

**Must We Know What We Say?**

Matthew Weiner

Department of Philosophy

Texas Tech University

Final Draft

For publication in the *Philosophical Review*

*Abstract*

The knowledge account of assertion holds that it is improper to assert that  $p$  unless the speaker knows that  $p$ . This paper argues against the knowledge account of assertion; there is no general norm that the speaker must know what she asserts. I argue that there are cases in which it can be entirely proper to assert something that you do not know. In addition, it is possible to explain the cases that motivate the knowledge account by postulating a general norm that assertions would be true, combined with conversational norms that govern all speech acts. A theory on which proper assertions must be true explains the data better than a theory on which proper assertions must be known to be true.

The knowledge account of assertion seems to provide a clear and simple connection between the key concept of epistemology, knowledge, and a key concept of philosophy of language, assertion. According to the knowledge account (presented by Williamson (1996, 2000), DeRose (1996, 2002), and Unger (1975)), assertion is governed by the norm of knowledge, that we should assert only what we know. The knowledge account draws support from cases in which we judge assertion without knowledge to be impermissible, including lottery cases in particular. It is a surprising and powerful hypothesis; it takes two concepts that are defined and motivated independently and posits a strong connection between them that appears to be confirmed by linguistic data.

I will argue, however, that the knowledge account cannot succeed in unifying independently important conceptions of knowledge and assertion. On the most obvious accounts of knowledge and assertion, the knowledge account is inferior to a truth

account, on which the primary norm on assertion is that we should assert only what is true. The knowledge account simply gets some cases wrong, because assertion without knowledge is sometimes permissible. When we do judge assertion without knowledge to be impermissible, that judgment can be derived from the truth norm and other norms of conversation.

It would be hasty to declare the knowledge account simply false. It may be salvaged by relying on accounts of knowledge or assertion other than the most obvious ones. However, the salvaged knowledge account would no longer connect conceptions of knowledge and assertion that had been motivated independently; to save the knowledge account, we would have to tailor our conception of knowledge to our conception of assertion or vice versa. Furthermore, these tailored conceptions would not be as important as the original untailored conceptions, on which the knowledge account fails. A conception of knowledge that tracks our judgments of assertability will not do everything we want a conception of knowledge to do. Finally, the salvaged knowledge account may fail to explain the lottery cases that originally motivated it.

Accordingly, the knowledge account, if it survives, does so only as a shadow of its former self. After the account has been saved, the concepts it deals with are not as important as they seemed at first, and its truth appears to be stipulative rather than conceptual. The truth account, by contrast, explains the data in terms of more intuitive and important concepts of knowledge and assertion, and it can explain why we should have an act of assertion that is governed by the truth norm. Indeed, the fact that knowledge is not the norm of assertion suggests that knowledge may not be as important as epistemologists think.

## 1. Preliminaries

To begin with, we need to specify the initial targets of our analysis: the obvious accounts of assertion, knowledge, and conformity to norms. These accounts will be rough sketches, for there are no obvious necessary and sufficient conditions on assertion or knowledge. Nevertheless, these rough sketches yield some intuitive judgments concerning what counts as knowledge and what counts as assertion, and those judgments will supply a starting point for our analysis. These judgments, we will see in section 2, cannot be reconciled with the knowledge account of assertion, so that saving the knowledge account would require rejecting even the roughly sketched accounts I will describe.

Williamson suggests the obvious account of assertion when he remarks that “the default use of declarative sentences is to make assertions” (2000, 258). The difficulty here is to distinguish default uses from non-default uses. Some uses seem easy to exclude. Non-literal and fictional uses of declarative sentences will not in general carry out the speech act that the same declarative sentence would carry out when uttered *in propria persona*. Declarative sentences can be used to establish definitions, legal propositions, or bets only given some initial preparation, such as a statement that the following sentence is a definition or the investiture of legal authority in the speaker. They do not describe the world, as we might expect assertions to do, but rather change it.<sup>1</sup>

---

<sup>1</sup> In the terminology of (Searle and Vanderveken 1985), these acts have a double direction of fit rather than the word-to-world direction of fit that assertives have. (Thanks to Adam Leite for suggesting these speech acts.) Unger classifies at least some of these acts as acts of saying that *p* without asserting that *p*, attributing to Gilbert Harman the example of someone who bets that *p* by saying *p* (Unger 1975, 267n7); as Unger points out, the knowledge norm clearly does not apply to such bets. Uttering *p* as a declarative sentence will usually not be enough to bet that *p* unless you have been asked “What’s your bet?”; this is the sort of special setting that plausibly means that bets are not default uses of declarative sentences.

This remark also suggests some factors that will not be important in determining whether a certain speech act is an assertion. Williamson cites declarative mood rather than the tense of the sentence; accordingly, we should expect that future, past, and present tense sentences can all be used to make assertions. Furthermore, assertion is a broad category, which will not be restricted to acts that are based (or purport to be based) on certain kinds of evidence. Assertions will be more than just reports, for instance; the same declarative sentence can generally be used to assert something known firsthand, heard about, inferred, or arrived at through speculation.

Hence we will begin with a conception of assertion as a genus that comprises species such as reports, predictions, arguments, reminders, and speculations. This conception of assertion is, I think, the most obvious one. It is also a conception on which it is important to study the norms of assertion; a norm that applies to all these acts is more important than a norm that applies only to one of the species.

Our starting point for the analysis of knowledge will rest on intuitions that certain true beliefs do not count as knowledge, even aside from skeptical worries. This will not beg the question against contextualist or sensitive invariantist theories, on which the standards for knowledge depend in part on the ascriber's or the putative knower's context.<sup>2</sup> The

---

Williamson (2000, 249), citing Unger, suggests that the act of saying obeys the truth rule, but this seems mistaken with respect to Unger's examples. Uttering a falsehood as a bet may be expensive, but it is not improper. Nor is it clear that Williamson's own example, guessing the answers to the quiz, obeys the truth rule; it seems entirely proper to guess wildly on a quiz even if your answer turns out to be false. Indeed, it is not absolutely clear what Williamson means by 'saying' here, unless it is uttering a declarative sentence meant literally in any setting whatsoever. Since saying so defined includes many disparate illocutionary acts, we should not expect a homogenous account of its norms.

