

Knowledge and other norms for assertion, action, and belief: A teleological account

by Neil Mehta

Forthcoming in *Philosophy and Phenomenological Research*. Please cite the final published version.

ABSTRACT: Here I advance a unified account of the structure of the epistemic normativity of assertion, action, and belief. According to my **Teleological Account**, all of these are epistemically successful just in case they fulfill the primary aim of knowledgeability, an aim which in turn generates a host of secondary epistemic norms. The central features of the Teleological Account are these: it is *compact* in its reliance on a single central explanatory posit, *knowledge-centered* in its insistence that knowledge sets the fundamental epistemic norm, and yet *fiercely pluralistic* in its acknowledgment of the legitimacy and value of a rich range of epistemic norms distinct from knowledge. Largely in virtue of this pluralist character, I argue, the Teleological Account is far superior to extant knowledge-centered accounts.

Here I advance a unified account of the structure of the epistemic normativity of assertion, action, and belief. According to my **Teleological Account**, all of these are epistemically successful just in case they fulfill the primary aim of knowledgeability, an aim which in turn generates a host of secondary epistemic norms. The central features of the Teleological Account are these: it is *compact* in its reliance on a single central explanatory posit, *knowledge-centered* in its insistence that knowledge sets the fundamental epistemic norm, and yet *fiercely pluralistic* in its acknowledgment of the legitimacy and value of a rich range of epistemic norms distinct from knowledge.

I do not argue for this view from scratch, however. From the outset, I assume the truth of the **Knowledge Account**, according to which an assertion, action, or belief has a certain positive epistemic status central to ordinary practice, which I noncommittally call **okayness**, iff it is appropriately related to knowledge. Others have assembled impressive evidence in support of this core idea, evidence which I rehearse only in the interest of bringing relevant information to the forefront of the reader's mind (§1). Here I focus solely on the following guiding question: what is the best version of the Knowledge Account, i.e., what is the best understanding of okayness? The reader ought not be fooled by the apparent narrowness of this question, for in answering it I offer an ambitious and innovative account of the structure of epistemic normativity.

Toward the end of finding the best understanding of okayness, I begin with the simplest approaches, moving to more complex ones only if these prove inadequate. In the case at hand, the simplest approach is to identify okayness with some well-understood positive normative status such as *permissibility*, *goodness*, or *success*. The identification of okayness with permissibility yields a broadly deontological theory, but I show that, given a minimal assumption, this theory has the bitter consequence that a belief is epistemically justified only if it is knowledge (§2). The identification of okayness with goodness, meanwhile, yields a broadly axiological theory, but I show that this theory rules out plausible claims about epistemic goods besides knowledge – goods like truth and epistemic justification – while a more permissive variant of the theory cannot explain the distinctive normative salience of knowledge (§3).

I instead endorse the identification of okayness with success, which yields the Teleological Account. I argue for this account by inference to the best explanation, taking as *explananda* both the original data motivating the Knowledge Account and the data yielded by the failures of the deontological and axiological approaches (§4–§5). After elaborating the Teleological Account by locating two possible sources of epistemic teleology (§6), I conclude that it is the best version of the Knowledge Account (§7).

1. The Knowledge Account

A primary motivation for the Knowledge Account is *the pervasiveness of references to knowledge in our ordinary epistemic assessments*, so I begin by programmatically reviewing a sample of these references *en route* to the Knowledge Account itself. In addition, while many have offered one piece or another of the account, I find it most compelling as a whole and will present it as such.

The following data relate knowledge to assertion:

- (A1) ‘How do you know that?’ standardly insinuates a negative assessment of an assertion. ‘You don’t know that!’ standardly gives outright expression to a negative assessment of an assertion.
- (A2) Any assertion of the form ‘*p* and I don’t know that *p*’ is infelicitous, even when such an assertion is true and supported by substantial evidence.
- (A3) Suppose I have purchased a single ticket for a fair lottery with a thousand tickets, each of which has an equal chance of winning. Unbeknownst to me, the winning ticket has already been drawn, and it is not mine. Call this the **Lottery Case**. It is infelicitous for me to assert, ‘My ticket has lost the lottery,’ even though this assertion is not only true, but also overwhelmingly probable given my evidence.
- (A4) ‘Do you know . . . ?’ is a standard prompt for assertion. For instance, if I wish to prompt you to make an assertion about where the train station is, I might ask, ‘Do you know where the train station is?’¹

Roughly parallel data relate knowledge to *reasons for action*²:

- (R1) Suppose that on the way home from work, I drive towards the interstate, which often becomes congested with traffic. ‘Do you know that this is the best way home?’ standardly insinuates a negative assessment. ‘You know this isn’t the best way home!’ standardly gives outright expression to a negative assessment.
- (R2) It is always problematic to treat *p* and *I don’t know that p* as a reason for action.
- (R3) In the Lottery Case, it is problematic for me to treat the proposition *that my ticket has lost* as a reason for action, even though this proposition is both true and overwhelmingly probable given my evidence.
- (R4) When we deliberate about what to do, a standard starting point is to review what we know.³

A third strand of roughly parallel data relates knowledge to *belief*:

¹ Williamson (2000) calls attention to (A1)–(A3) and Turri (2010) calls attention to (A4).

² For brevity, I will usually omit the phrase “reasons for.”

³ Hawthorne and Stanley (2008) call attention to (R1)–(R3).

- (B1) ‘How do you know that p ?’ standardly insinuates a negative assessment of the assertion, ‘I believe that p .’⁴ ‘You don’t know that p !’ standardly gives outright expression to a negative assessment.
- (B2) It is always problematic to believe that p and *I don’t know that p* , even if this belief is true and well-supported by evidence.
- (B3) Suppose that in the Lottery Case I truly believe that my ticket has lost. My belief, though supported by substantial evidence, is problematic; the boldest non-problematic belief would be that my ticket has *very likely* lost.
- (B4) There is something problematic about true justified beliefs that fall short of knowledge, as in Gettier cases.⁵

These data may be explained by the hypothesis that knowledge is a norm of assertion, action, and belief. On this hypothesis, if I know that p , then my assertion that p , my treatment of p as a reason for action, or my belief that p satisfies this norm and is thus (let’s noncommittally say) **okay**, and if I do not know that p , then my assertion that p , my treatment of p as a reason for action, or my belief that p violates this norm and is thus **not okay**. I stipulate that the assessment ‘okay’ is always positive, but whether the assessment ‘not okay’ is sometimes neutral or always negative I leave as a substantive question.

