

Is the Norm on Belief Evaluative? A Response to McHugh

Alexander Greenberg and Christopher Cowie

Published in *Pacific Philosophical Quarterly* 98 (S1): 128-145 (2017). Please cite the final version, which is available at <http://doi.org/10.1111/papq.12158>.

Abstract: We respond to Conor McHugh's claim that an evaluative account of the normative relation between belief and truth is preferable to a prescriptive account. We claim that his arguments fail to establish this. We then draw a more general sceptical conclusion: we take our arguments to put pressure on *any* attempt to show that an evaluative account will fare better than a prescriptive account. We briefly express scepticism about whether McHugh's more recent 'fitting attitude' account fares better.

It is often claimed that belief is subject to a norm of truth.¹ This is taken to explain intuitive data about the nature of belief and the assessment of beliefs and believers.² Amongst those who accept this view however, there is little agreement as to how the norm should be understood. In a recent paper 'The Truth Norm of Belief' Conor McHugh has argued that the norm should be understood in *evaluative* terms rather than in *deontic* or 'prescriptive' terms.³

This marks a real departure from the existing literature: existing attempts to clarify the sense in which belief is subject to a norm of truth have worked with (a variety of) deontic formulations. We have some sympathy with McHugh's ambitions. And many of his arguments are initially appealing. But in this paper we argue against him. Our argument takes the following form. Firstly, we note that McHugh's case for the evaluative account is premised on the claim that the problems facing the prescriptive account are better dealt with by the evaluative account. Secondly, we argue against McHugh that these problems are not better dealt with by the evaluative account at all. We conclude that McHugh's case for the evaluative account is – as it stands - under-motivated. Although our arguments are targeted primarily against McHugh, we take them to generalise to

many – we hope, *any* - attempts to show that evaluative accounts are preferable to prescriptive accounts. We are also optimistic that these arguments have force against McHugh's more recent attempt to explain the normative relation between belief and truth in terms of 'fittingness'.

We begin in section 1 by summarising the prescriptive or deontic approach and some of the problems that McHugh identifies for it. In section 2 we present the evaluative account as an alternative. In sections 3-7 we show that – despite McHugh's claims - the evaluative account does not fare better as regards these problems. In section 8 we claim that our arguments generalise to pressure any evaluative account, and 'fitting-attitude' accounts as well.

1. Problems for the Prescriptive Account.

It is often claimed that there is a single, primitive norm relating belief to truth. This is taken to explain (for example) why there is something 'wrong' or 'defective' about a false belief, why there is something 'wrong' or 'defective' about irrational belief, why the 'right kind' of reasons for belief are truth-directed or evidential, why beliefs tend to actually be responsive to truth-relevant considerations, and why various 'paradoxes of belief' (including 'Moore-paradoxical' sentences) arise.⁴ We will suppose for sake of argument that these are good motivations for thinking that there must be some such norm. We are interested in how that norm should be understood. Arguably the simplest way of understanding the norm is to read the relationship between belief and truth in prescriptive terms. Prescriptive properties include the property of being what one *ought* to believe, what one *may* believe, or what one *may not* believe. The simplest, and one of the most popular, ways of understanding the norm is as follows:

(1) For any S, p : if p is true then S ought to believe p , and if p is false then S ought not believe p .⁵

This is taken to express the primitive norm on belief. Other norms on belief – such as norms governing *rational* belief – are then explicable in terms on this primitive norm.⁶ Despite some initial appeal however, this strategy for

understanding the norm on belief has been widely criticised, most influentially by Bykvist and Hattiangadi (2007). McHugh takes up and develops these criticisms. They will form the basis of his case for the evaluative account. His claim will be that 'the evaluative account is well placed to avoid all of the problems... raised for the... prescriptive account' (McHugh 2012, 19). So what are the problems facing the prescriptive account? We briefly sketch the four McHugh focuses on below.

Firstly, it entails – the falsehood - that one ought to believe arbitrarily long conjunctions. The argument is as follows. There are true conjunctions so long that we cannot believe them – take the proposition that is a conjunction of all the truths, for example. Because such propositions are, *ex hypothesi*, true, (1) entails that we ought to believe them. But we cannot believe such arbitrarily long conjunctions (owing to our cognitive limitations). But if we cannot believe them, then it is not the case that we ought to. So it is not the case that we ought to. So (1) is false.⁷

Secondly, (1) entails – the falsehood - that one ought to believe true 'blindspot propositions'. Consider the proposition that it is raining but no-one believes that it is raining. This proposition may be true. If so, (1) entails that one ought to believe it. But if one believes it then it is false. So if one believes such a proposition, it becomes the case that one ought not to believe it. So (1) entails that one ought to believe a proposition that – by its own lights - one ought not to believe.

Thirdly, (1) entails – the falsehood - that one ought to believe self-fulfilling propositions. An example illustrates. Suppose that I am playing tennis. At present, the evidence indicates that I will not win the game. However, if I form the belief that I will win, then – because of the extra confidence – I will win. (1) would appear to entail that I ought to believe that I will win. But it doesn't seem that I ought not to believe this (at present) evidentially insufficiently supported proposition. So (1) is false.