<sup>2</sup> DeRose (1995, 1996) argues for contextualism about knowledge, and in (DeRose 2002) argues that the knowledge account is implausible unless the standards vary with the ascriber's context. Williamson (2004) argues against such contextualism, although his earlier discussion of the knowledge account does leave room for contextualism (Williamson 2000, 254-5). Hawthorne (2004) holds to the knowledge account of assertion while pointing out the disadvantages of contextualism relative to subject-sensitive invariantism, on which the standards for knowledge vary with the putative knower's context rather than

intuition is that the cases I will describe (in section 2) do not count as knowledge under any standard. The evidence that the speakers have in these cases is too weak to yield knowledge, and the resulting beliefs would be unstable in a way that knowledge should not be. The resulting concept of knowledge will be important not only because it respects intuitions but also because it is important that our conception of knowledge comprise stability.<sup>3</sup>

The last preliminary is an initial account of what it is to conform to a norm. Here again I will begin by relying on messy intuitions: *Prima facie*, an assertion violates a norm if it somehow sounds wrong. There is a complication here, however, because of what DeRose calls the “secondary propriety/impropriety” (2002,180) that generally arises from norms.<sup>4</sup> Someone who reasonably believes that she is complying with a norm is in some sense acting properly, even if she is in fact violating the norm; and *vice versa*.<sup>5</sup> If I have every reason to think that I know that Alice is in her office, when in fact she has slipped out through the window, you may not condemn me for asserting that Alice is in her office, even though in so doing I violate the knowledge norm or the truth norm (whichever applies). This is because my assertion was secondarily proper even if it

---

with the ascriber’s. Leite (2004) argues that the knowledge account does not entail contextualism or subject-sensitive invariantism.

<sup>3</sup> This sketch of knowledge relies on claims about our intuitions; Weinberg, Nichols, and Stich (2001) (also in Nichols, Stich, and Weinberg 2003) have called into question any such use of intuition in epistemology. Nevertheless, the intuitions I cite will be useful in specifying the concept of knowledge that will initially be at issue, before we consider alternative conceptions of knowledge. In sections 4 and 5 I will argue that the resulting concept of knowledge is important for reasons above and beyond its conformity to intuition. This will take some of the philosophical weight off our intuitions about when a subject has knowledge, though I fear the intuitions still bear some weight in the argument.

<sup>4</sup> See also Williamson’s discussion of how it is reasonable to assert that *p* when it is reasonable to believe that you know that *p*, even though the assertion that *p* is not warranted (Williamson 2000, 257).

<sup>5</sup> This leaves open the question of what counts as reasonable belief that you are complying with a norm. One possibility is that the more important it is that a norm be complied with, the better supported your belief that you are complying must be in order to ensure secondary propriety. In the cases we discuss we can assume that the stakes are relatively low, so that primary and secondary propriety never do come apart.

violated the primary norm of assertion. We will deal largely with cases in which the distinction between primary and secondary propriety does not arise, because the speaker knows exactly what her epistemic situation is.

All three of these concepts, assertion, knowledge, and conformity to norms, are to some extent in play. We may decide that our initial conceptions need refinement or elaboration, and we may decide to revise one of the concepts in order to save the knowledge account. I will argue, however, that any revision that saves the knowledge account will leave the account much less significant than it seems at first glance.

## 2. The Data

The data that support the knowledge account fall into two kinds: lottery examples and sentences of the form “*p*, but I don’t know that *p*.” Each kind seems to support the knowledge account, given the accounts of the various concepts just sketched.

Suppose that Alice has bought a lottery ticket, with her odds of winning as low as you like. The drawing takes place at 9 a.m. and is publicly announced at noon.<sup>6</sup> At 10 a.m. Sarah says to Alice

(1) Your ticket didn’t win.

Sarah has no non-public information; her assertion is based on the overwhelming probabilities. Indeed, let us suppose that when the winning ticket is announced, it turns out that Alice’s has lost. Still, as Williamson notes (2000, 246), Alice is entitled to feel resentment against Sarah for asserting (1) on merely probabilistic grounds. *Prima facie*, this means that Sarah has violated a norm. The knowledge account explains this neatly:

---

<sup>6</sup> This presentation follows (Williamson 1996) (revised as chapter 11 of Williamson 2000), in particular in delaying the announcement of the lottery so that (1) is a past-tense statement. DeRose (1996) also cites the lottery cases in support of the knowledge account.

Intuitively, Sarah does not know that (1) is true, and so she has violated the knowledge norm.<sup>7</sup> By contrast, the truth norm does not seem to explain Sarah's impropriety. Her assertion fulfills the truth norm, nor is it secondarily improper, since she knows that her assertion is very probably true.<sup>8</sup>

Sentences of the form “*p*, but I don't know that *p*,” which I shall call *Moorean sentences*, generally seem paradoxical, and this also lends support to the knowledge account. Moore (1962, 277) observed that it seems odd to say

(2) Dogs bark, but I don't know that they do.

The knowledge account explains this oddity. If the second half of the sentence is true, then the speaker will not satisfy the norm of assertion for the first half. Either way, the norm of assertion is violated. Again, the truth norm seems to have trouble explaining why it always seems improper to assert (2), since (2) is sometimes true.<sup>9</sup>

---

<sup>7</sup> Note that there are accounts of knowledge on which Sarah does know (1). We do sometimes say to a lottery ticket holder, “Stop dreaming, you *know* your ticket won't win.” On the practical environment view of knowledge discussed by Hawthorne (2004, 179 and passim), a subject who is faced with certain decisions may know that her ticket won't win. I take it, however, that our intuition is that Sarah does not know, and that this case does lend initial plausibility to the knowledge account. (Compare Williamson's discussion of the jocular tone in which one might say “[Come off it—] Your ticket didn't win” (Williamson 2000, 246). Williamson takes it that this is acceptable but not a flat-out assertion; I think it would do less violence to intuition to say that it is an assertion but not flat-out acceptable, though my own view is that it is acceptable and a counterexample to the knowledge account. Hawthorne (2004, 18) argues that in this case the assertion is proper and the speaker does know what she asserts.)

<sup>8</sup> Secondary propriety requires that Sarah reasonably believe that she is conforming to the norm, and one might here object that merely probabilistic grounds are not enough for Sarah to reasonably believe that Alice's ticket did not win, no matter how high the probability. This, however, runs the risk of declaring that belief that one is obeying a norm is not secondarily proper unless one knows one is obeying the norm. On this account of secondary propriety, the truth account and the knowledge account would produce most of the same judgments concerning which assertions were improper; an assertion that violated the knowledge norm would be secondarily improper according to the truth norm. (An assertion of a proposition that the speaker knows but does not know she knows would be secondarily improper under the knowledge norm, but that will not hold of any of the examples under discussion.) This account of secondary propriety thus trivializes the truth norm. My objections to the knowledge account will also tell against this account of secondary propriety.

<sup>9</sup> I will argue, however, that whenever (2) is true the speaker lacks good reason to believe that it is true, and so the assertion of (2) involves a secondary impropriety with respect to the truth norm. See section 3.



