

BELIEF NORMS AND BLINDSPOTS

THOMAS RALEIGH

ABSTRACT: I defend the thesis that beliefs are constitutively normative from two kinds of objections. After clarifying what a “blindspot” proposition is and the different types of blindspots there can be, I show that the existence of such propositions does not undermine the thesis that beliefs are essentially governed by a *negative* truth norm. I argue that the “normative variance” exhibited by this norm is not a defect. I also argue that if we accept a distinction between subjective and objective norms, there need be no worrying tension between doxastic norms of truth and doxastic norms of evidence. I show how a similar approach applies to the attitude of guessing. I then suggest that if we distinguish between practical and theoretical rationality, we will prefer a negative form of norm that does not positively oblige us to form beliefs. I finish by considering an alternative possible subjunctive form of norm that would also avoid problems with blindspots, but I suggest this has a nonintuitive consequence.

A number of writers have claimed that beliefs must, in some interestingly normative sense, constitutively *aim* at truth.¹ In particular, Shah and Velleman (Shah 2003, Shah and Velleman 2005) have argued that rational agents must grasp that belief is “regulated by truth,” as this is the best explanation for the phenomenon of “transparency” in deliberation—deliberating whether to believe that *p* just is to deliberate whether *p* is true.² And Wedgwood (2002, 2009) has argued that in order to explain how a subject’s dispositions (both mental and behavioral) determine that some

Thomas Raleigh is an Assistant Professor in the philosophy department at Concordia University in Montreal. His research focuses primarily on philosophy of mind and epistemology, particularly the nature of conscious experience. He has held postdoctoral fellowships at the Instituto de Investigaciones Filosóficas at UNAM and the University of Essex, and he was a visiting fellow at Cornell University.

¹ For example, Shah (2010) notes that early versions of the thesis are found in Griffiths 1962–3 and Edgley 1969. Opponents of the thesis include Owens (2003) and Papineau (1999).

² Zalabardo (2010) argues that Shah and Velleman’s reasons do not establish that belief is essentially normative. See, e.g. Boghossian 2003 or Gibbard 2003, 2005. Vahid (2006) offers a “deflationary” account of belief’s aiming at truth.

attitude, or state, is a *belief*, these dispositions must be rational and so (according to Wedgwood) must essentially involve normativity. The normative dimension of such claims has been put in terms of prescriptions as to what a believer *ought* to do: “my proposal is that exercising the concept of belief involves accepting that in some sense one *ought* to believe that p only if p is true” (Shah 2003, 449) and “to put it roughly, the principle that a belief is correct if and only if the proposition believed is true” (Wedgwood 2002, 273).

An issue that immediately arises is how exactly we are to understand this ‘ought’. It is not, of course, the moral sense of ‘ought’, for we can easily imagine scenarios in which the morally best option is to have some false belief. And nor can ‘ought’ here mean “ought *given the aim* of believing only truths,” for this would make any such truth-norm vacuous and undermine the claim that belief is normative in any interesting sense. So, presumably it is some *sui generis*, distinctively doxastic sense of ought that is supposed to be at work here.

One obvious sort of worry about the idea that belief is essentially governed by a norm of *truth* is that the evidence available to a subject also seems to govern what they ought to believe. And there can clearly be cases in which a subject has very strong evidence for p though p is false. Or cases in which a subject does not have a good enough evidential basis for their belief that p (or has formed the belief that p by some unreliable process) though p is, in fact, true. So there seems to be another plausible kind of norm governing belief that mentions a subject’s evidence, or justification, rather than truth, and which could conflict with a truth-norm.

A different sort of objection comes from Bykvist and Hattiangadi (2007), who argue that the existence of Moorean “blindspot” propositions fatally undermines the idea that beliefs are essentially evaluable according to a truth-norm. An example of such a blindspot is:

(BS): It is raining and nobody believes that it is raining.

The aim of this paper is not to provide new arguments for the claim that belief is constitutively normative. My primary aim will merely be to show that the existence of blindspots does not undermine it. In section 3, I will show that a negative truth-norm (i.e., a falsity-norm) for beliefs is not threatened by the existence of true-belief blindspots. And in sections 4 and 5, I will argue that the “normative variance” displayed by this norm is not problematic. Then, in section 6, I will go on to consider another kind of blindspot: justified-true-belief blindspots. This will allow me to sketch a line of response to the first sort of worry. I will argue that accepting the common distinction between subjective and objective norms allows us to see that conflict between norms of truth and norms of justification need not undermine either kind of norm.

I will also argue, in section 8, that if we accept a further common distinction, between practical and theoretical, or intellectual, rationality, we will prefer a negative form of norm that does not positively oblige us to form beliefs. Along the way, I will also show how a similar approach applies to the attitude of guessing (section 7), and I will consider and then criticize an alternative possible approach based on a “subjunctive” form of norm (section 9). But to begin we must briefly consider the different kinds of propositions that might be called “blindspots.”

1. KINDS OF BLINDSPOTS

Bykvist and Hattiangadi provide the following explanation of what a “blindspot” is. “There are some propositions such that it is logically impossible to believe them truly: if they are true, then you don’t believe them, and if you believe them, then they are false. These ‘blindspots’ (Sorensen 1988) are not truly believable” (2007, 281).

As Bykvist and Hattiangadi here acknowledge, the term ‘blindspot’ was first pressed into such distinctively philosophical usage by Roy Sorensen (1988). It is important to note, however, that Sorensen allowed for a much wider application of the term.³ Sorensen first defines a “knowledge blindspot” as (roughly): a proposition that is possibly true but which cannot be *known* (or known by some subject S). Sorensen then goes on to define a blindspot relative to any given propositional attitude A (and subject S) as: a proposition that is possibly true but which cannot have attitude A taken towards it (by S).⁴ Now, with a factive attitude, like knowledge, one can only take the attitude towards the proposition if the proposition is true. So with factive attitudes there is no difference between the following two definitions of a blindspot:

Blindspot Def.(i): $\Diamond p \ \& \ \neg\Diamond A(p)$

Blindspot Def.(ii): $\Diamond p \ \& \ \neg\Diamond[p \ \& \ A(p)]$

But, of course, with nonfactive attitudes these are two different definitions of a blindspot. Consider the following proposition:

³ Strictly speaking, Bykvist and Hattiangadi only claim that a proposition that cannot be truly believed is a blindspot; they do not rule out the possibility of other kinds of propositions also counting as blindspots.

⁴ See Sorensen 1988, 52. We need not limit our generalized definition to propositional attitudes. A possible state of affairs, F, is a blindspot for a representing system, or medium, R, iff F cannot be true and represented by, or in, R. For example, F is the state of affairs “grass is green,” and there is no representation of the fact that grass is green. This state of affairs can obtain, and this state of affairs can be represented (I just did represent it in the previous sentence), but it cannot be both true and represented.

(P1): $q \ \& \ \neg S \text{ Bel}(q)$

Given some assumptions about the possibility of being wrong about one's own beliefs, it is possible for S to believe (P1) when (P1) is false—S believes q , perhaps unconsciously, and also falsely believes that she does not believe q . So according to Def.(i), (P1) would not be a belief-blindspot. But according to Def.(ii), (P1) *is* a belief blindspot—S cannot believe this conjunction when it is true, for if it were true it would both be the case that S believed q and not the case that S believed q . To avoid confusions over which definition of 'blindspot' is in play, we can say that (P1) is a *true-belief* blindspot. (As true-belief is also a factive attitude, the difference between Def.(i) and Def.(ii) becomes irrelevant.)

