

Normativism and Doxastic Deliberation

Conor McHugh

When we deliberate about what to believe, considerations relevant to what is *true* play a special role in motivating us to believe or not believe one thing or another. This special role should be understood, according to Nishi Shah, in terms of the phenomenon of *transparency*. Shah takes transparency as the base of an inference to the best explanation, whose conclusion is a *normativist* account of (the concept of) belief. In this paper I distinguish the alleged phenomenon of transparency from a range of other possible phenomena that might characterise truth's role in doxastic deliberation. I do not address the question which of these phenomena are genuine and which spurious. Rather, I argue that normativism does not explain transparency (if genuine), and furthermore that it explains at most a relatively weak claim about *rational* doxastic deliberation. Since this claim can also be explained in other ways, including by rival accounts of belief, doxastic deliberation gives no support for normativism.

1. Doxastic Deliberation

Doxastic deliberation is deliberation framed by the question what to believe about some subject matter (e.g. whether to believe p) and that aims to conclude in the formation of a belief about that subject matter (e.g. the belief that p or the belief that $\sim p$).¹ Although doxastic deliberation is deliberation about what to believe, it involves a special role for a seemingly different question, namely what is *true* (e.g. whether p is true; or simply whether p). Normally, the considerations on the basis of which a doxastic attitude is formed through deliberation are considerations seen as relevant to what is true with respect to the subject matter addressed (e.g. evidence about whether p). Truth thus has a special role in the

¹ Doxastic deliberation can also conclude in withholding belief (suspension of judgment) about the relevant subject matter. This is an unsuccessful outcome, in so far as doxastic deliberation involves the aim to make up one's mind.

deliberative motivation of belief.²

Nishi Shah characterises this special role of truth in doxastic deliberation in terms of *transparency*, which he describes thus:

“when asking oneself *whether to believe that p*, [one] must ... immediately recognize that this question is settled by the answer to the question *whether p is true*[.] ... Within the perspective of first-personal doxastic deliberation, that is, deliberation about what to believe, one cannot separate the two questions.”

(Shah 2003, 447)

Transparency is thus a relation between two questions. Question A is transparent to question B when a subject can and must answer question A simply by answering question B, and when the subject must immediately recognise this on asking herself question A. Shah's claim is that the question whether to believe *p* is, from the point of view of the deliberating subject, transparent to the question whether *p* is true (see also Shah and Velleman 2005, 499).

For Shah, truth thus excludes *all* other considerations from doxastic deliberation, and does so *immediately*. It is not that, when engaging in deliberation about what to believe, you reason that pragmatic (non-truth-directed) considerations are irrelevant, or motivationally inefficacious, and so you intentionally ignore them. Rather, on asking the question whether to believe *p*, your attention immediately focuses exclusively on evidence (truth-directed

² By using the term ‘deliberative motivation’ I do not intend to suggest that beliefs are formed voluntarily. Rather, I refer to cases in which there are reasons *for which* one believes, where these reasons are or can be taken up in consciousness and their probative force as reasons for belief explicitly or implicitly recognised, and where one believes directly on the basis of these reasons.

Motivation in this sense, while not necessarily involving voluntary control, is not mere causation either. Pragmatic considerations certainly can causally influence belief (as Shah is careful to point out), but this does not mean that they can be reasons for which one believes. The claims I discuss will all concern the kinds of considerations that can play the role of reasons for which one believes in the deliberative motivation of belief. On the relations between motivation, control and voluntariness, see McHugh (2011a, 2012a).

Note also that I am exclusively concerned with reasons for which one forms a *particular* belief—such as the belief that *p*—and not with reasons for which one goes about forming a belief about a particular subject-matter.

considerations) (Shah *loc cit.*). Considerations to do with whether it would be pleasant, advantageous or amusing to believe p , independent of whether p is true, do not occur in deliberation as bearing on whether to believe p . In this sense, transparency is supposed to be a feature of the “phenomenology of deliberation” (*ibid.*, 462)—of what it's subjectively like to deliberatively ask yourself whether to believe p . The claim is that it is subjectively equivalent to asking yourself whether p .

So:

Transparency: Pragmatic considerations cannot occur to a thinker, within doxastic deliberation, as relevant to what to believe.³

One might find the claim that doxastic deliberation is characterised by transparency too strong.⁴ Is it really an invariable feature of the phenomenology of doxastic deliberation that the thinker who asks herself whether to believe p cannot so much as be struck by pragmatic considerations as relevant to that question? One might wish instead to make the more modest claim that only (what the thinker sees as) evidence can deliberatively *motivate* the thinker within doxastic deliberation.⁵ Even if considerations of what is pleasant, advantageous or

³ What kind of impossibility is expressed by "cannot" in this formulation? For the purposes of the paper we can understand it as psychological impossibility, rather than something stronger like metaphysical or conceptual impossibility. There are two reasons for this.

First, transparency is supposed to be a datum, but, even if obviously true (which I doubt), it is certainly not *obviously* a metaphysical or conceptual truth. Anything stronger than a claim of psychological impossibility would have to be supported by philosophical theorising.

Second, since my target is the explanatory power of normativism, it is charitable to make the explanandum relatively weak.

For the rest of this paper, claims about what thinkers can or cannot do may be understood as expressing only psychological possibility or impossibility, unless otherwise indicated.

⁴ Shah's characterisations of transparency are not all equivalent. At times he seems to mean the putative feature I described; at others he seems to mean something weaker—what I will call strong exclusivity (e.g. Shah 2003: 453). I will reserve the term 'transparency' for the stronger feature.

Shah and Velleman (2005: 517-8) distinguish between deliberation about whether to “accord” a “pre-existing acceptance” that p the status of a belief, and deliberation about whether *tout court* to believe p . Transparency, they claim, is a feature of the latter sort of deliberation only. I restrict my attention to this.

⁵ By ‘deliberatively motivate’ I mean function within deliberation as reasons for which one (say) forms a belief, or as considerations that provide some motivation to (say) form a belief even if one does not actually

amusing to believe can occur within doxastic deliberation as relevant to whether to believe *p*, one might say, such considerations cannot move the thinker, in the guise of reasons, through doxastic deliberation. The thinker's doxastic state will be unaffected by them. Truth excludes all other considerations from doxastic deliberation, not phenomenologically, but at the point of motivational efficacy.

This putative feature of doxastic deliberation is *exclusivity*.⁶ It comes in stronger or weaker versions.⁷ The difference between the strong and weak versions reflects the fact that there are two ways in which doxastic deliberation can conclude: it can conclude in the forming of a belief, or it can conclude in the forming of a state of withholding belief (aka suspension of judgment). Thus, it might be claimed that pragmatic considerations cannot be motivationally efficacious at all in doxastic deliberation, either in providing motivation to believe or in providing motivation to withhold belief. Or, it might be claimed only that pragmatic considerations cannot provide deliberative motivation to believe, even if they can sometimes provide such motivation to withhold belief.⁸

Strong exclusivity: Pragmatic considerations cannot deliberatively motivate belief or withholding of belief.

