

An Evaluative Norm for Belief

MICHAEL-JOHN TURP *University of Canterbury*

ABSTRACT: It is often argued that belief is partly constituted by a norm of truth. Most recent discussions have assumed that the norm is deontic concerning what may or ought to be believed. I criticize two proposals, one canvassed by Krister Bykvist and Anandi Hattiangadi, and the other defended by Daniel Whiting. Instead, I argue in favour of an evaluative norm, according to which we would do well to believe the truth. I show that an evaluative norm fares better than its deontic competitors with respect to the demandingness of truth, the aim of truth, and epistemic blame.

RÉSUMÉ : On soutient souvent que la croyance est en partie constituée par une norme de vérité. La plupart des discussions récentes ont supposé que la norme est déontique concernant ce qui peut ou doit être cru. Je critique deux propositions, l'une mise de l'avant par Krister Bykvist et Anandi Hattiangadi, et l'autre défendue par Daniel Whiting. Je plaide plutôt en faveur d'une norme évaluative selon laquelle nous ferions bien de croire la vérité. Je montre qu'une norme évaluative réussit mieux que ses concurrentes déontiques en ce qui a trait au caractère exigeant de la vérité, au but de la vérité et au blâme épistémique.

Keywords: truth, epistemic value, epistemic normativity, aim of belief, nature of belief

Introduction

It is often argued that belief is essentially normative because partly constituted by a norm of truth. According to this view, sometimes referred to as 'doxastic normativism,' part of what it is to be an epistemic agent is to be subject to this norm when we form beliefs. This might be taken to be a conceptual necessity, a metaphysical necessity, or both. Most recent discussions have assumed that the norm is deontic, concerning what may or ought to be believed. I shall criticize

Dialogue 56 (2017), 227–239.

© Canadian Philosophical Association/Association canadienne de philosophie 2016

doi:10.1017/S0012217316000706



two such proposals, one canvassed by Krister Bykvist and Anandi Hattiangadi, and the other defended by Daniel Whiting. Instead, I shall argue in favour of an evaluative norm, according to which we would do well to believe the truth. I show that an evaluative norm fares better than its deontic competitors with respect to the demandingness of truth, the aim of truth, and epistemic blame.

1. Beliefs and Oughts

Bykvist and Hattiangadi criticize the following principle:

- (1) For any S, p : S ought to believe that p iff p is true.¹

Depending on the scope of ‘ought,’ (1) is ambiguous between:

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

(2) is only satisfied by S believing that p when p is true. (3) is also satisfied when p is false and S does not believe that p (see further §3). However, most of Bykvist and Hattiangadi’s discussion focuses on (2), i.e., the narrow scope interpretation. Subsequent commentators, including Whiting, also focus on (2). I shall do likewise initially. As Bykvist and Hattiangadi argue, (2) is too demanding.² For example, a proposition containing an infinite number of true conjuncts is true. However, we cannot entertain such a proposition and, therefore, cannot believe it. There are also ‘blindspots,’ which are propositions that might be true, but cannot be *believed* truly. For example, propositions of the form $\langle p$, but nobody believes that $p \rangle$ cannot be believed truly when p is true. According to (2), however, we ought to believe these true propositions. This violates the principle that ‘ought’ implies ‘can’ (OIC). Assuming that OIC is true, we should reject (2).

2. Belief, Truth, and Permissibility

In light of Bykvist and Hattiangadi’s arguments, Whiting proposes an alternative norm of truth stated in terms of what *may* be believed, rather than what *ought* to be believed. He endorses:

- (4) For any S, p : S may (believe that p) iff it is true that p .³

(4) is normative at least insofar as it rules out S believing that p when p is false. If S believes that p when p is false, then S has violated (4). However, (4) does not

¹ Bykvist and Hattiangadi (2007, 277).

² Bykvist and Hattiangadi (2007, 279).

