

MOORE'S PARADOX AND EPISTEMIC NORMS

Clayton Littlejohn

(Forthcoming in the *Australasian Journal of Philosophy*)

We shall evaluate two strategies for motivating the view that knowledge is the norm of belief. The first draws on observations concerning belief's aim and the parallels between belief and assertion. The second appeals to observations concerning Moore's Paradox. Neither of these strategies gives us good reason to accept the knowledge account. The considerations offered in support of this account motivate only the weaker account on which truth is the fundamental norm of belief.

I. Introduction

In the wake of Williamson's defence of the knowledge account of assertion, it is not surprising that an increasing number of authors are now defending the thesis that knowledge is the norm of belief.¹ If knowledge is the norm of belief, you should not believe p unless you know p . If you should not believe p unless you know p , it seems that a belief must constitute knowledge in order to be truly justified.²

Must we really refrain from believing what we do not know? I think not. I shall respond to arguments offered in support of the knowledge account, show that the view delivers the wrong verdicts, and show the considerations taken to support the view can be accommodated by alternative views such as the view that belief is governed by the truth norm.³ Because the knowledge account seems to have the implication that there are no false, justified beliefs, many

¹ Slote [1979], Williamson [2000], and Unger [1975] defend the thesis that knowledge is the norm of *assertion*. Adler [2002], Bird [2007], Huemer [2007], Sutton [2007], Unger [1975], and Williamson [2000] say that knowledge is the norm of *belief*. See Adler [2002], Kvanvig [forthcoming], and Sutton [2007] for arguments that seem to show that assertion and belief are governed by the same types of norms.

² At least, this seems to be an implication if we think of epistemic justification as a deontological notion and think that you can at most hope to be excused for violating epistemic norms. Sutton [2007] thinks this is a consequence. Bird [2007], Huemer [2007], and Williamson [2000] do not think knowledge that p is necessary for having a justified belief that p even though they think that knowledge is the norm of belief. Littlejohn [forthcoming b] argues that a belief's justification depends, *inter alia*, on whether that belief conforms to the norms governing belief. In this paper, I shall speak as if conforming to the norm or norms of belief is necessary for a belief to be justified. If the reader does not accept this thesis, the reader can take my remarks to pertain primarily to the norms of belief. None of the arguments in this paper assumes that justified beliefs conform to the norms governing belief.

³ They might also be accommodated by views on which the norms of belief all have to do with evidence. Conee and Feldman [2004] appear to defend a view on which belief is governed by evidential norms but not the truth norm.

will dismiss the view out of hand. Dismissing the view in this way does not help us see where the arguments for this view go wrong.⁴ Since the view seems to be an obvious consequence of views now being seriously defended, the view deserves serious consideration.⁵ Hopefully internalists will be interested to see why the view is unmotivated and can add my complaints to theirs. The externalists can approach this discussion with an eye towards determining whether the knowledge account is preferable to alternative externalist accounts of justified belief.

The aims of this paper are modest. It is not my intention to offer a positive account of the norms of belief or analysis of justified belief.⁶ My aim is to show that when it comes to describing the norms of belief the traditional truth-first approach is preferable to the knowledge-first approach. We shall first look at some of the arguments inspired by Williamson's remarks concerning assertion and the aim of belief. Next, we shall examine the arguments of those who claim that reflection on Moore's Paradox leads to the conclusion that knowledge is the norm of belief (e.g., Adler [2002]; Huemer [2007]). If I am right, none of the arguments considered in this paper support the claim that belief is governed by the knowledge norm but only the considerably weaker claim that the fundamental norm of belief is the truth norm.

II. Truth and the Aim of Belief

The first argument works from the assumption that all beliefs have a common aim and that our description of the aim of belief has a kind of normative significance. To say that beliefs aim at X seems to carry with it the commitment to the further claim that any belief that doesn't fulfil this aim isn't as it ought to be.⁷ Suppose that belief really does aim at truth (Velleman [2000];

⁴ See Brewer [1999], Littlejohn [forthcoming a], and Sutton [2007] for three different externalist responses to the objection that intuition shows that the falsity of a belief is no obstacle to its justification. In this paper, I am concerned with the showing that the truth account is preferable to the knowledge account. Questions about the epistemic status of false beliefs are bracketed. If someone can defend K without saying that knowledge is necessary for justified belief, someone can defend T without saying that truth is necessary for justified belief. If the defenders of K can get away with denying the possibility of false, justified belief, so can the defenders of T.

⁵ In addition to arguments from the norms of assertion to the norms of belief, Hawthorne [2004] and Hawthorne and Stanley [forthcoming] have argued that it is epistemically improper to rely on the premise p in reasoning whenever you do not know that p is true. In Littlejohn [MS a], I defend the idea that the norms of belief determine whether the belief that p is 'fit' for providing premises for the purpose of practical deliberation.

⁶ I defend a positive account in Littlejohn [forthcoming b; MS b].

⁷ For a defence of the view that claims about the aims of belief should be understood in normative terms, see Shah [2003]. Vahid [2006] seems to think that it is a mistake to try to draw normative conclusions from claims about what is constitutive about belief as a mental state. Vahid might reject the argument for the knowledge account right from the outset, but as I think the argument

Wedgwood [2002]; Williams [1973]). Let us also suppose that it has no independent aim which a false belief might fulfil that would dispose us to say that the belief is correct, successful, or (objectively) as it should be. Supposing this, it seems natural to say that even when something good comes of believing something false, such goods couldn't *justify* believing something false.⁸ If the norms of belief can be derived from the proper description of the aims of belief, it seems that you should not ever believe p unless p .

Williamson [2000: 241] suggests you can derive the knowledge account of assertion from the assumption that you ought not assert what is not true. If the derivation works for assertion, it seems a parallel derivation should work for belief. If there is such a derivation, it is far from straightforward. Consider the utilitarian view that the only thing that could justify an action is the fact that the action is optimific. Utilitarians do not think an action must be known to be optimific to be justified and no one has ever faulted them for this omission. It is not hard to see that you are courting disaster if you advance a view on which there are positive duties that depend upon more than just effort or good will and insist that all duties must be knowingly discharged. You will end up having to say that there are unknowable obligations that can be obligations insofar as they are knowable.

The difficulty is not so acute if we assume that the duties in question are primarily negative duties. If belief is governed by a truth norm, the norm tells us to refrain from believing what is not true:

T: You should not believe that p is true unless it is true.

The norm does not instruct us to believe everything that is true or anything that crosses our minds that happens to be true.

To say that the fundamental norm of belief is the truth norm is not to say that this is the sole norm that governs belief. Williamson seems to think that you cannot consistently say that there is a norm such as T without also saying that belief is governed by an evidential norm:

E: You should not believe that p is true unless you have
adequate evidence for believing that p is true.

In general, it seems that if it is wrong to Φ if condition C obtains, it is irresponsible to Φ unless you can reasonably assume that C does not obtain. Whether it is reasonable to assume this depends upon whether you have adequate evidence for believing C not to obtain. So, if we

fails at a later step I hope to bracket this issue for the purposes of this discussion. Thanks to an anonymous referee for pressing me on this point.

⁸ Instead, it seems an excuse is called for. While we might say there is some justification for something that is not permissible, it seems that something cannot be justified unless it is permissible.

assume that it is *prima facie* wrong to Φ irresponsibly, we can say that if T is true and it is wrong to believe p when p is not true, it is impermissible to believe p without adequate evidence of p 's truth because it would be irresponsible to so believe.⁹

On the assumption that belief is governed by the truth norm, we have to reject:

FAULT-1: All epistemic wrongs are fault-implicating wrongs (i.e., any condition that makes believing p wrongful is a condition that the believer can be faulted for failing to take account of if she believes p when that condition obtains).