Weak exclusivity: Pragmatic considerations cannot deliberatively motivate belief.⁹

form it. See n.2 above.

There may be an argument that if considerations of a certain kind cannot deliberatively motivate you to do something, then they are not normative reasons for you to do that thing (see Shah 2006). But this would not amount to transparency, which has to do with what the deliberating subject can *see* as reasons that count in favour of or against believing. In any case the point won't matter for my argument.

⁶ I take this label from Steglich-Petersen (2009).

⁷ I discuss this difference in McHugh (forthcoming a)

⁸ In what follows I will sometimes talk about considerations 'deliberatively motivating' or just 'motivating' belief, or withholding. This is meant to include cases where a consideration provides some motivational 'push', through deliberation, to believe (say), even if you do not in fact form the belief (e.g. because there are other considerations pushing you the opposite way).

⁹ The claim of weak exclusivity enjoys wide endorsement—for example from Bennett (1991), Walker (1996, 2001), Owens (2000, 2003), Kelly (2002) and Hieronymi (2008)—but is not often distinguished from strong exclusivity. I argue against both strong and weak exclusivity in McHugh (forthcoming a), as does Frankish

The claim here is that pragmatic considerations are *wholly* excluded from the deliberative motivation of belief (or of withholding, in the case of strong exclusivity). One might prefer to make a yet more modest claim. One might think that truth must play some significant motivating role in doxastic deliberation,¹⁰ even if other considerations can also play some role. For example, if offered a tremendous reward to believe a proposition that, by your lights, your evidence counts against, or for which your evidence is evenly balanced, you cannot nonetheless come to believe the proposition, for the reason that you will thereby win the reward, through the process of deliberation about whether to believe it. But even if a certain degree of evidential support is always necessary in the deliberative motivation of belief, it might be that, given the presence of that evidential support, pragmatic considerations, such as the prospect of a reward, can add some extra deliberative motivation, pushing the thinker to go ahead and form the belief.¹¹

What degree of evidential support might be required for the deliberative motivation of belief? Various answers might be proposed. For the purposes of this paper, we can consider the claim that the evidence must at least make the content of the belief more likely to be true than not, in the thinker's eyes.¹²

(2007). If strong exclusivity is an illusion, then so too is transparency, since the latter is stronger. For the purposes of this paper I do not question transparency or exclusivity, strong or weak. (See also Zalabardo 2010 for some doubts about transparency.) This paper and my argument in my (forthcoming a) constitute two independent attacks on the motivation for normativism.

¹⁰ When deliberation is motivationally efficacious at all. Perhaps it is sometimes not so. You can believe against your own assessment of the evidence. In such cases, truth arguably fails to play a significant motivating role. But such beliefs are formed and maintained not through doxastic deliberation, but through other psychological mechanisms—or so the claim would go.

¹¹ See McHugh (forthcoming a).

¹² Elsewhere (McHugh 2011b, forthcoming a) I have argued that there is indeed a phenomenon of efficacy, that the required degree of evidential support is much more than merely making the content of the belief more likely than not to be true, and that the required degree is in fact such as would be good enough for knowledge. I have suggested that this should be understood modally: the thinker would not easily have that evidence and yet the content be false. Similar views can be found in Owens (2000) and Adler (2002). These views are controversial, and it is more charitable in the present context to stick with the weaker claim of efficacy I have outlined.

Efficacy: A thinker cannot form a belief through deliberation on the basis of pragmatic considerations, unless there is also evidential support that, in the subject's eyes, makes the content of the belief more likely than not to be true.

I have described four possible ways of characterising the role of truth in doxastic deliberation, in descending order of strength: transparency, strong exclusivity, weak exclusivity, efficacy. My question is not which of these phenomena (if any) are genuine and which (if any) spurious, but rather what the explanation for these phenomena, if genuine, might be. The significance of this question lies in Shah's argument for his normativist account of belief, to which I now turn.

2. Normativism

Shah states his account of belief as follows:

“the idea that truth is the standard of correctness for belief ... is expressed in the *prescription* to believe that p only if p is true... . My hypothesis is that accepting this prescription is one of the conditions for possessing the concept of belief.”
(2003, 448-9)

According to this account, it is a conceptual truth about belief that a belief is correct if its propositional content is true, and incorrect if its propositional content is false. This standard of correctness is understood as *normative*. In particular, it is understood as a prescription: for any p the standard of correctness for belief *demands* or *requires* that you believe p only if p . Thus, the concept of belief is a normative concept: to grasp the concept of belief is to accept for any

p to which you apply the concept that you ought to believe p only if p ,¹³ in the sense that you apply to yourself and others the prescription to satisfy this standard.

There are various points at which one might disagree with this account. One might agree that truth is the norm of belief, but deny that this norm is prescriptive, holding that it is instead, for example, evaluative.¹⁴ One might hold that (the concept of) belief is not normative, but rather teleological: it is (the concept of) a state that has as its aim or function to be true, thus making truth not a norm but an aim or *telos*.¹⁵ One might hold a more deflationary view, on which there is a standard of correctness for belief, but this standard, like the standard of correctness for sudoku puzzles, has no implications for what anyone aims to do or ought to do.¹⁶ Or, one might even deny that truth is in any sense a unique ‘proper’ standard of correctness for belief.¹⁷

Why, then, should we accept normativism? Shah takes transparency as a datum, and his argument is an inference to the best explanation from this. Let's see then how normativism purports to explain transparency, assuming the latter is indeed a genuine feature of doxastic deliberation. Note that if it can explain transparency then it can also explain strong and weak exclusivity, and efficacy, since they all follow from transparency.

To deliberate about whether to believe p is to engage in deliberation that is framed by the question whether to believe p . Doxastic deliberation thus involves exercise of the concept of belief, Shah claims.¹⁸ Since the deliberating subject's grasp of the concept of belief involves

¹³ This ought has wide scope. You ought to: believe p only if p .

¹⁴ See McHugh (2012b).

¹⁵ Velleman (2000), Steglich-Petersen (2009), McHugh (2011b, forthcoming a, forthcoming b). This view can be developed as a version of the evaluative view, in so far as something's fulfilling its function *qua* K makes it a good K (McHugh 2012b). Shah's overall account involves teleological elements, but the normativist element is independent, and I focus on this.

¹⁶ Papineau (1999). See Rosen (2001) and Hattiangadi (2007) for the point that correctness need not be normative.

¹⁷ This view is considered by Zalabardo (2010).

