, on accepting a belief or attitude on the basis of which he ought to make accommodating changes in his other beliefs [...] he fails to make those changes. (Davidson 1982, p. 299)

This is right, but only partially so. We definitely see the evidence-insensitive believer as *epistemically* irrational; epistemic rationality requires susceptibility to new information about the world. This does not mean that we will see her as irrational *tout court*, however, for we may see her as appropriately gaining some *pragmatic* benefit with her belief; we may see her as epistemically irrational but judge her belief to be pragmatically appropriate.

I am presently concerned, however, not with judgments of doxastic rationality —that is, with what a believer *should* be doing— but with doxastic explanation —with what a believer *is* doing. And the point about explanation that Davidson’s observation misses is that the evidence-insensitive believer is to be *understood in a different light* from the evidence-sensitive believer. The former is not just being epistemically irrational; she has not simply made a mistake in seeking a truth. Her insensitivity to the counter-evidence before her shows that she is *not seeking truth at all*. A person in the grip of an *idée fixe* is not believing in order to gain some truth about the subject matter at hand; she is in the clutches of a belief that is not affected by other, clearly relevant, things that she holds true. Consequently, our account of her belief changes: we see her as being pulled to believe by some pragmatic benefit the belief is giving her, or as determined by some non-rationalizing cause.

When I epistemically explain Jones’s belief, I expect Jones to be sensitive to counter-evidence with respect to his belief; I expect him, roughly, to behave in such a way that he will maintain what he sees is a truth about the subject matter at hand. When I *pragmatically* explain someone’s belief, on the other hand, I have different, but analogous, expectations of her; I expect

the pragmatic believer to behave in such a way that she will continue reaping pragmatic benefit from her belief. Pragmatically rationalizing explanations recognize, as Richard Foley puts it, that ‘what one believes can dramatically affect one’s practical prospects [...]’.⁶ They explain by appealing to the belief’s contribution to the subject’s *well-being*, claiming that a belief is held because it brings the believer some utility other than truth. Accordingly, a pragmatically rationalizing explanation of someone’s belief that *p* involves an expectation that she will compensate in the face of challenges to her belief in order to maintain the pragmatic benefit that the belief confers.

Imagine that Smith has a daughter who is on trial, accused of a crime that he believes she did not commit. We explain his belief pragmatically: Smith believes that his accused daughter is innocent because he would feel enormous disappointment or pain were he to believe otherwise. In contrast to the epistemic case, in which we expect Jones to change his belief to fit his evidence, we expect Smith to respond to challenges to his belief *in order to continue to hold it*. It is only if Smith retains his belief that he will achieve the pragmatic benefit that he is gaining from this belief (e.g., comfort). Consequently, we expect to see Smith responding to challenges in such a way that he can continue to feel comfort or ease.

This would be most clearly demonstrated in the situation in which Smith were to give up the belief in his daughter’s innocence at the first sign of evidence for her guilt. In such a situation we would rightly grow doubtful about whether his belief really is the case of pragmatic believing that we initially thought it was. We would suspect that he was indeed after a truth of the matter about his daughter’s innocence, that he was following his evidence. We would suspect that he had always had an epistemic belief. As with epistemically rationalizing doxastic explanations, pragmatic explanations have expectations built into them.

⁶ Foley 1988, p. 131. I am ignoring one kind of pragmatic beliefs, those that Dion Scott-Kakures (2000) calls ‘unwelcome’. I think that the account of pragmatic believing I advocate in the final section of this paper applies to unwelcome believing as well, but I will not defend this here.

If Smith's belief that his daughter is innocent is indeed pragmatic, then there are an enormous number of ways in which Smith can respond to counter-evidence to his belief, evidence that his daughter is not innocent. He can undermine this counter-evidence, he can explain it away, or he can reinterpret it. Each of these strategies would allow Smith to retain his belief and continue to reap pragmatic benefit from his belief. This is precisely what we expect from those we see holding beliefs without seeking knowledge. We expect them to respond to our challenges in such a way as to *retain* the benefit-providing beliefs they have. Their aim is not to come to know whether p or not- p is true, but to retain the benefit of the beliefs they presently hold. They have a stake not in holding true beliefs, but in what they gain from the particular belief they hold.

This does not mean that Smith will *never* give the belief up. The extent to which Smith will compensate in order to retain his belief depends upon the features and context of Smith's particular belief (e.g., how strong is his desire for comfort) as well as how strong the challenges are to his belief (e.g., perhaps his daughter herself will tell him that she is guilty). The extent to which someone is susceptible to counter-evidence lies on a spectrum. At the epistemic end of the spectrum, the believer responds appropriately and sincerely in the face of counter-evidence; at the pragmatic end the believer will do everything in his power to reject counter-evidence that comes his way. If we think that Smith holds a pragmatically rational belief, we place him towards the latter end of the spectrum. We expect him to do what is *within his own doxastic integrity* to retain the belief, to ignore challenges or respond to them in such a way as to continue to retain the belief that is doing him some good.