³ Whiting (2010, 216).

provide further positive guidance with respect to belief formation. One advantage of (4), according to Whiting, is that, since ‘may’ does not logically imply ‘can,’ it does not impose impossible demands on us in the same way as (2). Thus, (4) is attractive partly because it is weaker than (2). Although we may not believe false propositions, we may either believe or not believe true propositions. To the extent that both belief and withholding belief are permissible, they are normatively on a par and the choice between them should be a matter of normative indifference. This seems wrong, however, because evidence in favour of the truth of a proposition is a *pro tanto* reason to believe it. To withhold belief on a true proposition that we entertain and for which we have good evidence is irrational in the sense that it is a failure to recognize and respond to the reason for belief given by the evidence. Because (4) does not capture this idea, it is insufficiently belief-guiding. Notice that the concern is not that (4) fails to *determine* what it is rational to believe. Indeed, the account that I defend below depends on the separateness of norms of truth and rationality. I shall argue that it is sometimes irrational to believe the truth. Rather, the concern is that (4) fails to weigh into the balance the fact that evidence in favour of a proposition being true is typically a defeasible reason in favour of believing that proposition. This does not show that (4) is false. It still could be one of the norms governing belief. The argument does, however, motivate the search for a stronger principle.

3. Aiming at True Belief

We can connect the limitations of (4) to a further concern about the wide scope interpretation of (1).⁴ Recall:

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

According to (3), S ought to ensure that the biconditional within the brackets is true for any p . Let ‘ B_{Sp} ’ read ‘ S believes that p .’ (3) is satisfied by either:

- (i) $B_{Sp} \ \& \ p$.
- (ii) $\neg B_{Sp} \ \& \ \neg p$.

Consider the positive requirement given by (i), i.e., that S ought to satisfy the *combination* of believing that p and p being true. This does not entail that S ought to believe that p if p is true. If p is true and S believes that $\neg p$, S can satisfy (3) by bringing it about that $\neg p$, i.e., by intervening in the world. Or, more precisely, S can then ameliorate her epistemic position by forming the new belief that $\neg p_1$, where p_1 has the same content as p , except that it is indexed to the time at which $\neg p_1$ is made true by S intervening in the world. The details depend

⁴ See Bykvist and Hattiangadi (2007, 284).

on the nature of propositions and the correct account of belief revision, but the significant point is that satisfying (3) gives *S* a reason to intervene in the world. However, the implication that we can do as we epistemically ought to do by bringing the world into line with our beliefs is implausible. I might believe that I shall raise my arm within the next two minutes, and ensure that the belief is true by actually raising my arm. But it seems that the fact that I believe I shall raise my arm gives me no *normative* reason to actually raise my arm. By contrast, the fact that I shall raise my arm, or intend to raise my arm, does give me a normative reason to believe that I shall do so. There is a normative reason to form a belief that represents the world as it is (or will be), but no normative reason to make the world conform to our beliefs. Thus, in terms of a common metaphor, (4), like (3), fails to capture the idea that belief ‘aims’ at truth.

It might be objected at this point that there is an important class of beliefs that *do* require us to intervene in the world if we are to be rational, namely moral or evaluative beliefs.⁵ If I am rational, then my beliefs concerning, say, the wrongness of causing animals to suffer unnecessarily will lead me to change my behaviour in order to reduce or eliminate that suffering. There are various possible replies here. One option would be to restrict the argument to non-evaluative beliefs. Another option would be to deny that evaluative beliefs are truth-apt and defend some form of expressivism. I think the best approach, however, is to distinguish between practical and theoretical rationality. Although failing to try to bring the world into line with evaluative beliefs may display a failing of practical reason, and perhaps even a moral failing, it is unclear that it displays a failing of theoretical reason. I might be content, for example, with determining the moral facts of the matter without having any interest in acting on those facts. Or, I might be suffering from *akrasia*. An *akratic* agent need not be a deficient epistemic agent.