A speaker's basis for assertion, however, can fall short of knowledge without relying on raw probability, and in these cases assertion without knowledge may be *prima facie* proper. This is particularly likely in the case of predictions and retrodictions (inferences to past unobserved happenings).<sup>10</sup> Take the following case of prediction: Captain Jack Aubrey has had long experience of naval combat against the French Navy. He and young Lieutenant Pullings have been watching French ships maneuver off Mauritius all day. At two p.m., Aubrey says to Pullings,

(3) The French will wait until nightfall to attack.<sup>11</sup>

Consider also the following case of retrodiction: Sherlock Holmes and Doctor Watson are brought to a crime scene. Holmes scans the scene and says (truthfully, as it turns out),

(4) This is the work of Professor Moriarty! It has the mark of his fiendish genius.

Holmes, at this point, has not found any evidence (in the criminal rather than epistemological sense) incriminating Professor Moriarty, but he is sticking his neck out based on his sense of what Moriarty's crimes are like.

Intuitively, Aubrey and Holmes do not know what they assert, even if their assertions turn out to be true. Contrast these cases with cases in which Aubrey's spy has intercepted the French orders, or in which Holmes has found some physical evidence linking Moriarty to the crime. In these cases it would be natural to say that they do have knowledge; in the cases originally presented, their beliefs do not rise to the same epistemic status. (At this point I am relying exclusively on an intuitive judgment that Aubrey and Holmes lack knowledge in the original cases. I will further discuss the significance of these intuitions in section 4.) Aubrey and Holmes also seem to be making

---

<sup>10</sup> The term 'retrodiction' comes from the list of assertives in (Searle and Vanderveken 1986).

<sup>11</sup> Readers who know the proper naval terms for 'two p.m.' and 'nightfall' are invited to supply them.

assertions; (3) and (4) are ordinary declarative sentences uttered in ordinary settings, and their evidential bases should not determine whether we count them as assertions (see section 1).

Nevertheless, the utterances of (3) and (4) do not seem to be improper. If the French do attack at nightfall, then it would be inappropriate for Pullings to feel resentment against Aubrey for asserting something that he did not know. If the French attack before nightfall, Pullings might be entitled to feel resentment against Aubrey for misleading him, but in such a case Aubrey would have violated the truth norm.<sup>12</sup> Similarly, Watson will not and should not resent Holmes' making an assertion without knowledge. Holmes has not claimed knowledge; Watson can tell that Holmes' assertion is based on a hunch.

These cases pose a problem for the knowledge account. They seem to be proper assertions in the absence of knowledge; the speakers' grounds for belief fall short of those necessary for knowledge, yet the assertions do not seem wrong or odd. The truth account, by contrast, predicts properly that these assertions will be permissible so long as they turn out to be true. If they turn out to be false, they will indeed have violated the truth norm; but if they turn out to be false, they do seem improper. The hearer may reproach the speaker for having told him something false (barring secondary propriety). So here the truth account yields better results than the knowledge account.

### 3. The Gricean Explanations

To show that the truth account is superior to the knowledge account, I will have to show not only that it can explain certain data that the knowledge account cannot explain,

---

<sup>12</sup> In some sense it may be inappropriate for Pullings to feel resentment in any case, because a lieutenant is not entitled to resent his captain. Abstracting from such considerations; we are concerned with whether Pullings *qua* epistemic being is entitled to feel resentment.

but also that it can explain the data that motivated the knowledge account: the impermissibility of assertion in the lottery case, and the general unacceptability of Moorean sentences. My explanation will be Gricean. The interaction of the truth norm with general norms on conversation explains why, in these cases, it is impermissible to assert without knowledge.

Briefly, the idea is that there is a norm, akin to Grice's Cooperative Principle (1989, 26), that one's utterances have some point. It would be pointless for Sarah to tell Alice that her ticket hadn't won unless Sarah had some inside information not obviously available to Alice. Accordingly, when Sarah asserts that Alice's ticket did not win, she implicates that she has inside information. Because she lacks inside information, Alice has the right to feel resentful. This resentment is grounded in the falsehood of Sarah's implicature rather than in the mere fact of Sarah's having asserted what she does not know.

Suppose Sarah's assertion

(1) Your ticket did not win

conveyed only its literal significance, that Alice's ticket did not win. On the truth account, this assertion would be proper if Alice's ticket did not win, which as it happens is true. But this assertion is also pointless. It violates Grice's second maxim of Quantity, "Do not make your contribution more informative than is required" (1989, 26). Indeed, the assertion of (1) is not required so long as the following three conditions hold:

(5a) Alice is already aware of the overwhelming likelihood that her ticket didn't win

(5b) Sarah is not presenting herself as knowing anything that could strengthen Alice's belief (cited in (5a)) that her ticket did not win

(5c) Alice does not need reminding that (most likely) her ticket didn't win.<sup>13</sup>

Contrapositively, for Sarah's assertion to have some point, one of (5a-c) must be false. Alice knows that (5a) and (5c) are true, and by the Cooperative Principle she is entitled to assume that Sarah's assertion does have a point; so she is entitled to assume that (5b) is false and Sarah is presenting herself as having inside information. Thus Sarah's assertion of (1) implicates that she has inside information.

Williamson considers and rejects this Gricean solution. His first objection is that if we use Gricean reasoning to explain why it is unacceptable for Sarah to assert (1), parallel Gricean reasoning commits us to the conclusion that it is unacceptable to assert

(6) Your ticket is almost certain not to have won.

Asserting (6) would be as pointless as asserting (1), since Sarah does not have anything to add to what Alice already knows. Yet, Williamson argues, asserting (6) is not unacceptable in the same way as asserting (1); "the worst to be said of [the] assertion [of (6)] is that it is banal and unkind" (2000, 147). So a parallel to our Gricean argument seems to yield the wrong result in the case of (6).

Examined more closely, however, the parallel breaks down. These are the conditions, parallel to (5a-c), under which asserting (6) is pointless:

(7a) Alice is already aware that her ticket is almost certain not to have won

(7b) Sarah is not presenting herself as knowing anything that could strengthen Alice's belief (cited in (7a)) that her ticket is almost certain not to have won

(7c) Alice does not need reminding that her ticket is almost certain not to have won.

---

<sup>13</sup> More precisely, asserting (1) is pointless if (5b) is true and (5a) and (5c) are mutually known to Alice and Sarah, but in the cases in question this mutual knowledge is present.

