



## Philosophical Review

---

Knowing and Asserting

Author(s): Timothy Williamson

Source: *The Philosophical Review*, Vol. 105, No. 4 (Oct., 1996), pp. 489-523

Published by: [Duke University Press](#) on behalf of [Philosophical Review](#)

Stable URL: <http://www.jstor.org/stable/2998423>

Accessed: 29/09/2011 12:31

---

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at  
<http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Duke University Press and *Philosophical Review* are collaborating with JSTOR to digitize, preserve and extend access to *The Philosophical Review*.

<http://www.jstor.org>

## Knowing and Asserting

Timothy Williamson

Assertions are praised as true, informative, relevant, sincere, warranted, well-phrased, or polite. They are criticized as false, uninformative, irrelevant, insincere, unwarranted, ill-phrased, or rude. Sometimes they deserve such praise or criticism. If any respect in which performances of an act can deserve praise or criticism is a *norm* for that act, then the speech act of assertion has many norms. So has almost any act; jumps can deserve praise as long or brave, criticism as short or cowardly. But it is natural to suppose that some norms are more intimately connected to the nature of asserting than any norm is to the nature of jumping. One might suppose, for example, that someone who knowingly asserts a falsehood has thereby broken a rule of assertion, much as if he had broken a rule of a game; he has cheated. On this view, the speech act, like a game and unlike the act of jumping, is constituted by rules. Thus, not all norms for assertion are on a par. Norms such as relevance, good phrasing, and politeness are just applications of more general cognitive or social norms to the specific act of assertion. Perhaps the norm of informativeness results from a more complex interaction between a general norm of cooperativeness and the nature of assertion as a source of information. But on this view, not all norms for assertion derive from more general norms; otherwise nothing would differentiate it from other speech acts.

This paper aims to identify the constitutive rule(s) of assertion, conceived by analogy with the rules of a game. That assertion has

---

Versions of this material were given in talks to a conference in Glasgow (Jim Edwards delivered a stimulating reply), a meeting of the Scots Philosophical Club, and at the Universities of Edinburgh, Canterbury (Christchurch), Waikato, and Wollongong, Monash University, the Australian National University, and the Instituto de Investigaciones Filosóficas of the Universidad Nacional Autónoma de México. I am grateful to the audiences there for helpful questions, to Jack Copeland, Matti Eklund, Jon Kvanvig, John O'Leary-Hawthorne, and Ralph Wedgwood for detailed criticism of the material, and to the University of Canterbury and the Research School of Social Sciences at the Australian National University for hospitality while the paper was written. Comments by a referee for the *Philosophical Review* have significantly improved the paper.

such rules is by no means obvious; perhaps it is more like a natural phenomenon than it seems. One way to find out is by supposing that it has such rules, in order to see where the hypothesis leads and what it explains. That will be done here. The hypothesis is not perfectly clear, of course, but we have at least a crude conception of constitutive rules, which we may refine as we elaborate the hypothesis. Although no attempt will be made here to *define* 'rule', some remarks on constitutive rules will focus the discussion.

Constitutive rules are not conventions. If it is a convention that one must  $\phi$ , then it is contingent that one must  $\phi$ ; conventions are arbitrary and can be replaced by alternative conventions. In contrast, if it is a constitutive rule that one must  $\phi$ , then it is necessary that one must  $\phi$ . More precisely, a rule will count as constitutive of an act only if it is essential to that act: necessarily, the rule governs every performance of the act. This idealizes the case of games, for, in the ordinary sense of 'game', games such as tennis gradually change their rules over time without losing their identity; the constitutive role of the rules is qualified by that of causal continuity. Similarly, in the ordinary sense of 'language', natural languages such as English gradually change their rules over time without losing their identity. Nevertheless, in a technical sense of 'language' that the philosophy of language has found fruitful, the semantic, syntactic, and phonetic rules of a language are essential to it (Lewis 1975). The richer ordinary sense of 'language' introduces needless complications. Linguistic conventions and the consequent possibility of linguistic change can then be accommodated at a different point in the theory: a population that at one time has the convention of speaking a language L may later change to a convention of speaking a distinct language L\*, constituted by slightly different rules. Likewise, in the present technical sense of 'speech act', the rules of a speech act are essential to it. A population that at one time has the convention of using a certain device to perform a speech act A may later change to a convention of using that device to perform a distinct speech act A\*, governed by slightly different rules. 'Game' can receive a similar sense. Henceforth, 'rule' will mean constitutive rule.