The first three problems concern the extensional adequacy of the prescriptive account. The fourth does not. The fourth problem is that the modifications that it

is necessary to make to (1) in order to resolve the above problems render it too complex to guide belief formation and revision. But the norm on belief must guide belief formation and revision (albeit at an abstract level). So, neither (1) nor a suitably modified version of it is true.

We propose to take these four problems for the prescriptive account at face value (though we discuss them in greater detail below). The question that we are interested in is whether – as McHugh claims - an evaluative account of the norm on belief will fare any better. We argue, against McHugh, that it will not. Having done this, we deal with a further objection that McHugh raises. He claims that the deontic or prescriptive account faces difficulties (that the evaluative account does not) with cases in which it is appropriate to *withhold judgment*. We discuss this separately in section 7.

2. The Evaluative Account

McHugh's proposal is to understand the basic norm relating belief to truth in evaluative rather than deontic or prescriptive terms. He works with the evaluative properties of goodness and badness. His basic proposal (which he then modifies) is the following:

(3) For any S, p : if S believes p , then that belief is good if p is true, and that belief is bad if p is false.

The motivation for this account is given entirely by (a) its ability to retain the initial motivations for thinking that there is some norm relating belief to truth, coupled with (b) the comparative lack of susceptibility of the evaluative account (compared to the prescriptive account) to the four problems noted above. We will simply grant McHugh (a). We will grant that evaluative accounts are able to retain the initial motivation for thinking that there is some norm relating belief to truth. Our focus will be on (b). We will focus on whether the evaluative account does fare better than the prescriptive account as regards the above objections. We will claim that it does not.

Before making our case below, it is important to be clear on a further detail of

the evaluative account. The evaluative account states the (purported) fundamental norm relating belief to truth. And it states that this norm should be understood in evaluative terms. What the evaluative norm *doesn't* do, however, is tell us how exactly that norm relates to any deontic or prescriptive norms on belief. A defender of the evaluative norm does not deny that there are such (prescriptive) norms. She simply claims that they are less fundamental than the evaluative norm. Specifically, the deontic or prescriptive norms are to be explained in terms of the evaluative norm and some general principles relating evaluative properties of belief to deontic or prescriptive properties of belief.

For McHugh's part, he does not go so far as to defend a particular account of what, exactly, the prescriptive norms on belief are or how they are derived from the more fundamental evaluative norm. Nevertheless, he suggests that these deontic norms will relate belief to truth in the kinds of ways that one might expect. He gives as an example: 'if your evidence for p is conclusive you ought to believe p rather than withhold or disbelieve it; you ought not believe q if you lack good evidence for it' (29, n.36). As McHugh notes, however, whether he is entitled to this will depend on whether there are independently defensible principles relating evaluative properties of beliefs to deontic properties of beliefs. We won't challenge McHugh's entitlement to the specific prescriptive norm mentioned in the preceding quotation. We do note, however, that there is a general issue here. There is a burden of proof on McHugh – or on a defender of the evaluative norm more generally - to show that the relevant prescriptive norms can be derived. McHugh does not claim to meet this burden ('how exactly does the evaluative account explain epistemic norms? I cannot deal with that question in detail' (21)).⁸ We return to related themes in some of the arguments below.

We now turn to the central question: does McHugh's evaluative account fare better than the prescriptive account as regards the four problems listed in section 1? We argue that it does not. We take each objection in turn. The general structure of our argument is as follows. We claim that the first three objections – all of which concern extensional adequacy - succeed in rendering the evaluative account preferable to the prescriptive account only if the fourth objection ('the

complexity objection') succeeds in rendering the evaluative account preferable to the prescriptive account. And we claim that the fourth objection does not succeed in this. So – contrary to McHugh's claims - the four objections do not succeed in rendering the evaluative account preferable to the prescriptive account.

3. First Problem: Believing Arbitrarily Long Conjunctions

The prescriptive account appeared to entail – implausibly - that one ought to believe arbitrarily long conjunctions. McHugh claims that the evaluative account does not have this implication. It does not have this implication because, in the first instance, the evaluative account does not entail *anything* about what one ought to believe. It only states which beliefs are good and which are bad. Furthermore, it would only entail that one ought to believe all the truths if it were coupled with an implausibly strong principle relating evaluative properties to prescriptive properties: a principle to the effect that 'whenever something is or would be good, you ought to do it or bring it about' (20). This principle, however, McHugh claims, is independently implausible. So, the evaluative account avoids the objection from believing arbitrarily long conjunctions.

We won't dispute McHugh's argument here. But we claim that it would be a mistake to take this as grounds to prefer the evaluative account over the prescriptive account. The primary reason for this is simple. It is that, *as McHugh himself acknowledges*, the prescriptive account can easily be modified to avoid the consequence that one ought to believe arbitrarily long conjunctions. The modification that McHugh proposes is to (1) that is sensitive to this. His proposal – which is a first modification of (1) – is:

(1*) For any S, p : if S has some doxastic attitude to p then [(if p is true then S ought to believe p) and (if p is false then S ought not believe p)].