The expectations that are built into rationalizing explanations are not precise and unambiguous. As are most aspects of our folk psychological practices, the constitutive characteristics of doxastic explanations are vague and inexact. However, anyone who has spent time arguing with other people will at

one time or another have experienced the feeling that expending any further effort in arguing would be futile. One feels that one's opponent is simply not interested in finding out the truth of the matter. She is not responding, appropriately and sensitively, to the considerations that one is introducing against her position. It will be at this point that one begins to suspect that one's opponent is (as we might say) pragmatically committed to her position, that she believes it because it is a position that is *important* to her. She believes it because she invented the position herself, she has spent so much time thinking about it, or she simply wants to beat you in an argument.

This sort of situation epitomizes the expectations that we have when we attempt to account for each other's beliefs. Argumentation—the presentation and consideration of counter-evidence to and defense of our beliefs—is a social practice in which we engage in order to epistemically improve our own and each other's beliefs. It is a central part of our epistemic process of attempting to get things right. Someone who aspires to get things right on a certain topic will be more sensitive to argumentation, and more affected by it, than someone whose believing is guided by other considerations. Both, though, will be responding to challenges in order to meet their aims.

We can see, then, that while epistemic and pragmatic explanations are both rationalizing, they are so in very different ways. Epistemic explanations work by showing that, from the believer's point of view, what she believes is likely to be true, while pragmatic explanations work by showing that the belief benefits the believer in some way. As such, epistemic and pragmatic doxastic explanations both carry expectations of the believer's future behavior. The expectation in both cases is that the believer will respond in order to either achieve a true belief—in the case of an epistemic belief—or receive the non-epistemic benefit—in the case of a pragmatic belief—that the believer is receiving from his belief.

3. *The Nature of Rationalizing Doxastic Explanations*

The fact that both kinds of rationalizing doxastic explanations bring with them expectations of the future behavior of the believer suggests that rationalizing doxastic explanations belong to one of the following three broad kinds:⁷

- (i) Explanations that cite an *occurrent dependence* between the belief and some other feature of the believer
- (ii) Explanations that portray the believer as *following a norm*
- (iii) Explanations that appeal to some *sub-intentional goal-oriented system* that the believer ‘possesses’ or is ‘part of’ the believer

Each of these strategies for accounting for a belief brings with it expectations about the future behavior of the phenomenon being explained. While there may be other possible accounts of rationalizing doxastic explanations, I do not know what they would be. In this section, I will discuss each of these three explanation strategies in turn. As we will see, the first sort of explanation is not a plausible candidate to account for the nature of rationalizing explanations of beliefs. In the final section of the paper, I will argue that the norm-following account is not a suitable candidate for explaining the nature of pragmatically rationalizing explanations.

(i) Occurrent Dependence Explanations

The case against conceiving of epistemic or pragmatic explanations of beliefs as being mere causal explanations cannot rest on the fact (established in the previous section) that offering an epistemic or pragmatic explanation involves having expectations

⁷The reader familiar with Daniel Dennett’s work will recognize an analogy between the three explanation strategies described in this section and Dennett’s physical, intentional, and design stances, respectively. See, e.g., his 1978. I will not in this paper explore any analogies or disanalogies that the claims I am making may have with Dennett’s work.

of the believer, for some mere causal explanations carry expectations with them. Occurrent dependence explanations account for the present occurrence of φ by citing the present occurrence of ψ , claiming that φ would not exist without ψ . Two examples:

The pan is on fire because the oil is too hot.

Jones's temperature is dropping because he is perspiring.

In each case, one fact or event is explained by citing the present occurrence of another fact or event. In the first example, the fact that the pan is on fire is explained by the fact that the oil is too hot.

In contrast to other brute causal explanations —those that cite what we might call 'triggering causes'— knowledge of occurrent causes lead to one's having certain expectations about the future behavior of the explanandum. So, for example, knowing that the pan is on fire because the oil is too hot can (indeed, should) lead one to expect that water will not put it out, because water can spread the burning oil rather than smothering it. The explanation of Jones's temperature drop in terms of his perspiring leads us to expect that his temperature will drop more quickly in less humid environments than it will in more humid environments. A triggering causal explanation, on the contrary, can lead to such expectation only on the assumption that the triggering cause is also the occurrent cause. The oil's being too hot may have been the triggering cause of the fire, but only if I also think that it is still the oil that is responsible for the fire in front of me will I have expectations about the behavior of the fire in response to my throwing water on it. If I think that hot oil triggered the fire but that the fire in front of me now is a burning kitchen towel that was too near the oil, then I will not have the same expectations of the behavior of this fire than I would were it to be burning oil.

While occurrent dependence explanations include the element of *expectation* in rationalizing doxastic explanations, they cannot, unfortunately, account for the fact that rationalizing doxastic explanations are *rationalizing*. As we saw above, rationalizing doxastic explanations essentially involve an attribution

of appropriateness to the believer in relation to his belief. This is not a contingent feature of rationalizing explanations. A rationalizing doxastic explanation has its explanatory power *in virtue of* the fact that it is rationalizing.

Take the following statement, which explains one belief by citing another:

Jones believes that his wife is home because he believes that today is Monday.