Notably, assertion⁶ is also normatively connected to items other than knowledge, such as belief, truth, evidence, reasons, and justification. For example, it is natural to express a negative assessment of an assertion by saying, ‘You don’t really believe that,’ ‘That’s not true,’ ‘You don’t have any evidence for that,’ etc. However, the present hypothesis predicts such patterns. After all, if I know that p , then I must believe that p ; p must be true; and I must have evidence, reasons, and justification for believing that p . Given the presupposition that knowledge requires belief, a negative assessment like ‘You don’t really believe that’ thus entails the negative assessment ‘You don’t know that.’ By contrast, it is difficult to explain the normative significance of knowledge via the claim that belief, truth, evidence, reasons, or justification set a norm of assertion, for none of these require knowledge.⁷

This hypothesis requires two refinements. First, assertions, actions, and beliefs are normatively assessable in a variety of ways that have little to do with knowledge: they may be kind or cruel, polite or rude, funny or dull. So knowledge is not the sole norm against which assertions are measured. Instead, norms come in distinctive clusters, one of which is **epistemic**. The difference between epistemic and non-epistemic norms can be brought out by well-worn examples, like the example of a patient who believes without evidence that she will recover from a serious illness and whose belief in fact helps her recover: her belief satisfies certain non-epistemic norms but violates certain epistemic norms. According to the present hypothesis, a central part of this epistemic cluster of norms is the knowledge norm, against which assertions, actions, and beliefs are measured.

Second, knowledge can be problematically disconnected from assertion, action, or belief. Consider Marvin, whose job is to direct tourists to the Big Burger Bistro. Marvin stands south of the Big Burger Bistro, pretending to give tourists directions when he in fact dishonestly

⁴ Sometimes we say ‘I believe that p ’ to make a hedged assertion that p , but I stipulate that this is not occurring here.

⁵ Gettier (1963).

⁶ Parallel remarks apply to action and belief.

⁷ Assertion is also normatively connected to certainty, as assertions are sometimes naturally met with responses like, ‘Are you sure?’ See Turri (2010) for a suggestion on how to square this with the Knowledge Account.

directs every tourist north. Thus, when a tourist asks Marvin how to get to the vegan Candle Café, which Marvin knows is to the north, Marvin replies with feigned cheer, ‘Just head north!’ with the sole intention of directing the tourist to the Big Burger Bistro. Even though Marvin asserts what he knows, his assertion is not epistemically okay because of its disconnection – causal, counterfactual, and otherwise – from his knowledge. Such disconnections between a subject’s assertion that p and her knowledge that p may occur for any number of reasons, including ulterior motives, carelessness, or a momentary lapse of memory.

This structure is familiar throughout the normative realm. While being ethically required to recycle, I might do so solely to keep up appearances; while having most reason to believe that my dog is more clever than yours, I might believe this merely because I love my dog. The solution, too, is familiar: to require an appropriate connection. For assertion and the treatment of items⁸ as reasons for action, the connection is *proper basing*. Accordingly, call an assertion that p or treatment of p as a reason for action **knowledge-based** just in case it is properly based on knowledge that p . On the present hypothesis, an assertion or treatment of a reason for action is epistemically okay just in case it is knowledge-based.⁹ For belief, the proper connection is simply identity: on the present hypothesis, a belief is epistemically okay just in case it *is* knowledge.¹⁰ For convenience, I will also disjunctively define the term **knowledgeable**: an item is knowledgeable iff it is a knowledge-based assertion or action *or* a belief that is knowledge.

We now arrive at the unified **Knowledge Account**:

Knowledge Account of Assertion: An assertion that p is epistemically okay iff this assertion is knowledge-based.^{11,12}

Knowledge Account of Action: A treatment of p as a reason for action is epistemically okay iff this treatment is knowledge-based.^{13,14,15}

⁸ I use the term ‘item’ as a catch-all sortal, covering states, events, objects, etc.

⁹ Turri (2011) makes a similar point.

¹⁰ Perhaps beliefs are never knowledge; perhaps beliefs are at best a proper part of knowledge, or perhaps there are belief types but no belief tokens. My core claims can be paraphrased to fit such views, though for simplicity I will continue to speak as though an identity obtains between every instance of knowledge and the associated belief.

¹¹ Williamson (2000, ch. 11) provides the now-canonical defense of this account. Related accounts appear in Moore (1962), Unger (1975), Slote (1979), and DeRose (1991).

¹² Perhaps knowledge is the *constitutive* norm of assertion. On this stronger hypothesis, something is an assertion wholly in virtue of being subject to the Knowledge Account of Assertion, just as something is a game of chess wholly in virtue of being subject to the rules of chess (Williamson (2000, ch. 11)).

¹³ Fantl and McGrath (2002) suggest that if I know that p , then it should not be a problem for me to act as if p , which approximates the right-to-left direction of the Knowledge Account of Action. Similarly, Hawthorne and Stanley (2008) defend something much like the Knowledge Account of Action, but their account is restricted to **p -dependent choices**. (Hawthorne (2004) contains a less-developed version of this account.) A choice between actions $x_1 \dots x_n$ is p -dependent just in case the most preferable of $x_1 \dots x_n$ given that p differs from the most preferable of $x_1 \dots x_n$ given that not- p . This restriction is designed to avoid cases where p is completely irrelevant to one’s choice of action. However, this restriction is unnecessary. To be sure, when p is completely irrelevant to one’s choice of action, it will typically be unhelpful and therefore misleading to say that one treated p as a reason for action, for one ought to discuss only *relevant* reasons for action. Nevertheless, when p is irrelevant, there is no harm in letting the treatment of p as a reason for action be trivially epistemically okay.

¹⁴ Note that the Knowledge Account of Action does not entail the Knowledge Account of Assertion. Admittedly, assertion is an action, so the Knowledge Account of Action entails that treating p as a reason for asserting that q is epistemically okay iff this treatment of p is knowledge-based. However, it does not follow that asserting that q is epistemically okay iff this assertion that q is knowledge-based. Brown (2012, 130-134) helpfully elaborates on this point.

Knowledge Account of Belief: A belief that p is epistemically okay iff this belief is knowledge.¹⁶

In short, an assertion, action, or belief is epistemically okay iff it is knowledgeable.

The Knowledge Account applies only to *existing* assertions, treatments of reasons for action, and beliefs, not to *merely potential* ones. Moving from a theory of the former to a theory of the latter is surprisingly difficult.¹⁷ To mention just one hurdle, even if I know that I have never mentioned the word ‘assertion,’ it is not epistemically okay for me to assert, ‘I have never mentioned the word “assertion.”’

The Knowledge Account cleanly explains the above data. Because these explanations run almost in parallel, here I will discuss only the assertion data:

- (KA1) ‘How do you know that?’ often implicates that you do not (or might not) know that. If you do not know that, then your assertion is not epistemically okay, so the question often implicates that your assertion is not (or might not be) epistemically okay and hence insinuates a negative assessment. ‘You don’t know that!’ entails that your assertion is not knowledge-based and thus not epistemically okay. That expresses a negative assessment outright.
- (KA2) For the assertion ‘ p and I don’t know that p ’ to be epistemically okay, I must know that p , and I must also know that I don’t know that p . But given the factivity of knowledge, that is impossible. So any assertion of the form ‘ p and I don’t know that p ’ is not epistemically okay.
- (KA3) In the Lottery Case, though I have strong evidence that my ticket has lost the lottery, I still do not *know* this. So my assertion is not epistemically okay.
- (KA4) An epistemically okay assertion that p requires the speaker to know that p , so by asserting that p , one implicates that one knows that p . Typically, when I prompt an assertion by asking, ‘Do you know whether p ?’, it is common ground (and arguably common *knowledge*) that I am not interested merely in whether or not you know whether p . Given that knowledge is the epistemic norm of assertion, a question like ‘Do you know where the train station is?’ implicates which speech act I would like you to perform: an assertion about the location of the train station.