We agree with McHugh that this formulation (or some other, similar formulation) is sensitive to the 'believing arbitrarily long conjunctions' objection: it doesn't require one to believe propositions that one couldn't hold a belief about. It is perhaps curious, then, that when McHugh lists the senses in which

‘the evaluative account is well placed to avoid all of the problems... raised for the... prescriptive account’ he lists the ‘believing arbitrarily long conjunctions’ objection amongst them. Given that a suitably modified version of the prescriptive account (i.e. (1*)) is not undermined by the ‘believing arbitrarily long conjunctions’ objection, why is the fact that the evaluative account isn’t undermined by this objection either presented as a comparative strength of the latter?

What McHugh has in mind here, we think, is that once the prescriptive account is modified to deal with the ‘believing arbitrarily long conjunctions’ objection all else ceases to be equal. In particular, a problem arises elsewhere concerning the *complexity* of the modified prescriptive norm, (1*). It is after all it is true that (1*) is more complex than (1). And recall McHugh’s worry (fourth objection, in section 1) that any formulation of the norm on belief must be simple enough to guide belief-formation. So this may be a sense in which all else is not equal and hence in which the evaluative account would fare better than the prescriptive account. We return to this when we consider the issue of complexity in section 6.

4. Second Problem: Blindspot Propositions.

Our response to the second problem for the prescriptive account (concerning blindspot propositions) is very similar to the response given above. We concede to McHugh that the evaluative account does not face the problem. But we note that a suitably modified version of the prescriptive account doesn’t either.

McHugh argues that the evaluative account does not face problem with blindspot propositions (20). Specifically, the evaluative account does not entail that it would be good to believe blindspot propositions. We agree with McHugh here. The evaluative account – as set out in (3) – does *not* entail that it is good to believe true propositions. So it does not entail that it would be good to believe true blindspot propositions. Instead, the evaluative account – as set out in (3) - states that whether it would be good to believe a proposition depends on whether the resulting belief would be good. And the belief that would result from believing a blindspot proposition would be a false belief. So the evaluative account entails that believing a true blindspot proposition would not be good (in

fact, it would be bad).

We won't dispute McHugh's argument here. But we claim that it would be a mistake to take this as grounds to prefer the evaluative account over the prescriptive account. This is because it is not the *evaluative* nature of the evaluative account that allows it to avoid the problem of blindspot propositions. Rather, it is a matter of how the evaluative account is *parsed* that allows it to avoid the problem of blindspot propositions. And, we claim, it is possible to parse the deontic or prescriptive account so that it too can avoid the problem of blindspot propositions.

Indeed, McHugh provides such a parsing himself. The parsing is as follows

(1**) For any S, p : if S has some doxastic attitude to p then [(S ought to believe p if S would thereby have a true belief that p) and (S ought not believe p if S would thereby have a false belief that p)].

This is a parsing of the prescriptive account that does not entail that one ought to believe true blindspot propositions. Rather, it entails that whether one ought to believe a proposition depends on whether the resulting belief would be true. And for blindspot propositions, the resulting belief would be false. So, (1**) does not entail that one ought to believe blindspot propositions. We take this to show that it is not the *evaluative* nature of the evaluative account that allows it to avoid the problem with blindspot propositions. Rather, it is the parsing of the view.

Given that McHugh provides us with (1**) it is again curious that when McHugh lists the senses in which 'the evaluative account is well placed to avoid all of the problems... raised for the two versions of the prescriptive account' he lists the problem with blindspot propositions. Why does he do this? Again, perhaps what McHugh has in mind here is that in re-interpreting the prescriptive account so as to make it sensitive to the objection from blindspot propositions, (1) must be made implausibly in some other respect: specifically, in respect of both its complexity and in respect of the problems that it creates for dealing with self-fulfilling beliefs, as we discuss below.

5. Third Problem: Self-Fulfilling Beliefs.

The third problem for the prescriptive account concerns self-fulfilling beliefs. Recall the basic problem. The basic problem is that there may be cases in which – on the prescriptive account – one ought to believe propositions for which one possesses insufficient evidence. The example of the tennis match given in section 1 illustrates this point. Suppose that my evidence indicates that I will not win the match. But if I believe – contrary to my evidence - that I will win then the extra confidence will cause me to win. So if I believe that I will win, I will thereby have a true belief. (1**) would appear to entail that I ought now to believe – contrary to my evidence – that I will win. This, McHugh claims, ‘seems unacceptable’ (13).