¹⁸ One might wonder in what sense doxastic deliberation need involve exercise of the concept of belief. Shah is clear that doxastic deliberation need not involve the thinker's explicitly posing to herself the question whether to believe p . But the question must be “in the background”, directing the thinker's deliberation (Shah, *loc. cit.*, 467). Deliberation “whose sole question is *whether p is true*” (*ibid.*)—what Zalabardo (*loc. cit.*) calls *inquiry* into whether p —does not count as doxastic deliberation. Shah is committed to this distinction

acceptance of the prescription to believe p only if p , the exercise of the concept activates a disposition to try to satisfy this norm, to the exclusion of competing norms. Only evidence is relevant to how one can satisfy this norm. Thus, the concept of belief and the very nature of deliberation together ensure that in doxastic deliberation the subject is disposed to have an eye only for evidence. As Shah puts it:

“a competent user of the concept of belief must accept the prescription to believe that p only if p is true for any activity that he conceives of as belief-formation. Because one accepts this prescription insofar as one is deliberating about *whether to believe that p* , determining *whether p is true* will be immediately imperative, to the exclusion of any other question, for anyone who entertains the deliberative question *whether to believe that p* .” (470)

In the next section I will ask whether this explanation succeeds.

3. What Can Normativism Explain?

I will argue that the above explanation of transparency fails. I will do this by arguing that normativism cannot explain strong or weak exclusivity, from which it follows that it cannot explain transparency (since transparency includes both forms of exclusivity). I will argue that normativism can at best explain efficacy, and even then only if the latter is understood only as a constraint on *rational* doxastic deliberation. Since such a constraint can easily be explained in other ways, doxastic deliberation provides no support for normativism.

between doxastic deliberation and inquiry, since inquiry into whether p surely need not involve any exercise of the concept of belief. Equally, in the case of inquiry there is no transparency of the question whether to believe p to some other question; the thinker simply begins by asking whether p (is true), rather than asking a question about what to believe. As Zalabardo (*loc. cit.*, 10) points out, once we recognise the distinction between inquiry and doxastic deliberation, the 'datum' of transparency becomes less compelling: perhaps we confuse the simple fact that we almost always ask ourselves, simply, whether p , and very rarely whether to believe p , with the transparency of the latter question to the former. (Thanks to Daniel Whiting for discussion here.)

3.1 Exclusivity and Transparency

I start with exclusivity. Shah says that the deliberating subject accepts a prescriptive norm enjoining her to believe p only if p , and that acceptance of this norm is what explains transparency, and hence exclusivity. The explanation thus appeals to the subject's *acceptance* of a norm in order to explain the subject's being *motivated* to satisfy that norm to the exclusion of any other consideration. It thus seems to assume that accepting a prescriptive norm guarantees that you will be motivated exclusively by considerations relevant to how to satisfy that norm.¹⁹ On the face of it, however, this isn't so. There is a difference between accepting that, in doing something, you ought to satisfy a certain condition, and being motivated exclusively by considerations relevant to that condition.²⁰ So Shah's explanation seems to involve a *non-sequitur*.

To avoid this, the normativist would have to invoke a principle to justify the step from attributing acceptance of a certain norm to attributing motivation by that norm to the exclusion of other considerations. This would, it seems, be a form of motivational internalism. 'Motivational internalism' is an umbrella term for a cluster of views according to which, roughly, acceptances of norms are necessarily motivating.²¹ Roughly, the idea is: necessarily if a subject accepts some norm governing her conduct (e.g. she judges that she ought to satisfy

¹⁹ It might be objected: Shah need not assume that this is generally true, but only that it is true with respect to the particular case of the norm of belief. That's correct, but as far as I can see Shah offers no explanation of why this particular instance of the claim, as opposed to any other, should be true, nor any reason to think that it is indeed so. The most promising way to find an explanation or justification for the particular claim is surely to look for a more general truth of which it is an instance—a general truth about norms of the relevant sort. I will be arguing that there is no such general truth. This not only leaves us with no justification for the particular claim; it also gives us reason to think it is false, in so far as the putative norm of belief should *ceteris paribus* behave like other, similar norms. Thus, if Shah's explanation of exclusivity rests on assuming that the particular claim is true, the explanation fails.

²⁰ Steglich-Petersen (2006, pp. 506ff.) makes this point, arguing that normativism cannot explain why we would be motivated (never mind *exclusively* motivated) to satisfy the alleged norm of belief, nor why our concern for that norm would be 'immediate' (this corresponds to what I have called 'transparency'). Some of the points made in what follows can be found in Steglich-Petersen's paper, but his treatment is brief; his main interest lies in defending the teleological account of belief, rather than in critiquing the normativist account.

²¹ Internalism comes in many varieties, as represented, for example, by Hare (1952), Smith (1994), Korsgaard (1996), McDowell (1998), Gibbard (2002) and Wedgwood (2007).

condition *C*) then the subject will be motivated to satisfy that norm (e.g. she intends to do what she thinks would satisfy *C*).

It might seem, then, that normativism could explain exclusivity (and perhaps transparency) when combined with a suitable form of internalism. The problem, however, is that no defensible form of internalism can play this role.

To see this, consider what the required internalist principles would be. To explain weak exclusivity in terms of normativism we would need:

Int₁: If *S* accepts a norm *N* of the relevant kind for act-type Φ , then *S* will be deliberatively motivated to Φ , if at all, only by *N*-directed considerations.

To explain strong exclusivity we would need:

Int₂: If *S* accepts a norm *N* of the relevant kind for act-type Φ , then *S* will be deliberatively motivated to Φ or to refrain from Φ ing, if at all, only by *N*-directed considerations.

Some clarifications: The 'relevant kind' of norm is the kind of which the putative norm of belief, the prescription to believe *p* only if *p*, is an instance. I use 'act-type' very broadly, to include things like beliefs. '*N*-directed considerations' means considerations that *S* sees (explicitly or otherwise) as bearing on how *S* can satisfy *N*.

The problem is that **Int₁** and **Int₂** are false. No sensible internalist would defend them. Acceptance of a norm *N* for a certain domain of conduct just doesn't have the consequence that only *N*-directed considerations can deliberatively motivate conduct in the domain. In particular, it doesn't have this consequence when *N* is of a certain form—the form that,

according to Shah, characterises the norm of belief.

The norm of belief, Shah tells us, prescribes that you believe p only if p . This norm has the conditional form: ' Φ only if C obtains'. To accept a prescription of this form is to accept that you ought not: Φ when C does not obtain. Crucially, this leaves it open whether or not to Φ when C obtains; the prescription is silent on that. So, if you *do* Φ for any reason, then your reason for doing so, if any, *cannot* be a reason that is directed towards the satisfaction of the norm ' Φ only if C obtains'. The fact that C obtains, or evidence that it obtains, is *not* a consideration that in any way favours Φ ing, as far as that norm is concerned, but that doesn't mean you can't go ahead and Φ for some reason. Given that (as far as you are concerned) C obtains, you might be able to go ahead and Φ for any reason you like.

