

PROCEEDINGS OF THE ARISTOTELIAN SOCIETY

ISSUE 3 | VOLUME CXIII | 2012 - 2013

The Russellian Retreat

CLAYTON LITTLEJOHN
KING'S COLLEGE LONDON



PROCEEDINGS OF THE ARISTOTELIAN SOCIETY
134TH SESSION

ISSUE NO. 3
VOLUME CXIII
2012 - 2013

THE RUSSELLIAN RETREAT

CLAYTON LITTLEJOHN
KING'S COLLEGE LONDON

MONDAY, 20 MAY 2013

17.30 - 19.15

THE WOBURN SUITE
SENATE HOUSE
UNIVERSITY OF LONDON
MALET STREET
LONDON WC1E 7HU
UNITED KINGDOM

*This event is catered, free of charge, &
open to the general public*

CONTACT
mail@aristoteliansociety.org.uk
www.aristoteliansociety.org.uk

B I O G R A P H Y

Clayton Littlejohn is Lecturer in the Department of Philosophy at King's College London. His publications include *Justification and the Truth-Connection* (Cambridge University Press, 2012), *This is Epistemology* (Wiley, Forthcoming), and *Epistemic Norms*, edited with John Turri (Oxford University Press, Forthcoming). His current research concerns the relation between theoretical and practical reason.

E D I T O R I A L N O T E

The following paper is a draft version that can only be cited with the author's permission. The final paper will be published in *Proceedings of the Aristotelian Society*, Issue No. 3, Volume CXIII (2013). Please visit the Society's website for subscription information: www.aristoteliansociety.org.uk.

THE RUSSELLIAN RETREAT

CLAYTON LITTLEJOHN

A standard approach to epistemic normativity starts from the idea that belief aims at the truth. On this truth-first approach, all epistemic norms are thought to be grounded in the norm of truth. I shall argue that this approach cannot explain some important features of epistemic assessment. One of the virtues of the knowledge-first approach to epistemic normativity is that it can explain why epistemic assessment has the inward looking character that it does.

A commonly held view about the relation between belief and truth is that the one aims at the other.¹ This talk of aims is metaphorical. The best way to interpret this metaphor is in normative terms. If your beliefs don't fit the facts, they are defective. They shouldn't be like that and you shouldn't have beliefs like that.

I haven't said much, but suppose that everything I've said is true. Is it important? Some think that it is very important. The dominant view in current epistemology seems to be that the fundamental epistemic norm is a truth-norm. This norm grounds all the other epistemic norms and explains why epistemic assessment has the concerns that it does.

The truth-first approach to epistemic normativity is undeniably attractive. It seems rather plausible that a belief is correct only if it is true. It also seems rather plausible that it is not a brute fact that epistemic assessment is concerned with facts about your evidence or the way you reason. As Michael Lynch puts it, "we take it to be correct to believe what is based on evidence *because* beliefs based on evidence are likely to be *true*, and thus the value of truth ... is more basic than the value of believing what is based on evidence".² The trouble with the truth-first approach is not that it commits you to saying lots of false things or prevents you from saying lots of true things. The problem is that it is wrong about one very important thing.

If the correct approach to epistemic normativity starts from the idea that the fundamental epistemic norm has to do with truth, it seems that we might be able to say everything that needs to be said about epistemic normativity without saying anything at all about knowledge. Although the view isn't universally held, it's widely thought that knowledge has no

¹ See Boghossian (2008), Littlejohn (2012), Millar (2004), Shah (2006), Shah and Velleman (2005), Wedgwood (2002b), Williams (1973), and Whiting (2012).

² Lynch (2009: 229).

deontic significance. There is no duty to know and no duty not to believe what you don't know. Richard Foley has long defended the view that there are two fundamental questions in epistemology. One question is a question about what we ought to believe. The other is a question about what we can know. He thinks that the failure of the Cartesian project shows that these questions have to be addressed independently. Since there is no method or procedure we can use that's guaranteed to provide us with knowledge, facts about what we can know tell us little if anything about what we should believe.³

Crispin Wright once suggested that we could live with the concession that we don't know some of the things that we believe provided that we can still say that we're justified in holding these beliefs.⁴ This is the 'Russellian Retreat', a stance Russell and Wright think you should adopt once you recognize that you have to settle for 'probability, defeasibility, and inconclusive justifications'. I think this is a mistake that's symptomatic of the failure to appreciate the normative significance of knowledge. None of us can take comfort in the thought that we're justified in holding our beliefs once we've been forced to concede that we shouldn't believe what we do, but this is just what this Russellian Retreat amounts to.

I.

The truth-first view says that the fundamental epistemic norm is a truth-norm. There are various ways of formulating such norms, but we shall focus on this one:

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

So stated, the truth-norm states only a prohibition. It tells you that you shouldn't believe falsehoods, not that you should believe truths. So far as this norm is concerned, we might have no positive epistemic duties to believe anything at all. I think this is a virtue of the present formulation. I don't think there are positive epistemic duties to believe. It also doesn't say whether it's permissible to believe anything at all. I also think that this is a virtue of the present formulation. It's plausible that you're permitted to believe any proposition you entertain, so long as you don't violate any epistemic norms. We don't need a norm to tell us what we may believe. Once we see that we've satisfied the norms that govern belief, we know we're in the clear.

Critics have criticized the truth-first approach to epistemic normativity on the grounds that the truth-norm is both too restrictive

³ See Foley (2001: 13; 2012: 127).

⁴ Wright (1991: 88).

too permissive. Those who claim that it is too restrictive argue that there is simply no obligation to refrain from believing falsehoods.⁵ Those who argue that it is too permissive think that it's a problem for the truth-first approach that there's more to meeting your epistemic obligations than simply fitting your beliefs to the facts.⁶ I don't think there's much to the first objection and I'll briefly explain why. Most of our discussion will focus on the second objection.

Some epistemologists find both objections compelling. They tend to embrace a kind of evidentialist view according to which the fundamental epistemic norm has to do with relations of 'proper fit' between your evidence and your beliefs. Feldman claims that things are not going terribly well for you if you irrationally believe lots of true propositions. He's right about that. It would be a mistake, however, to adopt his reductive approach to epistemic normativity and hold that the only normatively significant relations hold between your beliefs and your evidence for them. His proposal runs into two problems. The first has to do with the having-relation, the relation between you and your evidence. This relation has to be understood in normative terms. If you have something as part of your evidence, you have the right to rely on it as evidence. These rights do not always arise from relations of proper fit between your evidence and your beliefs. Whether you acquire a piece of evidence non-inferentially depends upon whether you're properly related to the facts and your evidence does not determine whether you stand in the proper relation.

