

*draft of 10 November 2008*  
*please cite only with permission*

**abstract:** In this paper, I sketch and defend a new account of the sort of propositional knowledge possessed by mature humans. The defense of this account comes mostly in pointing out its ability to (i) explain the normativity of knowledge, (ii) comport with three intuitions about knowledge (what I call: the anti-luck, value, and anti-skeptical intuitions), and (iii) make unified progress on the Gettier, Meno, and Cartesian problems. In short form, the proposed account is that knowledge is true belief that one ought to have. However, several refinements come in spelling out the precise sense in which one *ought* to have such beliefs. One of the keys to these refinements is deploying what I (following Finlay, forthcoming) call the “end-relational” theory of ‘ought’. This lets me articulate and defend the final version of the account, which is that knowledge is true belief that one ought to have in order to have a safe and rationally defensible belief.

## **Knowledge as True Belief that One Ought to Have**

Matthew Chrisman  
University of Edinburgh

“...in characterizing an episode or state as that of *knowing*, we are not giving an empirical description of that episode or state; we are placing it in the logical space of reasons, of justifying and being able to justify what one says.” Wilfrid Sellars, *Empiricism and the Philosophy of Mind*, sec. 36)

### **0. Introduction**

In this paper, I want to propose a new account of the sort of propositional knowledge possessed by mature humans or, if you will: rational inquirers. This account is motivated by the idea that epistemology is a properly

normative discipline and that we must understand our epistemic evaluations—especially our knowledge attributions—in a way that helps to integrate them with normative evaluations more generally. (This is the reason for the seductive but perplexing slogan from Sellars quoted above.) I am also motivated by a number of intuitions about knowledge that I think are widely held and that, at any rate, I think would be nice to validate insofar as possible. These are (i) the *anti-luck intuition* that knowledge is incompatible with getting things right merely by luck, (ii) the *value intuition* that knowledge is better than mere true belief, and (iii) the *anti-skeptical intuition* that there are many things that we know.

These intuitions lead me to be interested in proposing an account that goes some way towards meeting three diagnostic desiderata emerging from recent epistemological literature. First, I want an account that brings some resources to diagnosing the difficulty epistemologists have had in analyzing knowledge in light of the apparent failure of the classical JTB analysis to explain why luckily true justified belief doesn't seem to be knowledge. Or, to state this desideratum slightly differently, I want an account that explains the intuitions driving the Gettier problem. Second, I want an account that helps to diagnose the difficulty epistemologists have had in explaining why knowledge seems to be non-instrumentally better than epistemic states short of knowledge. Or, again to put it slightly differently, I want an account that explains the intuitions driving the Meno problem. Finally, I want an account that helps to diagnose the persuasiveness of skeptical arguments from counterpossibilities. Or, to put it differently again, I want an account that explains the intuitions driving the Cartesian problem.

This is a bold list of desiderata. Recent epistemology has understandably specialized considerably in response to one or another of these problems. Discussions of the Gettier problem and the viability of various reductive analyses of knowledge have turned to issues about the aetiology of our beliefs and the sorts of luck that are consistent and inconsistent with knowledge; and several epistemologists have come to question the analyzability of knowledge altogether. Discussions of the Meno problem and the value of knowledge have naturally turned to issues in value theory about topics such as the difference between intrinsic, extrinsic, instrumental, and final value; and epistemologists working on this problem have become interested in the additivity, transitivity, and location of epistemic value. Finally, discussions of the Cartesian problem and skepticism have turned to the nature of evidence, the plausibility of various closure principles, and the possibility that 'knows' is semantically context-sensitive. Despite this specialization, I think full and unified satisfaction of *all* three desiderata is what the theory of knowledge should aim for. Here, I will propose an account that I think makes at least some unified progress towards meeting each desideratum.

In §1, I'll present the basics of the account and discuss the way it fits into my general motivation to view epistemology as a properly normative discipline and epistemic evaluations as properly normative evaluations. Then, in §2, I'll

discuss the three desiderata and the theoretical resources I think my account brings to the Gettier problem, the Meno problem, and the Cartesian problem.

## 1. The Account and the Normativity of Knowledge

Almost all epistemologists agree that knowledge is factive, which means that, if some subject S knows some proposition p, then p must be true.<sup>1</sup> So, I will simply take it for granted that knowledge is true. Though it is perhaps slightly more controversial, I also want to take it for granted that knowledge requires belief. Whether a subject S could know some proposition p without believing that p depends on what we mean by 'believe'. For my purposes here, I am willing to accept whatever we have to mean by 'believe' such that believing that p *is* partially constitutive of knowing that p for rational inquirers like ourselves.<sup>2</sup> For the standard debate, the real rub comes, of course, in saying what else is required to turn true belief into knowledge.

Traditionally epistemologists have approached this problem by saying, schematically, that what else is required is *justification* and then attempting to give a positive characterization of what constitutes epistemic justification for a belief. For instance, internalists typically see this as a matter of the belief's rational relations (such as connections to reasons or evidence) to other things internal to the believer's mind. But then they face the Gettier problem arising from the fact that whatever this internal relation is, the internally justified belief could be luckily true and so apparently not knowledge, at least if we take the anti-luck intuition seriously.<sup>3</sup> This sometimes generates a search for an elusive further condition on knowledge that captures the sense that knowledge must be not only rationally defensible but also reliably formed. By contrast, externalists typically reject the internal conception of justification and seek to replace it with a conception articulated directly in terms of some modally robust relation (such as being formed by a reliable process) with the fact known.<sup>4</sup> But then they face difficulties in articulating this modally robust relation in a way that comports with

---

<sup>1</sup> For a dissenting voice, see Hazlett (forthcoming).

<sup>2</sup> This is by far the majority approach in contemporary epistemology, though it is not without its detractors who trace their inspiration at least back to Plato's *Republic*. There it is suggested that knowledge and belief are incompatible states of mind, the former an infallible apprehension of a necessary truth and the latter a fallible attempted apprehension of the world of appearances in flux. In endorsing the belief component in an account of knowledge, I think most contemporary epistemologists are simply using the word 'believe' in a different way from this tradition. Even Williamson, who argues that knowledge should not be analyzed in terms of belief, suggests that belief should be analyzed in terms of knowledge, where "Knowing is...the best kind of believing. Mere believing is a kind of botched knowing"(2000: 47).

<sup>3</sup> See Gettier (1963).

<sup>4</sup> It is of course also possible for externalists to work outside of the JTB framework by conceding the internal conception of justification but then insisting that it is not this but rather a modally robust relation to the fact known that is necessary for knowledge. This is, for example, the strategy pursued in Dretske (1981).

intuitions about a number of cases (such as BonJour's reliable clairvoyant<sup>5</sup>) where it looks like a belief can exhibit whatever modally robust, external relation to the fact one wants, but it lacks any rational relations to other things internal to the believer's mind.

Here is where I want to propose something different. I think we should understand the further requirement for knowledge not (directly) in terms of reasons, evidence, reliability, or any of the other proposals epistemologists have pursued in the past fifty years; nor do I think we should simply add some fourth condition like truth-tracking, safety, sensitivity, no false lemmas, etc. Rather, I want to suggest that we understand knowledge in terms of what the subject *ought to believe*. So, as a first pass, my suggestion is that S knows that p just in case S's true belief that p is one that S ought to have.

## 1.1 Initial Refinements

This will likely seem rather unpromising until we recognize several distinctions which in metaethics and deontic logic sometimes go under the heading of the "many senses of 'ought'". There are three initial distinctions necessary to narrow in on the sense of 'ought' I have in mind.

First, there is the distinction sometimes made between normative and nonnormative 'ought's. It's sometimes claimed that many of the ought-claims made in medicine, meteorology, and other areas of prognostics are in no way normative. For instance, when your doctor tells you that the broken foot he has patched up ought to be healed in six weeks, it's tempting to understand this as a mere prediction to the effect that, given certain background assumptions, it's very likely that your foot will be healed in six weeks.<sup>6</sup> By contrast, normative 'ought's are apparently not about likelihoods. Saying exactly what they are about if not likelihoods is a vexed question, but surely it has something to do with reasons and values.

"Suzy ought to clean up her room" seems to imply something about what Suzy has reason to do or what would be best for Suzy to do given her options. So, if this distinction between normative and nonnormative 'ought's is cogent, by 'ought' in my account, I mean the normative 'ought'. This ties in with the value intuition mentioned above. Nonnormative 'ought's don't have anything specifically to do with goodness and betterness.

---

<sup>5</sup> See BonJour (1985).

<sup>6</sup> Some philosophers would argue that even these are normative in that they are information relative epistemic 'ought's. The idea is that these predictions state what one ought to believe, given certain presupposed background information. In this case, it would be, roughly, that one ought to believe that it's highly likely that the foot will be healed in six weeks. See Wedgwood (2008: CH5), but see Finlay (forthcoming) for a contrary view. I don't want to adjudicate this debate here, and it suffices to note that the 'ought to believe' at the heart of my account is not meant as a claim about what beliefs we ought to believe are likely, given certain presupposed background information.

Second, there is the distinction sometimes made between categorical and hypothetical 'ought's. Hypothetical 'ought's are ones whose application to an agent is thought somehow to depend on some contingent personal end that this agent has set for herself, whereas categorical 'ought's are ones which do not so depend. For instance, having set as an end for myself that I finish this paper by the end of the month, we might say that I hypothetically ought to work diligently on it every day of the last week of the month. The idea is that if I hadn't set this end or if, realizing its futility, I drop this end, then the normative force of the 'ought' disappears. By contrast, categorical 'ought's do not depend in this way on contingent personal ends for their normative force. Sometimes it seems as if this distinction could be marked by whether the relevant ought-statements are conditional in their logical form, but this is widely agreed to be false. Some hypothetical 'ought's are not obviously conditional in their logical form—for example: "You ought to have children." And there are clearly conditionalized ought-statements which are categorical—for example: "If you have children, you ought to do what you can to provide for them." So, the distinction is not about logical form but rather about the dependency of the normative force of the ought-statement on the agent's contingent personal ends.