There is nothing wrong with rejecting this while accepting:

FAULT-2: Any condition that grounds the charge of epistemic fault is a condition that makes belief wrongful.¹⁰

It seems that Williamson's reason for thinking E is a consequence of T *might* be something along the lines of FAULT-2. He says:

[I]f one must not bury people when they are not dead, then one should not bury them when one lacks evidence that they are dead. It is at best negligent to bury someone without evidence that he is dead, even if he is in fact dead [2000: 245].

Let's suppose he is right. We do not have a full defence of the knowledge account:

K: You ought not believe p unless you know p .

We first have to determine what it takes to satisfy the evidential norm. Williamson [2000: 246] says that we do not satisfy the evidential norm unless we have evidence that puts us in a position to know that the belief in question is true. If we think about lottery propositions, it seems we do not have adequate evidence to believe or assert such propositions.¹¹ It seems the best explanation

⁹ This might be an overly complicated way of stating the obvious point that you should not be epistemically irresponsible and that assuming responsibility obligates refraining from believing without adequate evidence. The argument here is just restating Williamson's argument [2000: 245]. The success of this argument rests on the assumption that if one must refrain from Φ -ing unless C obtains, one should Φ only if one has evidence that C does obtain. If this assumption is unmotivated, then Williamson's argument that the knowledge account can be derived from the assumption that belief aims at the truth and is governed by the truth norm fails at an earlier step than I am claiming. Note that if knowledge is the norm of belief, this argument would show that it would be irresponsible to believe p unless you had adequate evidence for believing that your belief that p is true would constitute knowledge. If you think such evidence is unnecessary for permissible first-order belief, this is some reason to think that knowledge is not the norm of belief.

¹⁰ For an illuminating discussion of the relationship between wrongs and fault, see Gardner [2005].

¹¹ This is how Williamson explains the observation that you should not assert lottery propositions. It is not universally accepted that such assertions are epistemically improper. Weiner [2005]

as to why this rests on the observation that the evidence we have for believing lottery propositions without insider's information does not put us in a position to know that these claims are true.

There are two ways of reading Williamson's argument. On the first reading, his remarks concerning lottery propositions gives us a clue as to what he thinks it takes to satisfy FAULT-2. If you allow yourself to believe without first gathering evidence that puts you in a position to know that p , you seem to be at fault even if your belief turns out to be true. After all, you could have weakened your commitment by simply believing that p is likely or probable. On the second, we appeal directly to E and let intuition serve as our guide in determining what it takes to satisfy E rather than appeal to assumptions linking fault and justification.¹²

On the first reading, the argument amounts to this:

Belief aims at the truth. For this reason, you should not believe something unless it is true. If you do not know whether your belief is true, but you hold the belief anyway, this is irresponsible. We generally assume that if you do not know whether you would violate a strict prohibition by Φ -ing but Φ anyway, you are at fault for Φ -ing. It follows that you should not hold a belief about p if you do not know that p is true.

The argument rests on a further assumption about fault:

FAULT-3: You are irresponsible to Φ if you do not first know whether Φ -ing is permissible.

Without the assumption, you cannot derive anything stronger than the claim that you ought not Φ unless you have conformed to the truth norm and are not unreasonable to assume that you have given the evidence you have.

There are two reasons to reject FAULT-3. Here is the first. Suppose we say that knowledge is a condition necessary for permissible belief and also assume that FAULT-3 is correct. In saying this we would have to say either that mere knowledge of p 's truth is insufficient for permissibly believing p or we would have to endorse a KK thesis according to which you cannot know p unless you are in a position to know that you know that p is true.¹³ One of the more

argues that such assertions violate Gricean maxims which means he can say that there is adequate evidence for believing and asserting lottery propositions even assuming that the evidence needed for epistemically permissible assertion and belief are the same.

¹² Jonathan Sutton suggested this second reading in conversation.

¹³ If the point is not obvious, perhaps this helps. According to \mathbf{K} , it is not permissible to believe p unless you know p . According to Fault-3, it is not permissible to believe p unless you know that you satisfy the conditions necessary for permissibly believing and one such condition is that p is known. So, unless you thought that knowing that you know that p is a condition on having first-

persuasive objections to the KK thesis is this. It seems that knowledge of p 's truth requires that the means by which you arrived at the belief that p could not have easily led you to be mistaken about p . According to the KK thesis, you do not satisfy the conditions for first-order knowledge unless you are in a position to know that you satisfy these conditions. Second-order knowledge also requires that you would not easily be mistaken in the second-order belief (i.e., the belief that your first-order belief constitutes knowledge). Our ordinary knowledge ascriptions suggest that knowledge does not require being in a position to know that you know. We readily ascribe knowledge of p 's truth to someone knowing that she could have easily been mistaken in her second-order belief that she knew that p . We might know that the margin of error for second-order knowledge is slim but the margin of error for first-order knowledge is sufficiently wide so they couldn't have easily been mistaken about whether p but could have easily been mistaken about whether they knew p . Concerning such cases, not only does it seem we readily ascribe you knowledge, we do not think that it is wrong for you to believe p . This seems to disconfirm both the weak KK thesis and the thought that permissible belief involves more than just knowledge.¹⁴ If we reject both, however, we have to reject FAULT-3.

Here is the second reason to reject FAULT-3. If combined with the knowledge account, it commits us to the JTB analysis of knowledge and an infallibilist conception of justification. Epistemic irresponsibility can make an otherwise justifiable belief unjustified. According to FAULT-3, if you fail to know *for any reason*, you can be charged with epistemic irresponsibility. Thus, if you cannot be charged with epistemic irresponsibility because you are justified in believing p , the fact that you are justified in holding your belief is logically incompatible with (a) your belief being mistaken or (b) your belief being Gettiered. One consequence of this is that you cannot satisfy the justification condition if it is possible that someone should have just your reasons but be mistaken about whether p . Thus, your reasons must entail p if your belief that p is justified. But, no one seems to think that you must have entailing grounds to permissibly believe p . Second, it seems that FAULT-3 has the consequence that if someone does not know that p , they are not justified in believing p , in which case Gettier cases are impossible.¹⁵

order knowledge of p 's truth, you would have to say that knowledge of p 's truth is not sufficient for permissibly believing p . I think we can reject the idea that knowledge of p 's truth is insufficient for permissibly believing p on the grounds that if it were insufficient, the claim 'While p is true and he knows that it is, he should not believe it for epistemic reasons' could express a true proposition. It could not.

¹⁴ For further discussion, see Williamson's [2000: 145] remarks concerning iterations of knowledge and margins of error.

