To appear in *Thought: A Journal of Philosophy*

Zachary Mitchell Swindlehurst
zacharyswindlehurst@gmail.com

The Knowledge Norm of Belief

1. Doxastic Normativism

A sizeable number of philosophers hold that belief is constitutively normative (see, e.g., Shah 2003; Shah and Velleman 2005; Boghossian 2005; Whiting 2010; Fassio 2011; Gibbard 2012; Engel 2013; Wedgwood 2013; Raleigh 2013; Greenberg 2018). Let *doxastic normativism* ('normativism' hereafter) be the thesis that certain norms are constitutive of or essential to belief, such that no mental state not subject to those norms will count as a belief.

A common proposal is that a belief is *correct* just in case its propositional content is true. That is, normativists accept something like the following principle:

(C) ∀*p*: The belief that *p* is correct if and only if *p* is true.

Denying that 'correct' simply means 'true', the normativist insists that correctness is an essentially *normative* notion. Given this, normativists hold that correct beliefs are the beliefs subjects *ought* to have. Thus many normativists hold that belief is essentially governed by a 'truth norm':

(T) ∀*S*, *p*: *S* ought to believe that *p* if and only if *p* is true.

In recent influential work, Krister Bykvist and Anandi Hattiangadi raise a number of objections to (T) and its variants, primarily concerning the formulation of truth norms (Bykvist and Hattiangadi 2007, 2013). In this paper, I propose that we construe the fundamental norm of belief as a *knowledge* norm, rather than a truth norm. I will argue that a specific version of the knowledge norm – one that has a subject's obligation to believe that *p* depend on her *being in a position to know* that *p* – avoids these well-known difficulties with truth norms.

The goal of this paper is modest. I do not attempt to sketch a positive normativist theory, nor do I attempt to show that the knowledge norm I defend is constitutive of belief. Rather, I only want to show that a knowledge norm formulated with the help of the notion of being in a position to know is an option the normativist should explore further. For although the norm I defend doesn't seem to succumb to the difficulties Bykvist and Hattiangadi raise against normativism, there remain unanswered questions concerning the logic and semantics of 'being in a position to know' that demand careful attention.

2. Truth Norms and their Logical Forms

According to Bykvist and Hattiangadi ('B&H' hereafter), belief is not constitutively normative. They argue that truth norms for belief are either false or impose unsatisfiable requirements on believers. For one thing, (T) is *ambiguous*, since its *ought* operator can take either a narrow scope or a wide scope over the biconditional. The narrow scope reading (T_N) runs as follows:

(T_N) ∀*S*, *p*: *S* ought to (believe that *p*) if and only if *p* is true (B&H 2007, p. 278).

I will only be concerned with the narrow scope reading, which requires the subject to believe that *p* on the condition that *p* is true.[1] (T_N) is the conjunction of two conditionals:

(T_N*) ∀*S*, *p*: If *p* is true, then *S* ought to (believe that *p*).

(T_N**) ∀*S*, *p*: If *S* ought to (believe that *p*), then *p* is true.[2]

Since there are infinitely many true propositions (some of which are ungraspably complex), (T_N*) – according which the truth of *p* is a *sufficient* condition on *S*'s being such that she ought

[1] For discussion of the wide scope reading, see B&H (2007, pp. 283-84).
[2] Alternatively: '*S* ought to (believe that *p*) only if *p* is true'.

to believe that *p* – is clearly false. More precisely, it violates the widely accepted principle according to which 'ought' implies 'can':

('Ought' Implies 'Can') If *S* ought to φ, then it is possible for *S* to φ.[3]

On the other hand, (T_N**) – according to which the truth of *p* is a *necessary* condition on *S*'s being such that she ought to believe that *p* – is too weak. For if *p* is true, it doesn't follow from (T_N**) that *S* ought to believe that *p*. And if *p* is false, all that follows is that *S lacks* an obligation to believe that *p*; it doesn't follow that *S has* an obligation *not* to believe that *p*. Thus when *p* is false, it is permissible according to (T_N**) for *S* to believe that *p*. The norm is too weak.