Understood merely as an occurrent dependence explanation, this statement derives its explanatory power in virtue of a supposed *correlation* or *law* linking two types of mental states, to which (i) Jones's belief that today is Monday and (ii) Jones's belief that his wife is home, respectively, belong. Such an explanation accounts for (ii) by citing (i) and a law claiming that beliefs like (i) and (ii) tend to go together. However, there is nothing in the *correlation* between these two beliefs to tell us why it is *appropriate* or *fitting* for Jones to believe that his wife is home. A doxastic explanation that merely appeals to an occurrent dependence will thus be missing an element that allows it to qualify as rationalizing, because it explains in some manner other than by presenting the believer as believing because it is appropriate for her to do so. Its explanatory power comes from a very different source from that of a rationalizing explanation. As we will see below in my discussion of teleological explanations, occurrent dependence statements can be embedded in a theory in such a way that they become rationalizing. The point remains, however, that a *mere* occurrent dependence explanation is not rationalizing, and not a suitable candidate for the nature of rationalizing doxastic explanations.

(ii) Norm-Following Explanations

A second sort of explanation that brings with it expectations is one that appeals to a norm that a person is following, adhering to, or attempting to fulfill. Norm-following explanations characterize the agent as φ -ing because of her relation to a norm governing φ -ing. Accounts of norms vary, construing them as

(i) ideals or standards, (ii) rules, prescriptions, or imperatives, or (iii) both. All accounts of norms have in common an understanding of something to which actions, character, and mental states can *comply* (or fail to do so), something which qualifies the action, character, or mental state as *successful* (or not).

There are two elements in a norm-following doxastic explanation. First, the agent is said to *accept* or *subscribe* to the norm. An agent can be *judged* in virtue of a norm whether or not she accepts it, but we can only *explain* her believing in virtue of a norm if there is some sense in which she accepts it. As with accounts of norms themselves, accounts of what it means for an agent to *accept* a norm vary widely. Some accounts stress the agent's membership in a community, while others stress her individual possession of dispositions or commitments.

The second element of a norm-following explanation is that the agent's ϕ -ing is said to be a case of *following* or *adhering* to the accepted norm. It is characteristic of norms that they are the sort of thing that can be followed, either as ideals or as prescriptions. Norm-following doxastic explanations need not present the belief as resulting from the subject's *trying*, consciously or otherwise, to meet a norm. We often are engaged in certain investigative *actions*, in which we are *trying* to get at the truth. However, there is a difference between *doing* something —acting— in adherence with a norm, and *believing* something in adherence with a norm. The latter may be a case of norm-following, even though it is not itself an action.

In a norm-following doxastic explanation, an agent's believing that p is said to be a case of attempting to meet a particular ideal, or of abiding by a doxastic prescription. The agent's believing that p is dependent upon her thinking (although not necessarily consciously doing so) that she *should* believe that p . She believes that p because she accepts a norm that dictates her doing so. A norm-following explanation of, for example, Jones's perceptual belief that I am eating garlic will claim (i) that Jones accepts a norm that, say, vision or olfaction is a proper, successful, or reliable way to form beliefs about the world,

and (ii) that he believes that I am eating garlic because he is having a visual or olfactory experience as such. Jones has this belief because he adheres to a norm, and his believing is a case of following it. Evidential doxastic explanations, on the other hand, appeal to *epistemic reasons* for belief: Jones believes that he has the measles because he takes the red spots on his stomach to be a sign or indicator of his having measles, and he accepts a norm that says (something like) beliefs in the unobserved should follow observed signs and indicators.

In the final section of this paper, I will argue that there are no pragmatic doxastic norms, but if they were to exist, doxastic explanations in virtue of them would go something like the following: Smith believes that his daughter is innocent because (i) Smith accepts and adheres to a norm that says that one should or is allowed to form beliefs that alleviate one's pain or anxiety, and (ii) Smith is following this pragmatic norm in believing that his daughter is innocent. In both cases, the explanation cites both a norm that Smith accepts, and furthermore claims that his believing what he does is in virtue of following that norm.

In Section 1 of this paper, I claimed that most doxastic explanations explain by attributing a minimal amount of appropriateness to the believer's holding the belief that she does. What is this appropriateness, and why is it minimal, on the norm-following account of doxastic explanations? The appropriateness derives from the believer's believing because of a norm; Jones believes that p because Jones accepts some doxastic norm that dictates or suggests that Jones *should* believe that p . Jones is aiming for a truth about the matter at hand, and he accepts some norm that says that in certain conditions a belief is likely to be true. Jones's belief that p is appropriate is minimally justified, because it is a belief that, by Jones's own lights, he should hold. However, for me to recognize and explain Jones's believing that p in this way is, importantly, still short of my agreeing with or endorsing Jones's belief that p . I can think that Jones should not believe that p , all the while recognizing that Jones believes because *he* thinks that

he should believe that p . I can disagree with Jones, and still recognize that there is *some* rationale for his believing p . So, in contrast to occurrent dependence explanations of beliefs, norm-following explanations get the rationalizing element in doxastic explanations just right.