And the basic problem worsens when we examine what the prescriptive account appears to entail in more detail. For in some cases the prescriptive account appears to entail *not only* that one ought to hold a self-fulfilling belief, *but also* that one ought to believe an incompatible evidentially-supported belief. To see this, note that in the tennis example above there are two different strategies that I might take in order to form a true belief. One strategy is to believe, contrary to my evidence, that I will win. This will lead me to form a (self-fulfilling) true belief. The second strategy is to believe that I will *not* win. This will *also* lead me to form a true belief because – as McHugh stipulates the case - believing that I will not win will be a blow to my confidence that causes me not to win (13). The availability of these two strategies for forming a true belief is deeply problematic for the prescriptive account. It is deeply problematic because (1**) appears to require one to take both. It appears to require believing that I will win and to require believing that I will not win. Schematically, if (1**) is true, then there are cases in which I ought to believe that p and simultaneously I ought to believe that not-p. As McHugh puts it ‘that can’t be right’ (13).

McHugh argues that the evaluative account fares much better. The evaluative account does *not* entail that one ought to believe self-fulfilling propositions. And it does not entail both that one ought to believe p and that one ought to believe not-p. It may seem that the evaluative approach is committed to the claim that one ought to have self-fulfilling beliefs, because self-fulfilling beliefs are always

true, and therefore good. But McHugh claims that the evaluative approach is not committed to this. This is because whether one ought to believe that p , McHugh argues, depends not only on whether a belief that p would be *good*, but whether that belief would be better than other doxastic states one could be in (20). And in the tennis example, McHugh claims that there are other states – such as believing that you will not win – that would be equally good as believing that you will win. McHugh therefore concludes that defender of the evaluative approach is probably committed to no more than the claim that one *may* hold self-fulfilling propositions.

This may seem a persuasive case against the prescriptive account and in favour of the evaluative account. But we do not think that it is. As McHugh himself acknowledges, the problems with self-fulfilling beliefs only arise at all for the specific parsing of the prescriptive norm given in (1**). As McHugh himself acknowledges it is possible to parse the prescriptive in such a way that no such problem arises at all. Specifically, he claims that (1***) avoids the problem with self-fulfilling beliefs.

(1***) For any S, p : if S has some doxastic attitude to p then $\{[S$ ought to believe p if $(p$ is true and S would thereby have a true belief that $p)]$ and $[S$ ought not believe p if S would thereby have a false belief that $p)]\}$.

We agree with McHugh that (1***) allows a defender of the prescriptive approach to avoid any problems with self-fulfilling beliefs. This is because it allows a defender of the prescriptive account to avoid commitment to the claim that one ought to hold self-fulfilling beliefs at all. It does so because (1***) only requires one to hold a belief in p *both* when doing so would lead one to form a true belief *and* when p is true antecedently to one's holding it. And self-fulfilling beliefs do not fulfil the second conjunct. So (1***) does not entail that one ought to hold them.

It follows that – by McHugh's own lights - it is not obvious that the evaluative account fares better than the prescriptive account as regards the problem of self-fulfilling beliefs at all. In fairness to McHugh, however, he is explicit that the modification necessary to render the deontic or prescriptive account immune to

the objection from self-fulfilling beliefs comes at a cost. Specifically, the prescriptive norm that is immune to the objection from self-fulfilling beliefs - (1***) – is *too complex* to play the role that it must play in guiding belief-formation. This is a worry that we have encountered in each of the objections to the prescriptive accounts discussed above. We turn to it below.

6. Fourth Problem: Complexity.

The fourth worry, unlike the previous three, does not concern the extensional adequacy of the prescriptive account. The fourth worry is that any feasible version of the prescriptive norm (e.g. (1***)) is too complex to guide belief-formation in the way that it must. This is an important objection, because, as we have seen above, none of the three previous objections speaks in favour of the evaluative account over the prescriptive account *unless* this objection does. Specifically: each of the three previous objections are dealt with *at least* as well by the prescriptive account as by the evaluative account unless the modifications necessary for the prescriptive account to deal with the three previous objections renders the prescriptive norm too complex to guide belief.

Let's begin by thinking about McHugh's 'complexity' objection to the prescriptive account in more detail. According to McHugh the norm on belief must play two roles. Firstly, 'it must be the basis of our epistemic assessment of beliefs'. Secondly, 'it is supposed to ground the ways in which we are motivated to form, revise and extinguish our beliefs' (14). According to McHugh this 'requires the norm to do some psychological work, either explicitly or implicitly, in our thinking *about* beliefs, and in our thinking that leads *to* beliefs' (14-5). And, McHugh claims, (1***) is too complex to do this psychological work. The evaluative account, by contrast, is not (he claims) too complex to play this role: 'I found (1***) far too complex to be plausible as a norm that we actually deal with. (3) does not suffer from this problem' (20).

We will not directly dispute McHugh's claims about the roles that a norm on belief must play. But we will argue that he is wrong to think that evaluative norms fare better than prescriptive norms on this measure.

As some preliminary support for this, we note that (3) is not McHugh's final formulation of the evaluative norm. The version of the evaluative norm that McHugh ends up defending (for reasons that we neither discuss nor dispute here) is actually more complex than (3). It is:

(3*) For any S, p : if S believes p , then that belief is a good doxastic attitude to have to p if p is true, and that belief is a bad doxastic attitude to have to p if p is false.

This is, as McHugh concedes, more complex than (3). So some of the simplicity that (3) enjoys over (1***) is lost when (3) is modified to (3*).