So, should the 'ought' in an account of knowledge be hypothetical or categorical? To appreciate the answer to this question, notice that it is possible for S to set some contingent and personal end for herself, such as having a false belief about her prospects for overcoming cancer or not believing anything unpleasant about her children. In these cases, abstracting away from any other contingent personal ends and assuming the relevant proposition p to be true, it would not be the case that S ought to believe p *hypothetical on her contingent personal ends*. Yet, if, for instance, she has overwhelming evidence that p is true, it still seems that she ought, at least in some epistemic sense, to believe p. What we want to say is that, given her overwhelming evidence and the truth of p, she *categorically ought* to believe that p, whatever her contingent personal ends happen to be. Because of this, I think the 'ought' relevant to an account of knowledge must be categorical rather than hypothetical.<sup>7,8</sup>

If the first two distinctions needed to specify the sense of 'ought' relevant for my account make it seem like I am talking about a very robustly normative sense of 'ought', the third will deflate things a bit. This is the distinction sometimes made between what one ought to do *pro tanto* or in some respect and what one ought to do all things considered. For instance, it might be that,

---

<sup>7</sup> For a related argument, see Kelly (2003). I discuss Kelly's argument in my (forthcoming) and attempt to diagnose why epistemic 'ought's are categorical rather than hypothetical.

<sup>8</sup> There maybe a deeper sense in which these 'oughts' are hypothetical on the contingent aims of humans in epistemic communities. That is, perhaps, if humans had radically different cognitive capacities or if somehow our communal lives were better served by different epistemic norms, then the 'ought's to believe that are relevant to my account of knowledge would not hold. I don't know whether this idea is cogent, but the present point is that the relevant 'ought's don't seem to be hypothetical on the believer's contingent and personal ends, which is what I mean by 'hypothetical'.

with respect to my teaching, I ought to prepare my lectures for next semester, but, with respect to my research, I ought to finish this paper by the end of the month; yet, all things considered, perhaps I ought to quit being an academic altogether. Personally, I remain unsure whether there really is such a thing as the 'all-things-considered ought', but even if there is, I think it is unlikely that the 'ought to believe' relevant in epistemology is an all things considered ought. After all, some of the truths we might come to know could be horrendous and ruin the lives of ourselves and others if we were to believe them. For example, it's sometimes suggested that knowing the truth about what each of us think about each other would be horribly depressing and lead to the demise of human community. Similarly, it sometimes seems as if being a good parent requires not believing certain relatively obvious truths about your children.

If something like this possibility holds, then it would not be the case that, all things considered, we ought to believe the relevant truths. However, that doesn't change the fact that, as we might say, with respect to *epistemic* considerations, we ought to have these beliefs insofar as they are supported by our evidence. So, in light of this distinction, I think the 'ought' relevant to understanding my suggestion that knowledge is true belief that one ought to have is an 'ought' *with respect to epistemic considerations* rather than an *all things considered* 'ought'. This means that although the 'ought' is categorical (because its normative force doesn't depend on the agent's contingent personal ends), it is only categorical with respect to epistemic considerations rather than all things considered.

These three distinctions let us make a second pass at my proposal in more specific terms that hopefully make it seem more promising. The idea is that S knows that p just in case S's true belief that p is one that, with respect to epistemic considerations, S normatively and categorically ought to have.

## 1.2 More Refinements and the End-Relational Theory of 'Ought'

While perhaps more plausible, the second pass at my proposed account does raise an important issue that will be necessary to address before the proposal is complete enough to evaluate. The issue is about what it means to say that someone ought *with respect to epistemic considerations* to believe some proposition. This leads to a fourth and much more controversial sense in which 'ought' is sometimes said to have many senses. This is the idea that there may be 'ought's of prudence, 'ought's of morality, 'ought's of law, 'ought's of etiquette, and even 'ought's of baseball, 'ought's of iambic pentameter poetry, 'ought's of modern decorating, etc. It's surely wrong to think of 'ought' as massively ambiguous along these many fronts, but I think there is also sense to be made of the idea that different ought-statements derive their force from different kinds of norms in systematic ways. This is why, in one context, it can seem correct to make one ought-statement, while, in another context, it seems correct to make an apparently incompatible ought-statement. For instance, in a context where a fugitive's best means for escape from the law

are at issue, it can surely make sense to say, “The fugitive ought to mingle with the crowd.” While, in another context, where the law is at issue, it can surely make sense to say, “The fugitive ought to turn himself into the authorities.” Here, we sometimes want to say something like the following: He ought, for escape, to mingle with the crowd, while he ought, legally, to turn himself into the authorities.

But how are we to capture this distinction without implausibly positing massive ambiguity to the word ‘ought’? I think that an *end-relational* theory of these ‘ought’s is a promising way to do so.<sup>9</sup> The end-relational account says that these ‘ought’s are, if not explicitly, then implicitly related to an end. This means that we can more transparently render them with an ‘in order to...’ modifier. For instance, we can render the first ought-statement about the fugitive as: “In order to evade arrest, the fugitive ought to mingle with the crowd”; and we can render the second ought-statement about the fugitive as: “In order to conform to the law, the fugitive ought to turn himself into the authorities.” Clearly these two claims are compatible, and, given certain assumptions about the fugitive’s situation, they both seem true.

Doesn’t this make them merely hypothetical ‘ought’s? I don’t think so. The issue of the categoricity versus hypotheticality of an ‘ought’ seems to me to be orthogonal to the end-relativity of (at least some) ‘ought’s.<sup>10</sup> To see this, notice that either of the previous ought statements *could* be hypothetical if the agent has set the right sort of personal and contingent end. For example, if he intends to evade arrest, then we might say that, hypothetical on that end, he ought to mingle with the crowd. But if he intends to conform to the law, then we’d likewise say that, hypothetical on *that* end, he ought to turn himself into the authorities. However, if he intends to go down in a blaze of glory like a true rebel, then, although it seems true that he to mingle with the crowd in order to evade arrest and that he ought to turn himself into the authorities in order to conform to the law, neither of these ‘ought’s have any normative force for him.<sup>11</sup> That is to say, neither of these ‘ought’s correspond to reasons for acting in the way that the ‘ought’s prescribe, because the ends to which these ‘ought’s are related are optional ends and he has not opted for either of them.<sup>12</sup>

---

<sup>9</sup> A more worked out end-relational theory of ‘ought’ is defended in Finlay (forthcoming), from which I borrow the example of what the fugitive ought to do.

<sup>10</sup> Compare also Finlay (2008) for a similar point.

<sup>11</sup> This claim assumes that neither what he ought to do in order conform to the law nor what he ought to do in order to evade arrest are categorical oughts. If you think one of them is categorical, then you can substitute a different example where all of the ends relative to which he ought to do what he ought to do are optional.

<sup>12</sup> This is not meant to commit me to a particular view in the debate between Humeans and Kantians on the *ultimate source* of normativity or valid rules. The former will argue that all valid rules ultimately derive from our interests and inclinations, while the latter will argue that at least some valid rules can float free of our interests and inclinations. The end-relational theory of ‘ought’ rest on a much more conservative claim about the logical form of rules couched in terms of ‘ought’: these can always be more perspicuously stated in terms of ‘ought...in order to...’.

So, whether or not an ought-statement is end relative (either implicitly or explicitly) doesn't determine its status as hypothetical or categorical. Rather its status as hypothetical or categorical is determined by whether its normative force is thought to be hypothetical on the agent's contingent personal ends. Categorical end-relative 'ought's will be ones where the relevant end is not optional for the agent to which they apply. Moreover, it seems that some ends may be nonoptional, with respect to epistemic considerations, while they may nonetheless be optional all things considered. This is how it could be the case that you categorically ought, with respect to epistemic considerations, to believe p, while you hypothetically ought not, with respect to other considerations (such as surviving your cancer or being a loving parent), to believe that p.

A general end-relational theory of the normative 'ought' says that all normative 'ought's are, if not explicitly, then implicitly related to an end. When this end is not made explicit, it is features of the context of utterance that determine what the end is. I'm tempted by this theory, but all that I need for the purposes of my account of knowledge is a more modest claim. This is the thought that specifically *epistemic* 'ought's are end-relational, and that this feature can be used to distinguish them from other sorts of 'oughts', in order to make sense of the idea that knowledge is true belief that, with respect to epistemic considerations, one ought to have. The idea would be that claims of the form "S ought to believe that p" are context sensitive in virtue of their implicit relativization to a contextually determined end; some contexts determine a specifically epistemic end, thereby making the relevant 'ought' epistemic.

But, is it plausible to think that epistemic 'ought's are end-relational? That is, is there a plausible rendering of statements of the form "S epistemically ought to believe that p" in terms of an 'in order to...' modifier? What ends would make an 'ought' an *epistemic* ought? According the venerable tradition of truth monism in epistemology, specifically epistemic considerations are, in general, to be distinguished by their relation to the goal of truth. That is to say, roughly, that epistemic reasons, epistemic norms, epistemic virtues, epistemic evaluations, etc. are all to be distinguished from their cousins in other domains by appeal to the goal of truth. This suggests a strategy for isolating specifically epistemic 'ought's. The idea would be to say that "S epistemically ought to believe that p" basically means that S ought to believe that p, in order to believe truly.

Unfortunately, however, I think that is neither the right way to apply truth monism to the issue of isolating the specifically epistemic 'ought' nor a very plausible idea for understanding the 'ought' relevant to knowledge.