B&H also consider and reject a proposal from Ralph Wedgwood which seems, initially, to dissolve the foregoing objections to normativism (cf. Wedgwood 2013). Wedgwood suggests we restrict the true propositions one ought to believe to propositions one has *considered*:

(T*) ∀*S*, *p*: If *S* considers whether *p*, then *S* ought to (believe that *p*) if and only if *p* is true.

(T*) would seem to solve the problem, since it doesn't violate ('Ought' Implies 'Can'). But there are some propositions – Moore-paradoxical propositions, or 'blindspots' (Sorensen 1988) – which it is logically impossible to believe truly. For example:

It is raining and nobody believes that it is raining.

There are no believers.

Dogs bark but I don't believe that dogs bark.

[3] ('Ought' Implies 'Can') is not entirely uncontroversial. See, e.g., Mizrahi (2012).

And so on. If you believe that there are no believers, then it is false that there are no believers, in which case you lack an obligation (according to (T*)) to believe that there are no believers. Likewise, if it is true that there are no believers, then you don't believe that there are no believers. B&H's objection to (T*) is that in the case of true blindspot propositions, (T*) violates the following principle:

> ('Ought' Implies 'Can Satisfy') If *S* ought to φ, then it is possible for *S* to φ while still being obligated to φ (B&H 2007, p. 282).

The problem with (T*), then, is not that the propositions in question cannot be *believed*, but rather that the obligation to believe that *p* just in case *p* is true cannot be *satisfied*. Thus (T*) seems hopeless, and so B&H reject the narrow scope reading of the truth norm.

3. Knowledge Norms

I want to suggest that we construe the fundamental norm of belief as a *knowledge* norm, rather than a truth norm. A specific kind of knowledge norm, I will argue, does not succumb to some of these well-known difficulties faced by truth norms.[4]

As Littlejohn (2013, p. 294) notes, many of the difficulties Bykvist and Hattiangadi raise against truth norms don't seem to arise for the 'knowledge first' approach. This is correct, but some of these difficulties do arise for knowledge norms, and knowledge norms face some difficulties of their own.[5]

To begin with, suppose we have the following knowledge norm:

[4] Different knowledge norms are defended by, e.g., Williamson (2000), Smithies (2012a), Gibbons (2013), Engel (2013), and Littlejohn (2013).

[5] See, e.g., Hughes (2017) and Fassio (2018).

(K) ∀*S*, *p*: *S* ought to believe that *p* if and only if *S* knows that *p*.

Like (T), (K) is ambiguous, since its *ought* operator can take either a narrow scope or a wide scope over the biconditional. I will again only be concerned with the narrow scope reading:

(K_N) ∀*S*, *p*: *S* ought to (believe that *p*) if and only if *S* knows that *p*.[6]

(K_N) is the conjunction of two conditionals:

(K_N*) ∀*S*, *p*: If *S* knows that *p*, then *S* ought to (believe that *p*).

(K_N**) ∀*S*, *p*: If *S* ought to (believe that *p*), then *S* knows that *p*.[7]

(K_N*) does not, unlike its alethic counterpart (T_N*), violate ('Ought' Implies 'Can'). For we cannot know infinitely many truths, nor can we know ungraspably complex ones (they are, after all, ungraspable). But (K_N*) does seem too strong: it is not possible to *violate* the norm, since knowledge entails belief. Arguably, any plausible norm implies the possibility of violation: normative *force* goes hand-in-hand with normative *freedom* (Railton 2000). We might capture this thought with the following principle:

('Ought' Implies 'Can Not') If *S* ought to φ, then it is possible for *S* to fail to φ.