Taking for granted the truth of the Knowledge Account, I ask: What is the best version of it?¹⁸ Answering this question requires finding the best understanding of okayness, and the simplest approach is to identify this property with a familiar positive normative property such as permissibility, goodness, or success. In the next three sections, I examine these identifications in turn.

¹⁵ On a **credence-oriented approach**, there are substantial normative links between one’s rational **credences** (i.e., degrees of belief) and one’s reasons for action. The Knowledge Account of Action is silent about the credence-oriented approach, for it is silent about credences. But Moss (2013) and Weisberg (2013) suggest ways of combining the Knowledge Account of Action with the credence-oriented approach.

¹⁶ Williamson (2000, p. 11) expresses sympathy for this account, and Sutton (2005) endorses it.

¹⁷ As Whitcomb (forthcoming) shows.

¹⁸ For notable critiques of the Knowledge Account of Assertion, see Weiner (2005), Lackey (2007), Brown (2008a), Stanley (2008), Sosa (2009), Brown (2010), and Maitra and Weatherson (2010). For notable critiques of the Knowledge Account of Action, see Brown (2008b), Neta (2009), Lackey (2010), and Gerken (2011).

I will pay particular attention to the Knowledge Account of Belief, which is representative in that almost all of my remarks about it can be recast *mutatis mutandis* as remarks directly about the Knowledge Account of Assertion or Action. As it happens, I also regard the Knowledge Account of Belief as the linchpin of the Knowledge Account. For I take assertion to be the speech act with the function of expressing belief, and I further take belief in a proposition to underpin the treatment of that proposition as a reason for action, so I hold that assertion and the treatment of propositions as reasons for action inherit the epistemic norms applying to belief. But I will not rely on this further view.¹⁹

2. The Deontological Account

The **Deontological Account** identifies okayness with *permissibility*:

Assertion Deontology: An assertion that p is epistemically permissible iff this assertion is knowledge-based.

Action Deontology: A treatment of p as a reason for action is epistemically permissible iff this treatment is knowledge-based.

Belief Deontology: A belief that p is epistemically permissible iff this belief is knowledge.

The great popularity of this account is understandable.²⁰ For among the most common norms of ordinary life are rules expressing what is permissible, required, or forbidden – codes of law, of etiquette, of professional conduct. It is natural, then, to conceive of the Knowledge Account as expressing rules of a similar stripe in the epistemic domain.

But, focusing as promised on the case of belief, I will argue that the Deontological Account offers far too restrictive a conception of epistemic permissibility. Consider:

Belief-Belief Link: If a belief that p is epistemically justified, then this belief is epistemically permissible.

I adopt this as a partial stipulation of the meaning of the technical term ‘epistemic justification.’²¹ I prefer this more complete stipulation:

Biconditional Belief-Belief Link: A belief that p is epistemically permissible iff this belief is epistemically justified.

¹⁹ For skepticism about the unity of the epistemic norms of assertion, action, and belief, see Stanley (2008) and Brown (2012).

²⁰ Defenders of Assertion Deontology include Williamson (2000). Defenders of Action Deontology include Hawthorne (2004) and Hawthorne and Stanley (2008). Defenders of Belief Deontology include Sutton (2005) and, more tentatively, Williamson (2000).

²¹ For those who prefer a different stipulation, I can recast the argument below as follows while bypassing any talk of epistemic justification: if Belief Deontology is true, then it is epistemically permissible to have a false belief based on mountains of misleading evidence, and the Gettiered subject’s belief is epistemically impermissible. But both of these beliefs plainly are epistemically permissible, so Belief Deontology is false.

But my argument requires only the simple Belief-Belief Link.

This principle, when conjoined with the left-to-right direction of Belief Deontology, immediately entails that if a belief that p is epistemically justified, then this belief is knowledge; that is, a belief is epistemically justified only if it is knowledge. Yet this is manifestly false. For some false beliefs are epistemically justified, as when one has mountains of evidence which uniformly point in the wrong direction, and some epistemically justified true beliefs are not knowledge, as Gettier cases illustrate.²² To say that justification requires knowledge is to ignore some of our strongest and most stable epistemic intuitions.²³

Perhaps the deontologist will reply that beliefs falling short of knowledge are never *fully* epistemically justified but may still be *partially* epistemically justified.²⁴ Introducing graded epistemic justification alongside epistemic justification *simpliciter* does not help, however, for the Gettiered subject has epistemic justification *simpliciter*. Since she lacks full epistemic justification, epistemic justification *simpliciter* cannot require full epistemic justification; it must require only partial epistemic justification. Further, given the Belief-Belief Link, the Gettiered subject's belief is epistemically permissible. Belief Deontology then falsely entails that her belief is knowledge.

The deontologist may also suggest that I have confused excusability with permissibility: while the belief of the Gettiered subject is entirely excusable, it is also impermissible.²⁵ But this reply runs together two distinct epistemic statuses, as the delusive beliefs of a psychiatric patient may be excusable despite being epistemically impermissible. Nor is it fruitful for the deontologist to distinguish epistemic from non-epistemic excusability with the hope of identifying the former as the relevant status of Gettiered belief, since this approach similarly conflates two distinct epistemic statuses. For instance, Andrew Wiles' lengthy initial 'proof' of Fermat's Last Theorem contained a subtle gap which he was perfectly capable of recognizing but did not. Wiles' belief, before the discovery of the gap, that he proved Fermat's Last Theorem is a paradigm of an epistemically excusable but impermissible belief, the excuse being that the attempted proof was complex and the flaw subtle. But this is not the positive epistemic status we accord to the beliefs of Gettiered subjects.

Gettiered beliefs, unlike merely excusable beliefs, are formed by *epistemically proper internal processes*. I say little about what makes an internal process epistemically proper – perhaps it must fit the subject's evidence,²⁶ or reliably produce true beliefs in typical environments,²⁷ or be internally like a knowledge-producing process²⁸ – and I thus do not say whether all intrinsic duplicates must be alike with respect to the property of *using an epistemically proper internal process*. But I do say that this property suffices for epistemic justification and hence for epistemic permissibility. I would add that with respect to a Gettiered belief, the subject is not only exempt from epistemic blame but also merits epistemic praise, as from an epistemic perspective the subject's method of belief formation deserves active reinforcement.

²² Gettier (1963). Many who defend parts of the Knowledge Account explicitly rely on the view that Gettier cases exist – most prominently, Williamson (2000).

²³ For criticisms in the same vein, see Jackson (2012), Smithies (2012), Cohen and Comesaña (2013), and Turri (2014).

²⁴ Williamson (2013a) and (2013b) defends this view.

²⁵ For attempts to explain certain judgments of acceptability in terms of excusability rather than permissibility, see Hawthorne and Stanley (2008), Hellie (2011), Leite (2013), and Williamson (2013b).

²⁶ See Conee and Feldman (1985).

²⁷ See Goldman (1979).

²⁸ See Bird (2007).