To see this, notice that the truth monist is not someone who says that there are no other epistemic ends besides truth, rather he is someone who says that whatever is valuable epistemically is ultimately explicable in terms of

---

And, as I go on to discuss in the text, all that is important for my purposes here is that the logical form of specifically epistemic 'ought's can be made more perspicuous in this way.

the finally valuable goal of true belief. This means that the truth monist can recognize a distinction between what one ought to believe, (i) in order to believe truly, (ii) in order to believe in accord with one's evidence, (iii) in order to believe knowledgably, (iv) in order to believe in a rationally defensible way, (v) in order to believe responsibly, (vi) in order to believe safely, (vii) in order to believe nonfalsely, etc. The end-relational theory of 'ought' helps to see how these are all different ought-statements. The truth monist is not someone who denies that; rather he'll just argue that, whatever epistemic value attaches to these various ends, it is ultimately explicable in terms of the value of true belief.

In any case, there seems to be a distinctively epistemic sense of 'ought', where we want to say that someone believed as she ought to have believed, but because of factors beyond her ken, that belief turned out false. For example, consider a woman who has overwhelming evidence that her husband has died at war; perhaps she has a first-hand report from one of his fellow soldiers of a deadly attack. Here, we might say that, given the evidence, she ought to believe that her husband died at war. In fact, it demonstrates virtue not to languish in false hopes despite their comfort. But what if we know that the evidence was misleading or merely luckily indicative of the truth? Perhaps the fellow soldier was trying to deceive her to cover up something about the attack. Or perhaps he was trying to deceive her, but, unbeknownst to him, the husband really did die in some other battle. In all of these cases, there seems to be a sense in which she ought to believe that her husband died at war and a sense in which she ought not to believe that—and both senses seem to me to be *epistemic*.<sup>13</sup>

Two things are important here. First, if it's possible to say cogently that the widow ought not to believe that her husband died at war, even though he did, then, as long as this is an epistemic ought, it cannot be that all epistemic 'ought's are distinguished always by their implicit coupling with 'in order to believe truly'. In this case and on this sense of 'ought' she ought to believe something false. Second, adopting the end-relational account of 'ought' and allowing for several different possible epistemic ends provides a nice way to make sense of this variation of epistemic evaluation. For instance, in the case where the evidence was luckily indicative of the truth, we could say that the

---

<sup>13</sup> This echoes a point made by Kolodny that in saying that you ought to have some attitude, "we are saying one of two things. Sometimes we are saying that there is a *reason* for you to have the attitude, because of some feature of your *actual situation* ... Other times we are saying that it would be *irrational* of you not to have the attitude, because of some *other attitude* that you have" (2005: 510). In a footnote attached to the text I have just quoted, Kolodny claims further that, in the case of saying what one ought to believe, both things we may be saying in making such claims pull apart from what is required to believe truly. He apparently thinks that when we say that someone ought to believe p, we never mean that he ought to believe p because p is true. I think this is wrong. For instance, regardless of their evidence, students of geography ought to believe that London is the capital of the United Kingdom; this is because it is true (and perhaps relevant to their geography lessons). One value of the end-relational account is that it allows us to distinguish this third thing that we might be saying when we say that someone ought to believe p from the other two things we might be saying.

woman ought not to believe that her husband has died, in order to believe in a way that could not easily have been mistaken, though she ought to believe that he has died, in order to believe in the way best supported by the evidence to which she has access. If this is right, then some epistemic 'ought's are subjective in the sense that the end to which they are (implicitly) related involves factors only internal to the subject's ken, while other epistemic 'ought's are objective in the sense that the end to which they are (implicitly) related involves factors at least partially external to the subject's ken.

This raises an important question that I see as orthogonal to the monism vs. pluralism debate in epistemology: Is it then a subjective or objective epistemic 'ought' that is relevant for knowledge? This is a relative of the question that often seems to be at issue between internalist and externalist accounts of epistemic justification. That debate is often concerned with the issue of whether or not a belief's justification can outrun what is within the subject's ken. However, I think that the traditional way of framing the issue implausibly favors internalism because of the conceptual connections between 'justified' and 'reasonable'.<sup>14</sup> Framing the question instead in terms of what sort of 'ought' is relevant for knowledge doesn't import this prejudice, since, with the end-relational theory of 'ought's, there is clearly the conceptual space for both subjective and objective epistemic 'ought's; to get these, we just have to vary the relevant ends implicitly coupled to the 'ought's with the 'in order to...' modifier.

Although we clearly can evaluate what someone subjectively ought to believe, the anti-luck intuition about knowledge would suggest that this could not be the only sort of evaluation in play when we evaluate someone as having knowledge. If someone has gotten things right due to luck, as in the case of the widow who truly but luckily believes that her husband has died, we tend to think that she doesn't have knowledge, even if she believes as she *subjectively* ought to believe—that is to say, even if based on internally accessible considerations, she believes as she ought to believe. This suggests that the 'ought to believe' in an account of knowledge should be an objective 'ought'. However, the clairvoyant who has safe true belief but no idea why his belief is safely true is someone who believes as he objectively ought, but he also seems to fall short of knowledge. This, by contrast, suggests that there must be some subjective element to the 'ought to believe' relevant for knowledge. This is a recapitulation of the standard debate between internalists and externalists about knowledge.

Luckily, the end-relational theory of 'ought' provides a nice way transcend this debate. There are several options here, all of which are compatible with the general account of knowledge that I am proposing, but so as to have a specific version on the table, let me propose the following. Generically, we might begin by saying that the 'ought' relevant for knowledge is

---

<sup>14</sup> In fact, this is part of the reason that externalists often encourage us to jettison the notion of 'justification' in the theory of knowledge. See for example, Alston (2005) and for further discussion my (2007b).

what one ought to believe in order to believe knowledgeably. While true, this is clearly uninformative. A more informative answer comes from combining internalist and externalist ideas and saying that what one ought to believe in order to believe knowledgeably is what one ought to believe in order to have a safe and rationally defensible belief. By saying this, we respect externalist motivations by making *safety* crucial to the account of the epistemic 'ought' relevant to knowledge.<sup>15</sup> However, the account also respects internalist motivations by making *rationally defensibility* equally central. The end-relational theory of 'ought' provides a way to incorporate both of these elements into our account of knowledge as true belief that one ought to have.

How exactly are we to understand safety and rational defensibility? Again, there may be several options compatible with my general idea, but since it will prove useful later to have a more specific version of the account on the table, let me say a little more about how I am understanding each of these notions and the way they incorporate aspects of two recently prominent accounts of the semantics of knowledge attributions.

First, safety is often understood in terms of whether the belief is true in all nearby possible worlds. However, since some of our beliefs have necessarily true propositions as their contents, this understanding of safety can't work in general. For these beliefs are true in *all* possible worlds and thus true in all *nearby* possible worlds. Because of this, I think we should understand safety in terms of the reliability of the process by which the belief is formed. And this can, in turn, be understood in terms of whether the process generates true beliefs in all nearby possible worlds.<sup>16</sup>

Obviously, such epistemic safety comes in degrees. So, I think we should take the degree of safety required for knowledge to be a matter determined by context of attribution. In some contexts, the process by which a belief is formed must be reliable across a wider set of possible worlds to count as what one ought to believe in order to have a safe belief, whereas in other contexts a belief can be formed by a process that is reliable across a narrower set of possible worlds but still count as what one ought to believe in order to have a safe belief. This is a straightforwardly contextualist aspect of my account of knowledge.<sup>17</sup>

---

<sup>15</sup> Compare Sosa (2007) for a different attempt to combine internalist and externalist factors into an account of knowledge. In general, I suspect that many aspects of the account sketched here could be transposed into a virtue-theoretic key. However, I have found the notion of 'ought' to be more flexible and thus more theoretically useful than the notion of 'virtue' in developing an account of knowledge.

<sup>16</sup> In going this way, I think we bring the safety element much closer to the original idea of reliability operative in reliabilist accounts. As such, we also face the generality problem originally articulated in Goldman (1979) and forcefully argued in Conee and Feldman (1998). Here, I don't attempt to solve the generality problem.

<sup>17</sup> In this respect, the account is similar to the contextualist accounts of knowledge attributions proposed by DeRose (1995) and Lewis (1996), though I don't here commit to any particular account of the ways in which conversational dynamics might affect the degree of safety required for a belief to count as knowledge.

A common complaint against contextualist accounts of ‘knows’ is that they obscure the evaluative aspect of knowledge attributions. This is because they treat ‘knows’ as a verb referring not univocally to a two-place relation between a subject and a fact but rather variously to different three-place relations between a subject, a fact, and epistemic standards; the idea is that the epistemic standards are relevant for individual tokenings of a knowledge attribution and correlatively which three-place knowledge relation is the reference of the tokening of the verb ‘knows’ vary according to the context in which it is tokened (as opposed to the context of the subject). This commits the contextualist to the possibility that, from one context of attribution, ‘knows’ picks out a relation that the subject stands in to the fact and low standards, while from another context, ‘knows’ picks out a relation that the subject does *not* stand in to the fact and high standards. The reason critics complain that this obscures the evaluative aspect of knowledge attributions is that it appears to undermine the usefulness of knowledge attributions to criticize or legitimate actions based on the putative knowledge. For example, when half way out the door, if one has pretty good evidence for the belief that—but does not know that—the stove is off, one might be criticizable for not going back to check. Similarly, if one spends weeks digging about for a lost treasure on an island, this action might be legitimated by the fact that one knows that the treasure is buried somewhere on the island. The problem for straightforward versions of contextualism is that there is no univocal, context-invariant answer to the question whether one knows such things. As a result, it becomes unclear how, on the contextualist account, knowledge attributions can play this evaluative role of criticizing and legitimating actions based on the putative knowledge.<sup>18</sup>

Epistemologists who push this line of criticism against contextualism sometimes take to be a reason to favor a view of knowledge which treats it as a single relation, picked out by the verb ‘knows’ in all of its tokenings; a relation, however, that is sensitive to the subject’s practical context (as opposed to the ascriber’s epistemic context). The idea is that what it takes to stand in the knowledge relation to a fact varies depending on what is at stake in the context in which the potential knower finds herself. If a lot hangs on the line, it will be more difficult for the subject to know that *p* than if relatively little hangs on the line. Either way, from whatever context we ascribe knowledge, we pick out this single relation. This view is sometimes called *subject-sensitive invariantism*.<sup>19</sup>

---

<sup>18</sup> See Hawthorne (2005: 86) and Williamson (2005) for versions of this objection. I have tried to answer this challenge on the contextualist’s behalf in my (under review), where I argue that the objection works only if we assume that the relevant criticisms and legitimizations of actions are not themselves context-dependent in a way mirroring the contextualist’s view of knowledge attributions. I think neither side of the debate has argued effectively that they are or are not context dependent in this way. So, I don’t think this objection undermines straightforward versions of contextualism; nevertheless, I think there is something right in the vague complaint that standard versions of contextualism miss the normativity in knowledge attributions. See my (2007a) for an initial statement of this worry and an attempt to resolve it. The present account represents an elaboration of that idea in slightly different terms.