In Section 2, we saw that offering a rationalizing doxastic explanation brings with it expectations of the future behavior of the believer we are explaining. This is readily accounted for by the norm-following account. If I think that an agent is believing in adherence to a particular sort of norm, then I will expect her, in the face of challenges to her belief, to continue to do so. If, for example, I think that Jones is following a norm like ‘Believe in accordance with the evidence before you’, then I will expect his belief to be sensitive to deal with any counter-evidence that may come his way. We expect his beliefs to lessen in the face of counter-evidence to which he cannot respond, and strengthen in the face of further positive evidence. Similarly, a pragmatic norm-following explanation would present the believer as following a norm like ‘Believe so that you avoid the anxiety or pain that would be brought on were you to believe the opposite’. Such a norm would dictate that the believer believe so as to gain some pragmatic benefit from her belief, and as such we expect that she will either avoid dealing with counter-evidence, or deal with it in such a way that the belief—bringing her benefit—is still held. She will do what she can in order to retain her belief, so that she can persist in reaping its benefit. In both cases, once we posit—in our explanation—that the agent is following a norm, we thereby expect her to respond to any challenges that she comes across in such a way as to continue to follow it.

(iii) Sub-Intentional Teleological Explanations

In “Self-Deception and the Nature of Mind”, Mark Johnston (1988) draws a distinction between intentional and sub-intentional purposive processes, and argues that self-deception is best conceived in terms of the latter.

If we call mental processes that are purposive but not initiated for and from a reason *sub-intentional* processes, then we can say that [the] over-rationalization [by certain writers] of self-deception consists in assimilating sub-intentional processes to intentional acts, where an intentional act is a process initiated and directed by an agent because he recognizes that it serves a specific interest of his. (Johnston 1988)

Johnston opposes this ‘over-rationalization’, defending instead a sub-intentional —or as he also calls it, a *tropistic*— view of self-deception.⁸

I am going to adopt the distinction that Johnston draws between sub-intentional and intentional *processes* in order to introduce a type of doxastic *explanation* distinct from the norm-following strategy discussed in the previous sub-section. What I will here call *sub-intentional teleological explanations* are to be contrasted with norm-following explanations by the fact that they drop all reference to the agent’s relationship to a norm or rule. In a sub-intentional teleological account of a belief, Jones’s thinking that he should believe that *p* plays no role in his doing so. Instead, the teleological explanation models the belief as the outcome or product of a goal-oriented system.

Teleological systems involve what Arthur Collins calls a ‘two-sided causal relationship’. A disturbance to the system *causes* the system to *cause* goal maintenance. This capacity to *compensate*, Collins claims, is ‘the essential core’ of teleological systems: “A system [*S*] is teleologically organized only if threats to goal maintenance cause events in *S* that cause goal maintenance in spite of the threats.”⁹ A system is teleological only if it responds to perturbations in some way appropriate towards continuing to achieve the set goal of the system. It is notable that Collins (like other writers) prefers to speak of teleological ‘systems’, and not, say, ‘processes’. A system is teleological only if there are *various* processes by which that system can maintain

⁸ The claim (although not the argument) of the next section of this paper is in broad agreement with Johnston’s conception of motivated believing.

⁹ Collins 1978, p. 544. My discussion in this sub-section draws heavily upon Collins’s work on teleological processes.

or work towards its goal. A system is goal-oriented only if it has at its disposal multiple processes for maintaining its goal. This is not to say that a process cannot itself be considered teleological, but it will be only in so far as it is part of a goal-oriented system. A system is compensatory, and thus teleological, because of its possible utilization of various processes; it continues working, aiming, to reach an intended goal under different conditions.¹⁰

The relationship between the two properties “. . . is a compensatory system” and “. . . is a goal-oriented system” is a conceptual one. Goals are inherent to compensating systems, and these make little sense without them. Compensation just is compensation in response to stimulus *for the sake of* goal-maintenance. A waterfall and a ball working its way down a pinball machine are *directional* processes, but if one describes the various structures redirecting a drop of water or a pinball in terms of *compensation*, one must do so metaphorically. Waterfalls and pinballs do not have goals, and one cannot make sense of compensation without appreciating that there is a state that a compensatory system *works to* achieve or maintain. Compensation is compensation *on behalf of* some state, in virtue of which that state is the goal of the system.

Collins’s favored example of a compensatory system is that of homeostasis, the biological maintenance of body temperature via processes like perspiration, shivering, and blood vessel dilation and contraction. He writes:

Things happen in the environment of the human body that would, of themselves, cause significant changes in internal body temperature, but body temperature does not change much. Temperature stability is maintained in the face of destabilizing factors. (Collins 1984, p. 348)

¹⁰ Other writers prefer the notion of *plasticity* to that of *compensation*, but the concepts are essentially the same. See, for example, Enç and Adams 1992, at pp. 650ff. See also any of the writings associated with the abstract study of systems —called ‘cybernetics’— concerned with control and goal-maintenance. For an introduction and bibliography, see Heylighen and Joslyn 2001.

This physiological example allows us, first, to emphasize the fact that we identify systems as compensating, and thus teleological, without attributing to the agent any pro-attitude or norm-acceptance. Secondly, processes like homeostasis reveal that we can identify a system as teleological *before* we know the various causal pathways by which it compensates. We knew that constant temperature was maintained in mammals long before we knew *how* it was maintained. In this case we attributed a goal, that of maintaining the right temperature, to mammalian bodies, before we knew the particular causal processes by which temperature is maintained.