For example, you accept a prescription for all Φ to Φ only if you would not thereby treat another person as a mere means to an end. Suppose for some Φ you are satisfied that by Φ ing you would not treat anyone as a mere means to an end. That still leaves open whether to Φ ; after all, Φ ing might have nothing else positive to be said for it at all. If you go ahead and Φ , you might do so for any sort of reason at all—that it would be fun, or that it is a means to some important goal, or that it would help someone, or whatever. Indeed, if you don't have some such reason, then typically you won't go ahead and Φ , since you don't have any reason to.

Applied to belief, the point is that evidence for p is *not* a consideration that favours believing p as far as the prescription to believe p only if p is concerned. All that the norm entails is that evidence against p favours not believing p (and disfavors believing p).

We often believe for reasons, and our reasons for believing are usually evidential. (Of course Shah doesn't deny that; he sees it as part of his *explanandum*.) These reasons must derive from something other than concern for a norm that merely forbids false belief. If we do accept that norm, then this acceptance does *not* restrict our deliberative motivation to considerations that

are directed towards the satisfaction of the norm—and nor should we expect it to, given general considerations about norms of that form.

Since the acceptance of a norm does not entail any such restriction on our deliberative motivation, then in particular acceptance of the norm of belief does not entail that we cannot be motivated in doxastic deliberation by pragmatic considerations. **Int₁** and **Int₂** are false, and normativism cannot explain either strong or weak exclusivity.

It might be said that the solution to this problem is to reformulate the norm of belief. Rather than a prescription merely to believe p only if p , the norm should instead be understood as a prescription to believe p if and only if p .²² A norm of the form ' Φ iff C obtains' *does* tell you what to do if C obtains—it tells you to Φ . So the fact that C obtains, or evidence that it obtains, will constitute a reason to Φ , as far as such a norm goes.

In considering a norm of the merely conditional form, I follow Shah. It is no accident that Shah formulates the norm in this way. If the norm were a prescription, for any p , to believe p if and only if p , then it would enjoin you to believe any true proposition. It is not plausible that belief is governed by such a prescription; this would be implausibly demanding.²³ Nor is it plausible that grasp of the concept of belief involves acceptance of such an overdemanding norm. We don't behave as though we accept such a norm.

An alternative suggestion is that, even if the *norm* of belief merely forbids the holding of false and therefore incorrect beliefs, the activity of doxastic deliberation has as its constitutive *aim* that one arrive at a correct belief (and not merely that one avoid arriving at an incorrect one).

²² An alternative suggestion would be that the norm should be understood as a *permission* to (believe p) if and only if p (Whiting 2010). But this would not help: if condition C merely permits you to Φ , C 's obtaining still leaves it open whether to Φ , something that you would presumably settle by considering other reasons that favour or count against Φ ing (McHugh 2012b).

²³ See Bykvist and Hattiangadi (2007, forthcoming), Whiting (2010). In McHugh (2012b) I consider in detail what the character of the truth norm of belief might be. There are candidate prescriptive norms of belief that do enjoin certain true beliefs, and are not obviously overdemanding. For example, one might propose something like: if you consider whether p , or if you hold any doxastic attitude towards p , then you ought (believe p iff p) (see also Wedgwood, 2002). However, as I argue in that paper, there are other very serious problems facing the view that belief is governed by such a norm. Further formulations are considered and rejected by Bykvist and Hattiangadi (forthcoming). Space prevents me from discussing all these possibilities here.

This might promise to explain why, in doxastic deliberation, we are motivated to believe propositions by evidence for those propositions, but not (typically) by pragmatic reasons for believing those propositions.

In fact, Shah and Velleman seem to endorse something like this suggestion in the following passage:

“A norm of correctness forbids the holding of beliefs that would be incorrect, but it merely permits the holding of correct beliefs. One is not required to hold every belief that would be correct. In deliberating whether to believe that p , however, one is committed to forming the belief if it would be correct, and this commitment tends to supply the injunctive half of a biconditional norm, mandating a belief in p if and only if p is true.” (Shah and Velleman 2005: 519.)

It's not clear what Shah and Velleman mean by 'commitment' in this passage, but it's hard to see how it could amount to anything other than an aim or intention. After all, as we have seen, it cannot be the acceptance of a norm, since we do not accept a norm enjoining us to believe truths.

As far as I can see, Shah and Velleman don't offer anything to back up the claim that, in deliberating whether to believe p , you aim to form the belief if it would be correct.²⁴ But even if that claim is right, it does not entail exclusivity, strong or weak.²⁵ For you might have other aims too. You might also aim to arrive at a belief quickly, or entertainingly, or to arrive at a

²⁴ They claim that deliberation about whether to Φ is reasoning that is aimed at issuing or not issuing in a Φ ing, in accordance with norms for Φ ing (Shah and Velleman, *loc. cit.*, p. 502). This strikes me as false. You can deliberate about whether to make an assertion without caring about the norms for assertion. You can deliberate about whether to perform actions that are associated with no specific constitutive or proper norms. In any case, their claim about the nature of deliberation doesn't entail that, when deliberating about whether to Φ , you are committed to (Φ -ing if it is permissible to Φ), since doing everything permissible is not necessary for adhering to the relevant norms.

²⁵ As pointed out (in slightly different terms) by Owens (2003), whose argument I follow here. A closely related point is in fact made by Shah himself (Shah, *loc. cit.*). See McHugh (forthcoming a, forthcoming b) for discussion.

belief that is not too disturbing. In deliberating over aim-directed behaviours, we typically take into account considerations bearing on the satisfaction of the various relevant aims we have, and opt for a course of action that we expect to give the best overall outcome with respect to those aims. So, the mere fact (if it is one) that you aim to form a correct belief should not prevent you from taking into account, and being motivated by, considerations that are not relevant to the truth of p . Compare: if you are deliberating about whether to guess that p , or whether to conjecture that p , you will aim to guess or conjecture correctly and not incorrectly. But you might also be motivated by other considerations—speed, whether it would be more costly to incorrectly guess that p or to incorrectly guess that $\sim p$, and so on.

In any case, if Shah were to invoke an aim that is constitutive of doxastic deliberation in order to explain exclusivity (weak or strong), then this would undermine his argument for normativism. That argument claimed precisely that it is the norm of belief, and not some aim that is held by believers, that explains exclusivity (and transparency). If exclusivity is, after all, explained by an aim that is constitutive of doxastic deliberation, then the argument fails.²⁶ Thus, normativism cannot explain exclusivity, weak or strong. *A fortiori*, it cannot explain transparency.