As with (5a) and (5c), Alice knows (7a) and (7c).<sup>14</sup> The Gricean argument with respect to (1) was that, since Sarah's assertion is pointless if (5a-c) are untrue, Alice is entitled to assume that (5b) is false. With respect to (6), however, Alice also knows that (7b) is false. If (7a) and (7c) are true, then Alice already knows that her ticket is almost certain not to have won and does not need reminding of it. *A fortiori* she already has full belief in what Sarah is asserting, and Sarah cannot present herself as knowing anything that would strengthen this belief further.<sup>15</sup> So (7b) cannot be false either.

Hence the most charitable interpretation Alice can make of Sarah's assertion of (6) is that, contrary to appearances, Sarah thinks that (7a) or (7c) is false. Then Sarah would mean to inform or remind Alice that her ticket is almost certain not to have won. If this interpretation is implausible, the most plausible interpretation is that Sarah is violating the Cooperative Principle by making a pointless assertion. We judge it to be banal and unkind because of its pointlessness.

Williamson also argues that the impermissibility of lottery assertions cannot rest on a Gricean implicature, because the impermissibility cannot be removed by an utterance meant to cancel the implicature: "I have no more evidential authority to assert 'Your ticket did not win, but I do not mean to imply that I have inside information' than I have to assert the plain 'Your ticket did not win'" (Williamson 2000, 248). Williamson is correct that it would be improper for Sarah to assert

(10) Your ticket did not win, but I do not mean to imply that I have inside information,

---

<sup>14</sup> Indeed, they are mutually known; see note 13 above.

<sup>15</sup> What if Sarah had evidence that Alice's ticket was even less likely to have won than Alice thought? This would not strengthen Alice's belief that her ticket was almost certain not to have won. Rather it would give her grounds for a belief with a different content, "My ticket is even less likely to have won than I thought before." Hence this scenario would not falsify (7b).

but the impropriety of (10) would not rest on Sarah's lack of evidence. Rather, (10) is improper because it is pointless on its face. It violates a general norm of conversation rather than a specific norm of assertion. Hence (10) is reproachable in a different way than (1) is. If Sarah asserts (1), Alice can reproach her by saying "You didn't know that," once she discovers that Sarah lacked inside information. If Sarah asserts (10), it would be inappropriate for Alice to say "You didn't know that"; Sarah has already admitted as much. Rather, Alice can reproach her by saying, "Why say anything, if you lack inside information?" Sarah's assertion of (10) is banal and unkind rather than misleading.<sup>16</sup>

The third Williamsonian objection that we will consider raises the most profound issues. This objection is that the Gricean explanation fails to account for a case in which Alice does not know the odds against her ticket winning. In such a case Sarah would be telling Alice something new by asserting

(1) Your ticket did not win.

In this case (5a) is already known to be false: Alice is unaware of the overwhelming odds against her ticket winning. Hence (5b) need not be false in order for Sarah's assertion to have a point. Then it seems as though there is no way to generate the implicature that Sarah knows that Alice's ticket did not win; so on the Gricean account it should be acceptable for Sarah to assert (1). Yet, Williamson argues, this assertion is as unacceptable in this case as in the case in which Alice already knew that her ticket is unlikely to have won.

---

<sup>16</sup> Alternatively, Alice may conclude that Sarah does not mean her denial of inside information to be taken seriously. (This might be akin to "You didn't hear this from me, but your ticket didn't win.") Then the implicature is not canceled because the assertion that would cancel the implicature is itself not meant to be believed. The implicature that Sarah has inside information is still generated by Gricean principles in this case; Alice reconciles Sarah's utterance with the Cooperative Principle by concluding that "I do not have inside information" is not to be believed, so that Sarah's utterance does have some point.

Note first that there are circumstances in which it is acceptable for Sarah to assert (1) on probabilistic grounds, precisely because Alice does not know that her ticket is almost certain not to have won. Suppose that the drawing of the ticket is broadcast on television at a certain time. Some prankster has hooked up Sarah and Alice's television to a videotape of a simulated announcement that Alice's ticket has won. Sarah knows of the prank, Alice does not. After watching the simulated announcement, Alice believes that her ticket won the lottery. Then it is appropriate for Sarah to say, "Your ticket didn't win; that's a videotape of a fake broadcast."<sup>17</sup> Sarah's assertion, given her explanation, would not deceive Alice into thinking that Sarah had inside knowledge of the actual drawing; she needs to correct Alice's ill-grounded belief, and the overwhelming likelihood that Alice's ticket did not win is sufficient grounds for her assertion.

In the sort of case Williamson has in mind, however, it is unacceptable for Sarah to assert (1) even though Alice does not know that it is unlikely that her ticket has won. We imagine a case exactly like the original lottery case, in which the winning ticket has been drawn but not announced, except that Alice thinks there are very few tickets and Sarah knows there are millions. In this case, it is still inappropriate for Sarah to assert that Alice's ticket did not win, even though this (most probably) conveys a true proposition that Alice does not believe. Sarah should assert (6), that Alice's ticket almost certainly has not won.

Nevertheless, this case can be explained in terms of the truth norm and more general norms of conversation, and this explanation will generalize. It will allow us to explain why Moorean sentences generally seem paradoxical, and why Aubrey's and Holmes'

---

<sup>17</sup> Note that Sarah must give some explanation after saying "Your ticket didn't win"; otherwise her assertion will seem obviously false.

assertions can be proper in the absence of knowledge. The explanation is that, if assertions are governed by the truth norm, the hearer may reasonably expect that the speaker has some warrant for what she says. The most plausible warrants for assertions that are not predictions or retrodictions will generally, in conjunction with truth, be sufficient for knowledge.

Recall first DeRose's distinction between primary and secondary propriety/impropriety: If an act is governed by a norm, primary propriety is determined by whether the act conforms to the norm, and secondary propriety is determined by whether the agent has reason to believe that the act conforms to the norm. Thus, if assertion is governed by the truth norm, an assertion is secondarily improper if the speaker does not have reason to believe that it is true. When Sarah asserts (1), Alice is entitled to assume that Sarah has some justification not obviously available to Alice herself for believing that her ticket didn't win. Since the drawing has already taken place, it is plausible that this justification is that Sarah has somehow learned which ticket was drawn. Then Sarah would know that Alice's ticket did not win; so Alice is justified in inferring that Sarah knows that her ticket did not win, and has cause for resentment when she discovers the merely probabilistic grounds for Alice's assertion.

Why should Alice infer that Sarah's warrant is not merely probabilistic? The answer lies in Grice's maxim of Manner. If Sarah wants to make clear that her warrant is merely probabilistic, she has the option of saying

(6) Your ticket is almost certain not to have won.