¹⁵ To bring the knowledge account in line with epistemic intuition, we should distinguish between a believer that is justified in believing p and a belief that is justified. For discussion of this

In light of these problems, I think an alternative reading of the argument might be more charitable. On this reading, the argumentative burden is shouldered not entirely by assumptions about fault and epistemic responsibility, but also by cases such as those involving lottery propositions.¹⁶ The assumptions about fault are supposed to support the idea that *some* sort of evidential norm governs belief. Our intuitions about lottery propositions are supposed to help us see what it takes to satisfy this evidential norm. We start from the assumption that you should not assert or believe lottery propositions. It is then suggested that the natural explanation for this is that you do not have evidence for believing these propositions that would put you in a position to know that they are true. From here, the argument might go in one of two directions. If someone said that it followed from this alone that you ought not believe what you do not know, this would repeat the mistakes we have just discussed. All that follows is that you ought not believe p if you are in a position to appreciate that your evidence does not put you in a position to know p . In Gettier cases and in cases where you do not know that p is a lottery proposition, it seems p is not known, you do not know that you are not in a position to know p , but it is *not* obvious that you have violated E. If you think you *do* violate E in such cases, it seems you will once again be forced to accept the JTB analysis of knowledge and an infallibilist account of justification. If, however, someone were argue as follows, these difficulties might be avoided:

Intuition tells us that evidence is needed for permissible belief. Intuitions about lottery cases tell us that beliefs in lottery propositions are defective because there is not adequate evidence for believing them outright.¹⁷ The best

distinction, see Bach [1985], Engel [1992], Littlejohn [forthcoming a], and Lowy [1978]. Because the knowledge account tells us that justified beliefs constitute knowledge, if you accept the knowledge account, you can say that Gettier cases are possible if you think of such cases as cases in which the subject's belief is unjustified but the subject is justified in believing p . However, for a *believer* to be justified on Bach's view, the believer must be epistemically responsible. According to FAULT-3, however, it is impossible for the believer to be justified if the believer does not know that what she believes is true. If you accept FAULT-3, you have to say that it is never possible for the believer to be justified if the belief does not constitute knowledge, which means that Gettier cases are impossible on this view.

¹⁶ I am only suggesting that *one* way to take Williamson's remarks about fault would be to offer the sort of argument offered above. In their defence of the view that you must never reason from premises that you do not know to be true, Hawthorne and Stanley [forthcoming] appeal to the observation that it is '*prima facie* negligent' to reason from p when one does not know that p . The qualification '*prima facie*' makes it difficult to interpret this remark, but I see no obvious connection between fault, negligence, and knowledge. If they are saying that whenever someone does not know that p but reasons from the belief that p their ignorance counts towards our saying that they are negligent, this seems to call into question the very coherence of the notion of non-culpable ignorance.

¹⁷ It is important to stress that the adequacy of evidence is not simply a matter of degree of strength. If one were to say that the reason beliefs in lottery propositions are defective is that the

*explanation of the observation that you do not satisfy E unless your evidence puts you in a position to know that the relevant proposition is true is that K is the fundamental norm of belief.*¹⁸

This argument rests on a pair of assumptions. First, the argument assumes that beliefs in lottery propositions do not constitute knowledge. Second, it assumes that you ought not believe lottery propositions. Were we to reject the first assumption, we could not appeal to intuitive verdicts about lottery cases to motivate the knowledge account. If we were to reject the second assumption while accepting the first, lottery cases would provide counterexamples to the knowledge account. In response to this argument, I want to say two things. The knowledge account cannot give the best explanation if independent considerations show that **K** is not a norm that governs belief. The verdicts the knowledge account delivers for covert lottery beliefs and for Gettier cases are counterintuitive. In addition, cases involving covert lottery beliefs suggest that while the knowledge account delivers the right verdict about some familiar lottery cases, it gives the wrong reason for thinking this is the right verdict. Intuitions concerning cases of covert lottery beliefs suggest that the reason we ought not believe lottery propositions is not that they cannot constitute knowledge *per se*, but that subjects that believe lottery propositions are wrong to do so in light of considerations accessible to them (i.e., considerations about the kinds of grounds they have for believing lottery propositions). While such grounds might not put the subject in a position to know, the normative significance of this is not what the knowledge account takes it to be.

evidence we have for such beliefs is not strong enough, it seems one would be committed to a view on which nothing short of certainty warrants outright belief. Presumably, advocates of the knowledge account think that the inadequacy of the evidence available for believing lottery propositions is that such evidence is of the wrong kind to warrant outright belief. What gives us warrant to believe mundane propositions (e.g., that I have hands) is that we have more than purely statistical grounds for such beliefs or, more carefully, that we do not have to take ourselves to have nothing beyond such grounds for such beliefs. Previously I had said that the evidence we have for belief in lottery propositions was ‘insufficient’, and this misleadingly suggested that what was needed for the right to believe a proposition is a ground for the belief stronger than that available for belief in a lottery proposition. Finally, it is important to remember, as an anonymous referee reminded me, that there are situations in which it seems proper for a subject to assert and believe propositions about the outcomes of lotteries (e.g., when the subject has insider’s information). For discussion, see DeRose [1996].

¹⁸ An anonymous referee suggested that this rationale might be useless for Williamson’s purposes. I do not disagree. I think it is worth considering this rationale because of the role it has played in recent defences of the knowledge account. Sutton [2007: 53] maintains that beliefs in lottery propositions fail to fulfil **E** and that this fact can only be explained by appeal to **K** on the grounds that the reason we are disposed to think that such beliefs fail to fulfil **E** is that they fail to constitute knowledge. I think this locates the explanation in the wrong place since the absence of knowledge is not accessible to the subject but the reason for refraining from believing in lottery propositions is. Also, I think we can use a variant of the standard lottery cases to directly challenge the claim that **K** governs belief much in the way that we have used Gettier cases to this effect.

The distinction between covert and overt lottery beliefs is a familiar one, but the terminology is not. Let us say that a *covert* lottery belief is a belief whose truth or falsity depends on the outcome of a lottery when the believer is not in a position to appreciate that this is so. Let us say that an *overt* lottery belief is a belief in a lottery proposition held by someone who has no insider's information. If you look at your bank statement and see that you are down to your last few dollars, you might reasonably believe that you will not be able to go on safari. If your mother has just purchased you a ticket for a lottery drawing being held later this afternoon without telling you, that belief is a covert lottery belief. Were you to believe that the ticket that your mother bought you will lose, that would be an overt lottery belief. (We are assuming that you know that you would be able to afford to go on safari if only you were to win the lottery drawing being held this afternoon.) It seems that overt and covert lottery beliefs will either both constitute knowledge or neither will. If you think that safety is necessary for knowledge, it will be just as easy for a covert lottery belief to turn out to be false as an overt one to turn out to be false. If you think that some suitably formulated closure principle holds true, someone will be in a position to know that a covert lottery belief is true only if this subject is in a position to know that an overt lottery belief is true. Assuming, as we are, that overt lottery beliefs fail to constitute knowledge, it seems we have two reasons for thinking that covert lottery beliefs similarly fail to constitute knowledge.

If this much is correct, the knowledge account commits us to saying that you should not hold or form covert lottery beliefs.¹⁹ I think this is bad news for the knowledge account. First, in defences of the knowledge account, the focus has been on the judgments that overt lottery beliefs should not be held and cannot constitute knowledge. No intuitive support has been offered to back the claim that neither type of lottery belief ought to be held. In fact, you might think that one of the reasons that the lottery paradox is so interesting is that we are not naturally disposed to think of covert lottery beliefs held by others as beliefs they should not continue to hold for reasons of which only we are aware (i.e., that unbeknownst to them the truth of their beliefs is contingent on the outcome of a lottery). Second, not only is the knowledge account's verdict about covert lottery cases not intuitive, it seems positively counterintuitive. To see this, consider a modified

¹⁹ If this much is *not* correct, I think this undermines the arguments from lottery propositions to the knowledge account. If someone were to say that overt lottery beliefs constituted knowledge, there would be no surprising facts for the knowledge account to explain. If someone were to say that such beliefs did not constitute knowledge but were permissibly held, we would have our counterexample to the knowledge account. If someone were to say that overt lottery beliefs fail to constitute knowledge and ought not be held, but held that covert lottery beliefs did constitute knowledge, this would bring the knowledge account in line with the intuition that it is not wrong to harbour covert lottery beliefs. However, that would require denying the closure principle and that safety is necessary for knowledge.

version of one of Hawthorne's examples. A friend writes you an email on Monday before a lottery is held, but you only read it Tuesday after the results of that lottery are known to you. It contains the following line of reasoning:

The ticket for tomorrow's lottery is a loser.