The brave-hearted theorist may consistently retain Belief Deontology simply by denying one or more premises of this argument,²⁹ but she should at least admit that her position is counterintuitive. For Gettier cases are among the *paradigms* of epistemic justification: as a matter of sociology, such cases have overwhelmingly persuaded the philosophical community that true justified belief does not suffice for knowledge. The appeal of Belief Deontology lies not in its stringent conception of justification, but in its ability to explain the compelling data about the epistemic centrality of knowledge. If we can explain these data without the deontologist's procrustean methods, as I will ultimately argue that we can, then Belief Deontology is obsolete.³⁰

3. Two axiological accounts

Knowledge is an epistemic good, so of course a knowledgeable belief is epistemically good. Arguably, so too are knowledge-based assertions and actions, for do they not shine with the borrowed light of the underlying knowledge? These observations motivate the **Strict Axiological Account**, which identifies okayness with *goodness*:

Strict Assertion Axiology: An assertion that p is epistemically good iff this assertion is knowledge-based.³¹

Strict Action Axiology: A treatment of p as a reason for action is epistemically good iff this treatment is knowledge-based.

Strict Belief Axiology: A belief that p is epistemically good iff this belief is knowledge.

I will focus on Strict Belief Axiology, which has one notable advantage. Recall the Belief-Belief Link, which says that if a belief that p is epistemically justified, then this belief is epistemically permissible. Belief Deontology, as we have seen, is incompatible with this principle, but Strict Belief Axiology can accommodate it. For permissibility in some domain does not require goodness in that domain; something may be *neutral* or *bad*, but still permissible, in a given domain,³² as is perhaps illustrated in the ethical domain by a failure to develop one's talents out of sheer laziness. The strict axiologist may therefore hold that a Gettiered belief that p , while perfectly epistemically permissible, is nonetheless not epistemically good because it is not knowledge. And indeed a Gettiered belief that p would plausibly have been epistemically better had it been knowledge that p .³³ The Strict Axiological Account thus appears to have much to recommend it.

²⁹ Sutton (2005) heroically defends Belief Deontology, and Williamson speculates that the left-to-right direction of Belief Deontology holds (2000, p. 11). Both simply deny that beliefs falling short of knowledge are ever epistemically permissible (Williamson (2013b, p. 92)). Meanwhile, Weatherson (2003) flirts with rejecting the existence of Gettier cases.

³⁰ Perhaps some even find it intuitive that justification requires knowledge. While I amicably admit that I cannot persuade such theorists, they may still wish to see how best to square the predominant intuition about justification, which they do not share, with the Knowledge Account.

³¹ Turri (2014) meticulously defends a view very similar to the left-to-right direction of Assertion Axiology. My arguments below will undermine even this weaker position.

³² As Turri (2014) nicely brings out.

³³ For similar reasons, Assertion Axiology is perfectly compatible with the Biconditional Belief-Assertion Linking Principle, and Action Axiology is perfectly compatible with the Biconditional Belief-Action Linking Principle.

Unfortunately, the Strict Axiological Account is difficult to reconcile with other facts about epistemic goods. *Truth* is a paradigmatic epistemic good, so true beliefs that fall short of knowledge may presumably be epistemically good. Similarly, *justification* is an epistemic good, and accordingly justified beliefs that fall short of knowledge may also be epistemically good. Perhaps such beliefs are not *as* epistemically good as beliefs that are knowledge, but they may surely be at least *somewhat* epistemically good.

The strict axiologist might reply that though truth and justification are epistemic goods, falling short of knowledge is epistemically bad enough to outweigh these epistemic goods, so true justified beliefs falling short of knowledge are not epistemically good *overall*. This parallels the ethical consequentialist's claim that due to the ethical goodness of pleasure and the epistemic badness of pain, taking mild sadistic pleasure in the great pain of another is ethically bad overall, as the situation involves a balance of pain over pleasure.

While this strategy might be plausible in particular cases, there is no reason to accept it in general, for there is no reason to think that falling short of knowledge is always epistemically bad. The strict axiologist might infer this from the platitude that having knowledge is epistemically good. But in general it is fallacious to infer from the premise that all items having a particular property are good to the conclusion that all items lacking that property are bad, e.g., from the premise that paintings by Monet are aesthetically good to the conclusion that paintings not by Monet are aesthetically bad. For an item may have other good-making properties, and lacking the relevant property may also be merely neutral rather than bad.

Further, even if falling short of knowledge is always epistemically bad, the objection is not yet answered. The strict axiologist must further hold that this epistemic badness always outweighs the epistemic goodness of truth and justification. For if not, then some true justified beliefs falling short of knowledge will still be epistemically good overall, despite the epistemic badness of falling short of knowledge. But there is no obvious reason why falling short of knowledge must always be *so* epistemically bad as to outweigh the epistemic goods of truth and justification.

The Strict Axiological Account recognizes the distinctive value of knowledge, but only at the expense of neglecting the value of other central epistemic goods. Perhaps the following **Permissive Axiological Account**, which identifies epistemic okayness with *a particular degree of epistemic goodness*, can do better:

Permissive Assertion Axiology: An assertion that *p* is epistemically good to a particular degree iff this assertion is knowledge-based.

Permissive Action Axiology: A treatment of *p* as a reason for action is epistemically good to a particular degree iff this treatment is knowledge-based.

Permissive Belief Axiology: A belief that *p* is epistemically good to a particular degree iff this belief is knowledge.

This account avoids the problems just raised for its stricter cousin, for it is compatible with the claim that true or justified beliefs which fall short of knowledge may be at least *somewhat* epistemically good overall. The Permissive Axiological Account entails only that such beliefs do not reach the particular degree of goodness of their knowledgeable counterparts.³⁴

³⁴ Recall that an item is **knowledgeable** iff it is a knowledge-based assertion or action *or* a belief that is knowledge.

The acceptability of even this consequence is unclear. For compare, on the one hand, my knowledgeable belief that my car is in the parking lot, and on the other hand, a Gettiered true belief supported by extensive observational data which unifies theoretical physics. Setting aside any antecedent commitments to a particular theory of epistemic goodness, I find it difficult to judge which of these beliefs is epistemically better, and I feel stirrings of doubt even about the idea that there exists only a single axis of epistemic goodness along which these beliefs are to be ranked. Aren't these beliefs simply epistemically good in distinct and incomparable respects?

But I will bracket these concerns for the moment, as the Permissive Axiological Account suffers from a more obvious defect: it is *so* permissive that it can no longer explain the distinctive normative salience of knowledge. By analogy, consider point-scoring in basketball. Scoring at least 73 points is always good to at least a particular degree, and always better than many other scores. But we do not single out this score, or any other for that matter, for special normative attention. To do so would be deeply arbitrary and therefore silly. Hence, the property of *being good to a particular degree* does not always warrant special normative emphasis.

Knowledge is not only normatively salient, but also *far more* normatively salient than countless other epistemic properties, as evidenced by the data in §1. But the Permissive Axiological Account cannot explain this special emphasis, given that knowledgeable beliefs fall in the middle of a long scale of epistemic goodness. On one side, true or justified but non-knowledgeable beliefs, while epistemically good, are not as epistemically good as knowledgeable beliefs, and on the other side, some knowledgeable beliefs are epistemically better than others: knowledge based on overwhelming evidence from many sources is epistemically better than knowledge based on merely solid evidence from a single source. But the Permissive Axiological Account alone cannot explain why we single out the particular degree of goodness reached by knowledgeable beliefs over the arbitrarily many other surrounding degrees of goodness.