As well as true-belief blindspots there can be *justified-belief* blindspots. But unlike true belief, justified belief is not a factive attitude—so we need to be clear as to whether we are defining “blindspot” according to Def.(i) or Def.(ii). For example, consider the proposition:

(P2): $q \ \& \ \neg S \text{ Justified Bel}(q)$

Matters get somewhat more complicated here given the potential for disagreement over the nature of justification. Suppose we grant some externalist assumptions about justification—namely, the possibility of one's belief that p being justified despite one's taking it to be unjustified. Then if we understand the notion of a blindspot according to Def.(i) (and assuming also a fallible notion of justification), (P2) is *not* a justified-belief blindspot. For S could have a justified belief in q and also have a false but justified belief that she does not have a justified belief that q . But if we define “blindspot” according to Def.(ii), then (P2) *is* a justified-belief blindspot, that is, it is a justified-true-belief blindspot. For S cannot have both a justified true belief that q and a justified *true* belief that S does not have a justified true belief that q .⁵ Notice also that (P2) is *not* a true-belief blindspot—S could truly but unjustifiedly believe q and also truly believe that she does not have justified belief that q .⁶ I will briefly return to the tricky issue of whether a subject can really count as having a well-evidenced belief in propositions such as (P2) in section 6, below.

⁵ Likewise, notice that according to Def.(i), (P1) is *not* a justified-belief blindspot, because S might have a justified (unconscious) belief in q and a justified but false belief that she does not believe q . But if we define “blindspot” according to Def.(ii), (P1) will count as a justified-belief blindspot because it is a justified-true-belief blindspot.

⁶ Sorensen (1988, ch. 4) offers some interesting examples of justified-belief blindspots that are not true-belief blindspots. For example, “The world was created, complete with memories, fossils, etc., 15 minutes ago.” One might allow that this is a possibly true proposition that one can truly believe, yet, there is no evidence for its truth or falsity and so no possibility that a belief in it could be justified.

The following four sections of this paper will be concerned with rejecting the idea that blindspot propositions *as defined by* Bykvist and Hattiangadi, that is, true-belief blindspots, rule out a viable truth-norm for beliefs. So for the following five sections of this paper:

p is a blindspot for S iff: $\Diamond p \ \& \ \neg\Diamond[p \ \& \ S \text{ Bel}(p)]$.

In section 6, however, I will return to considering justified-true-belief blindspots and suggest that these are also unproblematic.

2. THE TROUBLE WITH BLINDSPOTS

In this section, I briefly summarize Bykvist and Hattiangadi's argument. Bykvist and Hattiangadi formulate the claim that ought implies can as follows:

(OC): If I have an obligation to perform an action A, then it is possible for me to do A while being obligated to do A.

Bykvist and Hattiangadi want to hold onto (OC) and argue that (OC) is bound to conflict with the idea of a "truth-norm" for beliefs. So, they conclude, we should drop the idea of a truth-norm for beliefs. I will not be concerned with defending or justifying (OC). It looks a very plausible claim—though perhaps one could find reasons to doubt it.⁷ In any case, I want to argue that a suitably modified norm avoids conflicting with (OC). For the purposes of this paper, then, I will assume that (OC) is correct.

The simplest version of a truth-norm that Bykvist and Hattiangadi attack is the following:

(T): For any S, p : S ought to (believe that p) iff p is true.

It is fairly clear that (OC) and (T) are going to conflict. For surely it is not possible for a subject to believe every true proposition. Yet, (T) puts us under an obligation to believe with respect to *any* true proposition. Quite apart from the indefinitely many true propositions there must be in "logical space" concerning states of affairs we could never have the slightest inkling about, Bykvist and Hattiangadi point to two kinds of propositions that look especially bad for (T) given a commitment to (OC). The first kind are:

Overcomplex Propositions: Propositions that are too long and complicated for one to even entertain in thought.

⁷ Martin (2009) provides some tentative possible counterexamples to (OC).

Boghossian's (2003) response to this problem was to reject one direction of entailment in the biconditional in (T). So he would endorse:

(Ti) S ought to (believe that p) $\rightarrow p$ is true

But he would reject:

(Tii) p is true \rightarrow S ought to (believe that p)

As Bykvist and Hattiangadi point out, while retreating from (T) to (Ti) prevents imposing an impossibly difficult obligation, it leaves us with an unacceptably weak norm. (Ti) entails:

$\neg[p \text{ is true}] \rightarrow \neg[\text{S ought to (believe that } p)]$

i.e., p is false $\rightarrow \neg[\text{S ought to (believe that } p)]$.

Bykvist and Hattiangadi rightly draw attention to the important scope difference between:

p is false $\rightarrow \neg[\text{S ought to (believe that } p)]$

and:

p is false \rightarrow S ought to $\neg(\text{believe that } p)$.

The former states the lack of an obligation, the latter states an obligation to not do something.⁸ When p is false it is only the latter norm which forbids belief in p . All that (Ti) tells us, when p is false, is that there is not an obligation to believe p —which leaves it open that it is still quite alright to believe a false p . So for all (Ti) has said: “it does not follow, from the falsity of p , that S ought not to believe that p ” (Bykvist and Hattiangadi 2007, 280). And given that Boghossian has rejected (Tii), no obligation would follow from the truth of p either. So neither the truth nor the falsity of a proposition places us under any obligation whatsoever.

A different amendment to our initial norm, (T), which avoids the problem of overcomplex propositions, is suggested by Wedgwood:⁹

(T*): For any S, p : if S considers p , then S ought to (believe that p) iff p is true.

⁸ At the risk of stating the obvious we might also note the scope distinction between “S ought to $\neg(\text{believe that } p)$,” and “S ought to believe that $(\neg p)$.” The former tells us that we should not believe p , leaving open the option of being agnostic about p , whereas the latter positively obliges us to believe that p is false.

⁹ Bykvist and Hattiangadi report that Wedgwood suggested this revised norm to them in personal communication.

Bykvist and Hattiangadi allow that this amended version avoids obliging belief in impossibly complex propositions. However, they then raise their main objection with a second sort of problematic proposition:

Blindspot Propositions: A proposition that is possibly true but which cannot be both true and believed.¹⁰

According to (T*), such a blindspot proposition,¹¹ if true, ought to be believed, and yet by definition it is impossible for such a proposition to be both true and believed—which contradicts (OC). Bykvist and Hattiangadi conclude: “the hypothesis that belief is constitutively normative is false” (277).

3. A BETTER TRUTH-NORM

I want to propose a different truth-norm, one that Bykvist and Hattiangadi do not consider:

(F): For any S, p : S ought (not to believe that p) iff p is false.¹²

By imposing an obligation not to believe, (F) avoids both the problem of overcomplex propositions and the problem of blindspot propositions. For it is eminently possible not to believe an immensely complex (false) proposition, just as it is perfectly possible not to believe a (false) blindspot proposition.

(F) is a negative rather than a positive norm; it forbids belief rather than requiring belief. Still, it is a truth norm. And, it seems to me, it is a truth-norm with just as much intuitive appeal as (T). That we should avoid believing falsehoods seems no less basic or important than believing truths. Moreover, (F) does not need the slightly clunky restriction to propositions we actually

¹⁰ Note that overcomplex propositions will count as blindspots by this definition. If for some reason we want to keep these two types of proposition distinct, we should add the requirement that the proposition can be believed, as in the definition given at the end of section 1.

¹¹ Bykvist and Hattiangadi provide as an example of a blindspot: “It is raining and nobody believes that it is raining.” One might worry about this blindspot proposition that it does not explicitly specify a *time* at which it is raining and nobody believes it is raining. To make the time explicit we might modify the blindspot to: “During temporal interval t_1 – t_2 , it is raining and nobody believes that it is raining.” It is perfectly possible for this to be true and believed by S at any time outside of the interval t_1 – t_2 . Still, though, during the interval t_1 – t_2 , the truth of this modified blindspot would, according to (T*), create an obligation to believe it. And were one then to act on this obligation during t_1 – t_2 , the blindspot would be rendered false, so one would end up believing a falsehood. Making the temporal range of the claim explicit then will not rid (T*) of the blind-spot problem. For this paper, I assume, for simplicity, that a proposition’s truth-value can vary at different times.