Since Sarah did not say (6), Alice is entitled to assume that her warrant is not merely probabilistic. Since (1) is ambiguous between two kinds of warrant that Sarah may have,



and (6) is appropriate if and only if Sarah has a probabilistic warrant (by the first maxim of Quantity), the assertion of (1) implicates that Sarah has a non-probabilistic warrant.<sup>18</sup>

This account predicts that, the more reason Alice has to believe that Sarah cannot have non-probabilistic grounds for asserting (1), the more acceptable it will be for Sarah to assert (1) on probabilistic grounds. For Alice's reason to believe that Sarah lacks non-probabilistic warrant may override the implicature generated by the maxim of Manner.

This prediction is borne out. Suppose that Sarah says, before the drawing of the lottery,

(11) Your ticket won't win.

This sounds more acceptable than the assertion of (1) in the original set-up. It may be banal, unkind, and pointless to assert (11), but it is not likely to mislead. This is because it is much more unlikely that Sarah has inside information about the winning ticket before it is drawn than after it is drawn. So (11) carries at most a weak implicature that Sarah has non-probabilistic warrant for what she says, much weaker than (1).

Sarah's assertion of (11) is even more acceptable if her information is new to Alice.

Suppose that there is a puzzle contest with ties to be broken by lot. Alice thinks that she is the only one to have solved the puzzle, but Sarah knows that there have been millions of correct entries. Sarah can say

---

<sup>18</sup> An anonymous referee suggests that a parallel argument would entitle Alice to conclude that Sarah has a probabilistic warrant. Sarah did not assert (6\*):

(6\*) Your ticket is absolutely certain not to have won

which would entail that the speaker has non-probabilistic warrant. So, the suggestion goes, if asserting (1) instead of (6) implicates that Sarah has non-probabilistic warrant for her assertion, asserting (1) rather than (6\*) implicates that Sarah lacks non-probabilistic warrant for her assertion. Both implicatures cannot hold, so if the argument is truly parallel asserting (1) does not implicate non-probabilistic warrant.

I do not think that the argument is truly parallel. Though (6\*) entails that Sarah has non-probabilistic warrant, (6\*) cannot be felicitously asserted in every case in which she has non-probabilistic warrant. If Sarah has merely heard from a usually reliable source that Alice's ticket did not win, she cannot assert (6\*), because she is not *absolutely certain* that the ticket did not win. So in asserting (1) rather than (6\*), Sarah does not implicate that she has probabilistic warrant rather than a non-probabilistic warrant that falls short of absolute certainty. (I thank the referee for making me see that "if and only if" rather than "only if" was required in the text.)

(12) You won't win; there were millions of correct entries.

This seems acceptable, if it turns out to be true. It is not calculated to mislead, since Sarah has made her non-probabilistic grounds clear, and it does not even have the air of paradox. If the knowledge account were true, it should be self-contradictory for Sarah to simultaneously assert that Alice's entry won't win and to acknowledge that her grounds are probabilistic, hence presumably insufficient for knowledge.<sup>19</sup>

Similar reasoning explains why Aubrey's and Holmes' assertions are acceptable in the absence of knowledge. In each case, it is clear that the speaker's grounds are not the sort that would be sufficient for knowledge. For Aubrey to know what the French will do, he would need better of evidence than he can have gathered from watching them maneuver. Similarly, Holmes has only glanced at the crime scene; he has not had time to gather any evidence that would justify us in saying that he *knows* that Moriarty is the culprit.<sup>20</sup> The audience for these assertions can see how the asserters gathered their evidence and so will not be misled into thinking that they have better evidence than they do.

Indeed, the maxim of Manner has little bearing on these cases, because in them the grounds for assertion are non-probabilistic. Aubrey's evidence does not definitively establish that the French probably will not attack before nightfall; the reason he has to believe this is just the reason he has to believe that the French will attack before nightfall. Asserting "The French will probably attack before nightfall" rather than "The French will attack before nightfall" would convey a lack of confidence rather than the (already

---

<sup>19</sup> It is necessary for Sarah to make the non-probabilistic grounds explicit because she could have non-probabilistic grounds for her assertion; she might have known that Alice's solution was flawed.

<sup>20</sup> Again, I am for the moment relying on the intuitive judgment that Aubrey and Holmes lack knowledge; I will discuss the significance of these intuitions in section 4.

obvious) fact that Aubrey's grounds are insufficient for knowledge. Hence the listener can draw no conclusions from the fact that Aubrey chose the unhedged assertion.

The same basic line of reasoning explains why the Moorean utterance

(2) Dogs bark, but I don't know that they do

seems paradoxical. The person who asserts (2) must have some warrant for believing that dogs bark, or she has committed secondary impropriety with respect to the truth norm. This assertion is about the habits of dogs as they are now. By far the most likely warrant for such an assertion is having seen and heard dogs bark or having been told that dogs bark by a trustworthy informant, and either of these warrants is sufficient for knowledge. If the speaker's situation is so unusual that she has some other warrant for believing that dogs bark, she should say so. Similarly for sentences such as "It's raining" or "Albert won the batting title last year"; someone who asserts these while disclaiming knowledge is admitting that she does not have the most likely and most satisfactory warrant for the truth of her assertion.

Predictions and retrodictions are generally acceptable in the absence of knowledge precisely because the most likely and satisfactory warrant for believing in their truth is not sufficient for knowledge.<sup>21</sup> Indeed, predictions and retrodictions can be maintained

---

<sup>21</sup> This will not hold for every prediction. DeRose discusses the oddity of saying, of a ball whose future causal path is undetermined, "The ball won't veer to the left, but it might" (1998, 74); he argues that the first half represents the speaker as knowing that the ball won't veer to the left, and the second half entails that the speaker does not know this. As he has set up the case, the only possible grounds for asserting that the ball won't veer to the left are either probabilistic or dependent on access to God's foreknowledge, which access would certainly be sufficient for knowledge. Probabilistic grounds would be better expressed by saying "The ball probably won't veer to the left"; so if the speaker simply asserts "The ball won't veer to the left," it is implicated that she has non-probabilistic grounds, which in this case entail knowledge.

Most predictions, however, will admit of a warrant that is not merely probabilistic but that still falls short of knowledge. Aubrey's and similar warrants are non-probabilistic in that they are not grounded in knowledge of the probability of what they assert, but they are intuitively insufficient for knowledge; for discussion of reasons other than intuition why we should not count them as giving knowledge, see section

while knowledge is explicitly disclaimed. Suppose that after Aubrey asserts (3), Pullings asks, “How do you know that the French will attack at nightfall?” and Aubrey responds,

(13) I don’t *know* they’ll attack at nightfall—we haven’t intercepted their orders—but my prediction is that they will.