Thus (K_N*) looks implausible. What of (K_N**)? A similar problem arises for (K_N**) as for (T_N**). For if *S* knows that *p*, it doesn't follow from (K_N**) that *S* ought to believe that *p*. Moreover, if *S* doesn't know that *p*, it doesn't follow from (K_N**) that *S* ought not to believe that *p*; it

[6] One might wonder why we are focusing on narrow scope knowledge norms, given that some prominent defenders of these norms construe them with a wide scope *ought* (e.g. Williamson 2000). The reason is that it is difficult to see how a wide scope knowledge norm will avoid the difficulties Bykvist and Hattiangadi raise against wide scope truth norms. Their objection, in essence, is that when a truth norm's *ought* operator takes wide scope, we can't *detach* from the norm a subject's obligation to believe that *p*, even when *p* is true (cf. Broome 2000). But wide scope (K) would also seem to have detachment troubles: even when *S* knows that *p*, we couldn't detach from the norm *S*'s obligation to believe that *p*.

[7] Alternatively: '*S* ought to (believe that *p*) only if *S* knows that *p*'.

follows only that *it is not the case* that *S* ought to believe that *p*. Since this is compatible with it being permissible for *S* to believe that *p*, the norm is too weak.

A natural next step for the normativist might be to modify (K) with the proviso that subjects must have considered the relevant propositions:

(K*) ∀*S*, *p*: If *S* considers whether *p*, then *S* ought to (believe that *p*) if and only if *S* knows that *p*.

Unlike (T*), (K*) will not prescribe believing blindspot propositions, and so will not violate ('Ought' Implies 'Can Satisfy'). It will not prescribe believing blindspots because we cannot have knowledge of them: if you believe some blindspot *p*, then *p* is false, and so you don't know that *p* (knowledge entails truth). Moreover, if *p* is true, then you don't believe that *p*, which means you don't know that *p* (knowledge entails belief). (K*), therefore, will never prescribe believing blindspots because its consequent will never be satisfied in the case of blindspot propositions: if you consider some blindspot proposition *p*, then it is not the case that you know that *p*, so you will lack an obligation to believe that *p*.

But this is precisely the difficulty with (K*): when you do not know that *p*, all that follows from (K*) is that you *lack* an obligation to believe that *p*. It does not follow that you *have* an obligation *not* to believe that *p*. Thus by modifying the knowledge norm with the proviso that subjects consider the propositions in question, we simply push the problem raised for (K_N**) back a step.

Is there a way out for the normativist? Hattiangadi (2010, p. 431) considers and rejects the following narrow scope knowledge norm of belief:

(K**) ∀*S*, *p*: If *S* considers whether *p*, then *S* ought to (believe that *p*) if she is in a position to know that *p*; and *S* ought not to (believe that *p*) if she is not in a position to know that *p*.

We are to understand 'being in a position to know', Hattiangadi suggests, in terms of the following counterfactual:

(PK) If *S* were to believe that *p*, then she would know that *p*.[8]

Hattiangadi suggests that *S* is in a position to know that *p* at some possible world W iff (PK) is true at W. And (PK) is true at W iff *S* knows that *p* in all worlds W* close to W in which *S* believes that *p*. Thus Hattiangadi raises the following objection against (K**):

> The semantics of this counterfactual can be given in terms of possible worlds. We do not want to say that you are in a position to know that *p* just in case all possible worlds at which you believe that *p* are worlds at which you know that *p*. For, this condition would rarely be met, and you would rarely be in a position to know that *p*. Instead, it makes more sense to say that you are in a position to know that *p* just in case all of the closest worlds at which you believe that *p* are worlds at which you know that *p*. This means that even if you are in a position to know that *p*, it is still possible that you will believe that *p*, and yet fail to come to know that *p*. In such a situation, [K**] implies that you ought to believe that *p*. However, if when you believe that *p* you realize a world at which you do not know that *p*, then you will be in a state that is not epistemically ideal, although you will have formed your doxastic attitudes in accordance with [K**] (Hattiangadi 2010, p. 431).