4. Deontic Norms

(1) – (4) have deontic content in the sense that they concern what we may believe or ought to believe. In the present context, I shall understand the verb ‘ought’ as corresponding to the noun ‘obligation’ and as standing in certain logical relations to the ‘may’ of permissibility as described in §9 below. Although by no means idiosyncratic, there are common and legitimate uses of ‘ought’ that are inconsistent with this understanding. For instance, we sometimes use ‘ought’ as a term of recommendation, saying, perhaps, that ‘you ought to buy new shoes’ or ‘I ought to go on a diet.’ In these cases, there is no corresponding obligation. One possibility is that these are applications of deontic language into properly evaluative contexts for the sake of commendatory force. To say ‘you ought to buy new shoes’ often seems equivalent to ‘you would do well to buy new shoes, and I recommend that you do so.’ However this may be,

⁵ I thank an anonymous referee for encouraging me to address this point.

the focus of my discussion is the circumscribed, philosophical understanding of ‘ought’ sketched above and developed below. I shall argue, then, that the primary normative relation holding between belief and truth is not deontic in this sense, but evaluative. I say ‘primary’ because I shall not argue that there are no deontic norms governing belief. Perhaps, for instance, there are restricted deontic norms governing belief, such as that we ought to form beliefs consistently or in accordance with *modus ponens*. I leave this open in the present discussion.

5. An Evaluative Norm

Consider as a first pass at an evaluative norm of truth for belief:

- (5) For any S, p : S would do well to (believe that p) iff p is true.

‘ S would do well’ is ambiguous dependent on the bearer of the value (the belief? the believer? the overall state of affairs?) as well as, perhaps, the perspective from which the belief is evaluated. The interpretation I have in mind is as follows: if S forms a true belief, then the state of affairs in which S ’s belief is true has non-instrumental value. The end of holding true beliefs is partly constitutive of belief formation, and the state of affairs in which this end is satisfied with respect to some proposition is epistemically good *simpliciter*. In this way (5) is intended to be consistent with (i) the idea that true beliefs are intrinsically valuable, and (ii) teleological accounts of epistemic normativity defended by Michael Lynch, William Alston, Asbjørn Steglich-Petersen and others. The important point for present purposes, however, is that the normative content of (5) is evaluative rather than deontic.

As an evaluative norm, (5) has two immediate advantages over the deontic principles considered above. First, like (4) and unlike (2), (5) does not impose impossible epistemic demands on us. The evaluative term ‘would do well’ does not imply ‘can.’ The fact that S would do well to prove Goldbach’s conjecture does not imply that S can do so. By contrast, from OIC, the putative fact that S ought to prove Goldbach’s conjecture would imply that S can do so. Second, like (2) and unlike (4), (5) accounts for the thought that there is a positive reason that favours S forming the belief that p when p is true. In general, the fact that S would do well to ϕ is a reason that favours S ϕ -ing.

6. Epistemic Value and Trivia

It is at best contingent that to do well epistemically is to do well in an all-things-considered sense. For instance, there may be decisive moral or prudential reasons not to believe what is true. One simple way to handle this is to introduce a from-an-epistemic-standpoint operator (EP). For example, (2) becomes:

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

Of course, (6) is no more plausible than (2). An EP operator focuses our attention on the possibility that different types of norm or value can bear on the same instance of belief formation, but it does not resolve the problems identified in §1. Moreover, it looks false that we ought-EP to believe certain types of true propositions. For instance, it does not seem that we ought to believe trivial true propositions concerning numbers of blades of grass, contents of telephone directories, and the like (henceforth, trivia).

Trivia provide interesting test cases for putative norms of truth for belief. Consider in this light (5), which becomes:

(7) For any S , p : S would do well-EP to (believe that p) iff p is true.

The fact that S would do well-EP to believe that p is merely a *pro tanto* reason for S to believe that p . So, S is not obliged to believe trivia, which seems correct. Nevertheless, according to (7), S would do well-EP to believe trivia. Indeed, according to (7), S would *always* do well-EP to believe trivia. This may be a surprising implication and a possible target for criticism. I think, however, that it is defensible. For, an advantage of an evaluative account is that goodness comes in degrees. The fact that an action or a belief realizes some value does not mean that it realizes more value than its alternatives or that it is to be pursued. By contrast, actions and beliefs are permissible or impermissible, obligatory or not. This difference is indicated in English by the fact that ‘good’ has comparative and superlative forms, whereas the deontic adjectives ‘right’ and ‘permissible’ do not. So, although S would do well-EP to believe trivia, we need not conclude that S would do better-EP or best-EP to believe trivia, or that believing trivia is the course of action to be pursued given finite cognitive resources. The good realized by believing trivia will normally be outweighed by competing sources of value. One reason is that trivia tend to be peripheral to our belief systems and, therefore, unlikely to be conducive towards further true beliefs. Moreover, trivia are unlikely to promote other epistemic goods, such as wisdom and understanding. And, of course, the value of believing trivia will normally be outweighed by non-epistemic values from an all-things-considered perspective.