The second problem has to do with support relations. Arguably, evidence consists of facts or true propositions. Plausibly, your evidence consists of only things that you believe. In rejecting the truth-norm, the evidentialist has to accept the first evidential norm ('the evidential-norm' henceforth) but reject the second as spurious:

E: You should not believe p unless you have adequate evidence to believe p.

H: You should not believe p unless you have p as part of your evidence.⁷

The reason they have to reject H is that evidence consists of facts or true propositions, not falsehoods. If you can meet your epistemic obligations

⁵ See Alston (1989), Cohen (1984), Conee (2000), Feldman (2000), and Steup (2001).

⁶ See Conee (1992), Feldman (2002), Maitzen (1995), Vahid (2006), and Williamson (2000) for versions of this objection.

⁷ Some epistemologists seem to think that there's a 'high bar' to evidence possession (e.g., your evidence consists of certainties). I don't. The evidence you have consists of those facts that you have the right to treat as a reason for forming further beliefs. (Having said that, I think Schroeder's (2011) bar is set too low.)

whilst violating the truth-norm, you can meet your epistemic obligations whilst violating H. That means that H isn't a genuine epistemic norm if T isn't a genuine norm. By rejecting H, the evidentialist has to say that the normative standing of your beliefs depends upon the support it receives, but not upon whether it can provide support.

This is a very odd idea. Justification is closed under known entailment. So long as you justifiably believe *p*, you have sufficient justification to believe *p*'s known consequences. If you don't have sufficient justification to believe *p*'s known consequences, you must not have the right to believe *p* in the first place. This closure principle does not sit terribly well with the idea that the normative standing of a belief is determined entirely by what supports it but doesn't depend at all upon whether it can lend support of its own. Standing at the stop, Audrey seems to remember that the buses don't run past 7:00. Her watch tells her that it's 7:45. She infers that there won't be a bus. There are two ways of filling out the details of the case. Here's one. On the basis of mountains of evidence, she falsely believed that the buses don't run past 7:00. She's long since forgotten what those reasons were, but they were good enough according to the evidentialist. The positive standing of her belief can persist even when she's forgotten her original grounds. Let's suppose that that's happened. Since she (allegedly) justifiably believes now that the bus doesn't run past 7:00 and knows that it's 7:45, the closure principle says that she has sufficient justification to believe that there won't be a bus coming until tomorrow. She competently infers this. Surely she doesn't now believe for sufficient evidence that the bus won't come again until tomorrow. Neither the fact that she believes nor the fact that she seems to recall is among her reasons for believing that the bus won't come again until tomorrow. While these facts are known to her, they cannot explain the normative standing of her belief that the bus won't come until tomorrow because she doesn't believe for these reasons and these reasons aren't sufficient on their own. What Audrey takes to be a reason is that the bus doesn't run after 7:00. That's not a reason to believe anything at all.

We could have filled out the details of the case differently. We could say that Audrey remembers. She knowingly judges that the bus won't come again until tomorrow. In this version of the story, she does have adequate reason to believe that the bus won't come again tomorrow—it's that the bus doesn't run past 7:00. We can retain closure and do justice to our intuitions about the persistence of normative standing across time if we insist that the normative standing of a belief depends upon whether *it* provides reasons, not simply upon whether it was once supported by additional reasons. If we do this, we'd have to accept H. If

we accept H, we have to accept the truth-norm.⁸ The trouble with the truth-norm can't be that it's too demanding if there's no denying that this norm governs belief.⁹

The more pressing problem for the truth-first approach is that the truth-norm seems too permissive. There is more to meeting your epistemic obligations than simply fitting your beliefs to the facts. Advocates of the truth-first approach have to explain how this could be.

Let's consider two related explanatory challenges to the truth-first approach. The first has to do with the normative significance of your evidence. The trouble with evidentialism isn't with the idea that the evidence matters, but that the evidence matters to the exclusion of everything else. You cannot meet your epistemic obligations unless you have adequate reasons for your beliefs. Moreover, you haven't met your epistemic obligations unless you believe for good reasons. Audrey might believe that her father wasn't involved in what happened at the mill because she just cannot bear the thought that that was something he'd do. If her belief in her father's innocence isn't based on good reasons, it doesn't matter that she happens to have good reasons available to her. She hasn't met her epistemic obligations unless she puts things together in the right way and (thereby?) believes for good reasons. Finally, only considerations that bear on the truth of what you believe can be reasons to believe. The truth-norm is formulated as a prohibition against believing falsehoods. The truth-first approach has to explain why epistemic assessment has its inward looking focus. Why should it be concerned with the relation between good reasons to believe and the reasons for which you believe?

The second set of explanatory challenges has to do with understanding the relations between two kinds of epistemic norm. Some

⁸ Bird (2004) appeals to cases like this one to argue that some of our evidence is acquired via inference.

⁹ Some authors (e.g., Cohen (1984)) think that any reasonably held belief is justified. A number of authors also hold that any justified belief conforms to the norms that govern belief. Since it's possible to reasonably believe falsehoods, it might seem that it should be possible to believe falsehoods without violating any of the norms that govern belief. In Littlejohn (2009), I argued that it's a mistake to think of reasonably held beliefs as justified on the grounds that you could not be excused for failing to conform to a norm unless it would be reasonable for you to believe that you conformed to the norm. If reasonableness is necessary for excusable violations of norms, it cannot be the mark of justification. In the recent literature, a number of authors have argued that it is possible to violate the norms that govern belief whilst having justified beliefs. See Bird (2007), for example. Luke Sutton (2007), I think that this is a mistake. Nothing much turns on this in this paper, however. We can, following Bird, say that your belief is justified if it conforms to the norms governing belief or fails to do so for reasons that you were non-culpably ignorant of. All that matters for our purposes is that any failure to believe with justification is a failure to conform to a norm, not whether any failure to conform to a norm means a failure to believe with justification.

epistemic norms are formulated in such a way that they pertain directly to your beliefs (e.g., the truth-norm and the evidential-norm). Some epistemic norms have to do with theoretical reasoning or doxastic deliberation. Examples would include norms that tell us not to treat certain considerations as reasons to believe or prohibit certain kinds of inferential transitions. To introduce some terminology, let's call the norms that apply directly to beliefs 'doxastic norms' and the norms that determine whether you've deliberated properly 'deliberative norms'. There appears to be consensus that the normative standing of a belief depends, in part, upon whether you've conformed to deliberative norms in coming to believe. If the truth-norm truly is the fundamental epistemic norm, the truth-first approach takes the fundamental epistemic norm to be a doxastic norm. Since we can conform to this norm however we form our beliefs, it's not obvious how the advocates of the truth-first approach can account for the normativity of deliberative norms. If you violate some deliberative norms in forming your beliefs, you haven't met your epistemic obligations. If what matters fundamentally is the fit between belief and fact, why should facts about how you try to fit your beliefs to the facts have any further normative significance?