But it seems to me that this objection to (K**) is unsuccessful. For (K**) prescribes believing that *p* on the condition that (PK) is true, and prescribes disbelieving *p* (i.e., believing that ~*p*) on

[8] In what sense does a norm formulated with the help of the notion of being in a position to know count as a *knowledge* norm? Given (PK), what (K**) prescribes, in essence, is acquiring knowledge that *p*. For (K**) tells the subject to form the belief that is conducive to her knowing that *p*. So (K**) counts as a knowledge norm insofar as it prescribes forming this belief. There is certainly more to be said here, however. Perhaps we need a deeper discussion of what it is for a subject to be in a position to know. But for our present purposes, I will set aside the question of whether (PK) is the correct interpretation. For discussion, see, e.g., Williamson (2000, p. 95), Rosenkranz (2007, pp. 68-75), and Smithies (2012b, pp. 732-3).

the condition that (PK) is false. But (PK) is false at any world W when *S*, at W, forms the belief that *p* without coming to know that *p*. Hattiangadi claims that *S* can be in a position to know that *p* even if it is possible for *S* to believe that *p* and yet fail to come to know that *p*. If this were true, (K**) would imply in such a situation that *S* ought to believe that *p*, even though she is in a state that is not epistemically ideal. But this situation is impossible: if *S* is in a position to know that *p* at W iff (PK) is true at W, then it is impossible for *S* to believe that *p* and yet fail to come to know that *p* at W. Hattiangadi's claim that (PK) can be true even when *S* believes that *p* but fails to come to know that *p* gets things wrong: (PK) is false in any world where this happens, including the actual world. After all, forming a belief does not amount to switching worlds.[9] So contra Hattiangadi, *S*'s forming a belief that *p* without coming to know that *p* violates (K**) on the suggested interpretation of 'being in a position to know'. And if this is right, then Hattiangadi's objection to (K**) is unsuccessful.

Now, does (K**) succumb to the same difficulties as its alethic counterparts? First, (K**) won't prescribe believing blindspot propositions, and so won't violate ('Ought' Implies 'Can Satisfy'). For (PK) is false with respect to any blindspot proposition *p*: if *S* were to believe that *p*, then *p* would be false, which means she wouldn't know that *p*. Thus we are never in a position to know blindspot propositions, so (K**) won't prescribe believing them.

Second, (K**) doesn't seem to violate ('Ought' Implies 'Can'), at least on the face of it. The reason is that we are not in a position to know ungraspably complex truths, precisely because they are ungraspable. Nevertheless, we are left with unanswered questions. For suppose that *S* is

[9] Thanks to an anonymous referee for pointing this out, and for prompting further clarification in this section.

in a position to know that *p* (having considered this proposition), and that *p* entails some ungraspably complex proposition *q*. Since *q* is a logical consequence of *p*, '*S* is in a position to know that *p*' would seem to entail that '*S* is in a position to know that *q*'. But since *q* is ungraspably complex, *S* is intuitively not in a position to know that *q*. (K**) would then imply that *S* ought to believe that *p* but disbelieve *q*. And since she can't even grasp *q*, it seems implausible to say that *S* ought to disbelieve it.

Although I cannot provide a proper discussion here, I suspect the answer to this sort of worry turns on whether the normativist accepts something like an epistemic closure principle for 'being in a position to know'. I want to suggest that the normativist need not accept that '*S* is in a position to know that *p*' entails '*S* is in a position to know that *q*' if she denies that the notion of being in a position to know is closed under logical consequence.