To be clear, I suspect but have not tried to argue that the Permissive Axiological Account is false; my account below will even be perfectly compatible with its truth. But I object that, even if true, it cannot explain why prizing knowledge is not as arbitrary, and as silly, as prizing basketball scores of 73 points or more.

4. The Teleological Account

So far, I have argued that our epistemic assessments of belief, and by extension assertion and action, fall into at least four distinct categories.

First, as evidenced by the data in §1, we make the binary rather than degreed distinction between *knowledgeable* and *non-knowledgeable* beliefs. When a subject fails this normative test, the flaw is sometimes hers, as in cases of careless reasoning, and sometimes in her environment or her connection to that environment, as in Gettier cases. Regardless, we pay keen attention to knowledge and organize many of our normative assessments around it.

Second, we make the further binary distinction between *knowledgeable beliefs*, *Gettiered beliefs*, etc., and *beliefs generated by wishful thinking, fallacious inference, etc.* In so doing, we focus on whether or not the subject formed her belief via an epistemically proper internal process. This distinction differs from the previous one in that Gettiered beliefs are on the positive side of the

ledger here, since they are properly formed,³⁵ and in that when the subject fails this test, the flaw is always hers.

Third and relatedly, when the subject fails this second test, we look for extenuating circumstances – perhaps the fatigue of a long day, the complexity of the problem at hand, or the indoctrination of ignorant parents – for her failure. If she is in circumstances like these, we do not (in our reasonable moments at least) hold her fully accountable; otherwise we do. Either way, however, we say that the flaw is in her cognitive system.

Fourth, we mark salient positive epistemic properties which a particular belief might or might not reflect, including *truth, justification, knowledge, evidence, thoroughness in gathering evidence, etc.* We can readily rank two such properties whenever having one entails but is not entailed by having the other – more justification is always better, and knowledge is always better than mere justified belief. But some rankings are harder to come by: it is difficult to say whether a justified false belief is better or worse than an unjustified true belief. Hence our assessments here lie along a fluid and approximate scale.

Our normative theory should reflect this complex structure, admitting at least four groupings of normative properties matching these distinctive profiles. While we may hold that these four groupings are interestingly related to one another, we should not conflate them.

These data are best explained by the position I will develop below: the **Teleological Account**, which identifies okayness with *success*. For as I will soon argue, the Teleological Account not only respects these distinctions but even explains the relevance of epistemic goodness, permissibility, and excusability in terms of knowledge. This theory thus underwrites a compact account of epistemic normativity.

Consider Harry, who applies for Auror training. Harry's application includes glowing recommendations from his Defense against the Dark Arts teachers, a moving essay that discusses his battle against Lord Voldemort, and outstanding results from a practical examination in which he single-handedly subdued a dementor. However, the malicious Malfoys have bribed the admissions committee at the Ministry of Magic, so Harry's application is rejected. How does this application stand with respect to the norms of applying for Auror training?

The application is obviously permissible. It is complete and accurate, and Harry submitted it well before the deadline. From a deontological perspective, the application thus merits a positive assessment. The application is also very good; any Auror will tell you that it is the best of this year's batch! So from an axiological perspective, the application again merits a positive assessment. Still, it also merits a certain negative assessment, for it is a failure: Harry is not accepted for Auror training, and any application for Auror training, *qua* application for Auror training, has acceptance as its aim.

This case highlights another salient normative property: *success*. An item is successful in a domain iff it fulfills its aim in that domain.³⁶

Harry's application illustrates two facts about success: that neither permissibility nor goodness entails success in a domain. Suppose that Draco Malfoy submits a middling

³⁵ Of course, this claim would be resisted by epistemologists who identify proper belief-forming processes with processes that in fact produce knowledge. But I worry that such epistemologists ignore a distinction that we make as a matter of course.

³⁶ Some items have multiple aims within a domain and may be successful in some respects but not others within that domain.

application that is accepted because of further bribery. Draco's application illustrates a third fact: that an item may be successful without being maximally good.³⁷

Ordinary practice reflects this normative multidimensionality. Hermione may praise Harry from the deontological perspective ('Sending in that application was the right thing to do') and from the axiological perspective ('What an outstanding application!'), yet these positive remarks are perfectly compatible with negative assessment from the teleological perspective: 'Your application failed – you didn't get into Auror training!' Applying this framework to the epistemic domain, the **Teleological Account** holds that assertion, action, and belief all aim at knowledgeability – or, as I will say more pithily, at knowledge.³⁸ This account thus identifies okayness with *success*:

Assertion Teleology: An assertion that p is epistemically successful iff this assertion is knowledge-based.

Action Teleology: A treatment of p as a reason for action is epistemically successful iff this treatment is knowledge-based.

Belief Teleology: A belief that p is epistemically successful iff this belief is knowledge.³⁹

According to this theory, both the Gettiered subject's belief and Harry's application are permissible and good, indeed excellent, in their respective domains. Nevertheless, through no fault of the subject, each is a failure. Success requires the aid of a fickle world.

The Teleological Account explains the relevance not only of epistemic success (i.e., knowledge), but also of epistemic permissibility, excusability, and goodness. For suppose that I have an aim that is **ongoing**, in that I will repeatedly have aims of the same type; **external**, in that the achievement of such aims requires the cooperation of the environment; and **social**, in that I belong to a group that cares about aims of this type. For instance, when I play pool, I aim to sink a particular ball or balls, and to continue to do so with each shot I take during the game. Doing so depends not only on my skill, but also on external factors such as the levelness of the table, the lack of interference by spectators, and so on. And, as a pool player, I belong to a community of pool players, spectators, and so on who care about this aim. For any such aim, members of the relevant group will by and large benefit from assessing four clusters of properties.⁴⁰ First and foremost, they will of course benefit from attending carefully to the aim itself – in this case, whether I sink the ball(s) that I call. But that is not all.

They will benefit secondly from assessing, roughly, *whether the individual uses a proper internal process*. The cloth on this pool table might be worn, throwing my shots slightly but unpredictably off course, or a careless bystander might jostle my elbow. Still, though these

³⁷ In some domains, an item may even succeed despite being bad: even had Draco's application been awful, a large enough bribe would have ensured its success. Below, I will suggest that this cannot occur in the epistemic domain.

³⁸ For more on the view that knowledge is the aim of belief, see Williamson (2000, p.11), Bird (2007), McHugh (2011), Smithies (2012), and Jackson (2012).

³⁹ While Jackson (2012) and Smithies (2012) defend theories much like the Teleological Account, my arguments will be independent of theirs. Smithies argues from the principle that one has justification to act in a certain way iff one has justification to believe that it is correct to act in that way. I rely on no such iterated justification principle, and indeed Smith (2012) argues persuasively against it. Jackson focuses on explaining how we can rationally act on partial rather than full beliefs. Though I am sympathetic to much of what he says, I focus instead on predicting and unifying the many dimensions of our epistemic assessments.