II.

Given just the resources of the truth-account, how can advocates of this approach account for the fact that epistemic assessment has the inward looking focus that it does? One place to look might be Shah's defense of evidentialism.¹⁰ As he sees it, the fact that belief is governed by the truth-norm explains why only considerations that bear on the truth of what we believe can constitute a reason to believe. If he's right, this might help to explain why epistemic assessment is concerned with whether your beliefs are based on adequate evidence.

The starting point for his argument that only evidence can constitute a reason to believe is the idea that those who grasp the concept of belief grasp that it is governed by the truth-norm. Those who grasp the concept of belief understand that the truth-norm captures the standard of correctness for belief. If doxastic deliberation is framed by the question whether to believe *p*, an individual who engages in such deliberation will grasp that only truth-related considerations can have any bearing upon whether to believe *p*. For this reason, Shah says, only such considerations can figure in doxastic deliberation. After all, nothing can be a reason to *X* unless it can figure in reasoning that disposes you to *X*.

¹⁰ See Shah (2003) and Shah and Velleman (2005).

Every part of this explanation is controversial, but let's grant each of his assumptions for the purposes of this discussion.¹¹ While Shah's argument suggests that certain kinds of considerations cannot constitute reasons for belief, it doesn't seem to lend any support to the idea that you need good evidence for your beliefs to meet your epistemic obligations.¹² If, say, Audrey comes to believe correctly that her father could not have started that fire on the basis of wishful thinking, it's not clear on the truth-first approach what epistemic wrong she's committed. She doesn't violate the truth-norm and she didn't deliberate from any non-evidential reasons. If the argument cannot vindicate that evidential-norm, we don't have an explanation as to why epistemic assessment is concerned with your evidence and the way you've handled it.

Someone could say that Shah's argument shows that there are deliberative epistemic norms that require us to exclude practical considerations from doxastic deliberation. It might seem that it's a short step from this to the further claim that factors unrelated to the truth of your beliefs shouldn't influence your beliefs. Even this doesn't seem to follow. If they show anything, Shah's arguments seem to show that certain kinds of considerations *cannot* figure in doxastic deliberation and so *cannot* constitute reasons to believe. To derive any normative conclusions from this about what the right to believe requires (e.g., that it requires reasoning from considerations that bear on the truth of what you believe), we'd need the further assumption that the right to believe depends upon whether your beliefs were formed in response to adequate reasons. That's a very plausible assumption, but it doesn't receive any support from Shah's argument.

Let's try a different tack. In the course of addressing the worry that the truth-norm is too 'objective' to be a genuine norm, various writers have proposed that there are also 'subjective' norms that have to do with evidence and rationality. Wedgwood, for example, distinguishes between different readings of 'ought'.¹³ On an objective reading, 'You ought not believe p' is true if p is false. On another, 'You ought not believe p' is true if it would not be rational for you to believe p. If Wedgwood is right that it's rational to believe p only if it makes sense for someone in your position to believe p given the aim of believing what's true and what it makes sense for you to believe depends upon your evidence, we seem to have the makings of a truth-first vindication of the evidential-norm.¹⁴

¹¹ See Hieronymi (2006) for a critical discussion.

¹² Shah never explicitly claims that it does. Raz (2011: 40) seems inclined to say that Shah's argument for evidentialism does help to explain why evidence has the normative significance we ordinarily take it to.

¹³ Wedgwood (2002b).

¹⁴ Wedgwood (2002a).

One potential problem with this approach to motivating the evidential-norm is that it's not at all clear what we've accepted in accepting that there's a sense in which you 'ought' to conform to the evidential-norm. Does this view say that you ought-subjectively believe only with evidence but may-objectively believe without it? If that's the view, it must be mistaken. Coop thought that Audrey's father was responsible for what happened at the mill, but he's started to have some doubts. In asking himself whether he should think that Audrey's father was responsible, it seems there's one and only question he has in mind. He wants to know whether he ought-really believe Audrey's father is involved. We now have three readings of 'ought', an objective reading, a subjective reading, and the sense of 'ought' that the conscientious and reflective subject has in mind. It seems that Audrey has two perfectly good ways of showing that Coop ought-really not believe her father was responsible for what happened at the mill. She might show that Coop couldn't have had any evidence of her father's involvement. He should-really not believe Audrey's father was involved because he should-subjectively not believe. She might show that Coop ought-really not believe her father was responsible by showing that he wasn't involved. He should-really not believe that her father was involved because he should-objectively not believe. Introducing different senses of 'ought' doesn't explain why overall epistemic obligation depends upon objective and subjective conditions. We need to explain how overall obligation depends upon both sorts of conditions to understand the relevance of the different readings of the epistemic 'ought'.

I suppose that someone could argue that these different senses of 'ought' should be introduced to explain why it appears that objective and subjective conditions both appear to work together to determine overall obligation. Someone could then argue that the challenge to explain how both sorts of factors work together to determine overall obligation is mistaken. When it comes to what you should-objectively believe, truth is all that matters.¹⁵ To further motivate the thought that there must be more to your epistemic obligations than simply believing what's true, consider cases of Moorean absurdities. There doesn't seem to be any sense in which someone believes what she 'ought' or as she 'ought' if she believes the following:

- (i) Audrey's father was involved, but I have no reason to think that he was.

- (ii) Audrey's father set the fire, but I don't know that he did.

¹⁵ A number of people have suggested that it's a mistake to try to show that false *beliefs* are defective and so objectively not as they ought to be. At best, poor reasoning or bad evidence shows that the deliberation that precedes belief is defective. The defects aren't inherited by the belief itself.

It would seem that belief in (i) or (ii), if possible, must surely be defective regardless of whether the propositions believed are true. Advocates of the truth-first approach should try to explain why there's a gap between believing what's true and meeting your objective epistemic obligations.¹⁶

Like Wedgwood, Boghossian thinks that there are subjective obligations associated with the objective obligation to conform to the truth-norm. He seems to think that there is a general connection between norms that aren't 'transparent' and derivative norms that can be followed directly:

Traders on the stock markets are attempting to comply with the rule: Buy low, sell high. But there is no direct way to recognize when a stock's price is low ... So traders follow certain other rules as a means of attempting to comply with the non-transparent rule that truly captures the aim of their trading activity ... Just so, I think, with the 'objective' norm that one ought to believe only what's true. Once again, this is not a rule that can be followed directly, but one that can only be followed by following certain other rules, the so-called norms of rational belief. For example: that we ought to believe that which is supported by the evidence and not believe that which has no support ... But, just as before, our story would be incomplete if we left out the fact that our following of these rules is a means of following the norm that we ought to believe only what is true. All of these norms are grounded in the objective norm of truth.¹⁷

Is this more promising?