Let's set aside this preliminary point however. There is a second, much more serious (and philosophically interesting) basis for our disagreement with McHugh. It is that because the evaluative account is evaluative rather than prescriptive, it is not clear that it could, by itself, guide belief formation and revision *at all*. This is our core point. To bring it into view, note that prescriptive accounts clearly and explicitly offer guidance as to how to revise and form beliefs. They state how one *ought* to believe. But evaluative accounts do not state how one ought to believe. They merely make claims about when beliefs are good or bad. And this does not provide direct guidance as to what to believe.

In order to extract some guidance as to how to revise and form beliefs it is necessary to add an additional principle allowing one to derive claims about how one ought to believe from which beliefs are good or bad. And this reintroduces the complexity; only this time for the evaluative account. Crudely put, in order for agents to take guidance on how to form or revise beliefs from the evaluative account they must (a) apply the relevant evaluative norm (e.g. 3*), (b) apply, in addition to this, a principle linking the evaluative properties of beliefs to prescriptions about how one ought to believe, and (c) apply the resulting derived norm on how one ought to believe to particular cases. And it is far from clear that this process requires any less of ordinary believers than is required by the application of a basic deontic or prescriptive norm (e.g. (1***)).

We can make the same point stand out more clearly via an analogy with *practical*

philosophy. According to a popular family of views in practical philosophy – consequentialist views - what one ought to do is determined by which states of affairs it would be good to bring about. McHugh is, in many respects, in an analogous position to these practical consequentialists. Like them, he claims that what one ought to *believe* is determined by whether the resulting beliefs would be (epistemically) good. Now consequentialists in practical philosophy face well established difficulties regarding the complexity of their view – especially at the psychological level. And seeing this helps to make the difficulties that McHugh faces (as regards complexity) stand out.

One of the difficulties that practical consequentialists face concerns *how* much work is required to derive facts about what one ought to do (i.e. ‘the right’) from facts about what it would be good to. To see this, let’s suppose that we are practical consequentialists; that is to say, we think that what it is right to do is determined by the goodness that would result from our actions. And let’s suppose furthermore that we agree that relevant metric of goodness for our consequentialism is human welfare. Note just *how wide* our options remain as regards the standard of right action. For example, should the right action be that which has the *best* consequences (‘maximising consequentialism’), or that which has the *least bad* consequences (‘negative consequentialism’), or that which produces consequences that reach some threshold level of goodness (satisficing consequentialists), or some other option? There is also the question of *how* we should take consequences into account. Should we aim to choose a *policy* that will have the best consequences (‘rule consequentialism’)? Or should we choose what to do on a case-by-case basis (‘act consequentialism’)? What these questions should make clear is the extent to which merely accepting a consequentialist view leaves open the standard of rightness. An analogue of this problem applies to McHugh’s evaluative account of the norm on belief. For merely accepting an evaluative account of the relation between belief and truth leaves open a similarly wide range of options when it comes to determining what one ought to believe. And this will impose a heavy psychological cost on agents who are trying to determine what they ought to believe based on evaluative norms such as (3) or (3*). This undermines McHugh’s claim that the way in

which the evaluative norm provides us with guidance is less complex than the prescriptive norm.

We can deepen the analogy – and the problem for McHugh – by considering another closely related difficulty that practical consequentialists face. Practical consequentialism has historically been criticized for demanding that agents perform excessive calculation. The point is nicely put by Mill:

[D]efenders of utility often find themselves called upon again to reply to such objects as this – that there is not time, previous to action, for calculating and weighing the effects of any line of conduct on the general happiness. (Mill, 1861, 224)

It seems that the same charge could be made against McHugh in the epistemic case. As we've said, McHugh requires there to be a principle linking the evaluative properties of beliefs to what we ought to believe, and, most relevantly to the current point, believers will need to apply this derived norm in particular cases. With any proposition I have a doxastic attitude towards, this will involve considering how likely believing that proposition is to be good or bad. This looks like it will involve considerable calculation, i.e. considering how likely p is to be true, considering how likely it is that other alternatives to p are true, considering my beliefs about the evidence pertaining to p and how likely they are to be true, etc. Once again, this is simply another way of putting the basic point made in (a)-(c) above. We take this to provide further reason to be sceptical of McHugh's claim that the way in which the evaluative norm provides us with guidance is less complex than that of the prescriptive norm.

We take these points to provide a fairly decisive case against McHugh's claim. But let's stop to consider an objection on McHugh's behalf. The objection is actually based on the analogy with practical consequentialism. Practical consequentialists have a standard response to the charge that their view is psychologically over-demanding. It is that consequentialist norms should not be understood as a 'decision procedure', but rather as a 'criterion of rightness'. This is a point made famously by Sidgwick:

[T]he doctrine that Universal Happiness is the ultimate standard must not be understood to imply that Universal Benevolence is the only right or always best motive of action. For, as we

have before observed, it is not necessary that the end which gives the criterion of rightness should always be the end at which we consciously aim. (Sidgwick, 1907, 413)

Sidgwick's thought is that if we view consequentialism as providing a criterion of rightness rather than a decision procedure, then consequentialism needn't require us to do excessive calculation. Viewed as a criterion of rightness, consequentialism can end up recommending seemingly un-consequentialist and easily-followable motives for action such as following rules of thumb, acting on the basis of virtuous dispositions, or acting from duty.