⁴⁰ Here I take inspiration from Sosa (2007) and from correspondence with John Turri, though my account is quite different from either of theirs.

unfortunate external circumstances might stop me from achieving my aim, I ought not locate the flaw in myself. I should instead focus on using the appropriate type of internal process, a process which tends to result in the correct ball or balls being sunk, to do well in the long run. Distinguishing internal from external flaws in this fashion will equip me to refine my shooting. In short, then, pool players should assess what is *permitted, required, or forbidden*, relative to the aim of sinking pool shots.⁴¹

If someone fails the second assessment, then to fairly apportion blame members of the relevant group should thirdly assess, roughly, *whether she could reasonably have been expected to do better*. Suppose that you and I are partners in a pool match, but I have been shooting poorly. If I am a beginner, or if I am suffering a panic attack, then perhaps you should not blame me for my bad performance, but blame is warranted if I have no such excuse. In short, then, pool players should assess the *excusability* of what is forbidden relative to an aim.

Fourth, a standard heuristic for achieving an aim is to break it into more manageable parts. Whether or not I make my shot, pool players can observe whether I meet *salient conditions typically or necessarily required to achieve that aim*, like calling the proper pocket and hitting the correct ball at the correct angle. Relatedly, they can assess whether I have *salient properties which dispose me to achieve my aim*, like shooting with a still elbow, or *which dispose me to become disposed to achieve my aim*, like practicing difficult shots regularly. Such properties are numerous and scattered, and this diversity will be reflected in our assessments. In short, then, pool players should assess what is *good, neutral, or bad* relative to this aim.

The aim of knowledge is ongoing, since we aim to achieve it in our many future assertions, actions, and beliefs; external, since (as Gettier cases remind us) knowledgeability requires the cooperation of the environment; and social, since it is an object of collective concern. In the wake of knowledge, we should thus expect to find assessments of permissibility, excusability, and goodness fitting the pattern I have described.

We find exactly this. As I remarked at the beginning of this section, we distinguish *knowledgeable beliefs, Gettiered beliefs, etc.*, from *beliefs generated by wishful thinking, fallacious inferences, etc.*, focusing on whether the assessed individual uses an epistemically proper internal process – roughly, an internal process suited to generating knowledge. If this test is negative, we then consider extenuating circumstances, focusing on, roughly, whether we could reasonably have expected her to do better. Finally, we laud a diverse mix of salient properties related to knowledge, including necessary conditions on knowledge, such as truth, justification, evidence; dispositions likely to lead to knowledge, such as thoroughness in gathering evidence; dispositions to become disposed to achieve knowledge, such as willingness to revise one's epistemic practices; and of course knowledge itself. The Teleological Account thus not only respects these many dimensions of epistemic normativity, but further predicts their specific contours. It puts knowledge first, but without the procrustean treatment often associated with that slogan.

But the aim of knowledge is in one important respect not analogous to the aim of sinking a pool shot: in the epistemic case, the link between permissible processes and fulfillment of the aim is constitutive rather than merely causal. Though it is unlikely, I might sink a pool shot via an impermissible process such as careless but lucky shooting. Such possibilities exist because the aim of sinking a pool shot is metaphysically independent from the permissible processes for achieving that aim; in other words, the relationship between permissible methods

⁴¹ In the context of pool, the term “permissibility” may refer either to deontological acceptability given the *rules* of pool or to deontological acceptability given the *aim* of pool. My meaning is obviously the latter, which is also naturally expressed with terms like “should.”

and success is merely causal and contingent. In the case of knowledge, however, it is part of the aim itself that one's method be permissible, for a true belief based on insufficient evidence or formed via bad reasoning cannot qualify as knowledge. Here, then, the relationship between permissible processes and success is one of metaphysical necessity.

As a consequence, the normative structure generated by the aim of knowledge differs, in subtle but important ways, from the normative structure generated by the aim of sinking pool shots. Negative assessment of someone's processes in pool can co-exist with positive assessment of her results: "I admit that your fancy trick shot sunk the ball, but you just got lucky! You should have taken the safe and easy shot instead of showing off." But such situations cannot arise regarding the aim of knowledge, as success in the absence of a permissible process is impossible.

In addition, since the permissibility of a process of pool shooting depends on its causal relationship to the aim of sinking pool shots, a pool player may take shots on purely probabilistic grounds. She may take a particular shot rather than another because she knows that she is more likely to sink the first shot, for example, even if she in fact misses the first shot and would by chance have sunk the second. Similarly, she may take a low-probability shot simply because she has no better shot available. But a process of belief formation is permissible in virtue of its constitutive relationship to the aim of knowledge; a merely probabilistic relationship is not enough, and indeed not even possible in the same sense. For if I believe, even truly, that a particular ticket has not won a fair lottery on the grounds that its chance of being selected is only 10%, my improper method of belief formation disbars me from achieving knowledge.^{42,43}

Though the constitutive rather than causal relationship between permissibility and success is a special feature of the epistemic case, once it is taken into account our varied epistemic assessments fall neatly into the explanatory net of the Teleological Account.

5. The Teleological Account and its rivals

In letter, the Deontological, Strict Axiological, and Permissive Axiological Accounts are each compatible with the Teleological Account, since it is conceptually coherent to identify epistemic success with epistemic permissibility, epistemic goodness *simpliciter*, or a particular higher degree of epistemic goodness.⁴⁴ I have no quarrel with the combination of the Permissive Axiological Account and the Teleological Account, though I will show at the end of this section that unlike the bare Permissive Axiological Account, the pluralistic Teleological Account explains the distinctive normative salience of knowledge. But I resist any attempt to combine the Teleological Account with the Deontological or Strict Axiological Accounts, for I have argued that these accounts blur distinct aspects of our epistemic assessments. I endorse only the *fiercely pluralistic* version of the Teleological Account which sharply distinguishes at least four categories of epistemic assessment. In addition, I resist for the further reason that any such combination would expose the Teleological Account to the objections that I have raised against

⁴² Smith (2010) and (2014) makes this point clearly.

⁴³ Though I hold that knowledge cannot be *based on* merely probabilistic evidence, I am non-committal about the possibility (defended by Williamson (2011) and (forthcoming)) of "improbable knowing," i.e., of cases in which the subject knows that p even though the epistemic probability *that she knows that p* given her evidence is low.

⁴⁴ Williamson (2000) suggests combining parts of the Teleological Account with parts of the Deontological Account, and John Turri (personal correspondence) does the same for the Strict Axiological Account.

the Deontological or Strict Axiological Accounts. But, as I will next show, the pluralistic Teleological Account is not vulnerable to those objections.

Given the minimal assumption that some justified beliefs are not knowledge, I objected that the Deontological Account is incompatible with the Belief-Belief Link, according to which if a belief that p is epistemically justified, then it is epistemically permissible. But the Teleological Account is perfectly compatible with the Belief-Belief Link. To be sure, the Teleological Account entails that any non-knowlegeable belief is thereby epistemically unsuccessful. This verdict is reasonable, since we can appropriately meet any such belief with negative assessments like, ‘You don’t know that!’ but this in no way conflicts with the Belief-Belief Link, for an item may be permissible despite being unsuccessful in a domain. Thus, the teleologist may hold that some beliefs which fall short of knowledge are epistemically permissible despite being epistemically unsuccessful. On this score, the teleologist, like the strict axiologist, has a leg up on the deontologist.