This latter point, that we are able to identify teleological systems independently from the causal processes that such a system utilizes in maintaining goal-orientation, is important to understanding how a teleological explanation works. Contrast the straightforward *causal claim*

Jones's perspiring causes his body temperature to drop,

with the *teleological claim*

Jones's body maintains its temperature over a range of external temperatures.

Even though Jones's ability to perspire is part of his system to control his body temperature, the teleological claim is independent of the causal claim. The teleological claim can be true even though causal claim is not. Even if Jones does not, for some reason, have the capacity to perspire, he still must maintain homeostasis. In most animals, there will be any number of causal processes that orient the creature towards the goal of temperature-maintenance.

The converse of this is true as well. We can and do learn causal claims without at the same time being aware that they are parts of teleological systems. We can imagine discovering that Jones's perspiring causes his temperature to drop before discovering that his body maintains its temperature at a constant state. In this case we might think that the fact that perspiring

causes a drop in temperature is merely a brute causal fact about the human body, and not part of any goal-oriented system. It may be much later that we discover the teleological fact that the body maintains itself at a constant temperature via different means. This order of progression, from first discovering a process and then later attributing a function to that process, will be familiar to the biologist.

Discovering a teleological system, of course, changes our understanding of the causal aspects of that system. Once we know that perspiration is part of a teleological system, then the explanation 'Jones's temperature is dropping because he is perspiring' becomes not just an explanation in terms of occurrent dependence, but an explanation in terms of the means to a goal. This is the vital point of the recognition of teleology. Once we realize that perspiring-causing-temperature-drop contributes to a wider compensating system that maintains body temperature, then the explanation 'Jones's temperature dropped because he was perspiring' becomes an instrumental, or goal-oriented, explanation. Once we see that there is a system expressing compensating behavior, then we will see a causal contributor to that system as a part of a teleological process. Perspiring is no longer a *mere* causal dependence, but the means to a goal.

This change of view brings with it the susceptibility to *assessment* or *evaluation*. A mere causal process, like a body's being wet causing its temperature to drop, is not open to evaluation. In contrast, a goal-oriented causal process, like perspiring, is. Once we see perspiring as a part of a teleological process, we can compare it to other actual or possible means to the goal of temperature maintenance. We can assess its cost, its speed, its sensitivity, or its efficiency as a means to its goal. In isolation from the goal-oriented system of temperature maintenance, such assessment makes no sense. There is no evaluating the fact that perspiring causes one's temperature to drop until we switch to thinking of perspiring *as a means to* temperature maintenance. Once we make that switch, however, evaluation becomes possible.

All of these elements can be drawn together into a sub-intentional teleological account of rationalizing doxastic explanations. A teleological explanation of Jones's believing that p involves claiming that Jones has a doxastic goal, an epistemic or pragmatic goal that his belief is held in order to gain. The minimal element of *appropriateness* that, as we saw in Section 1, is central to rationalizing explanations takes this form in teleological doxastic explanations: Jones has a goal, and he believes that p in order to achieve that goal. Jones's belief that p is minimally justified as an attempt to achieve a said goal. For the explainer to see this is, importantly, still short of her agreeing with or endorsing Jones's belief that p ; in order to do that, the explainer would have to endorse the details of why Jones believes that p . But one can disagree with Jones, and still recognize that there is *some* rationale for his believing p , namely as an attempt to attain a goal that Jones has. So, in the positing of goals the teleological account gets the rationalizing element in explanations just right.

Since teleological doxastic explanations are built upon the attribution of a goal to the believer, they bring with them expectations, expectations deriving from the compensatory behavior that the teleologist says the believer will undertake in order to achieve that goal. Part of what it means to conceive of something as *working towards a goal* is to conceive of it as engaged in a practice in which it responds to challenges to its gaining that goal. When one sees a system as goal-seeking, one *expects* the system to make necessary changes to attain it. Without such expectations, the explainer would not be conceiving of the system as teleological at all. To see a belief as brought about by or as a part of a goal-oriented system would necessarily involve the expectation that believers respond to challenges to their beliefs in such a way as to continue to achieve the doxastic goals that we are assigning them in the explanation. This would be reflected in different ways in epistemic and pragmatic doxastic explanations, each of which have different goals and, thus, bring with them differing expectations.

4. *Pragmatic Believing As Sub-Intentional*

Of the three types of doxastic explanation strategies that lead to our having expectations about beliefs, we have seen that one strategy—that of occurrent dependence explanations—does not include any features that could constitute the rationalizing nature of (most) explanations of beliefs. The remaining two strategies, norm-following and teleological, are (as far as we have seen) suitable for capturing the rationalizing nature of doxastic explanations. In this final section, I will focus my attention upon *pragmatically* rationalizing explanations, regarding which I will cast doubt on the appropriateness of a norm-following account. My argument relies on a distinction between pragmatic and epistemic believing, and so it is not applicable to epistemic believing. I will not address the question of whether epistemic doxastic explanations are best conceived as teleological or norm-following.