Even if we cannot determine directly whether we conform to the truth-norm and so have to try to do so indirectly by assessing the evidence, it's not clear what right Boghossian has to say that there are additional *norms* that govern the means we use when we try to conform to the truth-norm. It might be true that whenever you ought to X, you ought to adopt the means necessary for X-ing, but the situation we're considering isn't one in which conforming to an additional norm (e.g., the evidential-norm) is necessary as a means to conforming to the truth-norm. You can believe what's true without believing for any good reason at all.

¹⁶ Even if it's not possible to believe (i) or (ii), it's possible to believe collections of propositions like this: (iii) the only reason I could have to think that Audrey's father was involved was Audrey's testimony, (iv) Audrey told me that her father was involved, but (v) Audrey isn't trustworthy and so nobody can know anything on the basis of her sayso. The set consisting of (iii)-(v) is consistent. It is a contingent matter whether the elements of the set are true. Still, nobody should believe (iii)-(v). Anyone who properly believes (v) ought to suspend judgment on (iii) or (iv).

¹⁷ Boghossian (2008: 101).

I don't think Boghossian's rationale for recognizing an evidential-norm is simply that we need to follow the evidence in order to believe the truth. We all know that it's possible to believe truths without believing on the basis of any evidence at all. Sometimes good things happen for bad thinkers. If I had to hazard a guess, I think he thinks that there are norms that govern the means we use to try to conform to the truth-norm that's grounded in the truth-norm because he thinks of norms as things that should be followed. There's more to following a norm than simply conforming to it. If there are additional steps that need to be taken to follow a norm that has non-transparent application conditions and norms are the sorts of things that ought to be followed, there will be normative constraints on the means we adopt to conform to a norm. If this is right, this should vindicate the evidential-norm and explain why there should be deliberative epistemic norms.

III.

This, then, is the proposal we have before us. The fundamental epistemic norm is a doxastic norm, the truth-norm. This norm's status as a fundamental norm does not turn on whether there are other epistemic norms. It depends upon whether it derives its authority from some more fundamental epistemic norm. Thus, advocates of the truth-first approach don't have to deny that there are additional epistemic norms that have to do with evidence and doxastic deliberation. This norm's status as *the* fundamental epistemic norm implies that all other epistemic norms derive their status from it. The substance or the content of the fundamental epistemic norm isn't supposed to explain on its own why there are derivative epistemic norms. Part of the explanation has to do with the point of norms. Norms are the sorts of things that you're supposed to follow. The reasons associated with them, the guiding reasons, are the sorts of things that you should be guided by.

As Boghossian and Wedgwood remind us, we cannot follow the truth-norm directly. To follow it, we need to reason. To reason as we ought to reason, we have to reason in ways that enable us to follow the truth-norm. Thus, it's not surprising that there are deliberative epistemic norms that tell us that there are ways we shouldn't reason and would tell us not to draw the conclusions arrived at by means of defective deliberation. As Shah has argued, if we seek to conform to the truth-norm, we cannot try to settle the question whether p by deliberating on the basis of considerations unrelated to the truth or falsity of p . Thus, it's not surprising that the epistemic norms would require us to believe only on the basis of evidence.

Everything seems to hang together quite nicely if we combine the truth-norm with a further claim about what norms are supposed to do

and what their reasons require of us. The question I want to take up in this section is whether this further claim about reasons and their demands is correct.

Let me contrast two ways of thinking about guiding reasons (i.e., the reasons associated with norms that apply to you and demand things from you).¹⁸ According to the first way of thinking about guiding reasons, guiding reasons should be thought of as guides. They are things that you should follow and be guided by. To be guided by them, you need to be cognizant of them and they need to be operative. On this approach, guiding reasons (typically) demand compliance. A reason to X wants you to X out of consideration of this very reason or for this very reason. Any failure to do what a reason requires constitutes a wrong, and so this account implies that it would be wrongful to, say, believe without being guided by the reasons that bear on whether to believe. Thus, irrationally believing truths would be wrongful just as rationally believing falsehoods would be.

According to a second approach, guiding reasons are thought of as guidelines. So far as they're concerned, guidelines don't want to be crossed. That's their sole concern. You don't need to be cognizant of them. Reasons cannot be operative if you're not cognizant of them. Since they don't demand your attention, they don't demand any role in deliberation. On this account, reasons (typically) demand nothing more than conformity. A reason to X wants you to X and is satisfied iff you X. Any failure to do what the reasons require is a wrong. There are no wrongs without a failure to do what the reasons demand.

If this conformity account is combined with the truth-first approach, the distinction between believing what's right and 'right believing' is lost. Thus, it might seem that the conformity account is implausibly weak. It's widely thought that evidence for p can be a reason to believe p. It's also widely thought that you aren't in a position to judge that p if you don't appreciate the force of the reasons you have to believe p. You shouldn't believe p unless the reasons for which you believe are themselves good reasons to do so. If this is so, it doesn't look like the mere conformity account can do justice to this.

The advocates the truth-first approach might have their reasons for preferring the compliance account to the conformity account, but are these reasons any good? I suppose that if there were good reasons for thinking that the truth-norm was the fundamental epistemic norm, the account would be very attractive. The account seems to explain various features of epistemic assessment that otherwise seem rather puzzling.

¹⁸ For a helpful discussion of these two ways of thinking about guiding reasons, see Raz (1975) and Gardner (2007).

Unfortunately, it looks like the compliance account is deeply problematic. I don't think guiding reasons are supposed to guide us in the way that the compliance account suggests.

If the compliance account captures an important truth about guiding reasons and their demands, it would presumably tell us something important about guiding reasons of all sorts. If the reason that epistemic assessment has the inward looking character that it does is down to the fact that guiding reasons *as such* are supposed to guide us and we ought to be guided by them, practical assessment should also have this inward focus. In the course of attacking the doctrine of double effect, Thomson attacks this very idea:

It is a very odd idea ... that a person's intentions play a role in fixing what he may or may not do ... Suppose a pilot comes to us with a request for advice: "See, we're at war with a villainous country called Bad, and my superiors have ordered me to drop some bombs at Placetown in Bad. Now there's a munitions factory at Placetown, but there's a children's hospital there too. Is it permissible for me to drop the bombs? And suppose we make the following reply: "Well, it all depends on what your intentions would be ... If you would be intending to destroy the munitions factory and thereby win the war, merely foreseeing, though not intending, the deaths of the children, then yes, you may drop the bombs. On the other hand, if you would be intending to destroy the children ... then no, you may not drop the bombs" ... Can anyone really think that the pilot should decide whether he may drop the bombs by looking inward for the intention with which he would be dropping them if he dropped them?¹⁹

Like Thomson, I can't believe that we'd need to scrutinize the pilot's intentions to determine whether he'd act permissibly in dropping the bombs. We should be focused on things like the children, the factory, and the cause. Unfortunately, not everyone sees the reductios that Thomson and I see. Some people think that we should look at the agent's intentions and should think that the agent's reasons are among the factors that determine whether the agent acted permissibly.²⁰ To show that the compliance account is wrong about guiding reasons in the practical case, I'll need more than this clever example. I'll need some arguments.