My claims here are in strong disagreement with the following remarks of Shah and Velleman, regarding their explanation of transparency:

“This explanation relies on a very weak form of internalism about normative thought. The relevant form of internalism does not require a positive disposition to obey any norm that one applies; what it requires is the lack of a disposition to obey a different norm instead. One cannot genuinely apply the norm of truth to an attitude while simultaneously trying only to make it conform to some other,

²⁶ Compare Steglich-Petersen (2006, p. 508).

unrelated norm. This form of internalism does not rule out obedience to additional norms compatible with the one applied. One can aim to arrive as quickly as possible at a true cognition with respect to p —in which case, one will deliberate in accordance with a norm of speed as well as the norm of truth. One can aim to arrive at a true cognition in a manner that, if it leads to error, is more likely to err in stopping short of true cognition than in arriving at a false one; one will then deliberate in accordance with norms of caution as well as truth. What one cannot do, according to our weak form of internalism, is to apply a norm in thought while hewing to a competing norm in practice.” (Shah and Velleman, *loc. cit.*, n. 40.)

I find this passage puzzling. I do not see how the weak form of internalism mentioned here could suffice for the normativist's purposes. Shah and Velleman distinguish between norms that are compatible with a norm that one 'applies', and norms that are incompatible with it. They say that if you apply a norm N then you will not follow (obey, try to satisfy) only norms incompatible with N . Maybe so. But this form of internalism, combined with normativism, entails at most that in doxastic deliberation you cannot be motivated only by considerations directed towards norms incompatible with truth! This consequence is weaker even than efficacy.

Exclusivity and transparency rule out *any* non-truth-directed consideration from playing the relevant roles in doxastic deliberation. They do not merely rule out considerations directed to *incompatible* norms. If you can be motivated by norms of speed and caution in doxastic deliberation, then why shouldn't this show up in the reasons that you can take into account and that can motivate you to believe or not to believe? Why shouldn't you believe things in part for speed- or caution-directed reasons? (For example, why can't you deliberate like this: “Shall I believe p ? Well, I don't have much evidence for it, but I'd like to have a settled

attitude about p as quickly as possible, so I'll go ahead and believe it.” – and conclude by forming the belief?) The form of internalism envisaged by Shah and Velleman in the quoted passage does nothing to rule this out.

On the other hand, if the normativist made the stronger claim that applying a norm N rules out following *any* other norm, compatible with N or not, then this would be implausible, as we have seen.²⁷

I turn next to efficacy. The problems I raise for the normativist's explanation of efficacy will also be further problems for his explanation of exclusivity and transparency, since the latter include efficacy.²⁸

3.2 Efficacy

As we saw, acceptance of a norm doesn't guarantee *exclusive* motivation by considerations relevant to that norm. But perhaps it guarantees *some* motivation by such considerations. It is a claim of this sort—more recognisably like a defensible kind of motivational internalism—that the normativist requires in order to explain efficacy. Thus:

Int₃: If S accepts a norm N of the relevant kind for act-type Φ , then whenever S deliberates about whether to Φ , S will be significantly motivated by N -directed considerations.

²⁷ Shah and Velleman's internalism is presumably connected to another element of their overall view, namely 'norm expressivism'. This is the idea that “accepting a norm is a conative attitude that, among other things, disposes one to follow the norm and inhibits one from following any alternative” (Shah and Velleman 2005, p. 510). Parallel remarks apply here to those made in the text. A spelled out version of norm expressivism would have to say what the relevant disposition amounts to. If the disposition is said to be so strong that it gives rise to phenomena like exclusivity or transparency, then this will be an extremely implausible version of norm expressivism. Expressivists do not typically make such strong claims.

²⁸ The account of belief in Shah and Velleman (2005) is not purely normativist, but also contains a teleological element—the view, in a slogan, is that belief is *both* causally *and* normatively governed by truth. The teleological element of this view may suffice to explain efficacy. So, even if normativism cannot explain efficacy, this is not itself an objection to Shah's overall view. However, it is important to investigate the explanatory power of normativism as such, since Shah's *argument* for his view depends crucially on this.

By 'significantly motivated by *N*-directed considerations' I mean motivated enough to satisfy the condition of efficacy—that is, *S* will not intentionally violate *N*, nor conclude deliberation by doing something that she thinks as likely as not to violate *N*. Note that **Int₃** does not merely attribute a general disposition to be motivated by norms that one accepts; it attributes specific significant motivation in any particular case to which the norm applies.²⁹ This is what is required if there is to be any hope of explaining efficacy, since the idea of efficacy is not merely that we are generally disposed to be significantly motivated by evidence in our doxastic deliberation, but rather that we are so motivated in every instance of doxastic deliberation.

If **Int₃** is true, then normativism seems to give a straightforward explanation of efficacy: the subject deliberating about whether to believe *p ipso facto* accepts the prescription to believe *p* only if *p*, thus by **Int₃** is significantly motivated to satisfy that norm and will not form a belief that she thinks as likely as not to be false.

But at first glance **Int₃** looks as implausible as **Int₁** and **Int₂**. It's not the case that we are invariably, in particular cases, significantly motivated by the relevant norms that we accept. On the contrary, we often intentionally violate norms that we accept.

Consider assertion. Many philosophers hold that assertion is subject to one or more constitutive norms. Some of these philosophers hold that assertion is subject to a constitutive epistemic norm of knowledge: that is, it is constitutive of assertion that you ought to assert *p* only if you know *p*.³⁰ Another obvious possibility is that assertion is subject to a constitutive norm of truth: you ought to assert *p* only if *p*.³¹ Supposing that some such view is true, it wouldn't follow that we cannot intentionally make assertions whose contents we don't know, or whose contents are false. If a view entailed such a thing, it would constitute a *reductio* of

²⁹ Contrast the internalist principle (2) given by Wedgwood (2007, 28).

³⁰ The most influential defence of this is Williamson (2000, Ch. 11).

³¹ Versions of this view have recently been defended by Weiner (2005) and Whiting (forthcoming). Note that the scope of the 'ought' is wide in each case.

the view, because it is obvious that we can and do assert things that we don't know, and that we are fully aware we don't know, and things we don't believe, and things that we know are false. What such a view of assertion entails is that the relevant speech-acts violate the constitutive norm of assertion, not that they are impossible.³²

Is the putative norm of assertion a norm of the 'relevant kind'? We can certainly imagine a view on which the norm of assertion is analogous to the putative truth norm of belief in salient respects. It is a prescription. We can suppose that it is constitutive of the concept of assertion.³³ It is (we can suppose) a norm of correctness. None of that will make the norm any more inviolable, nor will it thereby make the view obviously false.³⁴

The normativist might object that the norm of assertion is nonetheless not of the relevant kind because it is not 'heavy duty' enough. The norm of assertion is merely one norm among others that one typically takes into account when deliberating about what to assert. In a particular case one might conclude, because of considerations of politeness, prudence, or whatever, that what one all-things-considered ought to assert is something false, or something one doesn't know. By contrast, the normativist might say, the prescription to believe p only if p is not merely *a* norm, among others, that one takes into account when deliberating about whether to believe p . Rather, it somehow silences or defeats any norms that make opposing recommendations, when it comes to the question what to believe—as moral norms have been claimed to do when it comes to the question how to act. That is, in accepting the prescription to believe only what is true, one accepts when deliberating about whether to believe p that

³² Steglich-Petersen (2006) makes this point about the speech-act of promising. This speech-act is plausibly constituted in part by a norm along the lines: you ought to (promise to Φ only if you intend to Φ). That the promising subject must regard herself as subject to this norm does not ensure that she will care about it in a particular case. Indeed, the promising subject might be quite aware that she has no intention of Φ ing.