Here Aubrey is still asserting that the French will attack at nightfall; for he explicitly describes it as a prediction, which (recall from section 1) is a kind of assertion. (See section 4 for more on why it is important to count predictions as assertions.) Nor does (13) sound paradoxical.<sup>22</sup> Hence, even if the bald Moorean formulation “*p* but I don’t know that *p*” sounds odd, it is possible to conjoin an assertion with a denial of knowledge.<sup>23</sup>

The truth account, together with Gricean mechanisms, thus can explain the data that motivated the knowledge account. For our assertions to be proper, not only must they be true, we must have reason to believe them true. The hearer thus is entitled to conclude that the speaker has some warrant for her assertion. In many cases the most likely warrant combined with truth will be enough for knowledge. When it is obvious that the speaker’s warrant would not be enough for knowledge, assertion without knowledge is permissible.

---

4. It is important to recognize that there can be non-probabilistic warrants that are still insufficient for knowledge.

<sup>22</sup> It resembles an actual unprompted utterance made by a friend at a baseball game. The team I root for had loaded the bases with the score tied and one out in the top of the tenth inning, and she said, “Your team is going to win. That’s my prediction.” (They did.)

<sup>23</sup> A contextualist might argue that the standards for knowledge shift when Aubrey stresses ‘know’ at the beginning of (13), so that his denial of knowledge takes place in a higher-standards context than his original assertion of (3). DeRose gives an example of such mid-utterance context shift through stress on ‘know’: A student, confronted with the brain-in-a-vat scenario, who says, “Well, I know [that I am sitting in a classroom], even though I don’t *KNOW* it” (DeRose 1998, 71). This account would still have to explain why standards shift back down when Aubrey reiterates his prediction. In any case, intuitively (and as discussed in section 4) Aubrey’s assertion of (3) should not count as knowledge by *any* standard.

#### 4. Possible Defenses of the Knowledge Account

I have argued that the truth account can explain the data that the knowledge account explains; we should consider whether the knowledge account can accommodate the data that the truth account explains. There are three non-exclusive ways to fit cases such as (3) and (4) to the knowledge account. One is to say that the utterances in question are not really assertions; one is to say that the speakers really do know what they say; one is to say that the utterances are not really permissible. Each of these defenses, I will argue, weakens the knowledge account by substituting a less important conception for a more important one. In addition, some of them drain the account of its power to explain the original lottery case (1). The steps required to save the knowledge account make it no longer worth saving.

First, let us consider ruling that (3) and (4) are not flat-out assertions. ‘Assertion’ is largely a philosopher’s term, which we are free to define; we might define it so that predictions and retrodictions do not count. The question is what principled definition of ‘assertion’ would exclude predictions and retrodictions without trivializing the knowledge account. It would be possible, for instance, to define assertions as claims to knowledge, but then the knowledge account of assertion would be true simply by definition. Even if we stop short of this, we may drain the knowledge account of interest. In section 1 I suggested that on the obvious account, assertion is the genus of speech act typically performed by utterance of a declarative sentence, which includes reports, predictions, retrodictions, arguments, reminders, etc. On this account, assertion is ubiquitous, and it is important to discover if any norms apply across the whole genus. If we restrict the set of assertions arbitrarily, we risk demoting assertion to one species

among many, so that it would be hard to see why we should care about the norms of assertion any more than the norms of prediction.

In any case, if we restrict the set of assertions so as to exclude (3) and (4), we will also exclude the lottery case that motivates the knowledge account. (1) is a retrodiction, and more typical lottery assertions such as (11) are predictions; it is hard to see why they should count as assertions if (3) and (4) do not. So the knowledge account of assertion with a restricted domain of assertion fails to explain the data that the knowledge account originally appears to explain.

If we exclude predictions and retrodictions from the set of assertions, we will be left with firsthand reports, secondhand testimony based on others' testimony, and perhaps inferences based on well-established laws as opposed to speculation. These, however, cannot be properly asserted in the absence of knowledge, even on the truth account. In section 3, I argued that on the truth account many assertions cannot properly be made in the absence of knowledge, because the speaker's most likely warrant for the assertion is enough for knowledge. This holds for all the assertions in the restricted domain; if the speaker is reporting what she has observed, passing along what she has been told, or drawing an inference according to a well-established law, then she knows what she asserts. So the knowledge account and the truth account yield the same predictions in this restricted domain. Restricting the domain of assertions to the domain in which the knowledge account is true removes any advantages the knowledge account might have.

More promising is the attempt to save the knowledge account by arguing that Aubrey and Holmes do know the truth of their assertions (3) and (4), respectively. To explain all the data while saving the knowledge account in this way, we would need an account of

knowledge that also explains why knowledge fails in the lottery case. DeRose's analysis of knowledge based on the Subjunctive Conditionals Account [SCA] (DeRose 1995, 1996) can be made to yield this result.<sup>24</sup> On the SCA, by and large someone fails to know that  $p$  when she would believe that  $p$  even if  $p$  were false.<sup>25</sup> This excludes knowledge in the lottery case; even if Alice's ticket had won, Sarah would still believe that it had lost. Plausibly, it does not exclude knowledge in the Aubrey and Holmes cases; in the nearest possible world in which the French attack before nightfall, they maneuver differently at 2 p.m., and Aubrey does not conclude from watching these maneuvers that they will wait to nightfall to attack; similarly for Holmes.

The SCA, read thus, has the intuitive cost that knowledge becomes much easier to attain than it seemingly should be. On the SCA, a believer can obtain knowledge that  $p$  based on any bit of non-misleading evidence; were  $p$  false, the evidence would not have been produced, and the believer would not have come to believe  $p$ . Knowledge seems as though it should require more than just a scrap of evidence.

Indeed, the resulting concept of knowledge lacks one of the qualities that make knowledge important: stability. As Williamson points out, knowledge that  $p$  is more stable than Gettierized justified true belief that  $p$ , and this can mean that knowledge explains the knower's actions better than mere justified true belief would. He gives the example of a burglar who spends all night ransacking a house because he knows that

---

<sup>24</sup> I do not wish to claim that DeRose is committed to this result; his analysis could easily be modified to avoid it (if it does not avoid it already). But if the analysis does not yield the result, it cannot be used to save the knowledge account of assertion as described in the text. (Nor do I wish to suggest that this is how DeRose would want to save the knowledge account of assertion.)

<sup>25</sup> (DeRose 1995) elaborates this account in a contextualist fashion to avoid failures of deductive closure. The account might also be elaborated in a subject-sensitive invariantist way (see Hawthorne 2004 for the definition of subject-sensitive invariantism). The details of the elaboration do not matter for our present purposes. Williamson's criticisms of DeRose (Williamson 2000, 156-61) apply to any account based on SCA.

there is a diamond in the house. If the burglar had a justified true belief in the diamond based on a false lemma, he might discover the lemma's falsity and stop. For instance, if he had been told that there was a diamond under the bed, when in fact the diamond was in a drawer, he would stop searching once he discovered there was no diamond under the bed (Williamson 2000, 62).