¹² After completing the initial drafts of this paper, I discovered that Whiting (2010) proposes an equivalent norm based on permission to believe truths in response to the problems with blindspots that Bykvist and Hattiangadi raise.

entertain, which Wedgwood proposed in (F*). I will automatically fulfil my obligation not to believe falsehoods with respect to all the (false) propositions that I will never so much as consider.

(F) also maintains, unlike Boghossian's proposal (Ti), a biconditional formulation so that the truth or falsity of any proposition places us under an obligation with respect to our beliefs. While it is clear that, according to (F), the falsity of a proposition places us under an obligation, it might not be so immediately obvious that a true proposition also places us under an obligation. To see that it does, consider that (F) is equivalent to:

(Fa): For any S, p : S ought (not to believe that p) iff $\neg p$ is true

and so to:

(Fb): For any S, $\neg p$: S ought (not to believe that $\neg p$) iff p .

(Fb) tells us that when a proposition is true we ought not to believe the (false) negation of that proposition. Notice then that the truth of a blindspot proposition, or the truth of an overcomplex proposition, *does* still impose an obligation on us, but this obligation concerns the (false) negation of such propositions rather than the propositions themselves. And as these obligations created by true blindspot propositions, or true overcomplex propositions, are perfectly possible to fulfil, there is no conflict with (OC). Of course, neither (F), (Fa), nor (Fb) tell us anything about when we *should* positively believe a proposition.¹³ I will return to considering whether a norm for beliefs should positively oblige belief in some circumstances in section 8. For now, I will just repeat that a requirement to avoid believing falsehoods looks to be no less important or fundamental than a requirement to believe truths.

4. AN OBJECTION

One might worry that blindspot propositions will still create problems for (F) based on the following line of thought. In slightly more condensed form we can write (F) as:

Ought $[\neg\text{Bel}(p)]$ iff $\neg p$

which entails:

$\neg\text{Ought} [\neg\text{Bel}(p)]$ iff $\neg\neg p$,

¹³ Sutton (2007) suggests: "Philosophers who have claimed that there are epistemic obligations have most commonly taken those obligations to be negative," though Sutton also notes the not uncommon claim that one ought to believe what one has good evidence for. I reply to this latter thought in section 8, below.

which is trivially equivalent¹⁴ to:

$\neg\text{Ought} [\neg\text{Bel}(p)]$ iff p .

Now, it is standard in deontic logic to treat not being obliged not to do something as equivalent to being *permitted* to do it. In other words: $\neg\text{Ought} (\neg X) = \text{Permitted} (X)$. So then (F) will have entailed:

$\text{Permitted} [\text{Bel}(p)]$ iff p .

At first glance, permitting belief in truths might seem a pretty sensible consequence for a doxastic norm to have. However, when we come to consider blindspot propositions, it is a consequence of (F) that we are permitted to believe a currently true blindspot, and yet, if we were to believe the blindspot, (F) would forbid the belief. So it looks like (F) both allows and forbids belief in the blindspot proposition. A failing in a norm, or so you might think.

5. RESPONSES TO THE OBJECTION

5.1

One way of pinning down our unease at the way (F)'s normative behavior *varies* when faced with blindspots is the following thought. As forming a belief in a blindspot always results in believing a falsehood, a decent truth norm should *always* forbid belief in blindspots—whether they are currently true or currently false. And yet, a consequence of (F) is that there is no norm forbidding belief in blindspots that are currently true.

We need to move cautiously here when talking about “believing true blindspots.” There is a scope difference between:

$\diamond[\text{Bel}(p) \ \& \ p \text{ is true}]$

and

$\diamond[\text{Bel}(p)] \ \& \ p \text{ is true}$

Where p is a true-belief-blindspot, the former claim is false by definition; there is no possible world in which (BS) is both believed and true. The latter claim, however, can be true; (BS) can be true in the actual world while there is a possible world in which (BS) is believed. In other words, it is impossible to

¹⁴ I have assumed, for simplicity's sake, bivalence throughout this paper. If one allowed propositions to lack truth-values, one might worry that (F) permits belief in propositions that are not false and so permits belief in propositions that are neither true nor false. In this case, we could try amending (F) to something like: For any S, p : S ought (not to believe that p) iff p is not true.

believe a still-true blind-spot. But with a blindspot that happens currently to be true, the act of believing that proposition—truth or falsity not specified—is a perfectly possible act. Now, there seems little point in forbidding the impossible.¹⁵ We do not need a norm forbidding belief in a still-true blindspot. But one still might have the intuition that we need a norm forbidding belief in blindspots—truth or falsity not specified—which remains operative when the blindspots happen to be true.¹⁶

I have two things to say in response here. Firstly, it seems to me that as much as we may have an intuition that attempting to believe a blindspot that happens to be true should be discouraged, we also have a conflicting intuition that we *would like* to believe such a true proposition—if only, per impossible, we could believe it without affecting its truth value. A feature of my proposal is that there is no truth-norm forbidding belief in a currently nonbelieved blindspot, but there is a norm that kicks in the moment that one ventures to believe the blindspot. I think this can be seen as a virtue rather than a vice, in that it does justice, as far as possible, to *both* of these conflicting intuitions. While the blindspot is nonbelieved it is a truth and we feel, I suggest, that *so long as it remained a truth* there would be nothing wrong in believing it, so there should not be a rule preventing belief in it so long as it remains true—as is entailed by (F). Of course, any actual instance of belief in a blindspot will be an instance of false belief, and we want a norm condemning such beliefs, which is just what (F) will do the moment anyone ventures to believe a blindspot.

Secondly, if we are merely entertaining a blindspot (that happens to be true) without yet believing it, we can, with a modicum of inference on our part, come to realize that *were* we to believe the blindspot, it would be false, and, thus, realize that *were* we to believe it, we would fall foul of my truth-norm (F). So, although (F) entails that there is no norm forbidding belief in a blindspot that happens to be true, my proposal does speak to anyone who can identify a proposition as a potential blindspot and make an elementary inference or two. Obviously, we can never identify a proposition as a *true*

¹⁵ However, Smullyan (1983), crediting Hintikka, provides the following, possibly not entirely serious, argument. Assume it is forbidden to do anything that entails the destruction of the world. Doing the impossible entails destroying the world. Hence, it is forbidden to do the impossible. Thanks to an anonymous referee for drawing my attention to this.

¹⁶ Bykvist and Hattiangadi discuss the difference between wide-scope and narrow-scope truth-norms towards the end of their paper and criticise a wide-scope variant of (T*). Notice that my proposed norm, (F), is, in their terms, a narrow-scope norm. (F) does *not* make the following wide-scope claim: For any S, p : S ought (not to believe that p iff p is false). Rather, the belief-act prohibited by (F) is specified entirely on the left-hand side of the biconditional without mentioning truth or falsity—truth and falsity are only mentioned in specifying the conditions under which the prohibition applies (on the right-hand side of the biconditional).

blindspot, for identifying it as such would be to believe it and so render it false. But we can identify propositions as (true-belief) blindspots while remaining uncommitted as to their truth-value.

5.2

But perhaps the reason one feels unhappy about (F)'s handling of blindspots is based on a different, more general thought: if a norm's condemning or not condemning an action depends on whether the act has actually been committed, then the norm is thereby defective. (F) does not forbid belief in a blindspot unless the act of belief is actually committed. So just how dodgy is this? Is it a fatal flaw in a norm?