Perhaps McHugh could make a similar move on behalf of his evaluative account of the relationship between belief and truth. He could claim that the evaluative norm should be viewed as a criterion of rightness (or 'goodness'), rather than a decision procedure. If he did then the complexity we have claimed is involved in the evaluative norm providing guidance would perhaps be unproblematic.

We don't think, however, that this will help McHugh. This is because even though thinking of the evaluative norm as a criterion of rightness can deal with the issues of complexity that we have raised with the evaluative norm, an analogous move can be used to defend prescriptive views against McHugh's charge that *they* are too complex to provide guidance. Specifically, a defender of the prescriptive norm can claim that (1^{***}) should be thought of as a criterion of rightness rather than a decision procedure. So McHugh will not be entitled to his key claim that '(1^{***}) [is] far too complex to be plausible as a norm that we actually deal with [but] (3) does not suffer from this problem' (20).

Indeed, it is worth noting that, in other contexts, defenders of the prescriptive norm already do claim that we have to distinguish between the precise formulation of the truth norm and how it guides belief formation and revision. This is Engel:

[T]here is no reason to suppose that we can *read off* the kind of regulation or guidance which a norm gives from its nature and its expression. In other words, the norm and its form—whatever way one conceives of it—is one thing, its psychological realization in the mind of a subject is another. (Engel 2013, 42)

Therefore, it looks like whatever arguments McHugh can appeal to argue that the complexity of the evaluative norm is unproblematic can be used to show that the complexity of (1***) is unproblematic. The complexity of (1***) does not give us reason to reject it *in favour of* an evaluative norm. So the objection from ‘complexity’ fails to serve McHugh’s purposes.

7. Withholding Judgment

To this point we have argued that McHugh’s argument turns on the claim that the modified prescriptive account is unfeasibly complex in relation to the evaluative account. In this section we briefly discuss a different kind of argument that McHugh claims favours the evaluative account over the prescriptive account. It concerns withholding judgment.

McHugh argues that on the prescriptive account, it will never be correct to withhold judgment on a proposition. This is because withholding judgment is never conducive to believing the truth. So, on the prescriptive account, withholding judgment is never what one ought to do. And this is problematic because on some occasions, withholding judgment *is* what one ought to do. McHugh claims that the evaluative account fares better. This is because the evaluative account ‘says nothing about withholding’. And so, ‘withholding is neither good nor bad as far as this norm goes’ (20). So, for all that the evaluative account tells us, suspending judgment is the right thing to do.

We think that this objection misses its mark in a sense that was prefigured in the quotation from Engel at the conclusion of the preceding section. The point is really about *guidance*. The prescriptive account is apt to guide believers in the obvious sense that it is prescriptive: it tells believers what to believe. But there is another sense in which it is *not* apt to guide believers. It tells us (roughly) to believe truths and not believe falsehoods. But this isn’t very helpful as a guide. We can’t just believe the truth at will: we aren’t omnipotent. Cognitively limited creatures like us need to take *means* to this end. This is a point about which those who have defended the prescriptive norm are quite explicit.⁹ The prescriptive norm should be thought of as an ‘objective’ standard that grounds subsidiary norms that we *can* follow. This echoes the contrast between a

‘criterion of rightness’ and a ‘decision procedure’ mentioned in the previous section.

Why is this relevant to McHugh’s claim that the prescriptive norm can’t account for the correctness of withholding judgment in some cases? It is relevant because although the prescriptive norm itself (i.e. (1**)) does not entail that one ought to withhold judgment, the subsidiary norms that it grounds *that actually guide creatures like us* may well do so. They may do so in roughly the following way. In order to conform to the prescriptive norm one must not only believe truths, but one must also avoid believing falsehoods. And in some cases – cases in which one’s evidence neither indicates strongly that a proposition is true nor that it is false – the best way to do this may be to withhold judgment. That way one will at least avoid doing as one ought not (i.e. believing a falsehood). So a prescriptive norm may well ground subsidiary norms that recommend withholding judgment.

Now one might object that this response amounts to a *further* complication for the prescriptive account and as such, a reason for scepticism about it in comparison to the evaluate account. But this is not true. The evaluative account faces the very same complication. We have already shown (in the previous section) that if the evaluative norm is to guide believers, then it must ground some subsidiary (presumably prescriptive) norm too. Given that the evaluative account doesn’t tell us anything about withholding judgment, the rules for withholding judgment must be written into this (perhaps these) subsidiary norm(s). And a proponent of the evaluative norm owes us an account of how this comes to be: how does the evaluative norm ground subsidiary norms that recommend withholding judgment? This is, structurally, much the same burden that a proponent of the prescriptive account faces.