As we saw above, norms need be neither consciously considered nor consciously followed. I can adhere to a norm without ever having thought explicitly about what that norm is. Norms of language and concept-use are well-known examples. Jones's using a concept in a certain way at a certain time can be explained in virtue of her accepting and following a particular norm, even though Jones could not state the norm that she is following. Nor, *a fortiori*, need it be true that in order to have my behavior dictated by a norm, must I be thinking about the norm at the time of use. I can follow a norm without being aware of what I am doing, without, that is, knowing what norm I am behaving in accordance with. Norms play a pervasive role in our lives, but in doing so they often remain in the background, hidden from our conscious purview of what we are doing.

There is a difference, however, between claiming that a norm *need not* be consciously followed, and claiming that it *cannot* be consciously followed. The latter makes no sense. It is essential to something's being a norm that it be such that we can come to follow it. It must be possible for us to be consciously aware that we are doing something in order to be said to be adhering

to a norm in doing it. Even if not consciously followed, a norm must be consciously *followable*.

Norms play a central role in our lives as agents. Indeed, crucial to our identity as agents is our ability to allow reflective adherence to norms, to what we think we ought to do, to guide our lives. We form beliefs about what we should do, we form beliefs about what the right norms are, and those beliefs lead us to behave in accordance with them. It is because it is possible for me to act *because* and *in the light of* my reflection upon what I think is the right thing to do, that I am an agent. My agency is essentially tied up with the fact that norms can, in my full awareness, dictate what I do.

The essential followability of norms raises the question that motivates Christine Korsgaard's discussion in *The Sources of Normativity*.

[A] successful normative theory must meet a condition which is sometimes called 'transparency' [...]. If a theory's explanation of how morality motivates us essentially depends on the fact that the source or nature of our motives is concealed from us, or that we often act blindly or from habit, then it lacks transparency. The true nature of moral motives must not be concealed from the agent's point of view if those motives are to be efficacious.

Otherwise, Korsgaard claims, the theory would not be a *normative* theory at all: "A normative moral theory must be one that allows us to act in the full light of knowledge of what morality is and why we are susceptible to its influences, and at the same time to believe that our actions are justified and make sense" (Korsgaard 1996, p. 17). A non-transparent moral theory would ultimately amount to a portrayal of moral behavior as other than norm-following. It would present it as motivated by considerations that are hidden from us. "The normative word 'reason'", Korsgaard writes, "refers to a kind of reflective success" (Korsgaard 1996, p. 93). Behaving on the basis of an accepted norm is to behave on the basis of something that one can consciously endorse and follow.

If someone tells me that I ought to φ , and also tells me that I will not be able to φ in full consciousness, then it would seem that she is telling me not, strictly speaking, that I *ought* to φ , but rather that it would be a *good* thing if I were to φ . Imagine someone's saying to you, 'You ought to get over your illness'. She is not giving you a norm, she is not telling you that you should get over your illness. Rather, she is telling you that it would be beneficial, a good thing for you, if you were to get over your illness, that she hopes that you do get over your illness. There is no norm 'one ought to get over one's illness' because getting over one's illness is not something one simply does. There are, of course, norms *in the area* of health and well-being. One should allow oneself to get over one's illness (e.g., before going back to work, before holding the baby), and one can do things that may help one get over one's illness (e.g., taking antibiotics). However, simply getting well is not one of them. Getting well is not, *simpliciter*, an example of one's agential behavior.

The reason why there are no pragmatic doxastic norms is related to, but not precisely the same as, the reason why there is no norm of simply 'getting over one's illness'. Like getting well, wishful thinking is not an example of one's agential behavior. However, in contrast to getting well, it is reflection itself that prevents wishful thinking from being an agent-level behavior. Being aware of getting well does not prevent me from doing so. The same is not true, however, of wishful thought. We cannot, in full awareness, follow a norm that would tell us to believe something because doing so would bring us some benefit, or because it would be a good thing to do.

The reason for this is that we are restricted in how we see our beliefs. My believing that something is true depends upon my not accepting a non-epistemic explanation of my believing it. The stronger is my conviction that my believing that p rather than q , or my believing that p rather than suspending belief on the matter altogether, depends upon pragmatic concerns, the weaker will be my belief that p .

It is this feature of beliefs that differentiates them from guesses. We can in full awareness guess in order to achieve a pragmatic aim—as we do, say, in trivia games—but we cannot believe with the same aims in mind. By thinking, ‘I have to say *something* if I am going to win [. . .]’, the game contestant characterizes what she is about to say *as* a guess and not as a belief. I cannot see myself as believing that p (rather than that q) *so that* I will pass a test or win a game. The believer cannot see her belief—her commitment that something is true—as not held solely in order to gain a truth. The belief is one’s *take on* the truth, and not one’s, say, *pressurized stab* at the truth or one’s *idle speculation* of the truth. The difference is not merely one of degree. It is rather that one sees one’s take on the truth (as opposed to these other attitudes) as not being held with the aim of relieving pressure or filling idle time. As David Owens has written: “The guesser can exercise control over her guesses by reflecting on how best to strike a balance between the goal of truth and other goals her guessing serves; the believer cannot” (Owens 2002, p. 395). Both the believer and the guesser aim at truth, but while the latter can think that her guess is dependent upon her attempt to gain some pragmatic end, the believer must not. To think, say, that being in a hurry has affected my *inclination* towards p , that I have this inclination in order to have it *now*, is to see my commitment to p as something like a guess, something other than a belief.