7. Epistemic Value and Wellbeing

The fact that the value of having true beliefs can be outweighed by other non-epistemic values might lead us to consider another class of putative counterexamples. In particular, it might be argued that we sometimes do well to hold false beliefs and do badly when we hold the corresponding true belief. For example, it may be that we do well to hold a false belief with respect to the departure time of a train that we intend to catch, but which will, unbeknownst to us, crash. Given that (7) implies that each true belief we hold is valuable, it would not help to point out that we would do better still to have the further true belief that the train will crash. This shows once

more that it is contingent whether to do well epistemically is to do well in an all-things-considered sense.

A large psychological literature studying the relationship between true beliefs and wellbeing has grown over the past 30 or so years.⁶ Given this literature, it would be a mistake to assume that the former are always conducive to the latter. For example, there is evidence that certain positive illusions about the self are associated with higher levels of social, psychological, and physical wellbeing.⁷ Such studies point to the possibility that some true beliefs are detrimental to our overall wellbeing. This possibility, which cannot be ruled out from the armchair, might seem to be in tension with (7).

Based on data such as these, it might be argued that the correctness of (7) is simply irrelevant to the pressing, practical question of how we should lead our lives. If (7) is correct, but believing the truth is not conducive to wellbeing, why care about (7)? What we need is information about the likely harms and benefits of holding true beliefs. This information will come from empirical psychology, not conceptual analysis.

I would like to respond in two ways. First, we should notice that empirical studies of wellbeing presuppose a normative account of what wellbeing is. We cannot say that something promotes wellbeing unless we have such an account. Perhaps wellbeing is being close to God, or *ataraxia*, or a life in accordance with reason, or the exercise of our basic capabilities. Or, perhaps, wellbeing is a preponderance of feelings of contentment that can be assessed using Likert-type scales. Whichever answer we think best, it is a substantive, normative thesis and not a finding of empirical psychology. Similarly, (7) is a normative thesis and, in this sense, prior to empirical findings concerning the relationship between holding true beliefs and overall wellbeing.

Second, (7) is strictly and only an epistemic norm, rather than a general account of wellbeing. Of course, many philosophers have assumed that forming true beliefs promotes, or is partly constitutive of, wellbeing. In antiquity, Plato, Aristotle, the Stoics, and the Epicureans all defended this view. By contrast, it was rejected by the Sceptics and the Cynics. Such debates continue and they can be helpfully informed by empirical findings.⁸

For all I say in this paper, it is possible that doing well *qua* epistemic agent is in no way conducive, or even antithetical, to wellbeing. This could be important. For example, as mentioned above, it might affect how much we *care* about (7). If part of what it is to be an epistemic agent is to be governed by a norm of true belief, and if true beliefs do not contribute to our overall wellbeing, then this might encourage the anti-constitutivist cry of ‘agency, schmagency.’⁹

⁶ I thank an anonymous referee for encouraging me to discuss this important point.

⁷ Taylor and Armor (1996).

⁸ See, for example, Stich (1990).

⁹ See Enoch (2006).

While this is a large topic, notice finally that the question is parallel to the question of whether acting *morally* well promotes wellbeing.¹⁰ Even if we answer ‘no’ or ‘just contingently,’ perhaps partly on empirical grounds, it would be too fast to conclude that moral principles are not genuinely normative.

8. Useful Falsehoods

Another interesting possible counterexample concerns systematically false domains of discourse that are nevertheless responsive to how things are and that help to guide our actions successfully. For example, if we mistakenly accepted the caloric theory, we would hold systematically false beliefs regarding heat. Still, we would truly believe propositions such as ‘the kettle is hot.’ In this case, we would do so on the basis of the false belief that ‘the kettle is at high caloric pressure.’¹¹ Moreover, the false belief ‘the kettle is at high caloric pressure’ would be appropriately action-guiding due its connection with the true belief ‘the kettle is hot.’ It seems, then, that we would do well to hold the false beliefs at least to the extent that doing so would enable us to satisfy certain desires (such as not burning ourselves on the kettle).