Recall the Strict Axiological Account, on which non-knowlegeable beliefs, assertions, and actions are never epistemically good overall. This account must deny that true or justified beliefs falling short of knowledge, as well as assertions and actions properly based on them, can sometimes be epistemically good overall. On one version of the account, truth and justification are not epistemic goods at all, while on another, lack of knowledge is always epistemically bad, and indeed always epistemically bad enough to outweigh the epistemic goods of truth and justification. Either commitment is implausible.

The Teleological Account is burdened with no such commitments. For in a given domain, something may be good in some respect, and even good overall, while still being unsuccessful. Thus, to avoid the first commitment, the Success Theorist may hold that true or justified beliefs, and assertions and actions properly based on them, are epistemically good in some respect. To avoid the second commitment, the Success Theorist may further hold that lack of knowledge need not be epistemically bad enough to outweigh the epistemic goods provided by truth or justification, and she may even hold that lack of knowledge need not be epistemically bad at all. Thus, she may hold that some true or justified beliefs that fall short of knowledge, and assertions and actions properly based on them, are epistemically good overall.

Finally, recall my central objection to the Permissive Axiological Account, which holds that knowlegeable assertions, actions, and beliefs reach a particular degree of epistemic goodness which their non-knowlegeable counterparts never do. This theory falters at explaining why knowledge is so normatively salient, given that knowledge falls in the middle of the long scale of epistemic goodness. Beliefs that are barely knowledge are a fine example: while they are epistemically good, they are epistemically better than some epistemically good beliefs and epistemically worse than others. From the perspective of the permissive axiologist, our practice of setting knowledge as an epistemic standard is deeply arbitrary.

The Teleological Account is perfectly compatible with the Permissive Axiological Account, but it readily explains the distinctive normative salience of knowledge: knowledge is distinctively normatively salient *because it sets the standard of epistemic success*. Beliefs that are knowledge, and assertions/actions properly based on them, are epistemic successes, while all other beliefs, assertions, and actions are epistemic failures.⁴⁵ Analogously, winning games is the aim of a basketball team *qua* basketball team. Because a team must score more points than the opposition to achieve this aim, our assessments of basketball teams are especially sensitive to the property of *scoring at least one point more than the opposing team*. This property is not salient

⁴⁵ At least, they fail with respect to the epistemic aim of knowledge. Perhaps there are additional epistemic aims as well.

merely because it marks a particular degree of basketball goodness, for many nearby properties to which we pay far less attention – *scoring at least two points more than the opposing team*, *scoring at least three points more than the opposing team*, etc. – also mark particular degrees of basketball goodness. We instead single out the property of scoring at least one point more than the opposing team because to instantiate it is to achieve the aim of basketball.

In general, then, the aim of any item of a given kind is salient in assessing that item *qua* member of that kind. Hence our emphasis on knowledge, even when barely obtained, over many nearby epistemically good properties – not because knowledge marks a particular degree of goodness (though it does), but because knowledge is an epistemic aim.

6. The source of epistemic teleology

In virtue of what do assertion, action, and belief aim at knowledge? Many aims exist because of our individual intentions. When I intend to sit on a chair but land ignominiously on the floor, my action fails because it misses the aim that I have set. But the aim of knowledge cannot be like this, for we rarely intend to meet any knowledge norm when we assert, act, or believe. We are not such deliberate creatures.

Other aims, however, exist because of our communal intentions. Suppose that George Weasley applies for Auror training as part of a crude practical joke: he slips a stink bomb into the application envelope. Though he knows perfectly well that his application will be rejected and is indifferent to this result, George's rejected application still fails *qua* application, for the aim of this Auror application *qua* application is determined not by George's intentions but by the community of applicants, admission committees, and so on. Indeed, *qua* application, George's application aims at acceptance even though no individual or group of individuals ever intends for it to be accepted. It has this aim simply in virtue of being an application.

Perhaps the aim of knowledge is like this. On this **communal account**, our communal intentions for our assertions, actions, and beliefs to be appropriately related to knowledge set an aim for these items. This aim is tacit, manifested in our patterns of assessment but rarely stated outright. It extends even to those beliefs which no one specifically intends to be knowledge, and even to those assertions and actions which no one specifically intends to be knowledge-based.

Yet some aims are set, not by human intentions, but by nature; in particular, by evolution. My heart has the proper function of pumping blood, and my liver of filtering toxins; as a consequence, these items have corresponding aims. These aims do not depend on my intentions or those of my community, but rather on the fact that the hearts and livers of my ancestors pumped blood and filtered toxins, thereby helping my ancestors to survive and reproduce.⁴⁶ According to the **proper functional account**, evolution also sets proper functions, and hence aims, for our capacities to assert, act, and believe. Knowledge is one such aim, since our ancestors survived and reproduced partly because of those assertions, actions, and beliefs which were knowledgeable.

The proper functionalist should admit that assertion, action, and belief have aims distinct from, and sometimes even incompatible with, knowledge. For our evolved capacities are complex and often messy, including not only aims nested within other aims, but also

⁴⁶ Though for purposes of vivid presentation I have filled in the details with the etiological theory of proper function developed by Wright (1975), Millikan (1984) and (1993), and Neander (1991), most alternative accounts of proper function can reasonably be substituted in its place.

competing aims. The aim of knowledge is no different. It is nested within other aims: belief aims not only to be knowledge, but also to combine with desire in effective action. And the aim of knowledge sometimes competes with other aims: when my unjustified belief that I will recover from cancer in fact helps my recovery, my belief fulfills one such competing aim. None of this should distract from the fact that knowledge is an important node in the tangled web of aims for assertion, action, and belief.

Officially, I remain neutral between the communal and proper functional accounts of the source of the knowledge aim, and indeed I leave open the possibility of some combination of these accounts.⁴⁷ Regardless, the Teleological Account takes only one notion as primitive: the notion of an *aim*. It provides paradigmatic examples of and theoretical remarks about aims without independently defining them, and it defines *success* in terms of aims. Assertion, action, and belief have many aims, one of which is the distinctively epistemic (rather than, say, practical) aim of knowledge.⁴⁸

7. Conclusion

Compelling data underwrite the Knowledge Account, but extant deontological and axiological formulations of it are deeply dissatisfying. Their cardinal sin is one of fixation: while they rightly recognize the centrality of knowledge, they do not do justice to the entire complement of epistemic norms.

My own view of the structure of epistemic normativity has instead been full-throated in its pluralism. I have claimed that knowledge sets the norm of epistemic success, and with this claim my opponents may agree in spirit and perhaps even in letter. But my innovation has been to treat this norm of knowledge as a principle of organization, rather than one of reduction, for other epistemic norms. In particular, I have held that epistemic permissibility requires, roughly, the use of an internal process suited to generating knowledge; that the impermissible is still epistemically excusable if, roughly, the subject could not reasonably have been expected to do better; and that epistemic goodness attaches to a disunified range of properties saliently related to knowledge, including of course knowledge itself, but also necessary conditions on knowledge, dispositions likely to lead to knowledge, and dispositions to become disposed to achieve knowledge.