³³ I do not know whether the knowledge norm of assertion, or the other norms that have been proposed (see, e.g., Brown 2008), are supposed to be constitutive of the *concept* of assertion. But a view on which they are so constitutive is not obviously false.

³⁴ Of course, assertion, unlike belief, is an intentional action. So, even if we can intentionally violate the putative norm of assertion, arguably we cannot intentionally violate the norm of belief, if that means doing so by performing an intentional action. But that is beside the point. The point is that we seemingly cannot conclude deliberation about what to believe by forming a belief that we think will violate the putative norm of belief.

what one all-things-considered ought to do, or what one 'just plain' ought to do, is believe p only if p .³⁵

Int₃ certainly appears more plausible if the 'relevant kind' of norm must have this status.³⁶ It's one thing to intentionally violate the rules of assertion because other considerations outweigh them; it's another thing to intentionally act contrary to what you take it you all-things-considered ought not do.

It is worth noting that this move requires the normativist to make a very strong claim: namely, that it is a conceptual truth that, for any p , you all-things-considered or 'just plain' ought to believe p only if p . Thus, even if believing a falsehood would save the world from destruction, you nevertheless ought not believe it. Of course, if your evidence counts against it then you *epistemically* ought not believe it. But the claim to which the normativist is committed, if making this move, is that the norm of belief entails that you ought not believe it *simpliciter*. And this is a conceptual truth; anyone who denies it is confused about the concept of belief.³⁷ Shah makes some remarks that suggest he might be happy to endorse this strong claim,³⁸ but many philosophers may find it hard to swallow that application of the concept of belief, or of any other concept that isn't a central normative concept like *good*, *reason* or *virtue*, could involve such strong normative commitments.³⁹

³⁵ As suggested in this sentence, there is more than one way in which this sort of idea could be filled out. One way is in terms of the notion of all-things-considered ought; on the need for such an idea, also referred to as ought *tout court*, *sans phrase*, all-out, or without qualification, see for example: Davidson (2001), Wedgwood (2007, 24-25), Alvarez (2010, 15-16). Another way to fill out the idea is in terms of the notion of 'just plain ought' (see Feldman 2000); here, the thought might be that, when it comes to belief, what you 'just plain' ought to do coincides with what you epistemically ought to do. A third way to fill out the idea would be to say that epistemic reasons silence other reasons when it comes to belief; on the notion of silencing, see McDowell (1978). In any case, the essential idea is that by accepting the norm of belief, you commit to a judgment to the effect that you ought not believe p if p is false, and this judgment has the same status as the sort of judgment about what you ought to do that can conclude practical deliberation and on which you base an intention to act.

³⁶ Not so for **Int₁** or **Int₂**, however. My argument against them in sec. 3.1 is unaffected by whether the 'relevant kind' of norm is 'heavy duty' or not.

³⁷ Philosophers who deny it include Foley (1993), Papineau (1999) and Reisner (2009).

³⁸ See Shah (2003, pp. 454-5), where he claims that it is a conceptual truth that there are no pragmatic reasons for belief.

³⁹ The matter is complicated by the seeming fact that, if a proposition is evidently false, you will normally not be able to do the following: believe it anyway, for the reason that doing so will save the world from

In any case we can put this point aside, because **Int₃** is still false, even if we assume that the 'relevant kind' of norm is 'heavy duty'—that it determines what you all-things-considered or 'just plain' ought to do in any case of doxastic deliberation. This is shown by the possibility of akrasia—that is, intentionally acting contrary to your own judgment of what you all-things-considered 'just plain' ought to do.⁴⁰ Akrasia is a familiar phenomenon; that it occurs, though puzzling, is undeniable. No matter how 'heavy duty' a norm, there is the possibility that a subject who accepts it akratically violates it. Take moral norms. If any norms are 'heavy duty', moral norms surely are. It is a truism that we can intentionally violate moral norms that we accept. A zealous Kantian might accept a norm that forbids lying, or treating other persons as mere means; that won't make her incapable of intentionally lying or treating others as mere means. We all regularly do things that we think we morally ought not do, out of laziness, weakness, indifference, selfishness, spite, or whatever.

Regardless of the status of moral norms, there are many familiar cases of subjects who intentionally act contrary to their judgment of what they all-things-considered 'just plain' ought to do. You intentionally stay under the warm blankets despite your firm judgment that what you ought to do is get up immediately and go to the office; you intentionally choose to have one last beer despite judging that you ought to go home instead.

In sum, then, it looks implausible that **Int₃** is true of any kind of norm. And it can hardly be claimed that the norm of belief has a special kind of motivational inescapability not possessed

destruction (that is, it seems that something like efficacy is indeed a feature of doxastic deliberation). Some form of 'ought implies can' principle may then be invoked, to argue that it is not the case that you ought to believe it. Now, this would not yet show that it is the case that you ought not believe it. But in any case, we can imagine a scenario in which you have some mechanism that can directly and instantaneously implant a given belief in your mind, and over whose operation you have voluntary control (see Bennett, 1991). Given that you would save the world by believing *p*, and you can come to believe *p* by operating this mechanism, is it nevertheless the case that you ought not believe *p*, and that you must accept this on pain of conceptual confusion? It might be said that, in such a scenario, although you ought to implant in yourself the belief that *p*, nevertheless you ought not believe *p*. I find it difficult to make sense of this proposal: how can you be required to do something that logically entails doing something you are required not to do?

⁴⁰ The normativist account of belief, combined with **Int₃**, would allow for the possibility of a certain kind of akratic belief—namely, that in which a belief contrary to the subject's own evaluation of the evidence is held in place by mechanisms other than doxastic deliberation. What is ruled out by **Int₃** is akratically violating a norm you accept through deliberation in which that norm is applied.

by other norms, even moral norms. That would be implausible, *ad hoc* and devoid of explanatory power.