Ross' argument against the deontic relevance of motives can be turned into an argument against the compliance account.²¹ He focuses on the motive of duty, but nothing turns on whether we focus on this

¹⁹ Thomson (1991: 293).

²⁰ Hanser (2005).

²¹ Ross (1930). As Sverdlik notes, however, such arguments don't show that motives are deontically irrelevant, only that your duty isn't to act from them. You might have obligations not to act on certain motives. For a discussion, see Sverdlik (1996).

motive or another. He thinks that the claim that your duty is to return what you owe from the motive of duty conflicts with the categoricity of moral reasons. Choice cannot produce a motive, such as the motive of duty. You can only choose to perform an act of a certain type. You might be able to cultivate certain motives over time and bring it about that you have the appropriate motive at some later time, but that has no bearing on what your present obligation is. Your present obligation is to return what you owe. If you can do your duty without acting for an undefeated reason, your duty isn't to act for an undefeated reason that demands that you return what you owe. This is a strike against the compliance account.

Consider a second objection to the compliance account. To comply with a reason to X, it has to be the reason for which you X. For it to be the reason for which you X, you have to be cognizant of that reason and that reason has to be operative. The reason cannot be motivationally idle. Any failure to meet the demands the reasons place upon you renders your response (or non-response) wrongful. It might not be all things considered wrong, but wrong to some extent. A reason that you're aware of might be motivationally idle. A reason that you're not aware of must be. Think about the reasons that you're non-culpably ignorant of. You cannot follow them. They cannot be the reason for which you believe or do anything. If what reasons demand is that you follow them, it would be wrong for you to fail to follow them even when you're ignorant of them. If you're non-culpably ignorant of them, the wrong might be excused, but the thought that there's a wrong to excuse or to try to justify doesn't ring true. If anything, it rings false.

Here's a third objection to the compliance account. There can be cases of overdetermination in which there are two perfectly adequate undefeated reasons to X. Coop has two reasons to kiss. One reason is romantic. The other has to do with the greater good. As Austin reminds us

Of all pleasures bodily or mental, the pleasures of mutual love ... are the most enduring and varied. They therefore contribute largely to swell the sum of the well-being ... But, though he approves of love ... it was never contended or conceived, by a sound, orthodox utilitarian, that the lover should kiss his mistress with an eye to the common weal.²²

If Coop kisses for romantic reasons (and not because the kiss serves the greater good), the action isn't wrongful for that. If Coop hasn't failed to do what the reasons require, the reasons that lined up on the same side aren't disappointed when only one of them is operative. If the idle reasons don't demand compliance, reasons don't require compliance.

²² Austin (1954: 108) taken from Sverdlik (2011: 45).

Their demands don't diminish simply because they're lined up alongside further reasons.

Of course, someone could say that it would be overkill to act for all the reasons that apply to you. Why can't the compliance account say that your duty is always to act for some undefeated reason or other? My worry about this response is that it's obscure what the rationale would be for such a principle. On the compliance account, the principle couldn't be grounded in the demands of any of the individual reasons that apply to you because they each demand compliance. Someone could offer an instrumental justification of the principle on the grounds that we'll do better a better job conforming if we always act for a good reason, but I don't think the instrumental justification could go far enough. Audrey might be clever enough to work out that the reasons warranted her in attacking her old rival and she might do so in order to settle an old score. If we want to explain why her actions were wrongful, we couldn't appeal to the instrumental principle to do that because she was clever enough to wait to settle the score for there to be adequate reasons for her to act.

Here is a fourth and final objection to the compliance account. Think about reasons to render aid. You might have a reason to jump into the pond to pull a child to safety. You might not be alone in this. There might be a handful of people well positioned to pull the child to safety. If you don't move to act quickly and someone else pulls the child to safety, there's nothing that the reasons demanded from you that you've failed to do. If reasons required compliance, wouldn't they require you to be moved to act? (Upon seeing someone else start to swim to the child, you don't have reasons to jump in and outswim them!) It's hard to reconcile the observation the reason is indifferent to whether you were moved to act with the thought that the reasons demanded that you act out of respect for them. Someone might say that reasons only demand that *somebody* bring it about that the child is pulled free. I can't quite see how this would help the account, but let's modify the example. Why should it matter that it's someone rather than something? If a passing turtle or log brings the child to safety, there's nothing left that the reasons required from you that you've left undone. An action is one way amongst many of bringing about the state of affairs that there's reason to bring about. If something other than an action can be adequate, it's hard to see what's left of the view that a reason's demands include a demand that the agent's agency is exercised in some particular way.

IV.

The objections discussed in the previous section might not show that reasons demand conformity and nothing further, but they do show that

they do not invariably demand compliance. If they don't demand compliance, then it's not true that we ought to be guided by them. We shouldn't expect that there should invariably be an intimate connection between the norms that governing the reasoning that leads us to believe or to act and the norms governing beliefs and actions. It should not be surprising that, as Thomson notes, the normative appraisal of actions is outward looking and focuses on whether the agent's actions fit the situation.

The truth-first approach tries to get normative assessment of an individual's reasons and reasoning into the picture by arguing that reasons by their very nature demand a role in our reasoning. This picture is mistaken. Reason only has work to do when you're in danger of acting against reasons or violating a norm. Sometimes you just get lucky and there's no reason for you to worry about violating a norm. When this happens, reason doesn't have to remind you about the guidelines or help you steer between them. In doing nothing, it did all that it needed to. If this is right, what role, if any, does reason and reasoning play in helping you do what the reasons and norms would have you do? It would play a mere facilitating role. If reasoning gets you back into conformity, it's done all that it needs to do.

In light of this, it is striking that epistemic appraisal does focus on the relationship between explanatory and guiding reasons. If we reject the compliance account, we should be puzzled as to why this should be. If it doesn't matter in the practical case whether the reasons for which you act are among the good, undefeated reasons to so act, it is a surprising fact that epistemic assessment is concerned with the relation between guiding and explanatory reasons. It's surprising that acting for a reason that isn't an undefeated reason to so act isn't a kind of wrong but believing without believing for an undefeated reason is a kind of wrong.