So if I keep the ticket I will get nothing.

But if I sell the ticket I will get a penny.

So, I'd better sell the ticket [2004: 29].

You know now that the first premise was not known to be true because of the grounds the subject had for that belief *and* know that the belief turned out to be true. Retrospectively, it seems you would agree with Hawthorne that this reasoning is unacceptable and would likely further agree that its unacceptability is due to the speaker's belief in the argument's first premise. Assuming that you should not hold beliefs that should not be trusted for the purposes of practical deliberation, we would arrive at the view that the speaker should not have held the first belief. Even without that assumption, you might agree that the subject should not have held the first belief regardless of whether it was fit to figure in practical deliberation. Now, suppose a different friend writes you an email on Monday before a lottery is held, but you only read it Tuesday after the results of that lottery are known to you. You had purchased this friend a ticket for this lottery without telling them, but now know that the ticket was a loser. They had written:

I want nothing more than to go on safari.

If I were to go on safari, I would want nothing more than to buy a new elephant gun.

The gun will be useless, however, since I cannot afford to go on safari.

So I guess I will use that money instead to do some repairs around the house.

The subject's belief in the first premise is known to you to be a covert lottery belief. The lottery was held and the ticket lost. You know this, so you know that the speaker's belief in the first premise was not known by the speaker to be true *and* that the speaker was in no position to appreciate this fact (i.e., it was an "unknown unknown" in Sutton's terminology). I think you would not take this reasoning to be unacceptable.²⁰ However, the knowledge account regards

²⁰ It is true that if you had responded in time, you might have alerted them to the fact that you had purchased them a ticket. However, in doing so the belief in the first premise would be an overt lottery belief rather than a covert lottery belief and it seems far more intuitive to say in that altered case that they should have reasoned differently than it does in the case as written that they

both instances of reasoning as unacceptable and takes them to be unacceptable for the very same reason. It says neither piece of reasoning is acceptable because both bits of reasoning involve crucially beliefs not known to be true. That the knowledge account delivers the wrong verdict in the case of covert lottery beliefs suggests that **K** does not govern belief. Additionally, it suggests that the knowledge account gives the wrong explanation for the unacceptability of the first bit of reasoning. The knowledge account seeks to explain the unacceptability of this reasoning in terms of a fact that is not accessible to the individual engaged in this bit of reasoning (i.e., that one of the beliefs involved in the reasoning is not known to be true). However, if overt and covert lottery beliefs have different normative statuses (i.e., one ought never hold overt lottery beliefs but may permissibly hold some covert lottery beliefs), it seems that the proper explanation as to why you should not reason from overt lottery beliefs should be given in terms of features distinctive of overt lottery beliefs (e.g., the kinds of ground available for overt lottery belief) rather than ignorance, *per se*.

We have not found a route from the truth norm or the thesis that belief aims at the truth to the knowledge norm. It is not for a lack of trying. Williamson is right that anyone who thinks there is a truth norm should think there is an additional evidential norm governing belief, but we know from Gettier that there is more to conforming to the knowledge norm than conforming to these two.²¹

III. Knowledge and the Aim of Belief

Rather than try to derive the knowledge account from the truth norm, we might try a different approach:

Belief does not aim at just the truth. Belief aims at knowledge. Any belief that fails to constitute knowledge is wrongful precisely because there is no distinct aim a belief serves that could potentially provide a justification for believing without knowing.

Since the argument assumes nothing about justification and fault, it should not face the problems the previous argument did.

should have reasoned differently or should not have believed the premises that figured in this piece of reasoning.

²¹ The considerations offered in support of the knowledge account might not show that you should not believe *p* unless you know *p*, but they might be taken to show that you should not believe *p* if you are in a position to appreciate that you cannot know *p* given your epistemic position. We shall later see that we do not need to assume that knowledge is the norm of belief to explain why this is an epistemic requirement on permissible belief.

The argument assumes that the aim of belief is knowledge (Bird [2007: 93]; Sutton [2007]; Williamson [2000: 48]). The problem with the argument is simple: knowledge is not the aim of belief. To test proposals about aims, we should consult our intuitions to determine what an external observer would say if she knew that another's belief fails to constitute knowledge. We know that belief aims at the truth, for example, because we know that if someone knows that someone else's belief about p is not true, this outside observer has sufficient warrant for asserting that this belief is incorrect or mistaken. On the hypothesis that belief also aims at knowledge, we should expect that those who know we don't know that p for any reason will be disposed to say we have made a mistake, we were wrong to believe what we did, or that we should suspend judgment. This is not what we find. If this is how we evaluate claims about the epistemic aim and the epistemic ought, we not only fail to find support for the knowledge account, we find evidence for denying that belief is governed by the knowledge norm. If beliefs that fail to conform to no norms are justified, we find evidence for denying that knowledge of p 's truth is necessary for the justification of the belief that p .

Suppose you think you saw a barn. You did, but you did not realize that you were in the land of fake barns. Because the hills were filled with convincing fakes, we do not think your belief constitutes knowledge. Knowing this, however, I do not think that your belief failed to fulfil its aim. Knowing that you do not know and why your belief is not knowledge, I would not be disposed to tell you are wrong to believe what you do or that you have made a mistake by believing that you see a barn. If a belief such as this does not miss its mark, nothing is left of the view that belief aims at knowledge. The fakes prevent your belief from fulfilling its aim only when they fool you into believing a fake barn is genuine.²² Moreover, if we set aside the question about aims and focus on the normative question, it seems that if someone said that you should not believe it is a barn knowing that your belief fails to constitute knowledge *simply* because the belief is Gettiered, it seems that they have made the mistake, not you. If that is right, there is nothing left of the view that knowledge is what is necessary for permissible belief.

In saying that it is not epistemically wrong to believe p if that belief has been Gettiered, it might seem I am denying something Reynolds [2002: 150] says in his discussion of Gettier cases and warranted assertion. He says the locals who know that you have been driving through fake

²² I suppose you *could* have as an aim knowing whether you saw a barn, but it doesn't seem you *must* have this as an aim insofar as you're engaged in the kind of deliberation that results in the belief that you saw a barn. I might also have as an aim believing things that make me popular or lead me to have interesting thoughts, but those aims, which I adopt, are not the aims I must have insofar as I'm reasoning towards a conclusion that we identify as a belief.

barn country would not say that you should believe you saw a barn.²³ They know that you were reasonable in holding this belief, but they know that given the grounds on which your belief is based, you did not have the power to distinguish fake from genuine barns.

This may be true, but the knowledge account is not necessary for explaining why the locals would not (and ought not) say it is permissible for you to believe you saw a barn. Reynolds says the locals do not know whether the particular belief you have formed is true. They only know that your grounds are not effective for determining whether your belief is true. Because of this, it would be wrong for them to assert that you should believe what you do because they do not have a true and reasonable belief that your belief satisfies the truth norm. Once we give the locals the additional piece of information that your belief is correct and they know that the sole reason you fail to know has to do with factors beyond those that determine whether you are justified or you are right, speakers are not disposed to think you should revise your beliefs and suspend judgment. Without this information, however, we cannot use their responses to evaluate the respective merits of the knowledge account or the weaker truth account.