<sup>19</sup> See Hawthorne (2005), Stanley (2005), and Hawthorne and Stanley (forthcoming).

I doubt that this is the right moral to draw from the ways in which knowledge attributions can be used to criticize or legitimate action. This is because it seems to me that there could be situations where a whole lot hangs on the truth of some proposition  $p$  but very little hangs on the truth of another proposition  $q$  but also where it seems highly plausible that one knows that  $p$  just in case one knows that  $q$ .<sup>20</sup> For example, consider a situation where you are traveling to London for a day of touring the art galleries and you overhear someone in the train saying “It’s going to be a rainy day in London, too bad the art galleries will be closed.” Not much hangs for you the truth of the proposition that it will be a rainy day in London; you have an umbrella and are accustomed to enjoying rainy days in London. However, a whole lot hangs for you on the truth of the proposition that the art galleries will be closed; after all, a great way to enjoy a rainy day in London is doing what you have planned to do: tour the art galleries. In such a case, assuming you have no specific prior evidence for either proposition and assuming that overhearing someone in the train can, in general, be a way to come to know something at least sometimes, the subject-sensitive invariantist would have to allow that you might have come to know that it will be a rainy day in London but not that the art galleries will be closed on the basis of overhearing what the person said on the train. This is because a lot hangs for you on the former proposition, but not very much hangs on the latter. But this seems absurd. If you know one of these propositions on the basis of overhearing what the person on the train said, then you know the other; or if you don’t know one of the propositions on the basis of overhearing what the person on the train said, then you don’t know the other either. Because of this, I don’t think we should capture the idea that knowledge attributions are evaluative directly in terms of how the subject’s practical context affects how difficult it is to stand in the knowledge relation to some fact. Rather I think we should capture it in terms of what it takes for a subject’s belief in a proposition to be one he ought to have in order to have a rationally defensible belief.

So, second, rational defensibility: It is convenient to understand this notion in terms of procedures of rationally defending a belief—that is, in terms of what it would take to establish one’s epistemic right to hold the belief in the face of a legitimate challenge. I take both the notions of ‘establishing an epistemic right’ and ‘legitimate challenge’ to be sensitive to the subject’s *de facto* epistemic situation. This is a subject-sensitive aspect of my account of knowledge. And it is also a key place where Sellars’ idea that ascriptions of knowledge “place one in the logical space of reasons, of justifying and being able to justify what one says” has influenced my account.

When is a subject’s belief open to legitimate epistemic challenge? In most epistemic contexts in which we find ourselves, many sorts of epistemic challenges to a belief would not be legitimate unless there were some positive reason to think that the belief might be false. For example, if you are at the

---

<sup>20</sup> Neta (forthcoming) provides a more careful statement of the following sort of argument against subject-sensitive invariantism.

bank where someone asks you what day it is, and you consult your watch before reporting “It’s Monday March 2<sup>nd</sup>,” it would not normally be legitimate to challenge your belief by responding “But have you ruled out the possibility that you’re being deceived by an evil demon?” However, it may be legitimate to challenge your belief by responding “This year is leap year, does your watch correct for that?” In general, I think there are norms governing the legitimacy of epistemic challenges, and these norms vary from situation to situation. But in every case a challenge is legitimate not because of the mere metaphysical or logical possibility that one could have been mistaken but rather because there is some positive reason to think that one could have been mistaken.

So, my idea of rational defensibility is that one’s belief counts as rationally defensible if one can meet all legitimate epistemic challenges in one’s *de facto* epistemic situation.<sup>21</sup> What does it take to be able to meet an epistemic challenge? Again, there may be several ways to go here that are all consistent with the overall account of knowledge that I am sketching; however, it seems to me that there are both forward- and backward-looking elements to being able to meet an epistemic challenge. In most cases, one will not have consciously considered all of the counter-possibilities to one’s belief, but, if they were brought up, one could easily rebut them. That means that, looking forward from the belief, one has it within one’s ken to adequately respond to possible legitimate challenges. For example, in response to the challenge generated by the possibility that your watch hasn’t corrected for the leap year, you might be able to say “I changed it for leap year yesterday.” You might not consciously consider this when forming the belief, but your ability to cite this fact would, in most circumstances, mean that you could adequately meet this challenge. And, more generally, it doesn’t matter if this or any other challenge actually comes up, as long as you are able to respond and thereby establish, according to the epistemic norms of your *de facto* epistemic situation, your epistemic right to the belief. However, sometimes there are cases where you could establish this right but only, as it were, by reasoning post-hoc. For example, if you guess the answer to a multiple choice question on a test but could back it up in the face of legitimate challenge were you to think about it, we might want to say that your belief is not actually rationally defensible until you actually think about it. This suggests that there is a standing norm to the effect that mere guesses don’t count as knowledge. Even if one doesn’t have to be able to articulate the precise psychological source of the belief, looking backwards, one must be able to point to some positive source of the belief besides a mere guess. That is one must be able sincerely to say things such as “It seemed to me that p was the most plausible explanation” or “I saw with my own eyes” or “A reliable informant told me that q, which made me think that p”, etc.

---

<sup>21</sup> This account of rational defensibility owes much to Rosenberg’s (2002: 193-200) “proceduralist” account of justification, according to which what it is to be justified in a belief is to be able to justify the belief. Compare also the “Default-Challenge” conception of justification and knowledge defended by Williams (2001: 150-157).

So far, then, I have made five distinctions under the rubric of “the many senses of ‘ought’”. These are:

- (i) nonnormative vs. normative ‘ought’s
- (ii) hypothetical vs. categorical ‘ought’s
- (iii) all things considered vs. in some respect ‘ought’s
- (iv) moral, prudential, legal, etc. vs. epistemic ‘ought’s
- (v) subjective vs. objective ‘ought’s

With respect to (i)-(iii), I’ve claimed that the second member of each pair is the one relevant for understanding knowledge. With respect to (iv), I’ve argued that the end-relational theory provides a nice way to distinguish between different ‘ought’s deriving from different kinds of norms. Context determines which one is relevant for a particular ought-claim, yet I’ve warned against thinking that truth is the monistic end determining a specifically epistemic context. I think there are many different possible epistemic ends, but, with respect to (v), I’ve claimed that the epistemic ought relevant to a substantive account of knowledge is the mixed objective-subjective end of safe and rationally defensible belief, where I am understanding what it takes for a belief to be safe as a matter sensitive to the ascriber’s context and what it takes for a belief to be rationally defensible as a matter sensitive to the subject’s epistemic situation. This allows me to make a third and final pass at my account:

*The Account:* S knows that p just in case S’s true belief is one that S normatively and categorically ought to have, in order to have a safe and rationally defensible belief that p, (where what counts as ‘safe’ is determined by context of attribution and what counts as ‘rational defensible’ is determined by context of belief).

### 1.3 The Normativity of Knowledge

With the account now stated more precisely, I want to return to my initial motivation for this account stemming from the slogan quoted from Sellars at the outset. In my view, epistemology is a properly normative discipline and we must understand knowledge attributions in such a way that helps us to integrate them into a more general account of normative evaluations. The way in which my account manifests this motivation is perhaps already quite apparent. Alongside ‘reason’ and ‘value’, ‘ought’ is surely a paradigmatically normative notion. So, by construing knowledge as partially constituted by belief which the knower ought to have, I think the account I have sketched clearly treats knowledge as normative and so knowledge attributions as a species of normative evaluations.<sup>22</sup>

---

<sup>22</sup> This raises an important question that I won’t address here, which is about the appropriate metanormative account of the semantics for knowledge attributions. This is typically done in terms of the truth-conditions of the claims. However, by construing them as a species of ought-claim, I think we open up theoretical space to consider all of the metanormative accounts

Of course, many other accounts of knowledge will claim to do this too in virtue of the fact that they take knowledge to require justification, and justification is a normative concept. However, whether that is a compelling way to capture the normativity of epistemic claims depends on how we conceive of justification. For reliabilist accounts, which conceive of justification in terms of being reliably formed, I think it's far from clear that they really do conceive of justification as a normative notion. After all, the issue of whether a process is reliable is generally thought to be a nonnormative matter.<sup>23</sup> Moreover, some accounts that eschew the project of analyzing knowledge altogether to insist instead that knowledge is the most general factive mental state<sup>24</sup> will also have a difficult time explaining how the study of knowledge is a normative discipline and how knowledge attributions are a species of normative evaluations. For the issue of whether a person is in a particular mental state is generally thought to be a nonnormative matter.