Here's an explanation as to why epistemic assessment differs from practical assessment. Epistemic assessment is concerned with the relation between guiding and explanatory reasons because truth isn't the fundamental norm of belief. Knowledge is the fundamental norm of belief:

K: You should not believe p unless you know that p is true.²³

To conform to the knowledge norm, you have to believe only what you know. Whether you know depends upon whether the reasons for which you believe are among the genuine reasons there are to believe p . That's

²³ Sutton (2007) and Williamson (2000) defend the view that the knowledge norm is the fundamental epistemic norm. Although I argued in Littlejohn (2012) that their approach was mistaken, I've since changed my mind for reasons discussed here.

why epistemic assessment is concerned with the relation between explanatory and guiding reasons. The content of the fundamental epistemic norm explains why this should be. It shouldn't be explained in terms of some false claims about what norms are supposed to do or what reasons require.

v.

There are two ways to resist the argument for the knowledge norm. The first is to argue that I've overlooked some important asymmetry. One might say that the demands of practical and epistemic reasons differ. Practical reasons don't demand compliance, but epistemic reasons do. The second strategy for resisting the argument is to argue that reasons are more demanding than the conformity account would have us believe even if they're less demanding than the compliance account makes them out to be. So long as they demand more than conformity, the truth-first approach might have the resources for explaining why epistemic assessment is concerned with the relation between explanatory and guiding reasons.

The first response won't do. The cases that undermined the compliance account of practical reasons can easily be modified to undermine the compliance account of epistemic reasons.

Ross thought that the compliance account conflicted with the categoricity of practical reasons and I think there's a similar conflict in the epistemic case. There is presently a debate about whether it's permissible to believe lottery propositions. I say that it isn't. Others disagree. They have their reasons and they aren't terrible reasons, but they are wrong. (Let's suppose.) Coop thinks you can have sufficient grounds for believing lottery propositions and he has good grounds for thinking this. He holds a ticket for a drawing that took place last night and he doesn't know what the results were. He has evidence for his belief (i.e., he knows how lotteries work and so appreciates that it's very unlikely that his ticket won), but he still shouldn't hold this belief. He ought to get rid of this belief. Thus, there is decisive reason for him to do so.

This reason cannot play any sort of guiding role. To be guided by it, he would have to be cognizant of it and moved by it. To be cognizant of it, it seems that he would have to believe correctly that he has such a reason. To believe what he ought about this reason, this further belief would have to be reasonable in light of his evidence. It couldn't be. He's in the grips of a false theory that's supported by good reasons that mislead him. Maybe there's just no sound route of reasoning that would show him that he has decisive reason to suspend judgment on lottery propositions. We have a situation in which it seems (i) the truth-account

should say that he shouldn't take himself to have a decisive reason to get rid of this belief and (ii) he should get rid of the belief. It's unfortunate that there's no good route of reasoning available to him that would show him that he ought to rid himself of his lottery belief, but that's what duty requires. The categoricity of the evidential-norm that enjoins us not to believe lottery propositions seems to be at odds with the idea the reasons associated with this norm ought to guide us. All we need to do to conform to the evidential-norm is not believe the lottery proposition. We don't have to stop believing for the right reason (i.e., an appreciation of the evidential-norm and its application to the case of lottery propositions).

Someone might object to the use of this sort of case because the guiding reason at issue is not a piece of evidence. I want the case to tell us something important about epistemic reasons, but the relevant reason in the case is not a piece of evidence. Some authors seem to think that all epistemic reasons are pieces of evidence. As Raz puts it, "Epistemic reasons are reasons for believing in a proposition through being facts that are part of a case for (belief in) its truth" and any reason that's a reason against believing something is a reason to believe its negation.²⁴ My reason doesn't act like this. The fact that you don't have sufficient evidence to believe a lottery proposition is a reason not to believe the proposition but it isn't evidence for the proposition's negation. Because my reason doesn't play by these rules, what right do I have to claim that it is an epistemic reason? The fact that you don't have enough evidence or the right kind of evidence to believe *p* seems to explain why you shouldn't believe *p*. I'm assuming that Raz is right that whenever you ought or ought not *X*, there are reasons that explain why this should be.²⁵ I don't see how pieces of evidence could play this explanatory role. Someone might say that my reason is non-standard because it isn't a reason we can follow in reasoning. I don't think that's right. You cannot believe a proposition or its negation on the basis of the fact that there's not sufficient evidence, but you can certainly suspend judgment for this reason. When there are epistemic obligations to suspend, there are reasons that we need to explain why the suspension is mandatory and I don't think evidence is terribly well suited for this role.²⁶ Once we recognize that epistemic reasons aren't limited to pieces of evidence, their categoricity causes trouble for the thought that our duty is to be guided by the epistemic reasons that apply to us.

²⁴ Raz (2011: 36).

²⁵ Raz (2011: 25).

²⁶ As Adler (2002) and Gibbons (2013) observe, the acknowledged fact that nobody is in a position to know whether *p* is a decisive reason to refrain from believing *p* and from believing $\sim p$. That nobody can know whether *p* is a reason against believing *p* but it isn't a reason to believe $\sim p$.

Remember that one objection to the compliance account had to do with ignorance of the reasons. If you're ignorant of some reason, you can't be moved by that reason. And if reasons require compliance, you can't comply with the reasons that you're ignorant of. If a failure to meet the demands of a reason is a wrong and there's no wrong here to excuse or to justify, reasons don't demand compliance. Consider the proposition that God is a property. Having considered it, you might come to reject it. I have no idea what your grounds are, but I suspect that they're pretty good. There was a discussion of this proposition online a few years ago and there was a contest to identify the best argument against it. If you believed on the basis of any of the many sound arguments against this proposition, there's probably nothing wrong with your belief. The chances are good that you're not moved to disbelieve by some of the arguments considered there, but many of them provided good reasons to disbelieve. There's nothing wrong with your beliefs simply by virtue of the fact that there were good reasons that weren't operative. You only needed to be moved by one good reason.

A third objection had to do with overdetermination. You read the first argument in the discussion thread. You consider the steps in the reasoning and the argument's structure. You found it convincing. If you believe for the reasons this argument provided without believing for others that other arguments provided, there's nothing wrong with your belief but there are epistemic reasons that you haven't complied with.