In another recent paper, Bykvist (2007) himself argues that there is nothing essentially wrong with "normative variance," that is, "that the normative status of an action may depend counterfactually on its own performance" (Bykvist 2007, 99). Bykvist argues that normative variance is only a problem for a norm if it makes it a poor guide to action, in particular, if it makes it impossible for the agent to comply with the norm. So does the normative variance of (F) make it a poor guide to action or make it impossible to comply with?

The answer is no on both counts. As we have seen above, when faced with blindspot propositions, given a modicum of inference, (F) can provide a perfectly good guide to action—if we can grasp that belief in the blindspot *would* render it false, then we grasp that believing in the proposition *would* be condemned by (F) (even though it is not currently condemned). And even if we are not able to identify a proposition as a blindspot, it is certainly still possible to comply with (F) even in the case of blindspots that currently happen to be true. To repeat: with a blindspot that happens to be true, (F) tells us not to believe its (false) negation, an obligation that it is eminently possible to fulfill.

I think the spookiness of this normative variance—that the wrong-ness of an act of belief depends on whether it is actually committed—is somewhat dispelled when we examine the *basis* for the counterfactual dependence. Consider the following quotidian proposition:

RAIN: It is raining.

Let us assume, for simplicity, that RAIN is sometimes true and sometimes false (rather than being tacitly indexed to a time and so eternally true or false). It might on occasion happen that just at the very moment one comes to believe RAIN, RAIN switches from being true to false. Here, the act of believing RAIN is only condemned by (F) once it is actually committed. Surely nobody would find this a spooky or untoward feature of the truth-norm (F). And here,

of course, there is no counterfactual dependency on the belief-act actually being committed; it is not *because* the belief-act was actually committed that it became condemned. Condemnation of the belief-act depended only on the truth or falsity of the belief, which here just happens to coincide with the belief act actually being committed.

Now in the case of the blindspot proposition there *is* such a counterfactual dependency—whether or not the option of believing the blindspot is actually taken up *determines* whether the act is condemned by (F). But, we should notice that it is not the actual commission of the act *in itself* that is relevant to (F), but only in so far as this actual committing affects the truth-value of the proposition. It is still the truth or falsity that is explanatorily fundamental to which belief-acts (F) will condemn. It just so happens with blindspot propositions that their truth or falsity is intrinsically linked to their actually being believed. In other words, if there is any feeling of spookiness here it should be attributed to the nature of the blindspot proposition, not to the variable behavior of the truth-norm (F) when dealing with the blindspot. (F) has no explicit concern with whether an act of belief is actually committed or not, (F) is solely concerned with the truth-value of propositions believed. It is the strange quirk of a blindspot that its truth-value depends on whether it is actually believed or not. Confronted with such a strange beast as a blindspot, (F)'s varying pattern of normative-condemnation and normative-permission is just what we should want and expect of a truth norm;¹⁷ it is not really so different to the sort of variation that might be coincidentally exhibited with respect to propositions like RAIN.¹⁸

6. SUBJECTIVE AND OBJECTIVE NORMS

I have argued that the modified truth-norm, (F), which is concerned to forbid false beliefs, rather than recommending true beliefs, avoids the problem of true-belief blindspots. This allows us to endorse the plausible principle (OC)

¹⁷ I return to this point, providing further considerations in its favor, in section 9.

¹⁸ It might be thought that if my defense of (F)'s normative variance is successful, it should also exonerate the behavior of the truth-norms (T) and (T*) with respect to blindspots. But it is important to bear in mind Bykvist and Hattiangadi's formulation of (OC): "If I have an obligation to perform an action A, then it is possible for me to do A *while being obligated to do A*." The sort of variance displayed by the truth norms falls foul of (OC)—I am obliged to believe a true blindspot but I cannot do this while still being obliged to, because the exact moment of complying with the obligation is the exact moment the obligation lapses (as the blindspot turns false). In contrast, according to (F) my obligation is not to believe a false blindspot, an obligation I can comply with while it remains in force, for I am able to not believe a blindspot while the blindspot remains false. I am very grateful to Jonathan Ichikawa for pressing me on this point.

while still holding on to the idea that the truth or falsity of any proposition places one under an obligation with respect to one's beliefs. I have also defended my proposed truth-norm against objections that its variable normative behavior is unacceptable.

My argument so far leaves it open that there might be other objections to (F), or to the more general idea that belief is constitutively normative. In particular, I have not addressed the other sort of worry I mentioned at the start of the paper. For example, one might think the following principle is plausible:

Parity Principle: If evidence for a true p makes it permissible to believe p , then equally good evidence for a false q makes it permissible to believe q .

Such a principle looks like it would cause problems for any norm, such as (F), which forbids belief in falsehoods. We can put this in terms of blindspots by considering how knowledge blindspots, or justified-true-belief blindspots, might seem to cause problems for (F). For example:

p & there is no justification for believing p ,

or

q & all the available evidence points to $\neg q$.¹⁹

(F) entails that one is permitted to believe these propositions when they are true. But are we really permitted to believe a proposition for which we have no justification? Or which all the evidence suggests is false? I want to briefly indicate how such apparent problems for (F) disappear if we accept the common distinction between subjective and objective norms. Pollock glosses this distinction as follows: "The subjective/objective distinction can be regarded as a distinction between evaluating the person and evaluating his act"²⁰ (Pollock 1987, 77).

Consider moral norms: clearly, if I have excellent evidence that an act is permitted, but it turns out that the act is actually morally wrong, I am exonerated from having acted irresponsibly—I was just unlucky. Nevertheless, the act was still wrong—it was not *really* permitted, it only *seemed* to be permitted from my unlucky evidential position. From a particular subject's

¹⁹ As mentioned in section 1, notice these are *not* true-belief blindspots—one could form a true, though epistemically irresponsible, belief in either of them. And perhaps, given a strongly externalist view of justification, these are not justified-belief blindspots (as opposed to justified-true-belief blindspots) either (see the discussion at the end of this section, below).

²⁰ Pollock (1987) goes on to explicitly draw an analogy between subjective moral obligation and a subjective notion of epistemic justification, and an analogy between objective moral obligation and truth (77–78).

point of view there is no sensible way of trying to abide by the norm, “one ought not to act badly” other than by trying to obey the norm, “one ought not to act in a way that one’s total evidence indicates to be bad.” But from some more enlightened, third-person perspective these two principles can be seen to pull in conflicting directions for some unlucky subjects.

Such occasional conflicts between our subjective and our objective moral obligations hardly cast doubt on the existence or coherence of moral norms. And nor should such conflict between subjective and objective norms governing belief impugn the proposed truth-norm, (F). The only way an agent can *rationally try* to abide by a norm like (F) is by forming beliefs according to her available evidence for a proposition’s truth or falsity. We might say that for a rational agent, attempting to abide by a truth-norm for belief will be *transparent* to attempting to abide by an evidential-norm for belief. We can frame a subjective norm of evidence as follows:

(EV): For any S, p : S *subjectively* ought (not to believe that p) iff S does not have adequate evidence that p .²¹

The same sort of normative variance we saw with the truth norm (F), will occur in parallel fashion with the evidential norm (EV). For instance, suppose that you do not yet believe the (currently) true proposition: “It is raining and nobody believes it is raining,” but you have mounting evidence that it is true. According to (EV), adequate evidence means you are subjectively permitted to form a belief in this proposition, just as (F) objectively permitted belief given the current truth of the proposition. But if one were to go ahead and form the belief, that very belief would, assuming normal access to one’s own beliefs, constitute one’s having undermining evidence that the proposition is false, and so one would no longer be (subjectively) permitted to hold this belief. Assuming my argument about normative variance with respect to (F) in section 5.2 was correct, the same considerations should apply here, and we should find nothing untoward about such parallel normative variance with (EV).