Of course, in some very broad sense, we *evaluate* a thermometer or a clock as reliable, and of course we can *evaluate* various belief-forming mechanisms as reliable in the same sense. And, also in a similarly broad sense, we evaluate people with respect to what state of mind they are in. But it is far from clear that either of these sorts of "evaluations" has anything to do with normativity or the way in which I want to see knowledge attributions as normative judgments. To put the point in Sellars' terms, judgments of reliability or what mental states one is in seem to be "empirical descriptions of an episode or state", whereas an account of knowledge in terms of what subjects ought to believe does a better job of explaining the way in which attributions of knowledge to rational inquirers "place one in the logical space of reasons".

Those sympathetic to reliabilism or the idea that knowledge is the most general factive mental state can, of course, deny that epistemology is a normative discipline in the sense of "normative" that I am using. So, I mention the way in which my account nicely captures the normativity of knowledge and knowledge attributions not because I think that this provides a knock-down argument against these competitors, but rather because it helps to explain some of the motivation for the account. Ultimately, the question of whether and in what sense epistemology is normative is probably an issue where "spoils go to the victor".

---

available for ought-statements more generally. In a previous paper, I floated the idea that knowledge attributions might admit of a plausible expressivist alternative to truth-conditional accounts. See my (2007a). While I no longer think it is right to see knowledge attributions as getting their meaning from the attitude that they conventionally express, I retain my opposition to truth-conditionalism in the semantics of knowledge attributions (and normative statements more generally).

<sup>23</sup> Brandom (2000: 97-122) argues against this claim, but this is by way of pointing out the drawbacks of standard versions of reliabilism, which attempt to reduce out the normative element in knowledge attributions.

<sup>24</sup> This is the sort of account defended in Williamson (2002).

## 2 The Motivating Intuitions and the Three Desiderata

An account of knowledge should simply and elegantly explain and unify our intuitions about knowledge. There are general intuitions about knowledge and more specific intuitions about a number of hard cases. At the beginning of this paper, I separated these into three vague and general intuitions I'd like to be able to respect as far as possible and three correlative diagnostic desiderata which have emerged from recent epistemological debate focused on a number of specific cases. In this section, I'll discuss both how I think my account comports with each of the vague and general intuitions and what resources I think it bring to meeting the diagnostic desiderata.

### 2.1 The Anti-Luck Intuition and the Gettier Problem

The first vague and general intuition was the *anti-luck intuition* that knowledge is incompatible with getting things right merely by luck. There are at least two different kinds of luck that affect getting things right. There's the luck that might be involved in being in the right place at the right time to form the relevant belief. This can be called *opportunity luck*. And there's the luck that might be involved in one's belief being correct given nearby possible worlds where one would have formed the same belief but it would have been false. This can be called *modal luck*.<sup>25</sup> Is knowledge incompatible with both kinds of luck? No. It's often a matter of luck that we look where we look, hear what we hear, etc., but given that we do look, hear, etc., as long as we do it competently, it seems that we can thereby acquire knowledge of what we've experienced by (opportunity) luck. However, if that experience happens to be indicative of the truth only by (modal) luck—in that we could have easily had the same experience, formed the same belief, but ended up with a false belief—then we don't have knowledge. This was illustrated in the case of the widow above. So it is modal luck rather than opportunity luck that is incompatible with knowledge. By conceiving of knowledge as true belief that one ought to have in order to have a *safe* belief, my account validates the idea that knowledge is incompatible with modal luck. Epistemic safety is safety from modal luck.

The Gettier Problem is related to the anti-luck intuition. Gettier famously proposed cases where it looks like the subject has a justified and true belief but intuitively fails to have knowledge. Any adequate account of knowledge will be able to explain the intuitive data surrounding Gettier-like cases. These sorts of

---

<sup>25</sup> Pritchard (2005: 134-147) offers a much more fine-grained set of distinctions between various kinds of luck that may be involved in belief formation; however, he recognizes what I am calling "modal luck" as the key sort of luck that seems to be incompatible with knowledge in Gettier-like cases. He also discusses what he calls "reflective epistemic luck" which is the luck exhibited when it's a matter of luck that a subject's belief is true, given her reflectively accessible evidence. It's unclear whether this sort of luck is always incompatible with knowledge, but, when it is, the account sketched here would rule it out in the subjective element in knowledge.

cases can be divided into two types.<sup>26</sup> First there are cases like the ones Gettier himself actually discussed.<sup>27</sup> These are cases where, although a subject ends up with justified true belief, the reason he believes what he believes isn't connected to the truth in the correct way. For instance, Smith justifiably believes that Jones owns a Ford and, from this belief, competently deduces that either Jones owns a Ford or Brown is in Barcelona. As it turns out, Jones doesn't own a Ford, but coincidentally Brown happens to be in Barcelona. So, Smith's belief that Jones owns a Ford or Brown is in Barcelona is justified and true. But it's only luckily true, and so we tend not to call it knowledge. Second, there are cases like the barn-façade case discussed by Goldman.<sup>28</sup> In these cases, one just happens to get lucky that the evidence one has for one's belief is truth-indicative; it could have very easily not been truth-indicative. For instance, riding through barn façade county, Henry happens to look up at the one real barn in the whole county and form the true belief based on visual evidence that what he is looking at is a barn; if he had happened to look at any of the other barn-like structures, he would have formed the same belief but it would have been false. Henry's belief that what he is looking at is true and based on what normally counts as pretty good evidence, but because it could so easily have been false we demur at calling it knowledge.

In both cases, the anti-luck intuition that knowledge can't be a matter of getting something right by modal luck leads most epistemologists to say that despite the presence of a justified true belief—at least given an internalist conception of 'justified'—there is no knowledge. But this verdict is not universal. Some epistemologists are willing to bite the bullet and say that some of these intuitions are mistaken; and undergraduate students often waiver on these cases.<sup>29</sup> I think this is because it seems that the subjects (especially in the second type of cases) have achieved something—correct belief—through the legitimate and apparently proper exercise of their belief-forming faculties. And although something is surely amiss, it seems overly demanding or even disrespectful to deny them the status of knowers.

So any theory of knowledge should be able to explain this tension between the anti-luck intuition that these subjects are not knowers and the sense that these subjects have gotten things right and not done anything wrong and so are fully praiseworthy, epistemically speaking.

---

<sup>26</sup> Pritchard (2008) makes this distinction.

<sup>27</sup> See Gettier *op. cit.*

<sup>28</sup> See Goldman (1973).

<sup>29</sup> Weatherson (2003) argues that these countervailing intuitions are outweighed by the theoretical utility and simplicity of the JTB account. Ernest Sosa suggests that these intuitions are at least partially misguided in barn-façade sorts of cases. In these cases, he thinks the subjects have "animal" knowledge by virtue of their getting things right in a way that manifests epistemic skill. See his (2007). See also Weinberg, *et. al.* (forthcoming) for empirical evidence suggesting that Western subjects are more likely than Eastern subjects to see Gettier cases as failures to know.

My account of knowledge as true belief that one ought to have brings distinctive resources to explaining what is going on with our intuitions in these cases. This is because of the distinction between subjective and objective senses of 'ought' made available by my appeal to the end-relational theory of 'ought'. There are at least two possible judgments about what one ought to believe that could come into play in Gettier-like cases. First, there is a subjective 'ought to believe', where we want to say that the Gettierized subjects believe what they subjectively ought to believe: They believe what they ought to believe in order to have a rationally defensible belief. I think it is this true claim about what the subjects ought to believe that leads us to the feeling that it is somehow unfair to withhold the status of knower from them. They have both achieved something—true belief through competent exercise of their belief-forming faculties—and done so in a way that is fully praiseworthy relative to what is within their ken. However, there is a second, more objective 'ought to believe', where we want to say that the subjects are not believing as they objectively ought to believe: in order to have beliefs that could not have easily been false, they ought not to believe what they in fact believe.<sup>30</sup> That is, from an objective point of view, their beliefs are very risky: in nearby possible worlds where they have the same experiences and use the same procedure to form their beliefs, these beliefs are false.

I think this is what Gettier-like cases highlight. For instance, had Brown happened not to be in Barcelona, Smith would have surely formed the same belief that either Jones owns a Ford or Brown is in Barcelona, but it would have then been false, even if rationally defensible. Likewise, had Henry happened to look at any of the many other barn-like structures in his field of vision, he would have surely formed the same belief that what he was seeing was a barn, but this belief would have been false even if rationally defensible. I think the reason these cases work so well as challenges to the classical JTB account of knowledge is that the way they are set up makes the end of safe belief particularly salient, and it is clear that the subjects do not have safe belief since their beliefs could have so easily been false.

So, if I'm right that knowledge is true belief that one ought have in order to have a safe and rationally defensible belief, then it's no wonder that we tend to think that the Gettierized subjects lack knowledge—in order to believe safely, they ought not to believe what they in fact believe. The countervailing feeling that the subjects are nonetheless not properly blamed and have indeed achieved something epistemically praiseworthy is, I submit, a reflection of our tendency to judge (implicitly and truly) that the subjects do believe what they ought to believe in order to have a rationally defensible belief, that is, in order to achieve some subjective epistemic end. That is how my account of knowledge coupled with our reflections on the end-relational theory of 'ought's and the

---

<sup>30</sup> Of course, there's also a different objective 'ought'—namely, what one ought to believe in order to believe truly. However, I argued above that this is not directly what is at issue for the epistemic sense of ought relevant to knowledge.

distinction between subjective and objective 'ought's brings new resources to explaining the intuitive data surrounding Gettier-like cases.