At this point it might be natural to wonder what exactly is meant, in Shah's account, by the notion of 'accepting' or 'applying' a norm. I have been construing this notion in such a way that it can roughly be assimilated to an implicit or explicit judgment about what one ought to do, where 'ought' expresses a prescription. This, I have argued, leaves a gap between acceptance of a norm and being significantly motivated to satisfy it. Am I being uncharitable when I construe Shah's claim in such a way that this gap is left?⁴¹

I don't see an alternative way of interpreting Shah's claim, that doesn't deprive the account of all explanatory power, or turn it into a terminological variant on an account that Shah explicitly distinguishes as a rival to his own.

One way to close the gap I have been focusing on would be to stipulate that the term 'acceptance' is being used in such a way that acceptance of a norm *entails* being significantly motivated to satisfy it. Then, the normativist account of belief would entail efficacy. But this would make the kind of norm-acceptance that is, according to normativism, constitutive of possession of the concept of belief, wholly mysterious. It would be a kind of explanatory *virtus dormitiva*. We can see this when we compare it with moral norms. Consider possession of the concept of something's being morally wrong. Possession of this concept just doesn't entail, in any particular case, that a subject will not (perhaps akratically) do what she thinks is wrong.

Alternatively, one might stipulate that 'acceptance' of a norm just *is* being significantly motivated to satisfy the norm. But now the normativist account collapses into the bare claim that efficacy is a condition on acceptance of the concept of belief. It's no longer clear in what sense such an account would be talking about *norms*; and in any case it is no more

⁴¹ Thanks to Tony Booth for raising this question.

explanatory than the previous suggestion.

A third way to close the gap would be to claim that the subject engaged in doxastic deliberation must be *aiming* to satisfy the norm of belief. I already discussed a similar suggestion in sec. 3.1 above; as we saw, it simply concedes the territory to the rival teleological account, leaving normativism playing no role after all in the explanation of the role of truth in doxastic deliberation.

If there is an alternative understanding of 'acceptance' of a norm, on which the normativist account can explain efficacy, then it seems to me that the burden is on the defender of the account to provide it.

3.3 Efficacy as a Constraint on Rational Doxastic Deliberation

Akrasia is a counterexample to **Int₃**. While not impossible, akrasia is nonetheless irrational. This suggests that we might find a plausible internalist principle by limiting it to cases of rational conduct:

Int₄: If *S* accepts a norm *N* of the relevant 'heavy duty' kind for act-type Φ , then whenever *S* deliberates about whether to Φ , if *S* is rational she will be significantly motivated by *N*-directed considerations.

Given how I have defined 'heavy duty' and 'significantly motivated', and given that akrasia is indeed irrational, **Int₄** looks plausible.⁴²

Normativism, when understood as claiming that the concept of belief involves a 'heavy duty'

⁴² Could this same move—restricting the relevant principles to rational conduct—be made in response to my case in sec. 3.1 against **Int₁** and **Int₂**? No. Acceptance of a norm, 'heavy duty' or otherwise, enjoining you to Φ only if *C* obtains, does not make it irrational to Φ or not Φ for some reason irrelevant to whether *C* obtains. The Kantian who accepts a moral norm forbidding conduct that involves treating others as mere means to an end is not irrational when motivated to Φ by some consideration orthogonal to whether Φ ing would involve treating another as a mere means to an end.

norm, and **Int**₄, together seem to entail that rational doxastic deliberation will be characterised by efficacy. That is, it seems that they can explain why a subject who concludes doxastic deliberation by forming a belief that, by her own lights, the evidence counts against, is irrational. If my arguments to this point have been correct, this is the *most* that normativism can explain about doxastic deliberation.

I take no position in this paper on whether normativism *should* explain more than this.⁴³

Rather, I want to point out that there are readily available alternative explanations of why rational doxastic deliberation will be characterised by efficacy.

One explanation would appeal to the teleological account of belief, according to which it is in the nature of belief to be regulated by activities or processes that have as their aim or function that the beliefs thereby regulated be true. Doxastic deliberation about whether to believe *p*, on this view, is characterised by the aim of getting it right about the truth of *p* (otherwise what you're regulating is not really *belief*). Considerations relevant to what to believe, as far as this aim goes, will be truth-directed considerations. But you can't rationally pursue an aim through a certain activity, unless you are guided, in your conduct of that activity, by considerations that you take to be relevant to how to achieve that aim. And you can't rationally pursue that aim if you conclude the activity by selecting a course of action that you think unlikely to satisfy the aim, foregoing another course of action that you think would be likely to satisfy it.⁴⁴ Thus, the subject who engages in doxastic deliberation will, if rational, be significantly motivated by truth-directed considerations.⁴⁵

⁴³ The answer to this question depends in part on whether a species of so-called 'epistemic akrasia' is impossible (in which case normativism would have more to explain) or merely irrational. The possibility of epistemic akrasia is accepted by Engel (2006), and denied by Adler (2002) and Owens (2002).

⁴⁴ If your evidence makes *p* improbable, then believing *p* would be unlikely to satisfy the putative aim of getting it right about *p*, while disbelieving *p* would be likely to satisfy it. There is a question of why, if the teleological account is true, you can't rationally form the belief that *p* when your evidence is exactly evenly balanced. In such a case, believing *p* would seemingly be no less likely to satisfy the aim of getting it right about *p* than would disbelieving *p*. In fact, however, it depends on exactly what the putative aim amounts to. Different answers to this question are offered in McHugh (2011b) and in Whiting (2012).

⁴⁵ Arguably, pursuing an aim by taking a course of action you think unlikely to achieve the aim, while

A second approach, more deflationary in spirit, would be to say that *epistemic* rationality requires doxastic deliberation to exhibit efficacy (and perhaps even weak or strong exclusivity). Epistemic rationality *just is* a kind of rationality that is directed towards truth in a certain way. Thus, failures of efficacy are always irrational in a certain respect, or from a certain point of view. This point of view matters to us because we are generally concerned about truth. That's why we care about epistemic rationality, and tend to evaluate beliefs and doxastic activities from the epistemic point of view. But this concern for truth need not enter into an account of the nature of belief, according to this proposal—or at any rate, not in the way that the normativist proposes. Epistemic rationality need not derive from a constitutive prescriptive norm of belief.⁴⁶

So, although normativism may be able to explain why rational doxastic deliberation is characterised by efficacy, this provides no basis for an inference to the best explanation in support of the view.

4. Conclusion

Shah's argument for normativism about belief is an inference to the best explanation from transparency, which is claimed to be a feature of doxastic deliberation. I have argued, however, that normativism fails to explain transparency, as well as a range of weaker putative phenomena: strong exclusivity, weak exclusivity and efficacy. While normativism may be able to explain why *rational* doxastic deliberation is characterised by efficacy, this provides no support for normativism over rival views, since plausible alternative explanations are available. The argument for normativism is therefore fatally undermined.