This is just one example. There are numerous false theories that are instrumentally valuable, enable us to form other true beliefs, and are constrained by mind-independent reality. Let us call these ‘useful falsehoods.’ Do we do well to believe useful falsehoods? From a pragmatic perspective, the answer is sometimes ‘yes.’ It is better to believe useful falsehoods than to believe useless falsehoods, less useful falsehoods, or to have no beliefs at all concerning matters on which we must act. Nevertheless, the claim I am defending is that useful falsehoods are less valuable than useful truths just to the extent that they are false. If all else is equal, we would do better-EP to believe useful truths than useful falsehoods. This is neither to claim that useful truths are always cognitively available nor that we should never believe useful falsehoods. It is to claim that if we compare a false theory with a true theory, then, *ceteris paribus*, we have reason to favour the true theory just because it is true. It is not a matter of normative indifference and we would go wrong from an epistemic perspective if we simply tossed a coin in order to decide which theory to endorse.

9. A Modified Permissibility Norm

These important considerations and qualifications aside, I would now like to compare (7) with a modified version of Whiting’s permissibility principle. Consider:

(8) For any S , p : S may-EP (believe that p) iff p is true.

¹⁰ See papers in Bloomfield (2008) for further discussion.

¹¹ I thank one of the anonymous referees for this example and for encouraging me to clarify the point at issue.

(8) improves on (6) because it permits *S* to satisfy competing demands at the expense of not believing that *p* when *p* is true. It also permits *S* not to believe trivia, which may appear to be an advantage over (7). Having already defended (7) on this point, however, I want to develop a different sort of concern about (8).

Note that (8) entails the following conditional:

(9) For any *S, p*: if *S* may-EP (believe that *p*), then *p* is true.

From which, by contraposition:

(10) For any *S, p*: if it is not the case that *p* is true, then it is not the case that *S* may-EP (believe that *p*).

Furthermore, in accordance with standard deontic logics, we can interdefine obligation and permissibility as follows:

(OP) It is not the case that *S* may ϕ iff *S* ought not to ϕ .

From (10) and (OP) we have:

(11) For any *S, p*: if it is not the case that *p* is true, then *S* ought-EP not (to believe that *p*).

Thus, *S* ought not to believe false propositions. In fact, Whiting mentions this implication in defence of (4) and he would presumably see it as counting in favour of (8) too.¹² He argues that, although (4) may not capture the intuition that belief aims at truth, it does at least capture the related intuition that belief aims at avoiding falsehoods.

10. Norms of Truth and Rationality

One difficulty with (11) is that it introduces a problematic tension between epistemic obligation and rationality. This is because it is sometimes rational to believe falsehoods. Consider the following proposition:

<N> The Loch Ness Monster does not exist.

It is rational by most lights to believe that <N>. The best available evidence supports it. Yet, it is possible that <N> is false. Let us assume that this is so, i.e., that Nessie is lurking in the depths. If so, according to (11), *S* ought not to believe that <N>. So, *S* ought not to form a rational belief. Although it is commonplace that it is sometimes irrational to believe the truth, I shall argue that this is an unattractive result in the present context.

¹² Whiting (2010, 217).

Before considering a standard response to this type of tension, by distinguishing between subjective and objective oughts, it will be helpful to consider whether a similar concern might be raised with respect to (7). After all, according to (7), *S* does not do well to form the rational belief that $\langle N \rangle$. This sort of tension between rationality and goodness is also superficially unattractive. I shall argue, however, that this tension is consistent with ascriptions of blame, whereas the tension between rationality and obligation seems not to be.

In general, we are not blameworthy when we form rational beliefs that are false. We may be non-culpable victims of epistemic bad luck, as *S* appears to be with respect to $\langle N \rangle$. This is consistent with (7). The fact that *S* would do well to ϕ is not sufficient to make it the case that *S* is blameworthy for not ϕ -ing. For example, if *S* would do well to prove Goldbach's conjecture, it does not follow that *S* is blameworthy for not doing so. In the present case, if *S* would do well to believe that not- $\langle N \rangle$, it does not follow that *S* is blameworthy for forming the rational belief that $\langle N \rangle$.