## 2.2 The Value Intuition and the Meno Problem

The second vague and general intuition to respect insofar as possible was the *value intuition* that knowledge is better than mere true belief. This intuition must be refined in two ways in order to be plausible. First, the answers to some questions, such as how many times Rawls said the word 'the', are probably not, all things considered, worth knowing. So, when we say that knowledge is better than mere true belief, we can't be talking about all-things-considered betterness. Second, some items of possible knowledge could hurt ourselves or others in serious ways. For instance, if a burglar threatens to kill your children if you acquire knowledge of his identity, then it's better, at least with respect to the safety of your family, to lack knowledge of his identity. So, we also can't be talking about prudential or moral betterness, rather we must be talking about *epistemic* betterness. Thus, a refined version of the value intuition is that knowledge is epistemically better than mere true belief.

My account comports with this intuition by conceiving of knowledge as true belief that one *epistemically* ought to have, which I cashed out as true belief that one ought to have in order to have a safe and rationally defensible belief. This explains how knowledge is epistemically better than mere true belief. For having mere true belief is not having what one ought to have in order to have a safe and rationally defensible belief; but having knowledge is having what one ought to have in order to have a safe and rationally defensible belief.

One might worry that using an end-relational 'ought' to explain the way knowledge is better than mere true belief will make that betterness necessarily hypothetical, that is, better only if we've set ourselves the relevant end. However, I don't think this is problematic. For, as we've already seen, with respect to moral or prudential considerations, knowledge may not be better than mere true belief; and, presumably, these considerations could outweigh epistemic considerations and make it the case that sometimes knowledge is not better than mere true belief, all things considered. However, the debate shifts when we restrict ourselves to specifically epistemic considerations. Here, I think the end of safe and rationally defensible belief is partially constitutive of knowledge, and the way that this is captured in my account is by construing the 'ought to believe' implicit in knowledge attributions as an ought to believe in order to have safe and rationally defensible belief.<sup>31</sup> The idea is that, because

---

<sup>31</sup> Compare the sense in which Plato and Aristotle seem to have thought that one ought to act morally in order to achieve the end of eudemonia. The idea is that acting morally is partially constitutive of living well, not an external means to it. Likewise, we drink nice wine in order to enjoy it, but drinking is partially constitutive of enjoying wine, not an external means to it.

safe and rationally defensible belief is partially constitutive of knowledge, and safe rationally defensible belief is the end relative to which knowers believe what they categorically ought to believe, knowledge is epistemically better with respect to this end.

The Meno Problem is related to the value intuition. Meno and Socrates famously struggled to explain why knowledge seems is better than mere true belief. They noted that true belief about the way to Larissa may get you to Larissa just as well as knowledge about the way to Larissa. Moreover, even if we add something to true belief that makes it better than mere true belief but not yet knowledge, it still seems that knowledge is better than any epistemic state short of knowledge. Generally, we'd rather know the way to Larissa than have true belief that is (say) accidentally justified. Any adequate account of knowledge will be able to explain this.

Now, one way to approach this issue is to argue that although mere true belief may get you to Larissa just as well as knowledge, there are other instrumental benefits knowledge sometimes has over true belief. That is to say, knowledge can be (part of) a better means for obtaining one of our contingent and personal ends than mere true belief. For example, if you want to impart knowledge about the way to Larissa to your audience, then knowledge about the way is clearly instrumentally better than mere true belief; for you cannot impart knowledge which you do not have.

However, as far as the Meno Problem goes, I think we must abstract away from the ways that knowledge can be instrumentally better than mere true belief. This is because the instrumental value of knowledge is relatively clear and yet knowledge has also seemed to many to be *non-instrumentally* better than mere true belief. This is both in the sense that knowledge is sometimes seems to be worth pursuing for its own sake and in the sense that, in an ordering of epistemic states from better to worse, knowledge seems to rank higher than mere true belief. This must be explained. However, because even non-instrumentally valuable knowledge can have detrimental effects, I think the sort of non-instrumental value we should think knowledge can have is non-instrumental value *with respect to epistemic considerations* rather than with respect to moral considerations or all things considered.

A standard way to try to explain this value is by arguing that, in addition to whatever instrumental value knowledge has for achieving our contingent and personal ends, it also has value in its teleological relation to some distinctive or definitive epistemic good.<sup>32</sup> For example, truth monists typically argue that truth is the distinctive epistemic good. And knowledge is epistemically valuable in that it promotes this good.

---

<sup>32</sup> The only prominent alternative in the literature is the virtue theoretic idea that knowledge is an achievement (success through or because of or manifesting ability). See Grecco (2003, 2008), Sosa (2007). In my view this fetishizes achievements and I won't discuss it further here, but see my (ms).

The problem with this particular proposal is that the value that accrues to knowledge in virtue of its teleological relationship to truth seems to be swamped in cases where true belief is achieved. This is the “swamping problem” that has emerged in response to reliabilist accounts of knowledge.<sup>33</sup> The problem for such accounts is that it seems that the reliability of a mechanism that formed a true belief cannot make that true belief any better, epistemically speaking, than it already was in virtue of being true. Compare: a gastronomically good cup of coffee is one that tastes good; so two identically tasting cups of coffee are gastronomically on par, even if they are made by different methods, one that reliably gets things right and one that gets things right only by a fluke. Likewise, the argument goes, an epistemically good belief is one that is true; so two beliefs that p which are both true are epistemically on par, even if they are produced by different belief-forming mechanisms, one that reliably gets things right and one that gets things right only by a fluke.<sup>34</sup> This seems to show that the reliabilist account of knowledge cannot explain why knowledge is apparently epistemically better than merely true belief; the reliability of the mechanism by which a true belief is formed doesn’t seem capable of making that belief any better, epistemically, than it already was in virtue of being true.<sup>35</sup>

This is not, it should be noted, only a problem for reliabilist accounts of knowledge. It’s a problem for any account that sees knowledge as true belief with some further property whose epistemic value derives from its connection to the truth of the belief. For example, if you think that knowledge is true belief based on reasons, and what makes being based on reasons good for a belief is that being based on reasons makes the belief more likely to be true, then you too will have a hard time explaining why knowledge seems better than merely true belief. For whatever value reasons bring to a belief they support will be swamped by the value of the belief already being true.

The account of knowledge developed above provides a natural way to respond to this challenge. The swamping problem clearly trades on thinking that truth is the only epistemic end relative to which various properties of a belief could be epistemically evaluated. However, as I argued above, we don’t have to see truth as the sole epistemic end. If there are multiple epistemic ends, then, although a merely true belief won’t be any better for getting the truth than a reliably formed true belief or a true belief based on sufficient reasons, it could be better for getting one of the other epistemic ends. For instance, reliably formed beliefs are safer than merely true beliefs; and beliefs based on

---

<sup>33</sup> See especially Zagzebski (1996), Jones (1997), Swinburne (1999), Riggs (2002), and Kvanvig (2003) for examples of pressing the Meno problem against reliabilist accounts.

<sup>34</sup> This example is due to Zagzebski (1996).

<sup>35</sup> This is the sort of argument against reliabilism pursued in more detail in Kvanvig (2003: 44-60). There is a whole literature now on the “swamping problem”; see fn. 28 above for some relevant literature. It would take me too far a field to critically assess even the most promising of the proposed solutions. As will come out in the text to follow, I think much hangs on whether we endorse and how we understand the idea that truth is the definitive epistemic value.

reasons are more rationally defensible than mere true beliefs. So, if we think that S knows that p because S truly believes that p and S ought to believe that p in order to have a safe and rationally defensible belief, then clearly such knowledge will be better than mere true belief vis-à-vis the end of having a safe and rationally defensible belief.

What I think this reveals is that, as long as we recognize the ends of having a safe belief and of having a rationally defensible belief as parts of the epistemic end relevant to knowledge—as my account urges—it’s easy to see that knowledge is better with respect to these ends than merely true belief. Compare: A cup of coffee that tastes good by accident is obviously just as good gastronomically as a cup of coffee that tastes good because of the barista’s skill. However, that doesn’t imply that the latter cup of coffee is no better than the former with respect to some other ends. For instance, the cup of coffee that tastes good because of the barista’s skill is clearly better with respect to the end of having a cup of coffee that could not have easily been bad or with respect to the end of demonstrating the barista’s skill than the cup of coffee that tastes good by accident.

So, it is in this way that I think the end-relational theory of the distinctively epistemic ‘ought’ helps us to see why knowledge seems better than mere true belief. It seems better because it is better for certain ends, like having safe beliefs or having rationally defensible beliefs. Conveniently, a very similar strategy can be used to explain why knowledge seems better than any epistemic state short of knowledge. These states may be as good as knowledge for certain ends such as truth or justification, but they will not be as good as knowledge for other ends such as safety or rational defensibility. For example, the sort of accidentally justified true belief we see in Gettier cases is just as good as knowledge with respect to the end of truth and even with respect to the end of having a rationally defensible belief, but knowledge is clearly better with respect to the end of having a safe belief. For an accidentally justified true belief is not safe. This is how the account I’ve been defending brings new resources to the first aspect of the Meno problem.