We have covered a considerable amount of ground. Let's take stock. You cannot derive the knowledge account from descriptions of the aim of belief. If we start from only the assumption that belief aims at the truth, we get no further than the claim that you ought not believe p unless you faithfully and faultlessly represent the world by having that belief. That is to say, you should not believe p unless you satisfy T and E. Intuitions concerning Gettier cases and covert lotteries suggest that belief does not aim at knowledge. It does not seem that beliefs that fail to constitute knowledge for purely Gettierish reasons are incorrect, mistaken, erroneous, or fail to fulfil some essential aim. It also seems it's not necessarily the case that you ought not hold beliefs that fail to constitute knowledge for Gettierish reasons. The Gettier cases are

²³ There is a second argument in favour of the knowledge account of assertion worth considering. It is not uncommon to say things like 'If you didn't know where the restaurant was, you should not have said we should have taken that exit'. Why is it that learning that someone did not know p leads us to challenge the propriety of the assertion and belief expressed if knowledge is not the norm of assertion and belief? This is an interesting question. It is true that in *some* contexts when we learn that someone did not know p we say they had no right to assert p and say in a critical tone that they ought not have asserted p . I do not think, however, that we imply that someone lacked the authority *simply* because they did not know p . We might say that they should not have told us to get off the exit having learned that they were *mistaken* about the restaurant's location. We might say that they should not have told us to get off the exit having learned that they were *guessing* at the restaurant's location. I do not think we would be critical if we learned that someone who was correct in their belief that p and reasonable in assuming p but did not know p asserted p . It has been observed by Goldman [2002] that in some contexts, we use 'know' as something akin to firmly held true belief and I think these are contexts in which we use 'know' in such a way that considerations having to do with accidental connections to the truth are irrelevant.

counterexamples to not only the JTB analysis of knowledge, but also the analysis of permissible belief as knowledge.

IV. Moore's Paradox and the Norm of Belief

We shall now look at arguments that appeal to observations about Moore's Paradox in the hopes of motivating the knowledge account. Consider the statement 'Custer died at Little Big Horn, but I believe he did not'. It seems contradictory to assert this statement, but it easily could have been true. (I know next to nothing about the history of the United States.) What accounts for the *appearance* of contradiction in the absence of contradiction? One suggestion is that anyone who holds the beliefs associated with Moorean absurd statements holds beliefs that conflict with the rational commitments that come with those very beliefs (Adler [2002]; de Almeida [2001]; Huemer [2007]).²⁴ For example, it is thought that belief has as its aim the truth, and someone who holds the beliefs associated with 'Custer died at Little Big Horn, but I believe he did not' would be committed to denying the accuracy of the belief expressed by the first conjunct. That is akin to a contradiction. The appearance of contradiction is explained in terms of the conscious conflict between the beliefs of which the subject is aware and the rational commitments that come with the beliefs associated with the Moorean absurd statements (e.g., such as the fact that, by her lights, she has misrepresented how things stand by believing Custer died at Little Big Horn).

The arguments we are about to consider rest on a methodological assumption that I shall accept for the purposes of this discussion:

²⁴ What is a 'rational commitment'? In speaking of rational commitments, someone might have one of two notions in mind. On one reading, to say that by Φ -ing S has violated some rational commitment is to say that given further attitudes of S's, S violates what Broome [1999] refers to as a 'normative requirement', a requirement that can be stated by a statement where the 'ought' takes wide-scope and enjoins or forbids S to form or from holding certain attitudes in combination. For example, it seems irrational to believe that the world was created in under seven days while believing it must have taken at least eight. Anyone who holds such beliefs will violate the normative requirement that states, in effect, you must not: believe the world was created in under seven days and at least eight days. On another, to say that by Φ -ing S has violated some rational commitment is to say that S's Φ -ing violates some norm which can be expressed by a statement where the 'ought' takes narrow scope and enjoins or forbids S to form or refrain from forming attitudes under certain conditions. Think about the knowledge norm. It enjoins S to refrain from believing p if S does not know p . While it seems plausible to think that minimally rational agents can be charged with irrationality for violating normative requirements, it seems far less plausible to think that every failure to conform to a norm is less than fully rational. It seems more charitable to think of talk of rational requirements that come with belief in terms of normative requirements. The arguments for the knowledge norm, then, will be attempts to derive norms from normative requirements. In effect, we will be looking at attempts to show that if, say, there is a normative requirement that states 'You ought not both believe p and believe you do not know p ' there is a norm according to which 'If you do not know that p , you should not believe p '.

MA: Anyone who holds the beliefs associated with a Moorean absurd statement holds beliefs that conflict with the rational commitments that come with those very beliefs.²⁵

The first argument tries to establish that knowledge is the norm of belief on the basis of two assumptions, one of which is undeniable and one of which is thought to follow from MA. According to the second, our intuitions will lead us to classify types of statements as Moorean absurdities that could only have that status if knowledge is in fact the norm of belief. According to the third, alternatives to the knowledge account are too weak to explain why certain kinds of Moorean absurdities have that status.

Huemer [2007] defends the thesis that knowledge is the norm of belief as follows:

*Consciously believing p rationally commits you, upon reflection, to comprehensively, epistemically endorsing your belief that p (MCP). Knowledge attribution is the most comprehensive epistemic endorsement (ETK). If you believed p it would be wrong to have that belief without endorsing it as knowledge. But, you should not endorse that belief as knowledge unless it is knowledge. Therefore, you ought not believe p unless you know p .*²⁶

In defence of MCP, remember that MA tells us that whenever someone utters a Moorean absurdity, they have uttered something absurd because the beliefs associated with that statement conflict with the rational commitments that come with those beliefs and so those beliefs cannot be comprehensively endorsed.²⁷ As for ETK, there seems to be no more comprehensive epistemic endorsement of a belief than one that says the belief constitutes knowledge.²⁸

²⁵ Those who think that Moorean absurdity has everything to do with speech and nothing to do with thought will reject the argument for the knowledge account at a stage earlier than I will by rejecting the methodological assumption central to this line of argument. For arguments that the paradox is *not* about speech alone, see Shoemaker [1996].

²⁶ Huemer [2007] defends both the view that you should not believe what you take yourself not to know (i.e., that there is a normative requirement to the effect that you should not believe both that p is true and that p is not known to you) and that there is a sense in which a belief that fails to constitute knowledge for any reason is epistemically not as it ought to be (i.e., that there is a norm to the effect that you should not believe p unless you know p to be true).

²⁷ An anonymous referee suggested that Huemer's [2007] MCP is not sufficiently different from Foley's [1993] account of egocentric epistemic rationality to avoid the difficulties that Foley's view faces. This might be right, but even if there is nothing wrong with this approach to epistemic rationality it lends no support to the knowledge account. So, I hope to bracket more general concerns about either of these authors' approaches to epistemic rationality.