Indeed, it is more plausible that we are appropriately subject to epistemic blame not when we fail to believe the truth, but when we fail to form a belief warranted by the reasons available to us, i.e., when we transgress norms of rationality rather than a norm of truth. If so, *S* would both do well *and* be blameworthy for believing that not- $\langle N \rangle$. While this may seem surprising, the situation is quite common. It is the situation of a bumbling but lucky Inspector Clouseau type of character, whose blameworthiness is not mitigated by the fact that he succeeds despite himself. While Clouseau deserves no praise, he acts well with respect to the aim of catching diamond thieves. Moreover, for Clouseau, it is the aim of catching diamond thieves that determines the rational response to the reasons available to him. Nevertheless, rationality and goodness are distinct aims. Sometimes, the means-end relation between the two can un/luckily fail. In such cases, it seems to be the rationality of the action or belief that is the appropriate criterion for praise or blame.

11. Oughts and Blameworthiness

According to (11), *S* ought to not believe that $\langle N \rangle$. However, it seems plausible that *S* ought to ϕ only if the following two conditions are met:

- (OIC) *S* can ϕ .
- (BLAME) *S* would be blameworthy for not ϕ -ing.

This is not an analysis of 'ought,' but a set of two necessary conditions. For example, *S* ought to prove Goldbach's conjecture only if *S* can do so and would be blameworthy for failing to do so. Again, *S* ought to tell the truth only if *S* can do so and would be blameworthy for failing to do so.

So, according to (11), *S* ought to not believe that $\langle N \rangle$. Therefore, according to (BLAME), *S* would be blameworthy for forming the rational belief that $\langle N \rangle$. This seems to be the wrong result insofar as it prescribes one to be irrational.

Indeed, I have just suggested that the reverse is true, namely that *S* would be blameworthy for forming the true, but irrational, believe that not- $\langle N \rangle$. The fact that this account of blame is consistent with (7), but inconsistent with (11), favours (7). The problem with the tension between rationality and a deontic norm for belief is that culpability falls in the wrong place, i.e., satisfying the norm of truth, rather than the norm of rationality.

Perhaps, though, (BLAME) is open to challenge. In particular, it might be objected that (BLAME) fails to account for the phenomenon of *excuses*. Say that Sam ought to attend a meeting, but uncharacteristically misses it due to the pressures of parenthood. In this case, the excuse may make it the case that Sam does not deserve to be blamed. We might be inclined, then, to say that Sam is not really blameworthy. However, the fact that we should not *blame* Sam does not entail that Sam is not *blameworthy*. For, *blameworthiness* is necessary, but not sufficient, for acts of *blaming* to be justified. Put another way, acts of blaming are forms of sanction, but blameworthiness may not be sufficient for sanction to be justified. So, while mitigating excuses sometimes undermine the justification for sanctioning *S* by blaming *S*, it would be too hasty to conclude that mitigating excuses make it the case that *S* is not blameworthy.

Sometimes, however, excuses not only mitigate blame, but also *eliminate* blameworthiness. In these cases, I suggest that the excuse undermines a necessary condition for there to be an *actual* ought. Perhaps Sam missed the meeting due to a flat tire. After the puncture occurred, Sam *could* not have attended the meeting. Consequently, OIC was not met and it was no longer the case that Sam ought to attend the meeting. In this case, blameworthiness is eliminated because the ‘ought’ is eliminated. The ‘ought’ turns out to be *prima facie*, rather than actual. In short, if the excuse eliminates blameworthiness, it does so by eliminating an actual ought. If, however, the excuse merely mitigates, then it may eliminate the justification for blaming, but it does not thereby eliminate blameworthiness.

12. Subjective and Objective Oughts

Perhaps, though, it will be objected that I have proceeded without taking into account the distinction between subjective and objective oughts. Perhaps, *subjectively*, *S* ought to form the rational belief that $\langle N \rangle$. *Objectively*, however, *S* ought to believe that not- $\langle N \rangle$, given that $\langle N \rangle$ is, in fact, false. Although the distinction is a helpful one, I do not think that it helps a defender of a deontic norm in the present context.