The failure of reliabilist accounts and other accounts to explain the value of knowledge stems, I am suggesting, from being blinded to the possibility that knowledge may be no better with respect to the end of truth, all the while being better with respect to other ends, such as safety or rational defensibility. So, when combined with the end-relational conception of ‘ought’, the idea that knowledge is true belief that ought to be believed helps to explain the intuitions driving the Meno problem.<sup>36</sup>

---

<sup>36</sup> It’s worth mentioning that you don’t have to accept my view that there are multiple epistemic ends to accept this response to the Meno problem. I’ve claimed that knowledge *seems* better than epistemic states short of knowledge because it is better for certain ends. Those ends may themselves not be epistemically relevant, in which case my response to the Meno and swamping problems becomes a debunking response. On this line of thought, knowledge turns out to be no better epistemically than other forms of true belief; however, we’re misled in to the

## 2.3 The Anti-Skeptical Intuition and the Cartesian Problem

The third general intuition I wanted to respect as far as possible was the *anti-skeptical intuition* that there are many things that we know. My account comports quite easily with this intuition as long as many of the propositions we actually believe are propositions that we ought to believe in order to have a safe and rationally defensible belief. Since my account has it that what counts as ‘safe’ is a matter of context of attribution and what counts as ‘rationally defensible’ is a matter of context of belief, all that is required to validate the anti-skeptical intuition is the thought that there are many context-of-attribution-context-of-belief pairs for which the bar is set low enough in the context of attribution that many of the true beliefs we actually hold are beliefs formed by processes that are reliable in nearby possible worlds and that and the bar is set low enough in the context of belief that many of our true beliefs are ones for which we can meet legitimate challenges. I think this is generally true. Although there may be far off possible worlds where my vision fails to deliver true beliefs about my immediate surroundings or my geography lessons fail to deliver true beliefs about world capitals, these possible worlds are not typically relevant to whether I could have easily been mistaken in forming beliefs in these ways. Likewise, although I might not be able to defend such ordinary beliefs against epistemic challenges stemming from extraordinary counter-possibilities, such as brains in vats and evil demons, these challenges are not normally legitimate in ordinary contexts of belief and so the beliefs may be perfectly rationally defensible nonetheless.

The Cartesian Problem is related to the anti-skeptical Intuition. Much ink has been spilled about different kinds of skepticism. The kind I think we should be most concerned with for the theory of knowledge is the kind raised by skeptical counterpossibilities such as the possibility that I am being deceived by an evil demon. This is because this sort of skepticism threatens our own first-personal claims to know various things, which is a corollary of the more general threat to undermine our widespread practice of attributing knowledge fairly liberally, which, I would submit, is the best handle we have on what we mean by ‘knowledge’ in the first place.<sup>37</sup> Such skeptical problems are usually raised by taking almost any proposition *p* we might ordinarily claim to know and then

---

intuition that it is better epistemically because it is better in some other putatively nonepistemic respect.

<sup>37</sup> Many epistemologists think that a more serious skeptical challenge comes not at the level of whether we have knowledge, given skeptical counterpossibilities, but at the level of whether any belief could be justified at all, given the absence of self-justifying starting points and the unsoundness of reasoning in a circle or an infinite regress. Compare Fogelin (1994: ch. 6) and Williams (2001: 58-68). I agree that this is a different sort of skepticism and an important epistemological challenge. However, I don’t think that this undermines my claim that challenges to our knowledge from skeptical counterpossibilities is the sort of skepticism that we should be concerned about in the theory of knowledge. Perhaps, this is where the theory of knowledge comes apart from the project of responding to the skeptic.

pointing out that our evidence for  $p$  is consistent with another proposition  $q$  that implies  $\text{not-}p$ . For example, if we know that it's dark outside based on our visual impression as if it's dark outside, the skeptic points out that this evidence is consistent with our being deceived by an evil demon to have visual impressions as if it is dark outside even though it really isn't dark outside. Then, the challenge is this: all of our evidence is consistent with the possibility that we are deceived by an evil demon, so it doesn't seem like we know this possibility not to obtain; and, if that's right, then we should be skeptical of the claim that we know that it's dark outside. That example is first-personal, but the more general case is third-personal. For example, we could be discussing your mother's knowledge that you are not at home, and the skeptic challenges this attribution of knowledge by pointing out that your mother's evidence for this proposition is completely consistent with her being deceived by an evil demon. Since she doesn't know this possibility not to obtain, the skeptic urges us to recant our third-personal attribution of knowledge much like we were urged to recant our first-personal attribution of knowledge.

The seductiveness of this sort of skeptical reasoning must be explained in a way consistent with our validation of the anti-skeptical intuition. Initially, the diagnostic resources my account of knowledge brings to this issue are straightforwardly contextualist. This is because the account has it that  $S$  knows that  $p$  only if  $S$ 's belief that  $p$  is one  $S$  ought to have in order to have a safe belief, where what counts as 'safe' is determined by context of attribution. Above I suggested thinking of safety as a matter of whether the process by which the belief was formed delivers false beliefs in nearby possible worlds. The contextualist idea, then, is that, what counts as 'nearby' is determined by the context of attribution. This allows us to see skeptical arguments as implicitly attempting to create a context where the scope of what counts as 'nearby' is broadened to include a larger and larger class of possible worlds—large enough, in fact, that, in one of these worlds, the belief-forming process is unreliable because (say) one is deceived by an evil demon. Of course, in more ordinary contexts, it takes much less to have a safe belief.<sup>38</sup>

I think the contextualist idea embedded in my account helps to meet the Cartesian problem. However, some epistemologists have worried that the contextualist response to skeptical arguments from counter-possibilities is too concessive. They grant that it legitimates many ordinary knowledge attributions made in contexts where the range of possible worlds is relatively restricted. However, they will point out that contextualism allows for the possibility of high-standards contexts where the contextualist must admit that the skeptic is right: we do not know what we think we know.

I think the strength of this worry depends on how easy it is to raise the epistemic standards in a context of knowledge attribution. Early discussions of contextualism were dominated by Lewis' (1996) "rule of attention", according to which merely getting one's interlocutors to attend to the skeptical scenarios is

---

<sup>38</sup> This is roughly the sort of response to skeptical arguments defended in DeRose (1995).

enough to raise the standards. I think this cannot be right; there must be a way to refuse to go off on philosophical flights of fancy.

The idea that knowledge is, in part, a matter of what one ought to believe in order to have a rationally defensible belief points towards a response to this worry. Skeptical arguments from counter-possibilities are often addressed at our first-personal claims to know, and, when they are not, the fact that knowledge is factive means that we can not typically assert that some other person knows that *p* unless we take ourselves know that *p*. And, if we view the skeptic's appeal to counter-possibilities as epistemic challenges to our right to hold these beliefs, then we can ask: when are these challenges *legitimate* challenges that we must be able to answer in order to rationally defend our belief? Above I suggested that they must be challenges for which there is some positive reason to think that the relevant counter-possibility holds. My account is consistent with different views about when there might be a positive reason to think that a skeptical counter-possibility holds; however, the sorts of wild counter-possibilities adduced in text-book skeptical arguments do not seem to me to be ones where there is some positive reason to think that they hold. In these cases, it seems to me that the appropriate response to the epistemic challenge from a skeptical counter-possibility is: "You say that I could be deceived by a demon, but why think that?" In this way, we can refuse to let the epistemic standards raise in a way that is concessive to the skeptic.

You don't have to agree with me in this response to the skeptic, in order to see it as a response to the worry generated by the contextualist response to skeptical arguments from counter-possibility. All that is important is that my idea that one knows only if one believes what one ought to believe in order to have rationally defensible belief provides a way either to respond to the worry or to diagnose its source. If you think there are some contexts where there is positive reason to think that a wild skeptical counter-possibility holds, then you will have a concessive attitude towards the skeptic; the reason I and your opponents think this is wrong is that we think there are different norms for what it takes for there to be some positive reason to believe something. If you are right, my account helps us to see the worry as based on misguided optimism, and if we are right then my account helps us to see why we don't have to concede so much to the skeptic. Either way, my account of knowledge as true belief that one ought to have brings new resources to the Cartesian problem.

So, it is these considerations about the context sensitivity of safety and rational defensibility that provide resources for addressing the Cartesian problem. Most of our ordinary knowledge ascriptions are made in contexts where it's simply true that the relevant belief is safe, on the contextually determined sense of 'safe'. So, many of these knowledge attributions will, on the account I've been defending come out as true. Moreover, even in contexts where one is challenged with a skeptical counter-possibility, it is far from clear that this challenge is a legitimate one that one must be able to meet in order to count as a knower.

Of course, it's a tricky issue saying how exactly conversational dynamics determine the vague extensions of words whose extension expands and contracts depending on context. For my purposes here, I don't need to grant that or explain how the skeptical arguments are actually successful in creating a context where pretty far-fetched counterpossibilities are enough to make it the case that our beliefs could have easily been wrong. All I need is that they can tend to make it seem this way to us when we are taking skeptical arguments seriously. The skeptical arguments are persuasive, insofar as they are persuasive, because knowledge involves what one ought to believe in order to have a safe and rationally defensible belief. And the notions of 'safety' is sensitive to features of the context of attribution, whereas the notion of 'rational defensibility' is sensitive to the context of belief in ways that can be influenced in ways to make skeptical arguments seem persuasive. Perhaps there are contexts in which the skeptical arguments are indeed sound, or perhaps there are only contexts in which the skeptical arguments seem to us to be sound but actually are not sound. Either way, the context sensitivity of these notions can explain the apparent persuasiveness of skeptical arguments and help to respond to the Cartesian problem.

### 3 Conclusion

That completes my positive case for the account of knowledge as true belief that one ought to have. This case involved a careful articulation of the notion of 'ought to believe' in terms of the end-relational account of 'ought'. Specifically, I suggested that this is to be understood as 'normatively and categorically ought to believe in order to have a safe and rationally defensible belief'. The case for the account also involved showing how it explains and underwrites the idea that epistemology is a properly normative study. Finally, the positive case involved discussion of how the account comports with the anti-luck intuition, value intuition, and the anti-skeptical intuition, which results in new explanatory and diagnostic resources for addressing the Gettier problem, the Meno problem, and the Cartesian problem. I'd like to close by rebutting what may seem like the two most natural objections; both attach to the centerpiece of my account, namely, the role of 'ought' in explaining knowledge.