8. Generalising

We have presented an argument against McHugh specifically. But we think that our basic strategy will generalise: it will be effective against other arguments for evaluative accounts. To see this, let’s just focus on the strategy that we used in sections 3-6 (and not on the argument concerning withholding judgment).

McHugh presented three arguments against the *extensional adequacy* of the prescriptive account and one that concerns its *complexity*. By his own lights the three arguments based on extensional adequacy don't succeed unless the complexity-based argument succeeds. And we have shown that the complexity-based argument won't succeed. Now think about *any* attempt to argue that an evaluative account is preferable to a prescriptive account. As McHugh has shown, no such argument will succeed on grounds of extensional adequacy. And if we're right, it won't succeed on grounds of complexity either. This is because – as we argued in section 6 above – if *any* evaluative norm is to provide guidance it must do so via a subsidiary norm. And the process of deriving and implementing the subsidiary norm is bound to re-introduce the requisite complexity.

We suspect that our basic strategy is also somewhat effective against McHugh's more recent attempt to understand the relation between truth and belief in terms of the sui generis normative property of *fittingness* (2014). As McHugh's modifications of (1) have shown, the fittingness-based account won't be extensionally superior to the prescriptive account. And we think it won't fare better in terms of complexity either. To see this, simply note that a norm relating belief to truth via the property of fittingness doesn't by itself provide guidance for believers any more than an evaluative norm. The mere fact that all and only truths are fitting objects of belief doesn't tell one what (if anything) one ought to believe (or how to go about making one's beliefs conform to this standard). If a fittingness-based account is to provide guidance it must – like an evaluative account - proceed via subsidiary norms too. And we see no reason to think that this won't introduce complexity.

If one is to argue that prescriptive accounts should be rejected in favour of either evaluative or fittingness-based accounts, one must proceed along different lines. One must provide measures on which evaluative or fittingness-based accounts are superior to prescriptive accounts *other than the measures of extensional adequacy and complexity*. We have argued that consideration of withholding judgement won't obviously help here. In his recent paper however McHugh does provide some measures that might be thought useful to this end (though he doesn't justify his provision of these measures in this way). One measure is

coherence: any account of the relation between belief and truth must cohere with existing claims in the theory of normativity. A second measure is *dominance*: any account of the relation between belief and truth must explain why evidence, rather than pragmatic considerations, is 'dominant amongst reasons for belief' (2014, 168).

It is primarily in terms of dominance that McHugh thinks fittingness-based accounts fare better than prescriptive accounts. But – although we won't go into this in much detail here – his arguments don't convince. The thrust of his reasoning is as follows (2014, 170). Prescriptive accounts fail to be extensionally adequate for familiar reasons (concerning arbitrarily long conjunctions, blindspot propositions, self-fulfilling beliefs). To resolve this, prescriptive accounts should be understood as stating that one *may* believe all and only truths as opposed to that one *ought* to. But this kind of 'may-based' prescriptive account fails to account for the dominance of evidence amongst reasons for belief (in that it fails to entail that evidence yields any more than *permissions* to believe). We agree that may-based prescriptive accounts face this problem (see footnote 7 above). But, by McHugh's own lights, prescriptive accounts understood in terms of *ought* can avoid the charges of extensional inadequacy if suitably modified (as in (1***)). So considerations of extensional adequacy do not push one from an ought-based account to a may-based account in the first place. So prescriptive accounts don't face this obvious worry in accounting for dominance.

There is, of course, much more to be said here. We have been unable to treat McHugh's more recent arguments fully. But we hope to have shown why we are optimistic that our basic arguments generalise.

9. Conclusion

We have argued that McHugh's evaluative account fares no better than the prescriptive account. We are optimistic that these conclusions generalise in the sense suggested above: merely changing the *kind of normative property* that figures in the truth norm – e.g. to goodness or fittingness – will not *by itself* solve the fundamental problems that the prescriptive account faces.¹⁰

Alexander Greenberg
Trinity Hall
University of Cambridge

Christopher Cowie
Fitzwilliam College
University of Cambridge

¹ See e.g. Wedgwood 2002, 2007, Shah 2003. For defence of a truth norm as opposed to a norm based on knowledge (as in e.g. Williamson 2000, 47) see Whiting 2013.

² For more detailed discussion of the motivations for the truth norm view see McHugh and Whiting 2015.

³ McHugh has since changed his mind: see section 7. We focus on the evaluative account expressed in his 2012 paper ‘The Truth Norm of Belief’. For another defence of an evaluative truth norm, see Fassio 2011. Although explicit defences of an evaluative truth norm are rare, many philosophers of different stripes have made the claim that true beliefs are ‘good’ or ‘valuable’. See Whiting 2013b for a survey of philosophers making this kind of claim.

⁴ See footnote 2.

⁵ See e.g. Boghossian 2005, Shah 2003, Engel 2004, 2013, 2013a, and Gibbons 2013 for presentations of the truth norm in terms of what one *ought* to believe. The view is sometimes presented in terms of ‘correctness’, especially by Wedgwood (see 2002, 2007), though this requires interpretation; and one interpretation reads correctness in terms of ‘ought’.