Matters are more complicated with a proposition such as: “It is raining and nobody has adequate evidence that it is raining.” As mentioned back in section 1, in order for S to have adequate evidence that this proposition is true, S would have to have both adequate evidence that it is raining, and adequate evidence that S herself does *not* have adequate evidence that it’s raining. It is not immediately clear whether or how this could happen in

²¹ Equivalently: For any S, p : S is subjectively permitted (to believe that p) iff S has adequate evidence that p .

normal, rational humans. Those with a strongly internalist conception of evidence will presumably deny that such a thing is possible—if one was fully aware of one's evidence for p and aware of one's evidence that one does not have adequate evidence for p , the net result would be that the proposition was inadequately evidenced for one to believe. However, if one's conception of a subject's evidence allows for the subject to have limited or partial access to their own evidence,²² then this possibility might open up. For instance, perhaps if all of one's evidence for rain and one's (misleading but good) evidence that nobody has evidence of rain were unconscious, one could be said to have, simultaneously, adequate evidence both that it is raining and that there is no evidence that it is raining—and so one's belief then would be subjectively permitted, though objectively wrong. However, this requires a contentiously externalist view of evidence, on which a subject can form a belief that it is raining on an "adequate evidential basis" despite the subject currently lacking conscious access to this evidence and, furthermore, having good evidence *precisely that they lack good evidence that it is raining*. Such a subject seems at least to be abnormally "alienated" from her own evidence.

The main point here is that if the agent is unlucky and contravenes (F) by forming a false belief, despite forming this belief in accordance with (EV), then she has formed her belief just as she ought given her subjective position, but her belief is still objectively not what it ought to be. Conversely, if an agent irresponsibly comes to form the lucky true belief: " q and all the available evidence points to $\neg q$," then she has not formed her belief as she subjectively ought; but her belief is, fortuitously, just what it objectively ought to be. Echoing Pollock, a norm of evidence such as (EV) evaluates the believer (the belief forming process); a truth-norm such as (F) evaluates the belief.²³ Far from being rivals, the two are connected in that a rational subject could only try to abide by the latter sort of norm in so far as they try to abide by the former.

7. GUESSING

I have argued that a norm for belief, whether it is an objective truth-norm or a subjective evidential norm, is best formulated negatively—in terms of what we ought not to do (or, equivalently, in terms of what we are permitted to do).

²² For example, Williamson writes: "Whatever evidence is, one is not always in a position to know what one has of it" (2000, 178).

²³ See also Vahid's (2006) distinction between doxastic and epistemic goals. Vahid distinguishes between a Belief's having truth as its aim and a Subject's having true beliefs as her aim.

Another mental attitude for which the “aims at truth” slogan seems apt is guessing.²⁴ Indeed, Owens (2003) argues that guessing is a better candidate than believing to count as an attitude that “aims at truth.” In this section, after making a couple of brief criticisms of Owens’s way of defining the attitude, I want to show how the present approach extends to guessing.²⁵

As an initial clarification, we should distinguish between two notions of what a guess is. In one sense, a guess is an act of selecting or indicating some option or outcome that occurs in the right kind of “guessing context,” for instance, in a quiz. In this sense, one may count as having “guessed that p ” even if one actually knows full well that not- p . However, in another sense of “guess” this would not count as a “genuine” guess. In this second sense, a guess is an act or attitude that reflects one’s subjective probabilities for the range of possible options—or at least does not conflict with one’s subjective probabilities so far as one is aware. So one cannot genuinely guess option A if *one is aware that* one’s own subjective probability for some other option is higher than that for A.²⁶

Owens suggests that a guess might still be a genuine attempt to “aim at the truth” despite a subject’s awareness that it does not reflect her subjective probabilities.²⁷ For example, in a quiz show situation, I suspect that option B is more likely, but the reward if I choose A and it turns out to be correct is much greater (my subjective probabilities favor B but do not rule out A as hopelessly unlikely). Owens’s idea is that because the only way of winning the greater reward is if option A is, in fact, correct, then my choice would still be genuinely aiming at the truth. I think this suggestion just muddies the useful distinction drawn above. One is “aiming at truth” here *only insofar as* one is aiming at maximizing one’s reward, and choosing truly happens to be the means to this end. The situation Owens describes is essentially betting on a long shot one does not expect to win, but which has (what you consider to be) generous odds. Suppose, after betting on horse A, a 500–1 outsider whose chances of winning you estimate to be only 50–1, someone suggests,

²⁴ I am grateful to an anonymous referee for urging me to consider how and whether my approach applies to the attitude of guessing.

²⁵ I make no attempt here to reply to the aforementioned main contention of Owens (2003), that belief does not “aim at truth,” at least not in the way that Owens takes guessing to so aim. However, see my criticisms of Owens’s way of explaining the notion of “aiming at truth” below in this section. See also Vahid 2006 for criticism of Owens’s overall position.

²⁶ Notice the italicized qualification here: we must allow that subjective probabilities might be less than perfectly transparent, and that people can be less than perfect reasoners—so perhaps one can *genuinely* guess A even though there is another option one has temporarily forgotten about that one thinks more likely. The point here is that a genuine guess must genuinely *attempt* to be in line with one’s own subjective probabilities.

²⁷ Though he does not fully endorse the suggestion.

“so you’re guessing A is gonna win.” A perfectly reasonable response here would be, “I don’t *guess* that A will win, in fact I strongly expect A won’t win, but at 500–1 I’m sure as hell *betting* that A will win.”

Owens, who endorses the claim that guessing constitutively aims at truth, explicates the notion of “aiming at truth” in terms of a subject’s intentions: “ \emptyset -ing that p aims at the truth if and only if someone who \emptyset s that p does so with the purpose of \emptyset -ing that p only if p is true” (2003, 292). Owens, then, is committed to the idea that someone who guesses that p must do so with the *purpose* of: guessing that p iff p . But it is not clear to me that making a genuine guess really does require such a purpose. If someone asks me to guess what number they are thinking of, where I have absolutely no evidence favoring any particular number from the potential infinity of options, I might still make a *genuine* guess, even though my subjective probability that this option is correct (or indeed any other option) approaches zero. My *purpose* in guessing here may have nothing to do with hitting the truth—my only purpose might be to indulge my interlocutor’s whim.²⁸ Still, my guess is genuine so long as the number I choose is as good an option as any other—there is no other option I assign a higher subjective probability (so far as I am aware).

In any case, let us assume that Owens is at least correct that guessing is in some sense constitutively normative and that the relevant norm is connected to truth/falsity. It does not immediately follow that guessing must be subject to exactly the same kind of norm(s) as belief—there are, after all, differences as well as similarities between the two attitudes. However, I think that the norm(s) for guessing are very similar to those for belief and should likewise be formulated negatively so as to avoid problems with overcomplex propositions and guessing blindspots.

Having distinguished between subjective and objective norms in the previous section, I propose that the fundamental negative norms for guessing are:

Objective Guessing Norm: For any S, p : S ought (not to guess that p) iff p is false.²⁹

Subjective Guessing Norm: For any S, p : S ought (not to guess that p) iff p is less well-evidenced than some other guessable option.³⁰

²⁸ More generally, it is not clear to me that genuine guessing requires that the subject have *any* definite purpose in mind (apart, perhaps, from the trivial purpose of guessing), but I will not pursue this any further here.

²⁹ Equivalently: For any S, p : S is permitted (to guess that p) iff p is true.

³⁰ Equivalently: For any S, p : S is permitted (to guess that p) iff p is as well-evidenced as any other option.