There is in fact good reason to deny this. Let *K* express '*S* is in a position to know that', and assume that *S* knows *modus ponens* and how to apply it. The relevant closure principle might read as follows:

(CL) $K\varphi, K(\varphi \rightarrow \psi) \vdash K\psi$

As Heylen (2016) points out, (CL) fails. To see this, consider the K-axiom for *K*, from which (CL) can be derived:

(K_K) $K(\varphi \rightarrow \psi) \rightarrow (K\varphi \rightarrow K\psi)$

(K_K), too, fails (Rosenkranz 2016). Assume that $K(\varphi \rightarrow (\psi \rightarrow (\varphi \,\&\, \psi)))$. ($K_K$) then gives us $K\varphi \rightarrow K(\psi \rightarrow (\varphi \,\&\, \psi))$. From this and ($K_K$), we get $K\varphi \rightarrow (K\psi \rightarrow K(\varphi \,\&\, \psi))$. So we can derive:

(1) $K\varphi \,\&\, K\psi \rightarrow K(\varphi \,\&\, \psi)$

According to (1), *K* agglomerates over conjunction. Now let κ express '*S* knows that', and note that κφ implies 'Someone sometimes knows that φ'. κ validates the following:

(2) κφ → φ

(3) κ (φ & ψ) → κφ

Now suppose that ψ is the proposition 'No one ever knows that φ'. It follows that ~◊κ (φ & ψ) (cf. Fitch 1963). It is plausible that if *K*φ holds, then ◊κφ also holds. Thus we get ~*K* (φ & ψ), and (1) then implies ~(*K*φ & *K*ψ). Rosenkranz writes:

> And yet, where ψ is 'No one ever knows φ', there are φ for which this conclusion would seem clearly unacceptable. Thus let φ be a fleeting truth of little interest, for example, 'There are at present exactly seven blossoms on the bougainvillea', in a context where the plant is in full view, a storm is about to hit, one is the only one around, the latter facts are known, and one also knows by introspection – and *a fortiori* is in a position to know – that one is far too unconcerned ever to find out about the matter. In such a case, both *K*φ and *K*ψ should hold, contrary to what (1) predicts – at least there would seem to be nothing inherent in the notion of being in a position to know that would preclude this. Therefore, (1) fails, and hence so does K_K (Rosenkranz 2016, p. 68).

If the foregoing is correct, then the normativist can reject (CL): she can deny that we are in a position to know the logical consequences of what we are in a position to know. However, there are costs involved in rejecting (CL). For what kind of semantics and logic for 'being in a position to know' would result if '*S* is in a position to know that *p*' doesn't entail that '*S* is in a position to know that *q*'?

According to Rosenkranz (2016, p. 69; 2018, pp. 316-19), the failure of (K_K) implies that the logic for *K* cannot be a normal modal logic. For standard epistemic logics are normal modal logics that include (K_K). They also include:

(RN_K) If $\vdash \varphi$, then $\vdash K\varphi$

(RM_K) If $\vdash \varphi \rightarrow \psi$, then $\vdash K\varphi \rightarrow K\psi$

In standard epistemic logics, (RM_K) is a derived rule. But given that (K_K) fails, the logic for *K* must be non-normal. If we had both (K_K) and (RN_K), we could derive (1). But we have seen that (1) fails. The logic for 'being in a position to know' cannot then be a normal modal logic. This suggests that we cannot give a standard Kripkean semantics for *K*: according to Rosenkranz (2018, p. 318), something like a neighbourhood semantics might be more appropriate. Perhaps this is right, but I cannot flesh out such a semantics here.

So, there are costs involved in rejecting (CL). The normativist who accepts (K**) must be willing to accept a non-normal modal logic – along with a corresponding semantics – for 'being in a position to know'. What exactly this amounts to is a difficult and substantive question, and I cannot attempt to answer it here. But it at least seems that if the foregoing is correct, a norm like (K**) – one formulated with the help of the notion of being in a position to know – is an option the normativist should explore further: it might avoid the well-known formulation problems with its alethic counterparts.

Acknowledgements

I would like to thank Alex Miller for helpful commentary on previous drafts of this paper. I am also grateful to an anonymous referee for *Thought* for several useful suggestions. This publication was made possible by a University of Otago postgraduate publishing bursary.