This is not to say that knowledge will be unshakable. If the knower receives enough of the right kind of misleading counterevidence, then he should abandon his belief. If, for instance, the burglar hears someone leave a message on the answering machine saying that Mrs. X should not worry, because the diamond she lost at the party an hour ago was found by the host, then he should abandon his belief that the diamond is in the house, even if the call is a practical joke, the diamond really is in the house, and the burglar had knowledge before hearing the phone call.<sup>26</sup> But knowledge entitles the knower to disregard counterevidence up to a point. If the burglar knows that there is a diamond somewhere in the house, for instance because he has been told by a trustworthy informant, then he should remain confident that the diamond is there even after it fails to turn up in the first few likely locations.

Note that in this case the burglar with the Gettierized belief would also be entitled to resist counterevidence. If he sees that the house is shabby in a way that ordinarily would make him doubt that there is a diamond in the house, but a usually trustworthy informant has told him that there is a diamond under the bed, then he is entitled to disregard the shabbiness of the house until he looks under the bed and sees no diamond there. But, if the diamond had been under the bed, this burglar would have known it. The reason that

---

<sup>26</sup> Thanks to an anonymous referee for this example, and for pressing me on the sense in which knowledge is stable.



this belief can resist counterevidence is that it is based on a warrant that in the right circumstances would suffice for knowledge. So stability is characteristic of knowledge in this way: A believer is entitled to resist counterevidence if her belief is based on a warrant that will suffice for knowledge if the world cooperates (that is, if the belief is true and not Gettierized).

Aubrey's and Holmes' beliefs entirely lack this stability; they are not entitled to resist counterevidence at all. If the French navy's continued maneuvers provide any indication that they will attack before nightfall, Aubrey should weaken his belief that they will wait till nightfall. Contrast the case in which Aubrey's spy has intercepted the French orders. In that case Aubrey is entitled to ignore counterevidence up to a point. If he sees them maneuvering as if to attack sooner, he has reason to dismiss that maneuver as a feint; only when the counterevidence becomes overwhelming must he abandon his belief. In this case it is intuitive to say that Aubrey knows when the French will attack, and his belief is stable in the way that knowledge should be. Similarly, Holmes is not entitled to discount any evidence of Moriarty's innocence.

I have cited Williamson's invocation of stability as important for explaining the knower's actions, but it is also important for the knower herself. We have an obvious interest in the stability of our true beliefs. If we are carrying out a plan based on some true belief, we will need the belief to last until we have completed the plan, or we will abandon the plan midway. A belief that needed to be revised as soon as any counterevidence came in would have little chance of enduring long enough to be the basis of any substantial plan. So there is good reason not to count beliefs that entirely lack

stability as knowledge. We should reject the SCA-derived account that counts Aubrey's and Holmes' unstable beliefs as knowledge.

Taking stability as criterial of knowledge helps in systematizing the intuitions about knowledge that I have been invoking. The warrants I have cited as intuitively sufficient for knowledge produce beliefs that should be stable. Reports grounded on observation or trustworthy testimony give us grounds for discounting future contrary evidence; hunches based on experience should be revised as new evidence comes in. Hence the predictions and retrodictions that are clearly based on hunches are the assertions that turn out to be acceptable in the absence of knowledge, because they are based on warrants that yield no stability. And, since stability is important, this conception of knowledge as requiring stability is a conception worthy of philosophical attention.

A third way of defending the knowledge account is to say that (3) and (4) really do violate norms of assertion, even though they seem permissible. This will have to go beyond the invocation of secondary propriety (see section 1), or of the reasonableness of asserting something that you have reason to believe you know (Williamson 2000, 257). Holmes and Aubrey are well aware that they do not know what they assert, so they are not respecting the knowledge norm even secondarily. The idea is rather that (3) and (4) seem permissible because we do not reproach the speakers for their violation of the norm of assertion. Williamson suggests that in certain contexts, "lively seminar discussion, or gossip," enforcement of the knowledge rule is lax (2000, 258). Perhaps (3) and (4) seem acceptable because they are uttered in contexts of lax enforcement.

This suggestion, however, cannot explain our judgment that the lottery assertion (1) is unacceptable. The contexts of assertion of (3) and (4) seem no more conducive to lax

enforcement than the context of (1), so there is no reason to think that (1) would be treated as a serious breach of the norms while (3) and (4) are not. Then (1) should seem as acceptable as (3) and (4); but the knowledge account is motivated by the obvious unacceptability of (1). Indeed, Williamson's suggestion concerning what happens in contexts of lax enforcement is that "we feel entitled to assert  $p$  whenever we are not confident that we do not know  $p$ " (2000, 259). On this account, Aubrey and Holmes should not feel entitled to assert (3) and (4); so this suggestion cannot reconcile the knowledge account with these data.

In any case, this suggestion would make it much less urgent to discover what the norms of assertion might be. If we often do not care whether the norms of assertion are violated, they are not as important as they seemed. The truth account performs better here. When someone asserts a falsehood her assertion is wrong; the assertion itself is faulty even if the speaker is not to be criticized. Violation of the truth norm can be criticized except when it is excused by some other general norm; when the speaker reasonably believed that what she said was true (making her assertion secondarily proper), or when the speaker conveys what is true and most important about the matter at hand. The truth norm has real teeth.

##### 5. The Significance of the Concepts

I have argued that, on the most intuitive and most important conceptions of assertion, knowledge, and conformity to norms, the knowledge account of assertion does not hold. Asserting what one does not know can be permissible and in conformity with the norm of assertion. Nevertheless, as the last section has shown, there are several alternative ways

of reading the concepts involved, and it is worth asking why the particular conceptions at issue are important. I have already said that a broad conception of assertion is more interesting than a narrow one because it covers a genus rather than species of speech acts, and that the conception of knowledge we began with (which generates the intuitions that Aubrey and Holmes do not know what they assert) is more important than a more permissive one because it ensures that our knowledge is relatively stable. Still, it is worth asking: Why should assertion, one of our most common speech acts, require truth rather than knowledge?

This echoes Williamson's question, based on the idea that the norm of assertion is knowledge, why we should have a speech act with that norm in our repertoire (Williamson 2000, 267). Williamson answers this with an analogy between assertions and commands. If I am commanded not to do *A*, Williamson argues, it is not enough for obedience that I do not do *A*; I must ensure or bring it about that I do not do *A*. Similarly,

[t]o make an assertion is to confer a responsibility (on oneself) for the truth of its content; to satisfy the rule of assertion, by having the requisite knowledge, is to discharge that responsibility, by epistemically ensuring the truth of the content (Williamson 2000, 268-9).