A false guess is objectively faulty in just the same way as a false belief—you ought to have guessed differently. If S has formed her guess according to the available evidence but the guess turns out, unluckily, to be false, one might say, “You made the smart guess there with option A, but with the benefit of hindsight we can now see that you ought to have guessed B.” So far as subjective norms go, clearly the evidential standard for guessing should be less strict than that for belief, whatever we understand the precise amount of “adequate” evidence for belief-formation to be. One can be permitted to guess things that one is not permitted to believe. Still, there is a subjective norm for guessing, and, as with the norms for belief, a necessary condition on trying to abide by the objective guessing norm is that one tries to abide by the subjective guessing norm. Indeed, I have claimed that a (normal, rational) subject cannot *knowingly* flout the subjective guessing norm and still count as *genuinely* guessing.

When it comes to guessing blindspots, there are some issues concerning whether or how it is possible to *genuinely* guess a proposition such as: p and I guess that not- p . The issues here concern whether (the degree to which) one’s own current guessing states are “transparent” to one, and what sorts of irrationality or self-blindness would have to be involved for us to make sense of this proposition being the content of a *genuine* guess. Clearly for any normally rational subject capable of the most basic inferences, it should be manifest that guessing such a proposition cannot be correct, so it is hard to see how a normally rational subject could make a genuine guess with this content. However, when the guessing-blindspot proposition is less obviously Moore-paradoxical—for example, “it is raining and nobody guesses it is raining”—it is easier for us to make sense of a subject’s failing to notice the self-undermining nature of making such a guess so as to credit the guess as being *genuine*.

Once more, the same sort of normative variance discussed in sections 4, 5, and 6 will arise with true-guess-blindspots. The negative norm does not oblige us to guess currently true blindspots, though it does permit such a guess; which guess is then immediately forbidden if it is actually carried out, for it will have turned the blindspot false. Once more, I contend that this normative variance, when faced with a proposition whose truth-value varies depending on whether it is actually guessed, is unproblematic.

8. NEUTRALITY

A different worry someone might have about (F) is that in imposing only a negative norm, and not imposing any positive obligation to seek true beliefs, we can abide by (F) simply by having no beliefs at all. A thing to say here first

is that (F) will provide an interesting, substantive normative constraint on anyone who is an active believer and will be fulfilled trivially by nonbelievers. Compare: “Thou shalt not covet thy neighbors oxen.” This moral norm will be trivially complied with by inanimate objects, but it is still a perfectly good and substantive moral rule for us humans who are in the business of forming desires. Given that at least some beliefs are being formed (and given bivalence), then a believer seeking to abide by the falsity norm will thereby end up “aiming at truth” on those occasions that she does venture to believe. So on the assumption that a subject is actively forming beliefs, the negative norm is enough to ensure that she aims at truths, though not at *every truth*, which is just as well given the existence of true overcomplex propositions and true blindspots.

However, a related worry is that the negative evidential norm (EV) permits a subject to remain neutral in the face of overwhelmingly strong evidence for *p*. It is certainly true that a subject who fails to form a belief that it is raining despite a wealth of subjectively available evidence that it is raining (and none against) strikes us as somehow crazy or “irrational.” But I want to suggest that the *kind* of irrationality involved here is *practical irrationality* rather than *intellectual or theoretical irrationality*.³¹ The proposed norms, (F) and (EV), are meant to tell us what we ought not to do qua theoretically rational believers. That they are silent on matters of practical reason is to be expected. It would be unreasonable of us to expect a belief norm, in addition to stating what theoretical rationality demands of us, to state what practical reason demands as well. For the demands of practical and theoretical reason can come apart, just as the doxastic obligations created by truth/falsity can come apart from those obligations created by our evidential position. Thus, that (EV) permits neutrality despite excellent evidence should not be seen as a defect.

If a subject’s belief-forming processes do not abide by (EV) then she is guilty of an intellectual failing. The way the belief was formed was an intellectual error (even if it is luckily true) in that the subject did not have the inferential (or other kinds of justificatory) basis to license its formation. But when a subject simply draws no conclusion whatsoever, despite possessing an adequate justification base to do so, this is a quite different kind of error (if it is an error at all). Such neutrality is not an *intellectual or theoretical error*. It may be a practical error; after all, when a subject has all the available justification she needs to believe *p*, it *normally* “costs” essentially nothing to go ahead and form the belief, and so gain the benefits of an improved “mental map” of

³¹ I am taking this common distinction to be clear enough for our purposes and offer no further elucidation of it here.

the world. Indeed, one suspects it would normally take *more effort* (assuming it is possible at all) to remain neutral in the face of excellent evidence than to simply allow oneself to form the belief, quite apart from the subsequent practical benefits of an improved set of beliefs. Nevertheless, no error of theoretical reasoning has occurred, for there has been no theoretical reasoning—no inference was made, no belief has been formed.

It might be suggested at this point that there is a difference between a subject who simply draws no conclusion whatsoever on the basis of her evidence, and a subject who actively decides to remain neutral, or to “withhold” belief *as an intellectual response* to her evidence.³² Whereas the former case may involve no intellectual or theoretical error, this latter case, it might be suggested, *is* an instance of intellectual or theoretical error.³³

I think that this suggestion is correct, but that such an attitude of deliberate neutrality concerning some proposition, *p*, can only count as an intellectual response to the evidence insofar as it is an *attitude concerning whether the evidence supports p*, as well as concerning *p*. What is the difference between a subject who simply gets bored or distracted while deliberating over her evidence concerning *p*, and so forms no belief at all, and a subject who actively decides to withhold belief as an intellectual response to her evidence? For the neutral attitude to be an intellectual response to the evidence, and so be subject to the demands of theoretical rationality, the subject must be taking a view as to what the evidence supports. Without any such view concerning her evidence, the subject’s neutrality collapses into the former case, in which there simply is no intellectual or theoretical response to the evidence. In requiring a view about one’s own evidence (one’s own potential beliefs), the attitude of active withholding concerning *p* differs from the attitudes of actively believing or disbelieving *p*, which require only a view about *p*. Michael Bergmann (2005) writes:³⁴

³² I am very grateful to Berislav Marusic for helpful conversation and correspondence on this point.

³³ Pyrrho is said to have counselled us *never to form beliefs at all*. I confess it is not entirely clear to me whether the motivation for Pyrrhonian neutrality is: (i) because it is the path to happiness (practical considerations), or (ii) because of skeptical arguments that we can never have adequate justification for belief (theoretical considerations). If Pyrrhonism is motivated in the former way, it may or may not be providing good practical advice, depending on one’s view of how practical a course in complete apathy (*ataraxia*) is, but the position should not be condemned for embodying a theoretical or inferential error concerning whether or when our evidence can support belief. However, if it is motivated in the latter way, then it may embody a theoretical or inferential error to the extent that the motivating skeptical arguments are unsound.

³⁴ Once more, I am extremely grateful to Berislav Marusic for drawing this passage to my attention.

Withholding p , then, is a propositional attitude distinct from mere failure to take up any attitude towards p . Like believing or disbelieving, it is taking an attitude towards a proposition. What more can one say about with-holding? As I shall be using the term, withholding p involves resistance, voluntary or involuntary, to believing p and to disbelieving p . The only thing one must consider in order to believe p or to disbelieve p is p (or its denial). But to withhold p (in the sense I have in mind) one must, in addition, consider the prospect of one's believing p as well as the prospect of one's dis-believing p ; otherwise one will not be able to resist both believing p and disbelieving p . So withholding p involves not only an attitude towards p but also attitudes towards attitudes towards p . (2005, 421)

Thus, an attitude of *active withholding* concerning p , despite having excellent evidence for p , would indeed count as an intellectual or theoretical error—but this is because such an attitude embodies a *false belief about what the evidence supports*. And so such an attitude would be forbidden by the falsity norm (F). So while active neutrality, or withholding, with respect to p , even when the subject has excellent evidence for p , is permitted by the evidential norm (EV), it is forbidden by the falsity norm (F), for it embodies a false belief about the subject's own evidence. (We have already seen, in section 6, that there is nothing problematic about conflicting normative verdicts from the norms (F) and (EV).)