Here is the epistemic analogue of the fourth and final objection. When your memory is in good working order, you retain lots of your beliefs. When your forgettery is in good working order, you lose lots of your beliefs. When you believe something you shouldn't, you can lose that belief by exposing that belief to reasons and revising your beliefs accordingly. You can also lose the belief by forgetting. If you believed *p* on the basis of spurious reasons and then simply lost that belief, there's no failure on your part to respond to the epistemic reasons that spoke against believing *p*. If so, they didn't demand compliance. Sometimes a leak is as good as a revision.

The compliance account isn't right about practical reasons or epistemic reasons. If we're stuck with the conformity account, the truth-first approach is sunk. At this point, someone might argue that the conformity account is implausibly weak and urge us to adopt a kind of compromise view. The conformity account might rightly reject the idea that your duty is to act for the right reasons, but it wrongly rejects the idea that your duty might be to refrain from acting on sufficiently bad reasons. Consider the dual demand account. It says that reasons place upon us a pair of conceptually related demands. They demand conformity and they also demand that you don't fail to show due deference to their status as reasons. It seems rather plausible that if you have reasons not to harm someone, you have reasons that demand that

you don't try to harm them or expose them to the risk of harming them. While the conformity account cannot easily account for the idea that there's just one reason here grounded in one thing that has this complex demand, the dual demand account can.

The dual demand account tells us that we have to focus on the individual's reasons because we need to know whether someone has shown due deference to a reason's status as a reason. As I understand the account, you can't fail to show due deference to reasons unless you've shown a willingness to act or believe against it. Thus, you don't fail to show due deference simply by failing to take notice of a reason. Moreover, you don't fail to show due deference to reasons simply by responding to other reasons. What they object to is your failing to conform to them or your willingness to act against them.

There is something attractive to this account. If we combined this account with the truth-first approach, we'd be able to explain why epistemic assessment is inward looking. It says that we should look in to determine whether subjects have shown due deference to the reasons that apply to them. A subject that didn't believe on what they could take to be strong evidence would be showing a kind of willingness to violate the truth norm by failing to exercise due care. So, the dual demand account nicely explains why we mustn't believe without evidence.

While this seems like a step in the right direction, it won't save the truth-first approach. The dual demand account tells us to focus on the individual's reasons for a specific reason. We're supposed to determine whether the subject has shown due deference to the reasons that apply to her. While it explains what's wrong with believing falsehoods and what's wrong with epistemic irresponsibility, I don't think that this quite captures why we want to look at the relation between explanatory and guiding reasons.

The limitations of the account become clear when we think about cases of responsible but fallacious reasoning and cases of fortuitous connection between belief and fact. There's a difference between believing responsibly and believing what the evidence supports. A responsible believer might be ignorant of the objective support relations that hold between the believer's evidence and he beliefs. When there is this gap, say, when a careful student is misinformed about some logical rule, the subject has shown due deference to the truth-norm but would fail to believe as she should even if her belief is correct. Consider a second example. A Roman physician who read all the peer-reviewed literature might have believed on good evidence that left-handed children were sinister. Suppose she believes on the evidence and quite correctly that some left-handed child is sinister. While her belief is based on evidence and true, the evidence leads only by accident to the truth. I think she shouldn't hold this belief, but the truth account cannot register

what's wrong with her belief. Thus, the dual demands account is no help in these cases.²⁷

In summary, I think that the charge that the truth-first approach cannot do justice to some familiar features of epistemic assessment is just. I don't see how advocates of this approach can do justice to the idea that the normative standing of your beliefs depends upon the reasons for which they're held. The advantage of the knowledge account is that it is easy to see how derivative epistemic norms are derived. Belief is governed by the truth-norm because knowledge requires truth. Belief is governed by the evidential-norm because knowledge requires evidence. There are derivative deliberative norms because the way you deliberate can help to determine whether you're in a position to know that something is so.

VII.

To remind us that we value the truth, Lynch wheels out Nozick's experience machine. The horror of life in the machine is supposed to remind us that we care about the truth:

Other things being equal, I wouldn't trade my present life, with all its ups and downs, for a life lived permanently within a pleasure machine ... Neither would I wish to live in the fool's paradise, where people just pretend to like and respect me. These examples, and others like them, show that we value something more than experience—even just pleasurable experience. We want certain realities behind those experiences, and thus we want certain propositions to be true.²⁸

²⁷ On Marcus' (2012) account of believing for a reason, to believe *p* for the reason that *q* is to represent *p* as to be believed on the grounds that *q*. He thinks we express the state of mind of believing *p* for the reason that *q* by uttering the demonstration '*q*, so *p*'. If believing for a reason essentially involves the representation of a rational connection between *q* and *p*, it might be thought that this gives the truth-account the resources to explain what's wrong with our Roman's attitudes. She would put her position this way, 'Bob is left-handed, and so sinister'. Just as she cannot say this without committing to the proposition that Bob is left-handed or the proposition that Bob is sinister, she cannot avoid the commitment to representing the one as to be believed in light of the other. Since it's false that's one to be believed in light of the other, can't we say that *this* is what constitutes the epistemic wrong? I don't think so. A successful demonstration isn't a list of truths and the force of the 'so' isn't captured by entailment. (Compare 'He is a bachelor, and so unmarried' to 'He is a bachelor, and so a bachelor'.) Our Roman takes it that one thing shows the other to be true. The natural way to understand this is in terms of one thing putting you in a position to know something else.

²⁸ Lynch (2004: 138-39)

He might be right that we care about the truth and right that cases like this remind us that this is so, but examples like this need to be treated with care. In showing us that the truth matters, they show us why the truth-first approach is fundamentally misguided.

In Nozick's story, there are realities behind the experiences. On his telling, the facts don't fit the appearances and so there's a gap between appearance and reality. That gap isn't an essential feature of the case. We can get rid of the gap while maintaining complete independence between appearance and reality. That's a common feature of Gettier cases and it's something we can build into the story if we choose to do so.

Audrey is in the machine. Before stepping in she had hoped that her brother would graduate. One morning it seems to her that her brother is graduating. We can tell the story Nozick's way. Her brother is dead, but it seems to her that her brother is graduating because the machine's operators have set things up that way. A smile stretches across her face. Can the operator say that she's smiling for the reason that her brother is graduating? Can she say that Audrey is happy that her brother has finally earned his degree? No, not truthfully. He's dead. We can tell the story Gettier's way. Her brother isn't dead. On this morning it just so happens that he's receiving his diploma. The operators have no idea that this is so. This is just a happy accident. Audrey's beliefs about her brother are correct just as they appear to be, but they aren't tethered to reality. Now can the operator say that she's smiling for the reason that her brother is graduating? Can she say that Audrey is happy that her brother has finally earned his degree? No, not truthfully. The mere match between belief and fact doesn't put Audrey in a position to think, do, or feel anything in light of this fact. The realities are hidden behind her experiences and so unknown to her.