Of course the knower does not ensure the truth of the content by bringing it about, but on Williamson's view of knowledge she is in a broad mental state that guarantees the truth of the content.<sup>27</sup>

Commands, however, are a specific, unyielding kind of imperative, unlike requests and advice. Perhaps to obey a command you must ensure that what is commanded is done; but you may get credit for doing what has been requested if you try your best and succeed. Assertions are a genus of speech act, more analogous to imperatives in general

---

<sup>27</sup> See (Williamson 2000), chapter 1.

than to commands in particular. This suggests that we may discharge our responsibility as asserters when we succeed in telling the truth by doing our best to do so. If we have reason to believe what we assert, and it is true, it does not matter that we were not in a state that ensured its truth.

The knowledge norm is too stringent because we should sometimes speak whereof we do not know. We are imperfect enough that we will sometimes have some idea concerning the truth of the matter at hand without being able to guarantee its truth. Sometimes, when speed is more important than guaranteed accuracy, it will be better to risk a falsehood in order to communicate economically what is most likely true. Keeping to what we know would sometimes require adding uneconomical qualifiers. This lack of economy would be particularly egregious in contexts like Aubrey's and Holmes', where the speaker's lack of knowledge is evident. This does not deprive us of the power to assume the heavier burden of guaranteeing the truth of our assertions; we can always explicitly claim knowledge. But we will not always want to.

The foregoing discussion assumes that knowledge is relatively hard to attain, so that our beliefs may be well grounded enough to be worth communicating but not well grounded enough for knowledge. This will be so if knowledge justifies some resistance to contrary evidence, as I argued in section 4. Even when we recognize that we have no special right to discount contrary evidence, we may still want to communicate our best estimate of the facts, especially if our listener would not be competent to evaluate that evidence or to make a similarly well-grounded estimate. But we may ask again whether knowledge has to justify this resistance. We want our beliefs to be stable, but we also

want them to justify assertion. Why should the former outweigh the latter when we ascribe knowledge?

We might think of rescuing the knowledge account by ruling that proper asserters do have knowledge; not by invoking an independent theory of knowledge, as with the SCA (discussed in section 4), but by identifying the epistemic position required for knowledge with that required for assertability. Then we would have an assertability account of knowledge rather than a knowledge account of assertion. This assertability account could gain support from the practical environment view of knowledge discussed by Hawthorne (2004, 176 and *passim*). On the practical environment view, whether a belief counts as knowledge on a certain occasion depends on whether it can properly be used in practical reasoning on that occasion (given that it is relevant to the decision in question). If the practical issue is “Should I express my belief that  $p$ , given that I wish to assert the truth?” then you will count as knowing that  $p$  whenever you should indeed assert that  $p$ .

This line of thought shows that we cannot have everything we might want from the concept of knowledge. On the perhaps more intuitive view, knowledge is a substantial accomplishment that entitles us to resist counterevidence, but it is not the norm of assertion. On the view that makes knowledge the norm of assertion, it can be had based on any bit of evidence, but it vanishes with the next countervailing bit. As we saw in section 4, this latter view makes knowledge an unsatisfactory guide to extended planning, because even when you know something you cannot be confident that you will continue to believe it for long enough to put your plan through. Thus, there is no conception of knowledge on which it serves as a guide to both assertion and to planning.

In the end, this suggests that knowledge may be less important than it is usually taken to be. When we ask which beliefs count as knowledge, we hope to discover beliefs that are satisfactory in several different ways: suitable to be communicated, bases for sound planning, results of good inquiry, certain and secure, as well as according with our intuitive judgments of knowledge.<sup>28</sup> But we have discovered that there is no conception of knowledge that lines up with all these ways of evaluating beliefs.<sup>29</sup> Perhaps epistemology should be less concerned with the exact contours of knowledge and more with the individual ways in which beliefs can be evaluated.<sup>30</sup>

### Works Cited

- DeRose, Keith (1995). Solving the Skeptical Problem. *Philosophical Review* 104: 1-52.
- DeRose, Keith (1996). Knowledge, Assertion, and Lotteries. *Australasian Journal of Philosophy* 74: 568-580.
- DeRose, Keith (1998). Simple 'Might's, Indicative Possibilities and the Open Future. *Philosophical Quarterly* 48: 67-82.
- DeRose, Keith (2000). Now You Know It, Now You Don't. In *Proceedings of the Twentieth World Congress of Philosophy*, vol. 5, 91-106. Bowling Green, Ohio: Philosophy Documentation Center.
- DeRose, Keith (2002). Assertion, Knowledge, and Context. *Philosophical Review* 111: 167-203.
- Grice, Paul. (1989). Logic and Conversation. In *Studies in the Way of Words*, 22-40. Cambridge, Mass.: Harvard University Press.
- Hawthorne, John (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Kaplan, Mark (1985). It's Not What You Know That Counts. *Journal of Philosophy* 82: 350-363.

---

<sup>28</sup> Several of these items correspond roughly to items on the list of desiderata on knowledge given by Hawthorne (2004, 111-2).

<sup>29</sup> See also (Kaplan 1985) for an argument that knowledge is not useful as a guide to inquiry.

<sup>30</sup> Thanks for helpful comments to Adam Leite, Gary Kemp, John MacFarlane, Jeff Rubard, Andrea Westlund, and the referees for the *Philosophical Review*. I am especially grateful to Ram Neta, who read through two drafts of this paper and offered extensive comments.

Leite, Adam (2004). How to Link Assertion and Knowledge without Going Contextualist: A Reply to DeRose's "Assertion, Knowledge, and Context." Manuscript.

Moore, G.E. (1962). *Commonplace Book: 1919-1953*. London: Allen & Unwin.

Nichols, Shaun, Stephen Stich, and Jonathan M. Weinberg (2003). Meta-Skepticism: Meditations in Ethno-Epistemology. In *The Sceptics: Contemporary Essays*, edited by Steven Luper, 227-247. Aldershot: Ashgate Publishing.

Searle, John R., and Daniel Vanderveken (1985). *Foundations of Illocutionary Logic*. Cambridge: Cambridge University Press.

Unger, Peter (1975). *Ignorance: A Case for Skepticism*. Oxford: Oxford University Press.

Weinberg, Jonathan, Shaun Nichols, and Stephen Stich (2001). Normativity and Epistemic Intuitions. *Philosophical Topics* 29: 429-460.

Williamson, Timothy (1996). Knowing and Asserting. *Philosophical Review* 105: 489-523.

Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.

Williamson, Timothy (2004). Knowledge, Context, and the Agent's Point of View. Manuscript.