In support of my claim that we need to distinguish between practical and theoretical irrationality, consider *when* (in what circumstances) it would be “crazy” to remain neutral in this way. This will depend on contextual matters and on the subject's cognitive resources and interests. For example, subjects can vary widely in their inferential abilities. It might be an immediately obvious inference for S1 that p follows from her other beliefs, so for her to remain neutral seems inexplicable or crazy. But for S2, such an inference might require a huge expenditure of time and cognitive resources, or be beyond her inferential abilities altogether. That S2 does not believe p despite having conclusive evidence is perfectly understandable—presumably we are all in S2's position with respect to some propositions that are entailed, but not obviously so, by our current beliefs. Similarly, when the proposition in question is something that is very unlikely ever to have any relevance to our interests and concerns (e.g., the proposition “the average monthly rainfall, to the nearest millimeter, in the Antarctic circle, in months whose names begin with vowels, is a prime number”), it seems much less “irrational” to remain neutral despite possessing adequate evidence than when the proposition is something that might well be relevant to our projects and interests (e.g., the proposition “smoking causes cancer”). And then, of course, one can imagine evil-demon scenarios in which were one to form a belief that p , the demon will kill innocent people, etc. Here, the subject's not making an inference that she

is both entitled and easily able to make, concerning a proposition that may be relevant to her desires and plans, is precisely the (practically) “rational” option.

In contrast, determining whether an inference or act of belief formation, given some body of evidence, is theoretically or intellectually “irrational” has nothing to do with the subject’s context, concerns, or cognitive resources. Whether, or to what degree, a particular body of evidence does or does not support a particular belief is entirely independent of one’s inferential abilities and interests. Even if a demon has contrived to set the situation up so that forming a belief that p is overwhelmingly in one’s best interests, there is still a quite separate fact of the matter as to whether or not such a belief has been (theoretically or intellectually) “rationally” formed given one’s evidence.

The norms (F) and (EV) concern how a believer should form her beliefs *qua intellectual subject*. One should not hold it against these norms that they are silent on when a believer should form her beliefs *qua practical agent*. For it is clear that there are situations where the demands of theoretical rationality and of practical rationality are in conflict, so we should not expect a single norm to cover both kinds of rationality.

9. SUBJUNCTIVE NORMS AND “TRUE-TURNS”

In this final section, I want to consider an alternative possible approach to dealing with blindspots. We could frame a “subjunctive” falsity norm as follows:

(F-SUBJ): For any S, p : S ought (not to believe that p) iff p would be false, were S to believe that p .³⁵

Clearly this subjunctive norm will not exhibit the sort of normative variance we saw with (F)’s handling of blindspots. So if one is unconvinced by the arguments of section 5, one might prefer this subjunctive variant. On the other hand, one might feel that such a norm is somehow an *ad hoc* modification to the core idea of “aiming for truth” (or “aiming away from falsity”).

I think the charge of *ad-hoc*-ness would be unfair here. Beliefs have the role, *inter alia*, of allowing for successful action by providing the subject’s “map of the world.” To provide a fully up-to-date “map,” these beliefs must

³⁵ Equivalently: For any S, p : S is permitted (to believe that p) iff p would be true, were S to believe that p . And for a subjunctive version of the subjective, or evidential, norm, (EV-SUBJ): For any S, p : S ought (not to believe that p) iff p would be inadequately evidenced for S , were S to believe that p . Or equivalently, for any S, p : S is permitted (to believe that p) iff p would be adequately evidenced for S , were S to believe that p .

accurately represent how the world is *including the existence of those very beliefs*. A subject who wants a totally up-to-date “map” by which to navigate reality should aim to update her beliefs so that they end up in accord with reality *after the updating is completed*, rather than aiming for beliefs that accord with the state of play prior to the updating’s completion. The process of belief formation then should aim not at the current actual state of affairs, but at the counterfactual state of affairs that will be brought about by the existence of the very belief that is the end product of the belief-forming process. Or so a proponent of (F-SUBJ) could argue. (Of course, for the normal case in which a proposition’s truth value does not depend on its being believed or not, these two tasks collapse into one.)

However, I think it should be admitted that this subjunctive approach does amount to a modification of the idea that belief aims at truth. For now the slogan should be something like: “believers aim at true beliefs,” or perhaps, “belief aims at what will be the truth.” (Though (F-SUBJ) still holds that the objective standard of correctness for a fully-formed belief is simply truth/falsity.)

I want to finish by raising a possible problem for this subjunctive norm. Consider the opposite sort of proposition to a true-belief blindspot:

(TRT): There is no greatest prime number and somebody believes that there is no greatest prime number.

This proposition can be false, but it cannot be false and believed.³⁶ Let us call such a proposition a “true-turn.” The Cogito provides other obvious examples: “I am currently conscious” or “I exist” are both true-turns. (Notice, though, that while belief in these propositions is sufficient to turn them true, it is not necessary. In what follows, I will assume that we are talking about true-turns that *require* belief in order to be true, rather than requiring consciousness or the subject’s existence.)

According to the nonsubjunctive norm (F), we ought not to believe a currently false true-turn. Of course, if one did go ahead and believe it anyway, one would still end up with a normatively kosher status, as the proposition would turn true—a win-win situation, though one exhibiting normative variance. In contrast, the subjunctive (F-SUBJ) holds that one *is permitted* to believe a currently false true-turn, whose normative status remains unchanged if one does actually form the belief. Now I want to suggest that there *is* something dodgy about the *act of forming* a belief in a currently false true-turn, though, the resultant belief, once fully formed, will be just as it

³⁶ In the formalism of section 1: $\diamond\neg p$ & $\neg\diamond[\neg p \ \& \ A(p)]$.

should be. And so we might feel that (F-SUBJ) is dodgy insofar as it permits us to go ahead and start *forming* a belief in a true-turn that is still false.

Let us suppose that forming a belief that p involves raising one's level of confidence (subjective probability) that p is true—that is, currently and actually true—until some threshold is crossed, at which point one *believes* that p .³⁷ And let us suppose a subject is in the position of considering, but not yet believing, the *currently false* true-turn (TRT). Now bear in mind here that we are considering a subject gearing up to form a belief in *the whole proposition* (i.e., both sides of the conjunction), in one fell swoop as it were. We are *not* imagining a situation in which a subject first believes that there is no greatest prime and then afterward (via introspection and inference) comes to believe that: “somebody believes that there is no greatest prime.”

In order to form a belief in this true-turn, the subject must raise her confidence that it is true up to a certain level. But it seems that the subject *ought not* to be growing more and more confident that the proposition is true, for, until the threshold of confidence for belief is actually reached, the proposition is (still) false. And though, of course, the final state of belief will be as it should be, it is not obvious that this is any reason to retrospectively exonerate the preceding action of increasing one's confidence in a falsehood. (Suppose the act of increasing confidence had stopped short of the threshold for belief.)