I cannot see my own belief as not being solely dependent upon my aiming to grasp some truth about the subject matter at hand. If I were to explain my believing that p non-epistemically, then I would be seeing my belief as aiming for some goal other than truth, or for no goal at all. These are incompatible attitudes to take towards a belief. I cannot attribute the exclusive goal of truth to the belief (by continuing to consciously believe) while taking it away (by accepting a non-epistemic explanation). Therefore, it is impossible to hold a belief and non-epistemically explain my holding it.

Stated again: in reflecting that I have a belief that p , I characterize myself as being in a state *only* for the purpose

of being committed to a truth about the subject matter at hand. However, I could not do this and at the same time see myself as following a pragmatic doxastic norm. If in full awareness I were to believe that p because of a non-epistemic norm dictating that I should do so, then I would see my commitment as aiming for some goal other than truth. I would understand myself as believing that p (rather than that q) because it would be a good thing for me to believe that p . But this is not possible, for such a state would not be a belief at all.¹¹ To consider and retain one of my current doxastic states requires seeing it as truth-aimed, and this is inconsistent with seeing myself as seeking a pragmatic doxastic benefit. Therefore, holding a doxastic commitment is essentially incompatible with seeing oneself as following a pragmatic norm.

We all recognize, of course, that certain of our beliefs fulfill other goals of ours, that our having them brings us benefit. Our moral beliefs are prime examples. But the believer must see each of these gains as *derivative* from the goal of truth. When I reflect upon a belief, I must see the possession of a truth as the primary aim of my believing what I do. To do otherwise would be to undercut the belief, and essentially to see it as something other than a belief.

The fact that beliefs cannot survive first-person pragmatic explanations of them explains some of the dialectical phenomena that Richard Foley describes.

[W]hen people reflect upon what reasons they have to believe something [...], they rarely even consider the practical advantages that might accrue to them by believing it [...]. Likewise, when someone tries to convince another person that he has reasons to believe something, they rarely even mention the practical benefits that might result from believing it. (Foley 1988, p. 214)

¹¹ Compare the often-quoted line from Bernard Williams: "If in full consciousness I could will to acquire a 'belief' irrespective of its truth, it is unclear that before the event I could seriously think of it as a belief, i.e., as something purporting to represent reality." From his 1973, p. 148.

We do not consider the benefits we may gain from holding beliefs which support our values or placate our desires, even though we may at times see that believing such things *would* bring us happiness, comfort, and satisfaction, and would do so even if what we believed was false. Nor, *ipso facto*, do we attempt to share this satisfaction by pointing out these practical benefits to others. The explanation for both phenomena is the same. If we were to see, upon reflection, that we have a belief in order to bring us pragmatic benefit, then the belief would disappear.¹²

The upshot of this is that a pragmatic doxastic norm would not be consciously followable. We cannot, in full awareness, believe by following a norm the content of which is something like, 'You should believe that *p*, because doing so would be good for you, or would make the world a better place'. Such a norm would be unfollowable. However, an unfollowable norm is not a norm at all. It may describe a good state of affairs, something that it would be good for us to do, but it is not a norm, something that can *guide* our behavior. While it might be a good thing for us to engage (at times) in a belief because it will bring us comfort, make us better people,¹³ or make the world a better place, it is not something that we can do *in adherence to a norm*. It follows that there is no such thing as a pragmatic doxastic norm, and that pragmatically rationalizing doxastic explanations are sub-intentional in nature.¹⁴ Pragmatically rationalizing doxastic explanations properly present a believer as believing *appropriately*, but not as believing in the face of *reasons* to believe what she does.

This is not to say that there is no such thing as a practical *argument* for a belief; however, it does mean that such arguments must work, and be understood, in a particular way. The most

¹² The claim of the previous few paragraphs is defended at length in Jones 2002.

¹³ In their much-discussed 1988, Taylor and Brown suggest that motivated believing is an essential characteristic of mental health.

¹⁴ In Dennett's terminology, pragmatic believing is to be explained from the design, and not the intentional, stance.

well-known practical argument for a belief is Pascal's Wager. Pascal argued that we should believe that God exists, as well as the rest of the Christian creed, because the utility of doing so is infinitely greater than the utility of not holding these beliefs. Yet Pascal was well aware that we could not consciously follow his Wager. After introducing the Wager, he has his interlocutor (who represents his ideal agnostic reader) respond, "Yes, but my hands are tied and my lips are sealed; I am being forced to wager and I am not free; I am being held fast and I am so made that I cannot believe. What do you want me to do then?" (Pascal 1966, fragment 418).¹⁵ Pascal's suggestion to his reader is well-known:

Learn from those who were once bound like you and who now wager all they have. These are people who know the road you wish to follow, who have been cured of the affliction of which you wish to be cured: follow the way by which they began. They behaved just as if they did believe, taking holy water, having masses said, and so on. That will make you believe quite naturally, and will make you more docile. (Pascal 1966, fragment 418)

Be among believers and *act* like a believer, and you will come to believe. His advice is complex, indeed so complex that it occupies the bulk of the *Pensées*.¹⁶ However, the important point here is that Pascal's Wager, as a practical argument for a belief, initially targets —*as it must*— actions. Pascal knew that making evident the benefit of religious belief could not, just like that, consciously bring about belief. The potential convert must take an indirect route, following norms for action that will in turn bring about the desired belief.¹⁷

¹⁵ Lafuma numbering.