References

Boghossian, Paul. 2005 [2008]. "Is Meaning Normative?" In *Content and Justification: Philosophical Papers*, 95–107. Oxford: Clarendon Press.

Broome, John. 2000. “Normative Requirements.” In *Normativity,* edited by Johnathan Dancy, 78–99. Malden: Blackwell.
Bykvist, Krister, and Anandi Hattiangadi. 2007. “Does Thought Imply Ought?” *Analysis* 67 (4): 277–85.
———. 2013. “Belief, Truth and Blindspots.” In *The Aim of Belief*, edited by Timothy Chan, 100–122. Oxford: Oxford University Press.
Engel, Pascal. 2013. “In Defence of Normativism about the Aim of Belief.” In *The Aim of Belief,* edited by Timothy Chan, 31–62. Oxford: Oxford University Press.
Fassio, Davide. 2011. “Belief, Correctness and Normativity.” *Logique et Analyse* 54 (216): 471–86.
———. 2018. “On the Generality Argument for the Knowledge Norm.” *Synthese*. https://doi.org/10.1007/s11229-018-01896-x.
Fitch, Frederic B. 1963. “A Logical Analysis of Some Value Concepts.” *Journal of Symbolic Logic* 28 (1): 135–42.
Gibbard, Allan. 2012. *Meaning and Normativity*. Oxford: Oxford University Press.
Gibbons, John. 2013. *The Norm of Belief*. Oxford: Oxford University Press.
Greenberg, Alexander. 2019. “Should I Believe All the Truths?” *Synthese.* https://doi.org/10.1007/s11229-018-1882-9.
Hattiangadi, Anandi. 2010. “The Love of Truth.” *Studies in History and Philosophy of Science* 41 (1): 422–32.
Heylen, Jan. 2016. “Being in a Position to Know and Closure.” *Thought: A Journal of Philosophy* 5 (1): 63–67.
Hughes, Nick. 2017. “No Excuses: Against the Knowledge Norm of Belief.” *Thought: A Journal of Philosophy* 6 (3): 157–66.
Littlejohn, Clayton. 2013. “The Russellian Retreat.” *Proceedings of the Aristotelian Society, Supplementary Volume* 113 (3): 293–320.
Mizrahi, Moti. 2012. “Does ‘Ought’ Imply ‘Can’ from an Epistemic Point of View?” *Philosophia* 40 (4): 829–40.
Railton, Peter. 2000. “Normative Force and Normative Freedom: Hume and Kant.” In *Normativity,* edited by Johnathan Dancy, 1–33. Malden: Blackwell.
Raleigh, Thomas. 2013. “Belief Norms and Blindspots.” *The Southern Journal of Philosophy* 51 (2): 243–69.
Rosenkranz, Sven. 2007. “Agnosticism as a Third Stance.” *Mind* 116 (461): 55–104.
———. 2016. “Being in a Position to Know and Closure: Reply to Heylen.” *Thought: A Journal of Philosophy* 5 (1): 68–72.
———. 2018. “The Structure of Justification.” *Mind* 127 (506): 309–38.
Shah, Nishi. 2003. “How Truth Governs Belief.” *The Philosophical Review* 112 (4): 447–82.
Shah, Nishi, and J. David Velleman. 2005. “Doxastic Deliberation.” *The Philosophical Review* 114 (4): 497–534.
Smithies, Declan. 2012a. “The Normative Role of Knowledge.” *Noûs* 46 (2): 265–88.
———. 2012b. “Mentalism and Epistemic Transparency.” *Australasian Journal of Philosophy* 90 (4): 723–41.
Sorensen, Roy. 1988. *Blindspots*. Oxford: Clarendon Press.
Wedgwood, Ralph. 2013. “The Right Thing to Believe.” In *The Aim of Belief,* edited by Timothy Chan, 123–39. Oxford: Oxford University Press.
Whiting, Daniel. 2010. “Should I Believe the Truth?” *Dialectica* 64 (2): 213–24.

Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.