⁶ See e.g. Wedgwood 2002, 2007, Shah 2003, Engel 2013a.

⁷ McHugh notes that it is possible to modify (1) in several ways to avoid this objection. We return to this at greater length below. But we note for present – in order to set it aside - one way of modifying (1) that McHugh considers. It is to replace ‘ought’ with ‘may’ in the formulation of the norm on truth (Cf. Whiting 2010):

(2) For any S, p : if p is true then S may believe p , and if p is false then S may not believe p .

McHugh refers to this as the ‘weak prescriptive account’. Clearly, it does not entail that one ought to believe all of the truths. McHugh considers this version of the norm in some detail. We will not. We will work with the formulation of the prescriptive norm given in

terms of 'ought' rather than 'may', because the weak prescriptive account faces problems of its own – problems that McHugh identifies elsewhere. Most obviously (2) doesn't entail that one *ought* to believe *anything*. And this has struck many commentators as unsatisfactory. There are some propositions that we ought to believe in light of the evidence that we possess (this is most obviously the case if we think in terms of belief-revision when suitable evidence becomes available, rather than belief-formation). For Whiting's defence see Whiting 2013a. It is also useful for us to note here that (2) is roughly equivalent to the following norm, which is also appealed to escape the problem of arbitrarily long conjunctions:

(2*) For any S, p : S ought to believe that p only if p is true.

(2*) is sometimes defended because it doesn't entail that one ought to believe any arbitrary truth (see Boghossian 2005). But it also faces the problem that it doesn't entail that one ought to believe anything.

⁸ Though he does gesture in this direction in his brief discussion of the relationship between values and (pro tanto) reasons (2012, 21). As McHugh does not claim to defend this in any detail however it would be uncharitable for us to base our criticism of his overall project on any specific failings of the relationship (between values and (pro tanto) reasons) in this discussion. So we do not.

⁹ E.g. Wedgwood 2001, Engel 2013a.

¹⁰ We are grateful to an anonymous reviewer for this journal and to Tim Crane for helpful comments on an earlier draft. This work was supported by the Arts and Humanities Research Council.

Bibliography:

- Boghossian, Paul (2005), 'Is Meaning Normative?' in Nimtz and Beckermann (eds.), *Philosophy—Science—Scientific Philosophy. Main Lectures and Colloquia of GAP.5, Fifth International Congress of the Society for Analytical Philosophy, Bielefeld 2003* (Paderborn: Mentis), 205-218.
- Bykvist, Krister, and Hattiangadi, Anandi (2007), 'Does Thought Imply Ought?' *Analysis* 67 (296), 277–285.
- Engel, Pascal (2004), 'Is Truth the Aim of Belief?' in Gillies (ed.), *Laws and Models in Science* (London: King's College Publications), 77-97.
- Engel, Pascal (2013) 'In Defence of Normativism about the Aim of Belief', in Chan (ed.), *The Aim of Belief* (Oxford: Oxford University Press), 32-63.
- Engel, Pascal (2013a) 'Doxastic Correctness', *Proceedings of the Aristotelian Society Supplementary Volume* 87(1), 199-216.
- Fassio, Davide (2011) 'Belief, Correctness and Normativity', *Logique et Analyse* 54(216), 471-486.
- Gibbons, John (2013) *The Norm of Belief* (Oxford: Oxford University Press)
- McHugh, Conor (2012) 'The Truth Norm of Belief', *Pacific Philosophical Quarterly* 93(1), 8-30.
- McHugh, Conor (2014) 'Fitting Belief', *Proceedings of the Aristotelian Society* 114(2), 167-187.
- McHugh, Conor and Whiting, Daniel (2014) 'The Normativity of Belief', *Analysis* 74(4), 698-713.
- Mill, John Stuart (1861) *Utilitarianism*, in Vol. X of Robson (ed.) *Collected Works of John Stuart Mill*, 33 Volumes (Toronto: University of Toronto Press, 1965-91).
- Shah, Nishi (2003) 'How Truth Governs Belief', *Philosophical Review* 112(4), 447-482.
- Sidgwick, Henry (1907) *The Methods of Ethics*, 7th Edition (London: Macmillan).
- Wedgwood, Ralph (2002) 'The Aim of Belief', *Philosophical Perspectives* 16, 267-297.
- Wedgwood, Ralph (2007) *The Nature of Normativity* (Oxford: Oxford University Press).

- Whiting, Daniel (2010) 'Should I Believe the Truth?' *Dialectica* 64(2), 213-224.
- Whiting, Daniel (2013) 'Nothing But the Truth: On the Norms and Aims of Belief',
in Chan (ed.), *The Aim of Belief* (Oxford: Oxford University Press), 184-
203.
- Whiting, Daniel (2013a) 'Truth: The Aim and Norm of Belief', *Teorema* 32(3),
121-135.
- Whiting, Daniel (2013b) 'The Good and the True (or the Bad and the False)',
Philosophy 88(2), 219-242.
- Williamson, Timothy (2000) *Knowledge and its Limits* (Oxford: Oxford University
Press).