The reason that life in the machine is horrible isn't that the beliefs you form in the machine are all false. The reason that life in the machine is horrible is that you've lost contact with reality. The example shows that there's a difference between having true belief and being in touch with reality. It can be good to have the truth in view. It doesn't do you much good if the truth is just 'out there'.

I've often been pressed to explain why we should think that there's a norm that governs belief. Here's the sketch of an answer. There are things that beliefs are supposed to do and there's a difference between the beliefs that can do what they're supposed to and the beliefs that cannot do what they're supposed to. Your beliefs are supposed to provide you with guiding reasons so that you can think what you should think, feel what you should feel, and do what you should do. These reasons consist of facts. Beliefs that don't fit the facts can't do what beliefs are supposed to do. That's why it's not wrong to say that belief is governed by the truth-norm. The reason that it's wrong to say that the

truth-norm is fundamental is that true beliefs can't always do what beliefs are supposed to do. It's only when your beliefs constitute knowledge that they can provide you with reasons.²⁹ That's why knowledge is the norm of belief.

²⁹ A point that's been ably defended by Hyman (1999), Unger (1975), and Williamson (2000).

REFERENCES

- Alston, W. 1989. *Epistemic Justification*. Cornell University Press.
- Austin, J. 1954. *The Province of Jurisprudence Determined*. Noonday.
- Bird, A. 2004. Is Evidence Non-Inferential? *Philosophical Quarterly* 54: 253-65.
- _____. 2007. Justified Judging. *Philosophy and Phenomenological Research* 74: 81-110.
- Conee, E. 1992. The Truth Connection. *Philosophy and Phenomenological Research* 52: 657-69.
- Conee, E. and R. Feldman. 2008. Evidence. In Q. Smith (ed.), *Epistemology: New Essays*. Oxford University Press, pp. 83-104.
- Feldman, F. 1986. *Doing the Best We Can*. Reidel.
- Feldman, R. 1988. Epistemic Obligations. *Philosophical Perspectives* 2: 235-56.
- _____. 2002. Epistemological Duties. In P. Moser (ed.) *The Oxford Handbook of Epistemology*. Oxford University Press, pp. 362-85.
- _____. 2000. The Ethics of Belief. *Philosophy and Phenomenological Research* 60: 667-95.
- Foley, R. 2004. A Trial Separation between the Theory of Knowledge and the Theory of Justified Belief. In J. Greco (ed.), *Ernest Sosa and His Critics*. Blackwell, pp. 59-72.
- Gardner, J. 2007. *Offenses and Defenses*. Oxford University Press.
- Gibbons, J. 2013. *The Norm of Belief*. Oxford University Press.
- Hanser, M. 2005. Permissibility and Practical Inference. *Ethics* 115: 443-70.
- Hieronymi, P. 2006. Controlling Attitudes. *Pacific Philosophical Quarterly* 87: 45-74.
- Hyman, J. 1999. How Knowledge Works. *Philosophical Quarterly* 49: 433-51.
- Littlejohn, C. 2009. The Externalist's Demon. *Canadian Journal of Philosophy* 39: 399-434.
- _____. 2012. *Justification and the Truth-Connection*. Cambridge University Press.
- Lynch, M. 2004. *True to Life—Why Truth Matters*. MIT University Press.
- _____. 2009. Values of Truth and the Truth of Values. In A. Haddock, A. Millar, and D. Pritchard (ed.), *Epistemic Value*. Oxford University Press, pp. 225-41.
- Maitzen, S. 1995. Our Errant Epistemic Aim. *Philosophy and Phenomenological Research* 55: 869-76.
- Marcus, E. 2012. *Rational Causation*. Harvard University Press.
- Millar, A. 2004. *Understanding People*. Oxford University Press.
- Raz, J. 1975. *Practical Reason and Norms*. Oxford University Press.
- _____. 2011. *From Normativity to Responsibility*. Oxford University Press.
- Ross, D.

- Russell, B. 1970. The Elements of Ethics. In W. Sellars and J. Hospers (ed.), *Readings in Ethical Theory*. Prentice-Hall, pp. 3-28.
- Schroeder, M. 2011. What Does it Take to 'Have' a Reason? In A. Reisner and A. Steglich-Peterson (ed.), *Reasons for Belief*. Cambridge University Press, pp. 201-22.
- Shah, N. 2003. A New Argument for Evidentialism. *The Philosophical Quarterly* 56: 481-98.
- Shah, N. and D. Velleman. 2005. Doxastic Deliberation. *The Philosophical Review* 114: 497-535.
- Steup, M. 2001. Epistemic Duty, Evidence, and Internality. In M. Steup (ed.), *Knowledge, Truth, and Duty*. Oxford University Press, pp. 134-51.
- Sutton, J. 2007. *Without Justification*. MIT University Press.
- Sverdlik, S. 1996. Motive and Rightness. *Ethics* 106: 327-49.
- _____. 2011. *Motive and Rightness*. Oxford University Press.
- Thomson, J. 1991. Self-Defense. *Philosophy and Public Affairs* 20: 283-310.
- Vahid, H. 2006. Aiming at Truth: Doxastic vs. Epistemic Goals. *Philosophical Studies* 131: 303-35.
- Wedgwood, R. 2002a. Internalism Explained. *Philosophy and Phenomenological Research* 65: 349-69.
- _____. 2002b. The Aim of Belief. *Philosophical Perspectives* 16: 267-97.
- _____. 2007. *The Nature of Normativity*. Oxford University Press.
- Whiting, D. 2012. Does Belief Aim (Only) at the Truth? *Pacific Philosophical Quarterly* 93: 279-300.
- _____. Forthcoming. Stick to the Facts: On the Norms of Assertion. *Erkenntnis*.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford University Press.
- Zimmerman, M. 1996. *The Concept of Moral Obligation*. Cambridge University Press.



THE ARISTOTELIAN SOCIETY

PRESIDENT: Sarah Broadie (St. Andrews)

PRESIDENT-ELECT: E.J. Lowe (Durham)

HONORARY DIRECTOR: Lucy O'Brien (UCL)

EDITOR: Matthew Soteriou (Warwick)

LINES OF THOUGHT SERIES EDITOR: Scott Sturgeon (Oxford)

EXECUTIVE COMMITTEE: Ben Colburn (Glasgow) / Alison Hills (Oxford) / Rosanna Keefe (Sheffield)
Marie McGinn (UEA) / Samir Okasha (Bristol) / Ian Rumfitt (Birkbeck) / Robert Stern (Sheffield)

EXECUTIVE ADMINISTRATOR: Mark Cortes Favis

ASSISTANT EDITOR: David Harris

EDITORIAL ASSISTANT: Lea Salje