This does not mean that it is not a conceptual truth that true belief is correct belief and false

foregoing another course of action you think would be likely to satisfy it, is incoherent in some more profound way than the mere irrationality of akrasia—if not flat-out impossible. Thus, if epistemic akrasia is not just irrational, but more profoundly incoherent than that, this may be explained by the teleological view of belief. See McHugh (forthcoming b) for discussion of what a teleological view can explain.

⁴⁶ Something like this appears to be Papineau's view (Papineau 1999, forthcoming).

belief is incorrect belief. Rather, it means that we have little reason to view this standard of correctness as a prescriptive norm.

Shah has provided a normativist account of intention that parallels his account of belief, and he offers an analogous inference to the best explanation to support it.⁴⁷ It seems to me that the points I have made in the present paper would apply, *mutatis mutandis*, to that analogous argument.⁴⁸

References

- Adler, J. (2002). *Belief's Own Ethics*. Cambridge, MA: MIT Press.
- Alvarez, M. (2010). *Kinds of Reasons: An Essay in the Philosophy of Action*. Oxford: OUP.
- Bennett, J. (1991). Why Is Belief Involuntary? *Analysis*, 50, 87-107.
- Brown, J. (2008). The Knowledge Norm for Assertion. *Philosophical Issues*, 89-103.
- Bykvist, K. and Hattiangadi, A. (2007). Does Thought Imply Ought? *Analysis*, 67, 277-85.
- Bykvist, K. and Hattiangadi, A. (forthcoming). Blindspots and Doxastic Oughts: A Reply to our Critics. In T. Chan, *The Aim of Belief*. Oxford: OUP.
- Davidson, D. (2001). How is Weakness of Will Possible? In *Essays on Actions and Events*. Oxford: OUP.
- Engel, P. (2006). Epistemic Norms and Rationality. In M. G. W. Strawinski, *Mysli o Języku, nauce I wartosciach, Ksiegaofiarowana Jackowi Jadakiemu*. Warsaw: Semper.
- Feldman, R. (2000). The Ethics of Belief. *Philosophy and Phenomenological Research*, 60, 667-95.
- Foley, R. (1993). *Working Without a Net: A Study of Egocentric Epistemology*. Oxford: OUP.
- Frankish, K. (2007). Deciding to Believe Again. *Mind*, 116, 523-47.
- Gibbard, A. (2002). Normative and Recognitional Concepts. *Philosophy and Phenomenological Research*, 64, 151-67.
- Hare, R. (1952). *The Language of Morals*. Oxford: Clarendon.
- Hattiangadi, A. (2007). *Oughts and Thoughts: Rule-Following and the Normativity of Content*. Oxford: OUP.
- Hieronymi, P. (2008). Responsibility for Believing. *Synthese*, 151, 353-73.
- Kelly, T. (2002). The Rationality of Belief and Some Other Propositional Attitudes. *Journal of Philosophical Research*, 110, 163-96.
- Korsgaard, C. (1996). *The Sources of Normativity*. Cambridge: CUP.
- McDowell, J. (1978). Are Moral Requirements Hypothetical Imperatives? *Proceedings of the Aristotelian Society, Supplementary Volumes*, 52, 13-29.
- McDowell, J. (1998). Virtue and Reason. In *Mind, Value and Reality*. Cambridge, MA: Harvard UP.
- McHugh, C. (2011a). Judging as a Non-Voluntary Action. *Philosophical Studies*, 152, 245-69.
- McHugh, C. (2011b). What Do We Aim At When We Believe? *Dialectica*, 65, 369-92.
- McHugh, C. (2012a). Epistemic Deontology and Voluntariness. *Erkenntnis*, 77, 65-94.

⁴⁷ Shah (2008).

⁴⁸ Thanks to Tony Booth, Daniel Whiting and Jonathan Way for comments on earlier drafts, and to an anonymous referee.

- McHugh, C. (2012b). The Truth Norm of Belief. *Pacific Philosophical Quarterly*, 93, 8-30.
- McHugh, C. (forthcoming a). The Illusion of Exclusivity. *European Journal of Philosophy*.
- McHugh, C. (forthcoming b). Belief and Aims. *Philosophical Studies*.
- Owens, D. (2000). *Reason Without Freedom: The Problem of Epistemic Normativity*. London: Routledge.
- Owens, D. (2002). Epistemic Akrasia. *The Monist*, 85, 381-97.
- Owens, D. (2003). Does Belief Have an Aim? *Philosophical Studies*, 115, 283-305.
- Papineau, D. (1999). Normativity and Judgment. *Proceedings of the Aristotelian Society*, 73, 16-43.
- Papineau, D. (forthcoming). There Are No Norms of Belief. In T. Chan (ed.), *The Aim of Belief*. Oxford: OUP.
- Reisner, A. (2009). The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem. *Philosophical Studies*, 145, 257-72.
- Rosen, G. (2001). Brandom on Modality, Normativity and Intentionality. *Philosophy and Phenomenological Research*, 63, 611-23.
- Shah, N. (2003). How Truth Governs Belief. *Philosophical Review*, 112, 447-82.
- Shah, N. (2006). A New Argument for Evidentialism. *The Philosophical Quarterly*, 56, 481-98.
- Shah, N. (2008). How Action Governs Intention. *Philosophers' Imprint*, 8, 1-19.
- Shah, N. and Velleman, D. (2005). Doxastic Deliberation. *Philosophical Review*, 114, 497-534.
- Smith, M. (1994). *The Moral Problem*. Oxford: Blackwell.
- Steglich-Petersen, A. (2006). No Norm Needed: On the Aim of Belief. *The Philosophical Quarterly*, 56, 499-516.
- Steglich-Petersen, A. (2009). Weighing the Aim of Belief. *Philosophical Studies*, 145, 395-405.
- Velleman, J. (2000). On the Aim of Belief. In J. Velleman, *The Possibility of Practical Reason*. Oxford: OUP.
- Walker, M. (1996). The Voluntariness of Judgement. *Inquiry*, 39, 97-119.
- Walker, M. (2001). Williams, Truth-Aimedness and the Voluntariness of Judgement. *Ratio*, 14, 68-83.
- Wedgwood, R. (2002). The Aim of Belief. *Philosophical Perspectives*, 16, 267-97.
- Wedgwood, R. (2007). *The Nature of Normativity*. Oxford: OUP.
- Weiner, M. (2005). Must We Know What We Say? *Philosophical Review*, 114, 227-51.
- Whiting, D. (2010). Should I Believe the Truth? *Dialectica*, 64, 213-24.
- Whiting, D. (2012). Does Belief Aim (Only) at the Truth? *Pacific Philosophical Quarterly*, 93, 279-300.
- Whiting, D. (forthcoming). Stick to the Facts: On the Norms of Assertion. *Erkenntnis*.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: OUP.
- Zalabardo, J. (2010). Why Believe the Truth? Shah and Velleman on the Aim of Belief. *Philosophical Explorations*, 13, 1-21.