²⁸ An anonymous referee suggested that it was unclear how Huemer's [2007] account could explain the absurdity of assertions such as, "You know that it is raining but I do not agree with you". I think he could do so as follows. The speaker knows that by asserting that another knows

Suppose we were to grant ETK. Here is an initial worry about Huemer's strategy. His argument implicitly assumes that each of the conditions that figure in a comprehensive epistemic evaluation pertains to the permissibility of belief.²⁹ This does not seem right. Consider moral evaluation. A moral evaluation that focused on just the permissibility of some action would *not* be a comprehensive evaluation. A comprehensive evaluation of an action should not only tell us whether the act was permissible, but also whether the act had moral worth. That my act lacks moral worth does not show that I have acted impermissibly and so does not show that my act lacked justification. Similarly, that my action was not supererogatory similarly does not show that my action calls for a justification. Huemer offers no reason to think that a comprehensive evaluation of belief would concern only a belief's deontic status. Thus, even if we assume ETK, it would not be surprising if some of the conditions necessary for knowledge were not necessary for permissible belief. We need some reason for thinking that a comprehensive epistemic endorsement is only concerned with properties that are of deontic significance, and he gives us no such reason. We have already seen some reason to think that the permissibility of believing *p* does not turn on that belief's being properly endorsed along all lines of epistemic evaluation. Suppose you thought that if a belief constitutes knowledge, that belief is more valuable from the epistemic point of view than a belief that fails to so constitute knowledge. Along one line of evaluation, a true belief someone is justified in holding might be less valuable than an item of knowledge. (Think of Gettier cases and covert lottery beliefs.) The additional value that attaches to beliefs that constitute knowledge would only be necessary for permissible belief if we were to say that you should never harbour covert lottery beliefs or hold beliefs in Gettier cases. So, we either have to say that a comprehensive epistemic endorsement concerns *more* than just that which bears on the permissibility of belief or deny the evaluative claim that items of knowledge are epistemically more valuable than beliefs that fail to do so.³⁰ If a comprehensive epistemic endorsement

p, the speaker is committed to, *inter alia*, the truth of *p*. By asserting that the speaker does not agree with what the other subject believes, the speaker is asserting that the speaker does not believe *p*. Thus, this case reduces to the omissive version of Moore's Paradox, "It is raining out but I do not believe it". I do not see why Huemer cannot offer the same treatment of the omissive version of Moore's Paradox that I do below.

²⁹ Huemer [2007: 149] says that a comprehensive epistemic endorsement is one that endorses a belief along all dimensions of epistemic evaluation and adds that any belief that cannot be so endorsed is such that there is no reason that the subject should not hold the belief.

³⁰ Many assume that knowledge is epistemically more valuable than mere true belief. For discussion, see Kvanvig [2003]. The assumption here is slightly different. I am assuming that a subject's having knowledge that something is true is more valuable than a proper subset of the 'parts' of this subject's knowledge. For a helpful discussion of strategies designed to explain why this is so, see Pritchard [2007].

concerns more than just those properties of a belief that are of deontic relevance, a belief might not be one we can endorse and not amount to knowledge even if it is permissibly held.

Maybe this worry is relatively minor. This next worry is much more serious. While the knowledge account does seem to follow from MCP and ETK, you *also* get the result that you should not believe p unless you know that you know that you know (etc...) that p . Clearly, the argument needs revision. It is not hard to find the needed fix. We simply have to rewrite MCP and say that consciously believing p requires that you should not believe yourself to fail to satisfy the standards of a comprehensive epistemic endorsement while also consciously believing p . This revision is independently motivated. It is simply not true that if you permissibly believe p , you ought to believe that your belief satisfies a comprehensive epistemic endorsement. Failing to have any belief about whether you know p when you happen to believe p (and happen not to be wrong to do so) is no sin at all. What rationality requires is that you revise your beliefs if you believe that you cannot comprehensively endorse them (i.e., by judging that they are false, that they don't amount to knowledge, etc...).

We should replace MCP with:

MCP-: Consciously believing p rationally commits you, upon reflection, to refrain from believing both p and that p cannot satisfy the standards of a comprehensive epistemic endorsement.³¹

Once we replace MCP with this, we have undermined the argument for the knowledge account. According to ETK and MCP-, all that permissibly believing p requires is that you do not both believe that you do not know that p while holding the belief that p , and that is a requirement you could easily satisfy even if you did not know that p . All it takes to satisfy this requirement is not conceding that you do not know p .

Adler [2002] offers a similar argument for the knowledge account, but it seems not to suffer from the same difficulties. He thinks we can use our intuitive sense of which combinations of attitudes would constitute Moorean absurd combinations to tell us something about the norms governing those attitudes and is as follows:

³¹ MCP- should do for the purposes of a normative solution to Moore's Paradox. Remember that the justification for MCP was the thought that someone who would believe the beliefs associated with a Moorean absurd statement would be conscious of violating some fundamental norm of belief, not that they hadn't bothered to affirm that every proposition they believed was in fact as it should be. Violations of MCP- are sins of commission, whereas violations of MCP, such as they are, can be sins of omission. As noted above, if you simply fail to believe you know p and believe p , there is nothing irrational, odd, contradictory, or untoward about that.

According to MA, if any judgment that expresses the belief that p is coupled with the acknowledgement that some condition C does not obtain, you are cognizant that, by holding these beliefs, you violate rational requirements on holding those beliefs. The judgment expressed by ‘ p but I do not know it’ is incoherent in the way Moorean absurdities are. Therefore, in judging that you do not know p , you know there is something wrongful in believing p —namely, that you do not know p .

This is an advance because the argument does not assume the Metacoherence Principle. Instead, it relies on the incoherence test:

IT: If it is incoherent in the way Moorean absurd statements are to Φ while acknowledging that C does not obtain, in acknowledging that C does not obtain, you are cognizant of something that makes it wrong to Φ —namely, that C does not obtain.

Given our methodological assumption concerning the proper resolution of Moore’s Paradox (i.e., that Moorean absurdity is due to the way in which a subject would consciously violate the rational requirements by consciously affirming the attitudes that constitute the absurd thought), it seems we can test proposals concerning the norms of belief as follows. If it is incoherent to simultaneously believe p while believing C does not obtain in the way a Moorean absurd thought is (i.e., apparently contradictory without being a belief in a contradiction), there is a norm that enjoins us to refrain from believing p when C does not obtain. For example, it seems incoherent to believe the following:

(1) I believe Custer died at Little Big Horn, but he did not.

In representing the belief about Custer as being false, it seems we have an incoherent combination of attitudes without a contradiction. Perhaps this is due to the fact that belief is governed by the truth norm. It seems similarly incoherent to believe:

(2) I believe Custer died at Little Big Horn, but there is no reason for me to think that.

On the assumption that the only thing that could be a reason for me *to* believe is a piece of evidence, it seems we can infer from the fact that (2) is a Moorean absurd thought that the belief about Custer is governed by an evidential norm. It has been observed that it is incoherent to believe both that p is true and that this belief fails to satisfy one of the conditions necessary for knowledge. It is also incoherent to believe that p but that p is not known to be true. So, by similar reasoning, it seems that we ought to accept that knowledge is the norm of belief.

Unfortunately, this test is insufficiently discriminating and might be useful only for uncovering normative requirements governing combinations of belief rather than useful for uncovering the norms that govern those beliefs individually. To appreciate the first problem, consider an example:

(3) God hates my atheism and it is raining outside.³²

This is a Moorean absurd thought. However, there is no norm that enjoins us to refrain from believing that it is raining outside unless God forgives the non-believer. It is no mystery as to why (3) is incoherent. It is incoherent because the belief that God hates my atheism is a Moorean absurdity in its own right. At the very least, IT needs to be reformulated to avoid these sorts of example. Even if we grant that for every incoherent pair of attitudes there is *something* you are cognizant of that makes one of the attitudes you are conscious of wrongful, we could say what makes it wrongful is *precisely* that it is held in combination with the belief that C does not obtain. It might be that this belief alone is absurd, irrational, or contravenes an epistemic norm. We might say, as it were, that whenever believing p while believing C does not obtain constitutes a Moorean absurdity, all that follows is that:

(4) You should not believe: p and that C does not obtain.

That is different from:

(5) If C does not obtain you should not believe p .