Recall the principle currently under discussion:

- (11) For any *S*, *p*: if *p* is not true, then *S* ought-EP not (to believe that *p*).

I have, in effect, been assuming an objective reading of the ‘ought’ in (11). This is natural, given that the antecedent refers to the truth-value of *p*, not the evidence

that *S* has for the truth-value of *p*. Consider, however, possible subjective and objective variants of (11):

- (12) For any *S*, *p*: if *p* is not true, then *S* objectively ought-EP not (to believe that *p*).
- (13) For any *S*, *p*: if it is rational for *S* to believe that *p* is not true, then *S* subjectively ought-EP not (to believe that *p*).

Now, in order to satisfy (12), *S* objectively ought-EP to believe that not- $\langle N \rangle$. However, in order to satisfy (13) *S* subjectively ought-EP to believe that $\langle N \rangle$. Given that *S* cannot satisfy both (12) and (13) with respect to $\langle N \rangle$, we might ask what *S* ought to believe both-oughts-considered? Put another way, which ought has practical, normative authority? It seems that the obvious candidate is the subjective ought (13). For, if (12) has practical authority and *S* ought, both-oughts-considered, to believe that not- $\langle N \rangle$, despite all of the available evidence and despite what it is rational for *S* to believe, then, it seems that *S* is blameworthy for forming a rational belief. I have argued that this is an implausible outcome.

The reason, then, for distinguishing between subjective and objective oughts in this context would most likely be the thought that (13), or a similar principle, has practical authority over *S*'s belief formation. If so, I am happy to agree. But, as I observed in §10, belief is governed by norms of rationality as well as a norm of truth. Indeed, I relied on separate norms of rationality and truth in order to explain why we should withhold praise from the irrational Inspector Clouseau (who, nonetheless, does *well*). If (13) is a plausible candidate for a deontic norm of *rationality* that governs belief, then the central thesis of this paper, namely that the norm of *truth* that governs belief is evaluative, remains unaffected.

13. Conclusion

When it is argued that belief is partly constituted by a norm of truth, it is not only meant that we ought to believe what we have most reason to believe is true. The basic intuition underpinning (1) and its various refinements is that the way the world is, independently of how we think about it, places normative constraints on our processes of belief formation, and that this fact is part of what it is to form beliefs at all (as opposed, perhaps, to having mere representations or being programmed with beliefs).

I have argued that the relevant norm of truth is best understood in terms of what we would do well to believe. This preserves the non-contingent, normative connection between belief formation and truth. Our beliefs are necessarily evaluable against the standard of how things are. Moreover, unlike its deontic alternatives, an evaluative norm of truth neither places impossible demands upon us nor entails that we are blameworthy for failing to believe what is true when the reasons available to us suggest what is false.

Acknowledgements: Many thanks to Doug Campbell, Carolyn Mason, Alex Miller, and an audience at the University of Otago for helpful comments and discussion.

References

- Alston, William
2005 *Beyond Justification: Dimensions of Epistemic Evaluation*, Ithaca: Cornell University Press.
- Bloomfield, Paul (ed.)
2008 *Morality and Self-Interest*, Oxford: Oxford University Press.
- Bykvist, Krister, and Anandi Hattiangadi
2007 “Does Thought Imply Ought?,” *Analysis* 67 (4): 277–285.
- Enoch, David
2006 “Agency, Shmagency: Why Normativity Won’t Come from What is Constitutive of Action,” *The Philosophical Review* 115 (2): 169–198.
- Lynch, Michael
2004 *True to Life: Why Truth Matters*, Cambridge, Mass.: MIT Press.
- Steglich-Petersen, Asbjørn
2011 “How to Be a Teleologist about Epistemic Reasons,” in *Reasons for Belief*. Asbjørn Steglich-Petersen and Andrew Reisner, eds., Cambridge: Cambridge University Press, 13–33.
- Stich, Stephen
1990 *The Fragmentation of Reason*, Cambridge, Mass.: MIT Press.
- Taylor, Shelley, and David Armor
1996 “Positive Illusions and Coping with Adversity,” *Journal of Personality* 64 (4): 873–898.
- Whiting, Daniel
2010 “Should I Believe the Truth?,” *Dialectica* 64 (2): 213–224.