¹⁶ For a discussion of the rationality of Pascal's suggestions to the non-believer, and in particular of whether it involves self-deception, see Jones 1998.

¹⁷ A number of writers are not aware of this. In his 1993 (at pp. 85–86), for example, Robert Nozick advocates a pragmatic rule for believing that we could not in full awareness follow.

Conclusion

Both epistemically and pragmatically rationalizing doxastic explanations bring with them expectations of the future behavior of believers, and I have taken this to suggest that rationalizing explanations are either teleological or norm-following. However, the norm-following account is inappropriate in the case of pragmatic beliefs. Pragmatic norms of believing could not be followed, and, thus, simply do not make any sense. Pragmatic believing must be a sub-intentional, and not a norm-following, process. Nothing I have said counts against the possibility that epistemic believing is best conceived as a norm-following process. If it is, then the correct accounts of epistemic and pragmatic believing differ. Motivated, or pragmatically induced, believing is the result of our sub-intentional systems, while epistemic believing is an expression of our behaving as agents, of our believing because we think we should.¹⁸

REFERENCES

- Collins, Arthur, 1984, "Action, Causality, and Teleological Explanation", in Peter French *et al.* (eds.), *Midwest Studies in Philosophy IX*, University of Minnesota Press, Minneapolis.
- , 1978, "Teleological Reasoning", *Journal of Philosophy*, vol. 75, pp. 540–550.
- Davidson, Donald, 1982, "Paradoxes of Irrationality", in R. Wollheim and J. Hopkins (eds.), *Philosophical Essays on Freud*, Cambridge University Press, Cambridge.
- Dennett, Daniel, 1978, "Intentional Systems", *Brainstorms*, Bradford Books, Cambridge, Mass.
- Enç, Berent and Fred Adams, 1992, "Functions and Goal Directedness", *Philosophy of Science*, vol. 59, pp. 635–654.

¹⁸ Thanks to Chris Allsobrook, Michaela Baker, Quassim Cassam, Nick Ferreira, Dylan Futter, Elisa Galgut, John Greco, Franco Huambo, Sharron Kraus, Tom Martin, Eusebius McKaiser, Brian McLaughlin, Don Ross, David Ryan, Nikolai Viedge, Heath Waller, and Francis Williamson. I worked on this paper during a visit to the Philosophy Unit at the University of Dar es Salaam, and I thank the University for the invitation and generous support during my stay.

- Foley, Richard, 1988, "Some Different Conceptions of Rationality", in E. McMullin (ed.), *Construction and Constraint*, Notre Dame University Press, Notre Dame.
- Heylighen, Francis and Clif Joslyn, 2001, "Cybernetics and Second-Order Cybernetics", in *Encyclopedia of Physical Science and Technology*, ed. R.A. Meyers, 3rd. ed., Academic Press, New York.
- Hursthouse, Rosalind, 1991, "Arational Actions", *Journal of Philosophy*, vol. 88, pp. 57–68.
- Johnston, Mark, 1988, "Self-Deception and the Nature of Mind", in B.P. McLaughlin and A.O. Rorty (eds.), *Perspectives on Self-Deception*, University of California Press, Berkeley, pp. 63–91.
- Jones, Ward E., 2002, "Explaining Our Own Beliefs: Nonepistemic Believing and Doxastic Instability", *Philosophical Studies*, vol. 111, no. 3, pp. 217–249.
- , 1998, "Religious Conversion, Self-Deception, and Pascal's Wager", *Journal of the History of Philosophy*, vol. 36, no. 2, pp. 167–188.
- Korsgaard, Christine, 1996, *The Sources of Normativity*, Cambridge University Press, Cambridge.
- Longino, Helen, 1979, "Evidence and Hypothesis: An Analysis of Evidential Relations", *Philosophy of Science*, vol. 46, pp. 35–56.
- Nozick, Robert, 1993, *The Nature of Rationality*, Princeton University Press, Princeton.
- Owens, David, 2002, "Epistemic Akrasia", *The Monist*, vol. 85, no. 3, pp. 381–397.
- Pascal, Blaise, 1966, *Pensées*, trans. A.J. Krailshiemer, Penguin, London.
- Pettit, Philip, 1979, "Rationalization and the Art of Explaining Action", in Neil Bolton (ed.), *Philosophical Problems in Psychology*, Methuen, London, pp. 3–19.
- Scott-Kakures, Dion, 2000, "Motivated Believing: Wishful and Unwelcome", *Noûs*, vol. 34, no. 3, pp. 348–375.
- Taylor, Shelley E. and Jonathon D. Brown, 1988, "Illusion and Well-Being: A Social Psychological Perspective on Mental Health", *Psychological Bulletin*, vol. 103, no. 2, pp. 193–210.
- Williams, Bernard, 1973, "Deciding To Believe", *Problems of the Self*, Cambridge University Press, Cambridge.

Received January 5, 2004; revised September 21, 2004; accepted October 6, 2004.