The former is a normative requirement and the ‘ought’ takes wide scope. The latter is a norm in which the ‘ought’ takes narrow scope. The former tells us what combinations of attitudes we ought to avoid. The latter tells us what sorts of conditions bear on whether to hold the belief in question. As we are trying to derive norms such as the knowledge *norm* (i.e., if you do not know p you must not believe p) from judgments about rational combinations of attitudes (i.e., it is irrational to believe both that p is true and not known to be true), we need some reason to think that we can proceed from intuitive judgments about irrational combinations of attitude to judgments about attitudes we have reason to refrain from holding when certain non-mental conditions obtain (i.e., that we have reason to refrain from believing falsehoods or those beliefs not known to be true even when we have no clue that our beliefs are false or fail to constitute knowledge).

Maybe these problems are not insuperable. To deal with the first, we can revise IT as follows:

IT2: If it is incoherent in the way Moorean absurd statements are to Φ while acknowledging that C does not obtain *and* the

³² The example is inspired by one of Sorensen’s [1988].

belief that C does not obtain is not itself incoherent, in acknowledging that C does not obtain, you are cognizant of something that makes it wrong to Φ —namely, that C does not obtain.

To deal with the second and more fundamental problem, we might say this. The reason that it is irrational to believe both that p is true and that C does not obtain is that in representing your present situation as one in which C does not obtain you thereby appreciate that if *that* belief is correct, you should expect there to be reason not to believe p . Moreover, if that belief is incorrect it is still by your lights a situation where there is reason not to believe p . To believe against what you take to be good reasons is itself a kind of epistemic wrong.³³

Perhaps this suffices to address the difficulties that arose for IT, but while we might have saved the test we have revised the test in such a way that it no longer lends support to the knowledge account. To appreciate this, note the following. The test only applies when the belief that C does not obtain is a belief that is not incoherent taken on its own. While we might grant that if C is a condition necessary for knowledge it is incoherent to believe both that p is true and that C does not obtain, we only find confirmation of the knowledge account if we assume *also* that for any condition C such that C is a condition necessary for knowledge it is *coherent* to believe on its own that C does not obtain. This is not what we find.

Let me explain. It is not incoherent to believe that you do not know p because p is false. So, according to IT2 there is a norm that enjoins us to refrain from believing the false. It is not incoherent to believe that you do not know p because your evidence does not put you in a position to know whether p . So, according to IT2, there is a norm that enjoins us to refrain from believing without evidence.³⁴ What of the conditions beyond this necessary for knowledge? So far, we have only confirmed that you should not believe the false and not believe without evidence. The belief that my belief about p is Gettiered, like the belief that God will not forgive my atheism, is incoherent taken on its own. There are many ways to Gettier a belief, but I shall focus on two. In the first sort of case, your evidence for believing p is undermined thanks to true propositions of which you are unaware. In the second sort of case, your evidence for believing p only accidentally leads you to the correct judgment concerning p .

³³ The principle that you have some reason to refrain from believing what you not unreasonably take yourself to have reason to refrain from believing seems intuitively more plausible than the principle that allows for bootstrapping and posits a reason *to* believe associated with any not unreasonable judgment that you have reason to believe.

³⁴ Adler [2002], de Almeida [2001], and Williams [1994] offer explanations along these lines and do not rely on the knowledge norm to explain the absurdity of the judgments associated with ‘ p but I do not believe p ’, ‘I believe p , but $\sim p$ ’, and ‘ p but there is no reason to believe p ’.

It is reasonably clear why you cannot coherently and correctly believe yourself to be in the first sort of case. To believe yourself to be in such a case, you have to believe of some piece of evidence (i) that you are unaware of it and (ii) that it would, if combined with your present evidence, undermine the justification you have for believing p . However, you would have this evidence in mind if you were to believe this. The case is not possible.³⁵ If you were aware of the evidence you would no longer be justified in believing what you did. This means this would not be a Gettier case.

What about the second sort of Gettier case in which your failure to know that p is due to the accidental connection between your grounds and the truth? You cannot coherently believe that you fail to know that p simply because you are only accidentally related to the truth about p . To believe that your connection to the truth is accidental is (roughly) to believe that if p turns out to be true, this is not to be expected. The judgment that you are in this sort of Gettier case amounts to the complex judgment that p is true but that you are not in an epistemic position to expect that it is true. This is a Moorean absurdity in its own right. We can account for its absurdity by noting that someone with such attitudes would take their belief that p is true to fail to fulfil the evidential norm since what is to be expected is a function of the evidence available. The subject's taking it to be the case that the truth of the relevant belief is not to be expected is acknowledging that by her lights she cannot expect to be right given her evidence. IT2 fails to confirm that this condition is necessary for permissible belief.

The challenging case for those who do not think knowledge is the norm of belief is this one:

(6) p but I do not know p .³⁶

³⁵ Michael Huemer suggested that I might be stacking the deck in describing the case as one in which you are both aware and unaware of a piece of evidence. Perhaps we should ask whether you might coherently believe that, say, you see a barn but there is evidence of which you're unaware which would undermine the evidence you now have for that belief. Note that if you had a belief about your evidence to this effect, your evidence would be undermined. This would not be a Gettier case because you would not be justified in believing that you saw a barn. It is impossible to coherently and correctly believe that you fail to know because you are in the first sort of Gettier case, in which case the incoherence test fails to confirm that not being in such a case is necessary for permissibly believing p .

³⁶ DeRose [2002] for example seems to think that assertions of the form of (6) are just the sort of case we need the knowledge norm of assertion to address. His solution is that the assertion of ' p ' represents the speaker as knowing p although what the assertion expresses is correct iff p . While it seems this account might have the resources for explaining why asserting (6) is improper, it does not seem to give us any explanation as to why the judgment that (6) is true is also incoherent or absurd.

If you take yourself not to know that p because you take yourself not to believe p , say, then we have a situation where your second-order belief is transparently falsified by a fact about your own mind. If you take yourself not to know that p because you take it to be that p is false, we can explain the incoherence by appeal to the truth norm. If you take yourself not to know that p because you have insufficient evidence by your own lights, we can explain this in terms of the evidential norm. That norm in turn can be derived from the truth norm in just the way Williamson suggested earlier. And, if you take yourself not to know because you take yourself to be in a Gettier case, we have already seen how the evidential norm can explain the incoherence of that attitude. It seems that we have our bases covered.

If we were to assume that the solution of Moore's Paradox should be given in normative terms and assume that MA is true, while the knowledge account might have the resources to explain the absurdity of the thoughts we have thus considered it seems the knowledge account is unnecessary for explaining why Moorean absurd thoughts strike us as contradictory. It seems we can explain the same data using either the truth norm combined with the evidential norm or the evidential norm taken on its own. So, while some might think that MA is a dubious assumption and dispense with the very idea of using Moore's Paradox as a way of uncovering the norms of belief, MA properly understood does not lead us to endorse the idea that the reason that the concept of knowledge is significant is that it plays an essential role in the formulation of the norms of belief.

V. Conclusion

Because people want to draw parallels between assertion and belief and because the knowledge account of assertion is so popular, it is not surprising that people are taking seriously the suggestion that knowledge is the norm of belief. Those who are willing to seriously entertain the suggestion that knowledge is necessary for permissible belief should be willing to seriously entertain the suggestion that permissible belief is a matter of faultlessly and faithfully representing the world. This view, which takes the truth norm to be the fundamental norm of belief, has no less intuitive support than the knowledge account. It does no worse than the knowledge account does in addressing Moore's Paradox. As an added bonus, it is motivated by reasonably uncontroversial claims about the aim of belief and does not deliver the wrong verdicts concerning cases of covert lottery beliefs and Gettier cases.