A proponent of the subjunctive form of norm might here protest that (F-SUBJ) is concerned with whether a proposition *would be* true or false if believed (rather than actual truth or falsity), and that what a subject *can* permissibly grow more confident about is that the proposition “there is no greatest prime and somebody believes that there is no greatest prime” *would be true if she believed it*. But this does not defuse the problem, for increasing one's confidence that “ p would be true-if-believed,” will only result in a belief that “ p would be true-if-believed,” *not* a belief in p itself. So increasing one's confidence that (TRT) would be true-if-believed only results in one believing that (TRT) would be true-if-believed, not in an actual belief in (TRT) itself. Thus TRT remains unbelieved. And what we were interested in was the act of believing TRT.

³⁷ This supposition is made for ease of exposition. I do not mean to suggest that belief formation could not be instantaneous. But as any candidate belief norm is presumably meant to provide an essential aim for any/all instances of belief, assuming we can even sometimes form beliefs by gradually increasing our confidence in the truth of a proposition—if this is so much as possible—then a genuine belief norm must be able to handle such gradual cases. And so the issues I raise in this section are relevant even if our belief formation is usually instantaneous, on the assumption that any candidate norm must still aspire to deal with possible cases of gradual formation. Thanks to Conor McHugh for helpful correspondence on this point.

Alternatively, the subject *can* permissibly grow more and more confident in the truth of a *different* proposition: “there is no greatest prime and somebody is *going to* believe that there is no greatest prime.” Once a certain level of confidence is reached, the subject will believe this different proposition, and so the original true-turn, TRT, will have turned true. But again, what the subject here ends up believing is simply a different proposition to (TRT). In order for (TRT) *itself* to be believed, a subject must raise her levels of confidence that (TRT) *is actually* true. And so the problem remains that as this action of confidence adjustment is going on, the subject’s confidence seems to be increasing as it ought not to.

There is, then, a kind of “leap into the dark” when forming a belief in an as yet false true-turn. Until the belief is fully formed, the process of belief formation seems to be aiming in the wrong direction, at falsity. Thus, I contend that the normative variance displayed by (F) with respect to a true-turn is *just what we want* from a belief-norm.³⁸ Whereas the normative invariance displayed by (F-SUBJ) has the anti-intuitive consequence that it is okay to be increasing one’s level of confidence in a (currently) false proposition. I suppose a proponent of the subjunctive form of norm might just deny the intuition that there is something amiss here—after all, though the proposition is currently false, it would be true if believed. And it is this latter condition that matters according to (F-SUBJ). But I take this to be an unintuitive cost of the subjunctive version of the norm—raising one’s level of confidence that a false proposition is *actually* true, does not seem okay just because that proposition will soon be true or would be true in different circumstances.

If for some reason one felt that blindspots and true-turns should receive differing treatments from a norm, it is possible to combine (F) and (F-SUBJ) as follows:

(F-SUBJ*): For any S, p : S ought (not to believe that p) iff p is false or p would be false if S believed it.³⁹

Like (F-SUBJ), (F-SUBJ*) does not exhibit normative variance with blindspots—it forbids belief in a currently true blindspot as well as a currently false one. But unlike (F-SUBJ), (F-SUBJ*) *does* exhibit normative variance with true-turns—it forbids belief in a currently unbelieved (so false) true-turn, but permits actually formed beliefs in (true) true-turns.

³⁸ Compare my comments on normative variance in section 5.2.

³⁹ Equivalently: For any S, p : S is permitted (to believe that p) iff p is true and p would be true if S believed it.

And if, despite my argument in section 8, one still yearns for a norm positively obliging one to believe truths, we could try amending Wedgwood's proposal, (T*), in the following way:

(T-SUBJ*): For any S, p : if S considers p , then S ought to (believe that p) iff p is true and S believing p would not render p false.⁴⁰

These are, admittedly, much more complicated and clunky norms for belief, but they seem to provide another, rather less elegant, way of avoiding problems with blindspots. Given the “flat-footed” obviousness of the sort of amendments made in (F-SUBJ), (T-SUBJ*), and (F-SUBJ*), it is perhaps surprising that Bykvist and Hattiangadi do not consider this method of avoiding trouble with blindspots. Perhaps they would consider such norms too egregiously ad hoc and clunky to warrant mentioning—and I do agree that these strike me as ugly formulations. Still, “clunkiness” is in the eye of the beholder and someone committed to finding a truth-norm/falsity-norm that is essential to belief would presumably want more of an argument against (T-SUBJ*) and (F-SUBJ*) than mere aesthetic distaste.

Between the simpler norm, (F), and its uglier cousins, (T-SUBJ*) and (F-SUBJ*), those sympathetic to the idea that belief is constitutively normative have nothing to fear from blindspots.⁴¹

REFERENCES

- Bergmann, M. 2005. Defeaters and higher-level requirements. *Philosophical Quarterly* 55: 419–36.
- Boghossian, P. 2003. The normativity of content. *Philosophical Issues* 13: 31–45.
- Bykvist, K. 2007. Violations of normative invariance: Some thoughts on shifty oughts. *Theoria* 73: 98–120.
- Bykvist, K. and A. Hattiangadi. 2007. Does thought imply ought? *Analysis* 67: 277–85.
- Edgley, R. 1969. *Reason in theory and practice*. London: Hutchinson.
- Gibbard, A. 2003. Thoughts and norms. *Philosophical Issues* 13: 83–98.
- 2005. Truth and correct belief. *Philosophical Issues* 15: 338–50.

⁴⁰ Of course, one could still trivially comply with (T-SUBJ*) by never even considering any propositions whatsoever.

⁴¹ I am extremely grateful to Krister Bykvist and Anandi Hattiangadi for their many helpful criticisms and comments. This paper grew out of a correspondence with them, and although I very much doubt that they will be convinced by my arguments, they must certainly take credit for preventing this paper from being much worse than it currently is. A version of this paper was presented at the 2010 symposium of the Canadian Society for Epistemology at the University of Sherbrooke. I am particularly grateful to Conor McHugh, David Hunter, Jonathan Ichikawa, and Berislav Marusic for their questions and comments on that occasion. Thanks also to David Papineau, Roy Sorensen, Ralph Wedgwood, and to this journal's anonymous referees for helpful written comments.

- Griffiths, A. P. 1962–63. On belief. *Proceedings of the Aristotelian Society* 63: 167–86.
- Martin, W. 2009. Ought but cannot. *Proceedings of the Aristotelian Society* 109: 103–28.
- Owens, D. J. 2003. Does belief have an aim? *Philosophical Studies* 115: 283–305.
- Papineau, D. 1999. Normativity and judgment. *Proceedings of the Aristotelian Society, Supplementary Volume* 73: 16–43.
- Pollock, J. L. 1987. Epistemic norms. *Synthese* 71: 61–95.
- Shah, N. 2003. How truth governs belief. *Philosophical Review* 112: 447–82.
- Shah, N. and D. Velleman. 2005. Doxastic deliberation. *Philosophical Review* 114: 497–534.
- Smullyan, R. 1983. *5000 B.C. and other philosophical fantasies*. New York: St. Martin's Press.
- Sorensen, R. 1988. *Blindspots*. Oxford: Clarendon Press.
- Sutton, J. 2007. *Without justification*. Cambridge, MA: MIT Press.
- Vahid, H. 2006. Aiming at truth: Doxastic vs. epistemic goals. *Philosophical Studies* 131: 303–35.
- Wedgwood, R. 2002. The aim of belief. *Philosophical Perspectives* 16: 267–97.
- 2009. The normativity of the intentional. In *The Oxford handbook of philosophy of mind*, 421–36. Oxford: Clarendon Press.
- Whiting, D. 2010. Should I believe the truth? *Dialectica* 64: 213–24
- Williamson, T. 2000. *Knowledge and its limits*. Oxford: Oxford University Press.
- Zalabardo, J. L. 2010. Why believe the truth? Shah and Velleman on the aim of belief. *Philosophical Explorations* 13: 1–21.