This Teleological Account has none of the shortcomings of its rivals but all of their merits and more. It explains the importance, and even predicts the specific contours, of epistemic permissibility, excusability, and goodness, and it further organizes these predictions in a compact theoretical framework in which knowledge has pride of place. I conclude that it is the best version of the Knowledge Account.

⁴⁷ Unofficially, I incline towards a pure proper functional account for belief and action and a combined proper functional and communal account for assertion.

⁴⁸ Per the discussion of epistemic and non-epistemic norms in §1.

ACKNOWLEDGMENTS: For their remarkably extensive comments on this paper, I must specially thank Nico Silins and John Turri. I have also profited greatly from discussion or correspondence on these topics with Andrew Bailey, Amber Carpenter, Jay Garfield, Tamar Gendler, Alexander Jackson, Peter Kung, Sam Liao, Cathay Liu, Michael Pelczar, Martin Smith, Matt Walker, and participants in the Yale-NUS reading group. This work was supported by Yale-NUS College through grant number R-607-263-010-121.

REFERENCES

- Bird, A. (2007). "Justified judging." *Philosophy and Phenomenological Research* 74 (1): 81-110.
- Brown, J. (2008a). "The knowledge norm for assertion." *Philosophical Issues* 18: 89-103.
- Brown, J. (2008b). "Subject-sensitive invariantism and the knowledge norm for practical reasoning." *Noûs* 42 (2): 167-189.
- Brown, J. (2010). "Knowledge and assertion." *Philosophy and Phenomenological Research* 81 (3): 549-566.
- Brown, J. (2012). "Assertion and practical reasoning: common or divergent epistemic standards?" *Philosophy and Phenomenological Research* 84 (1): 123-157.
- Cohen, S. and Comesaña, J. (2013). "Williamson on Gettier cases and epistemic logic." *Inquiry* 56 (1): 15-29.
- Conee, E. and Feldman, R. (1985). "Evidentialism." *Philosophical Studies* 48: 15-34.
- DeRose, K. (1991). "Epistemic possibilities." *Philosophical Review* 100 (4): 581-605.
- Fantl, J. and McGrath, M. (2002). "Evidence, pragmatics, and justification." *Philosophical Review* 111 (1): 67-94.
- Gerken, M. (2011). "Warrant and action." *Synthese* 178: 529-547.
- Gettier, E. (1963). "Is justified true belief knowledge?" *Analysis* 23 (6): 121-123.
- Goldman, A. (1976). "Discrimination and perceptual knowledge." *Journal of Philosophy* 73 (20): 771-791.
- Goldman, Alvin I. (1979). "What is justified belief?" In *Justification and Knowledge*, ed. G. Pappas (Dordrecht: Reidel Publishing Company).
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, J. and Stanley, J. (2008). "Knowledge and action." *Journal of Philosophy* 105 (10): 571-590.
- Hellie, B. (2011). "There it is." *Philosophical Issues* 21 (1): 110-164.
- Jackson, A. (2012). "Two ways to put knowledge first." *Australasian Journal of Philosophy* 90 (2): 353-369.
- Lackey, J. (2007). "Norms of assertion." *Noûs* 41 (4): 594-626.
- Lackey, J. (2010). "Acting on knowledge." *Philosophical Perspectives* 24 (1): 361-382.
- Leite, A. (2013). "But that's not evidence; it's not even true!" *Philosophical Quarterly* 63 (250): 81-104.
- Maitra, I. and Weatherson, B. (2010). "Assertion, knowledge, and action." *Philosophical Studies* 149 (1): 99-118.
- McHugh, C. (2011). "What do we aim at when we believe?" *dialectica* 65 (3): 369-392.
- Millikan, R. (1984). *Language, Thought and Other Biological Categories*. MIT Press.
- Millikan, R. (1993). *White Queen Psychology and Other Essays for Alice*. Cambridge, Massachusetts: MIT Press.
- Moore, G. E. (1962). *Commonplace Book: 1919-1953*. London: Allen and Unwin.
- Moss, S. (2013). "Epistemology formalized." *Philosophical Review* 122 (1): 1-43.
- Neander, K. (1991). "Functions as selected effects: the conceptual analyst's defense." *Philosophy of Science* 58 (2): 168-184.
- Neta, R. (2009). "Treating something as a reason for action." *Noûs* 43 (4): 684-699.
- Slote, M. (1979). "Assertion and belief." In *Papers on Language and Logic*, ed. J. Dancy (Keele: Keele University Library).
- Smith, M. (2010). "What else justification could be." *Noûs* 44 (1): 10-31.
- Smith, M. (2012). "Some thoughts on the JK-rule." *Noûs* 46 (4): 791-802.

- Smith, M. (2014). "Knowledge, justification and normative coincidence." *Philosophy and Phenomenological Research* 89 (2): 273-295.
- Smithies, D. (2012). "The normative role of knowledge." *Noûs* 46 (2): 265-288.
- Sosa, D. (2009). "Dubious assertions." *Philosophical Studies* 146: 269-272.
- Sosa, E. (2007). *A Virtue Epistemology*. Oxford: Oxford University Press.
- Stanley, J. (2008). "Knowledge and certainty." *Philosophical Issues* 18 (1): 35-57.
- Sutton, J. (2005). "Stick to what you know." *Noûs* 39 (3): 359-396.
- Turri, J. (2010). "Prompting challenges." *Analysis* 70 (3): 456-462.
- Turri, J. (2011). "The express knowledge account of assertion." *Australasian Journal of Philosophy* 89 (1): 37-45.
- Turri, J. (2014). "Knowledge and suberogatory assertion." *Philosophical Studies*.
- Unger, P. (1975). *Ignorance: A Case for Scepticism*. Oxford: Oxford University Press.
- Weatherson, B. (2003). "What good are counterexamples?" *Philosophical Studies* 115 (1): 1-31.
- Weiner, M. (2005). "Must we know what we say?" *Philosophical Review* 114 (2): 227-251.
- Weisberg, J. (2013). "Knowledge in action." *Philosophers' Imprint* 13 (22): 1-23.
- Whitcomb, D. (forthcoming). "Can there be a knowledge-first ethics of belief?" In *The Ethics of Belief: Individual and Social*, eds. J. Matheson and R. Vitz (Oxford: Oxford University Press).
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Williamson, T. (2011). "Improbable knowing." In *Evidentialism and its Discontents*, ed. T. Dougherty (Oxford: Oxford University Press).
- Williamson, T. (2013a). "Gettier cases in epistemic logic." *Inquiry* 56 (1): 1-14.
- Williamson, T. (2013b). "Response to Cohen, Comesaña, Goodman, Nagel, and Weatherson on Gettier cases an epistemic logic." *Inquiry* 56 (1): 77-96.
- Williamson, T. (forthcoming). "Very improbable knowing." *Erkenntnis*.
- Wright, L. (1973). "Functions." *Philosophical Review* 82 (2): 139-168.