There is something puzzling about the knowledge account. It is not hard to understand why the internal conditions necessary for knowledge could be taken to have normative significance. Someone who believes p without satisfying the internal conditions necessary for

knowledge has failed to exercise proper responsibility over her beliefs. It is not difficult to understand why the truth of a belief could be taken to have normative significance. What a belief is supposed to do is offer an accurate picture of the world, so any inaccurate beliefs is thereby defective. It seems that in offering a justification for belief what the justification is supposed to do is address charges that the belief is wrongful, which is a matter of showing that the belief is not defective and showing that you have not failed to exercise proper authority in coming to hold it.³⁷ The further conditions that distinguish a belief that faithfully and faultlessly represents the world from a belief that constitutes knowledge does not intuitively have much by way of normative significance. That you have faultlessly but accidentally managed to get what you want suggests that you have fulfilled your goals and done so in a way that does not speak badly of you. How could that not be enough for permissible belief?

If we look at some of the traditional moral doctrines, such as those associated with the doctrine of double effect, there is this thought that there is always two ways to do the wrong thing, either by bringing about that which ought to be avoided or by carrying on in a way that shows a lack of proper respect for the good.³⁸ The truth account captures this idea. We take as our foundational premise the suggestion that the norm of belief is truth. Because of this, you ought not believe without adequate evidence for believing without such evidence shows a lack of proper respect for the good. Because of this, you ought not believe what is false because such beliefs fail to promote what matters fundamentally. It is hard to grasp the suggestion that there is always some third way of going wrong. Perhaps you will agree that there is something to be said for the hypothesis that justified belief just is a belief that faithfully, faultlessly represents the world. It might not be much, but it might be more than there is to be said for the idea that only knowledge can earn you the right to believe.³⁹

³⁷ This is not the time to address those with ‘responsibilist’ sympathies who think that the justification of a belief cannot be threatened by conditions the subject cannot be held responsible for failing to take account of. Someone could say that while belief is governed by the truth norm, the mere failure to conform to that norm does not show that a belief is unjustified. Following Bird [2007], someone might say that a belief fails to be justified only if a subject’s failure to conform to the norms governing belief is something for which the subject can properly be held accountable. This might enable them to say that belief is governed by the truth norm while maintaining that the notion of justified belief is an internalist notion.

³⁸ Mark van Roojen thinks Philippa Foot might have said something to this effect. If she did not, someone should.

³⁹ I would like to thank Matt Burstein, John Gibbons, Sandy Goldberg, Mark van Roojen, Jonathan Sutton, as well as members of the SMU philosophy department works in progress group and the audience members at the CSPA for discussing the issues addressed in this paper. I also would like to thank Robert Howell, Mike Huemer, Todd Stuart, and two anonymous referees for this journal for incredibly useful written comments. Robert Howell and an anonymous referee for

References

- Adler, Jonathan 2002. *Belief's Own Ethics*, Cambridge, MA: MIT University Press.
- Bach, Kent 1985. A Rationale for Reliabilism, *The Monist* 68/2: 246-63.
- Bird, Alexander 2007. Justified Judging, *Philosophy and Phenomenological Research* 74/1: 81-110.
- Brewer, Bill 1999. *Perception and Reason*, New York: Oxford University Press.
- Broome, John 1999. Normative Requirements, *Ratio* 12/4: 398-419.
- Conee, Earl and Richard Feldman 2004. *Evidentialism*, New York: Oxford University Press.
- de Almeida, Claudio 2001. What Moore's Paradox is About, *Philosophy and Phenomenological Research* 62/1: 33-58.
- DeRose, Keith 1996. Knowledge, Assertion, and Lotteries, *Australasian Journal of Philosophy* 74/4: 568-80.
- DeRose, Keith 2002. Assertion, Knowledge, and Context, *Philosophical Review* 111/2: 167-203.
- Engel, Mylan 1992. Personal and Doxastic Justification in Epistemology. *Philosophical Studies* 67/2: 133-51.
- Foley, Richard 1993. *Working Without a Net*, New York: Oxford University Press.
- Gardner, John 2005. Wrongs and Faults, in *Appraising Strict Liability*, ed. Andrew Simester, New York: Oxford University Press: 51-80.
- Goldman, Alvin 2002. *Pathways to Knowledge: Private and Public*, New York: Oxford University Press.
- Hawthorne, John 2004. *Knowledge and Lotteries*, New York: Oxford University Press.
- Hawthorne, John and Jason Stanley. Forthcoming. Knowledge and Action, *Journal of Philosophy*.
- Huemer, Michael 2007. Moore's Paradox and the Norm of Belief, in *Themes from G.E. Moore*, ed. Susan Nuccetelli and Gary Seay, New York: Oxford University Press: 142-58.
- Kvanvig, Jonathan 2003. *The Value of Knowledge and the Pursuit of Understanding*, New York: Cambridge University Press.
- Kvanvig, Jonathan. Forthcoming. Knowledge, Assertion, and Lotteries, in *Williamson and his Critics*, New York: Oxford University Press.
- Littlejohn, Clayton. Forthcoming a. The Externalist's Demon, *Canadian Journal of Philosophy*.
- Littlejohn, Clayton. Forthcoming b. Epistemic Reasons and the Justification of Belief, in *Reasons for Belief*, ed. Andrew Reisner and Asbjorn Steglich-Petersen, Synthese Library.
- Littlejohn, Clayton. MS a. Belief's Aim and its Justification.

this journal were kind enough to read multiple versions of this paper and offer dozens of pages in comments between the two of them. I am exceptionally grateful to them for their hard work.

- Littlejohn, Clayton. MS b. The Myth of the False, Justified Belief.
- Lowy, Catherine 1978. Gettier's Notion of Justification, *Mind* 87/345: 105-8.
- Pritchard, Duncan 2007. Recent Work on Epistemic Value, *American Philosophical Quarterly* 44/2: 85-110.
- Reynolds, Stephen 2002. Testimony, Knowledge, and Epistemic Goals. *Philosophical Studies* 110/2: 139-61.
- Slote, Michael 1979. Assertion and Belief, in *Papers on Language and Logic*, ed. Jonathan Dancy, Keele: Keele University Library: 177-90.
- Sorensen, Roy 1988. *Blindspots*, New York: Oxford University Press.
- Shah, Nishi 2003. How Truth Governs Belief, *Philosophical Review* 112/4: 447-82.
- Shoemaker, Sydney 1996. *The First-Person Perspective and Other Essays*, New York: Cambridge University Press.
- Sutton, Jonathan 2007. *Without Justification*, Cambridge, MA: MIT University Press.
- Unger, Peter 1975. *Ignorance: A Case for Scepticism*, New York: Oxford University Press.
- Vahid, Hamid 2006. Aiming at Truth: Doxastic vs. Epistemic Goals, *Philosophical Studies* 131/2: 303-35.
- Velleman, David 2000. *The Possibility of Practical Reason*, New York: Oxford University Press.
- Wedgwood, Ralph 2002. The Aim of Belief, *Philosophical Perspectives* 16: 267-97.
- Weiner, Matt 2005. Must We Know What We Say? *Philosophical Review* 114/2: 227-51.
- Williams, Bernard 1973. *Problems of the Self*, New York: Cambridge University Press.
- Williams, Jon 1994. Moorean Absurdity and the Intentional 'Structure' of Assertion, *Analysis* 54/3: 160-66.
- Williamson, Timothy 2000. *Knowledge and its Limits*, New York: Oxford University Press.