First, one might wonder whether we couldn't drop the part of the account about what one *ought* to believe and just say that knowledge is true belief that is safe and rationally defensible. What does the 'ought' and the end-relational theory of 'ought's add? I think the answer to this depends on how we understand the notion of a rationally defensible belief. If we conceive of it as a nonnormative notion, then the account pared of its appeal to what beliefs one ought to hold loses the Sellarsian insight that knowledge ascriptions are not mere empirical descriptions of a state but evaluations of a state which place it into the logical space of reasons. However, if we understand rational defensibility as itself an irreducibly normative notion, then I suspect that my

account would be extensionally equivalent to the pared down account. However, extensional equivalence is not enough for a good account. To see this, notice that there may be some description of what we morally ought to do in purely psychological terms. However, this psychologistic description will lack the resources to explain many of the intuitions about what we ought to do. This is not necessarily because it will be intractably complex. Maybe “acts that maximize overall happiness” is extensionally equivalent to “what one morally ought to do”. However, the former concept will lack the explanatory force in ethics had by the latter concept. By incorporating the notion of ‘ought’ into the account of knowledge we gain certain explanatory resources even if there is some other account which is extensionally equivalent. It does so, in effect, by making explicit the fact that the inferential role of the concept of knowledge is caught up with inferences characteristic of normative inferences more generally; this is why knowledge attributions can be used as a species of normative evaluations. Other concepts may pick out the same class of knowers, even in all possible worlds, but they will not have the same inferential role in our cognitive economies.

Second, it may seem that the account I have sketched rests on a deontological conception of epistemic justification, and, as Alston has famously argued, deontological conceptions of knowledge illicitly import ethical notions such as duty, obligation, and—most crucially here—ought into a domain where they don’t belong.<sup>39</sup> According to him, the reason he thinks they don’t belong in the domain of epistemology is that epistemology is concerned with beliefs, and beliefs are things over which we don’t exercise the sort of control necessary to make sense of ideas like performing our epistemic duties, fulfilling our epistemic obligations, or believing as we ought to believe. Or, to put the point differently and specifically in terms of ‘ought’: one of the few venerable ideas in philosophy is that ‘ought’ implies ‘can’ in the sense of ‘voluntarily can choose to’; but we cannot voluntarily choose to believe, and so epistemic deontology is hopeless. And, as a result, my account of knowledge as true belief that one ought to have is radically mistaken.

I think this objection is misguided in at least two ways. The first is that I don’t think it is right to think of the account of knowledge that I have articulated and defended here as resting on a deontological conception of epistemic justification. I haven’t proposed an account of epistemic justification is, and I certainly haven’t used the terms ‘epistemic duty’ and ‘epistemic obligations’. That’s not because I think we don’t have epistemic duties and obligations; rather it’s because I think these are not what we are talking about when we talk about what beliefs someone ought to have. In my view, talk of epistemic duties and obligations is best thought of in terms of what we ought to do as good inquirers: we ought to do things like consider lots of evidence, talk to experts, think outside the box, etc. Now, in my account of rational defensibility, I have appealed to the notion of defending one’s ‘epistemic right’ to hold a belief. But

---

<sup>39</sup> See Alston (1988); see my (forthcoming and 2007b) for more discussion of his arguments.

the notion of a right to hold a belief much like the notion of a right to an honorary title doesn't presuppose that one can voluntarily and directly choose to have the thing to which one has a right. Of course, the idea of defending one's right may presuppose that one can voluntarily choose to engage in actions of defense, but I see nothing untoward about that presupposition of my account. Likewise, actions such as considering lots of evidence, talking to experts, and thinking outside the box are also actions that we can voluntarily choose to do. But these "epistemic actions" don't, in my view, exhaust the domain of epistemic 'ought's.

The only reason why you might think that they do—and this is the second thing misguided with the objection—is if you take the venerable idea that 'ought' implies 'can' in a much too strong way. I think it's undeniable that clocks ought to strike 12 at noon, classrooms ought to be kept tidy, dogs ought to be walked often, children ought to learn to count, parents ought to be more compassionate than they generally are, and compulsive drunks ought to abstain from the next drink. But none of these 'ought's are 'ought's that imply 'voluntarily can choose to'. What we need, as I have urged in another paper<sup>40</sup>, is to recognize a distinction between 'ought-to-do's applying narrowly to actions within an agent's voluntary control and 'ought-to-be's applying to the state of a thing regardless of whether it is an agent. If we do that, then we can treat 'ought-to-believe's as a species of 'ought-to-be's in a way that makes it perfectly compatible with Alston's observation that we don't typically exercise direct voluntary control over our beliefs.

If we do that, then I think it makes perfect sense to talk about what one ought to believe in order to have a safe rationally defensible belief. And my suggestion here has been that, by thinking of knowledge as true belief that one ought to have in this sense, we accomplish three very important things in the theory of knowledge. First, we make sense of the normativity of epistemology and begin to integrate epistemic evaluations into our understanding of normative evaluations more generally. Second, we provide a way to respect plausible versions of the anti-luck, value, and anti-skeptical intuitions. Third, we bring new unified theoretical resources to the diagnostic demands of the Gettier, Meno, and Cartesian problems.<sup>41</sup>

## Works Cited

[Insert author's papers]

Alston, William (1988) "The Deontological Conception of Epistemic Justification." *Philosophical Perspectives* 2: 257-299

----- (2005) *Beyond Justification* (Ithaca, NY: Cornell University Press).

---

<sup>40</sup> See my (forthcoming).

<sup>41</sup> [Insert acknowledgments]

- BonJour, Laurence (1985) *The Structure of Empirical Knowledge* (Cambridge, MA: Harvard University Press).
- Brandom, Robert (2000) *Articulating Reasons* (Cambridge, MA: Harvard University Press).
- Conee, Earl, and Feldman, Richard (1997) "The generality problem for reliabilism" *Philosophical Studies*, 89(1): 1–29.
- DeRose, Keith (1995) "Solving the Skeptical Problem" *Philosophical Review* 104(1): 1-52.
- Dretske, Fred (1981) *Knowledge and the Flow of Information* (Cambridge, MA: MIT Press).
- Finlay, Stephen (2008) "The Error in Error Theory" *Australasian Journal of Philosophy*.
- (forthcoming) "Oughts and Ends" *Philosophical Studies*.
- Fogelin, Robert (1994) *Pyrrhonian Reflections on Knowledge and Justification* (Oxford, New York: Oxford University Press).
- Gettier, Edmund (1963) "Is Justified True Belief Knowledge" *Analysis* 23.
- Goldman, Alvin (1973) "Discrimination and Perceptual Knowledge" *The Journal of Philosophy* 73.
- Greco, John (forthcoming) "Problem for Pritchard's Safety," *Synthese*.
- (2008) "The Value Problem," in A. Haddock, A. Millar, & D. H. Pritchard, eds., *The Value of Knowledge* (Oxford University Press).
- (2003) "Knowledge as Credit for True Belief," in Michael DePaul and Linda Zagzebski, eds. *Intellectual Virtue: Perspectives from Ethics and Epistemology* (Oxford University Press).
- Hawthorne, John (2005) *Knowledge and Lotteries* (Oxford: Oxford University Press)
- Hawthorne, John and Stanley, Jason (forthcoming) "Knowledge and Action" *The Journal of Philosophy*.
- Hazlett, Allan (forthcoming) "The Myth of Factive Verbs" *Philosophy and Phenomenological Research*.
- Hiller, Avram and Neta, Ram (2007) "Safety and Epistemic Luck" *Synthese* 158: 303-313.
- Kelly, Thomas (2003) "Epistemic Rationality as Instrumental Rationality: A Critique", *Philosophy and Phenomenological Research* 66(3): 612-640.
- Kvanvig, Jonathan (2003) *The Value of Knowledge and the Pursuit of Understanding* (Cambridge: Cambridge University Press).
- Kolodny, Niko (2003) "Why Be Rational" *Mind* 114 (455): 509–63.
- Jones, Ward (1997) "Why Do We Value Knowledge?" *American Philosophical Quarterly* 34: 423-40.
- Lewis, David (1996) "Elusive Knowledge" *The Australasian Journal of Philosophy* 74(4): 549-567.

- Moore, G. E. (1903) *Principia Ethica* (Cambridge: Cambridge University Press).
- Neta, Ram (forthcoming) "Anti-intellectualism and the Knowledge-Action Principle" *Philosophy and Phenomenological Research*.
- Pritchard, Duncan (forthcoming) "Sensitivity, Safety, and Anti-Luck Epistemology," Greco, ed. *Oxford Handbook of Skepticism*.
- (2008) "Radical Scepticism, Epistemic Luck and Epistemic Value" *Proceedings and Addresses of the Aristotelian Society* (suppl. vol.) 82 (2008)
- (2005) *Epistemic Luck* (Oxford University Press).
- Riggs, Wayne (2002) "Reliability and the Value of Knowledge" *Philosophy and Phenomenological Research* 64: 79-96.
- Rosenberg, Jay (2002) *Thinking about Knowing* (Oxford: Oxford University Press).
- Sosa, Ernest (2007) *A Virtue Epistemology: Apt Belief and Reflective Knowledge* (Oxford: Clarendon Press).
- Stanley, Jason (2005) *Knowledge and Practical Interests* (Oxford: Oxford University Press)
- Swinburne, Richard (1999) *Providence and the Problem of Evil* (Oxford: Oxford University Press).
- Weatherson, Brian (2003) "What Good Are Counterexamples" *Philosophical Studies* 115.
- Weinberg, Jonathan; Nichols, Shaun; and Stich, Stephen (forthcoming) "Normativity and Epistemic Intuitions" *Philosophical Topics*.
- Williams, Michael (2001) *Problems of Knowledge* (Oxford; New York: Oxford University Press).
- Williamson, Timothy (2002) *Knowledge and Its Limits* (Oxford: Oxford University Press).
- (2005) "Knowledge, Context and the Agent's Point of View" in *Contextualism in Philosophy*, edited by P. Preyer. Oxford; New York: Clarendon Press; Oxford University Press.
- (2005b)
- Zagzebski, Linda (1996) *Virtues of the Mind* (Cambridge: Cambridge University Press).