Given a game G, one can ask, "What are the rules of G?" Given an answer, one can ask the more ambitious question "What are noncircular necessary and sufficient conditions for a population to play a game with those rules?" Competent unphilosophical um-

pires know the answer to the former question but not to the latter. Given a language *L*, one can ask, "What are the rules of *L*?" Given an answer, one can ask, "What are noncircular necessary and sufficient conditions for a population to speak a language with those rules?" Given a speech act *A*, one can ask, "What are the rules of *A*?" Given an answer, one can ask, "What are noncircular necessary and sufficient conditions for a population to perform a speech act with those rules?" This paper asks the former question about assertion, not the latter. It cannot wholly ignore the latter, for assertion is presented to us in the first instance as a speech act that we perform, whose rules are not obvious. In order to test the hypothesis that a given rule is a rule of assertion, we need some idea of the conditions for a population to perform a speech act with that rule; otherwise we could not tell whether we satisfy those conditions. Fortunately, we need much less than a full answer to the second question for these purposes. Our task is like that of articulating for the first time the rules of a traditional game that we play; that does not require a full philosophy of games.

Constitutive rules do not lay down necessary conditions for performing the constituted act. When one breaks a rule of a game, one does not thereby cease to be playing that game. When one breaks a rule of a language, one does not thereby cease to be speaking that language; speaking English ungrammatically is speaking English. Likewise, presumably, for a speech act: when one breaks a rule of assertion, one does not thereby fail to make an assertion. One is subject to criticism precisely because one has performed an act for which the rule is constitutive. Breaches of the rules of a game, language, or speech act may even be common. Nevertheless, some sensitivity to the difference—both in oneself and in others—between conforming to the rule and breaking it presumably is a necessary condition of playing the game, speaking the language, or performing the speech act. The important task of elucidating the nature of this sensitivity will not be undertaken here.

The normativity of a constitutive rule is not moral or teleological. The criticism that one has broken a rule of a speech act is no more a moral criticism than is the criticism that one has broken a rule of a game or language. Although someone who knowingly asserts a falsehood may incur moral criticism, perhaps for having betrayed the hearers or inflicted false beliefs on them, such faults

are made possible only by the specific nature of assertion, which is not itself constituted by moral norms. Cheating at a game is likewise not a morally neutral act, but it is made possible only by the nonmoral rules that constitute the game. Nor is the criticism that one has broken a constitutive rule of an institution the criticism that one has used it in a way incompatible with its aim, whether the aim be internal or external. Consider a game with the internal aim of scoring more goals than the opposition and the external aim of exercising players or entertaining spectators. Breaking the rules can serve both internal and external aims. Conversely, lazy play can give away goals to the opposition, bore spectators, and fail to exercise players, without breaking the rules. Within the practice constituted by the rules, their authority does not require the backing of moral or teleological considerations.

What are the rules of assertion? An attractively simple suggestion is this: There is just one rule. Where 'P' is a schematic sentence letter and 'C(P)' expresses a condition, the rule says:

*The C(P) rule* One must: assert that P only if C(P).

In the imperative, assert that P only if C(P). As used here, 'must' expresses the kind of obligation characteristic of constitutive rules. The rule is to be parsed as 'One must ((assert that P) only if C(P))', with 'only if C(P)' inside the scope of 'One must' but outside the scope of 'assert'. The rule unconditionally forbids this combination: one asserts that P when not C(P). The combination is possible, or else it would be pointless to forbid it. The condition that C(P) may concern the content of the potential assertion (that P), contextual features (for example, speaker and time), or both. The C(P) rule is constitutive of the speech act: necessarily, assertion is a speech act A whose unique rule is "One must: perform A with the content that P only if C(P)." Furthermore, the envisaged account takes the C(P) rule to be individuating: necessarily, assertion is the *unique* speech act A whose unique rule is that one. In mastering the speech act of assertion, one implicitly grasps the C(P) rule, in whatever sense one implicitly grasps the rules of a game or language in mastering it. As already noted, this requires some sensitivity to the difference, both in oneself and in others, between conforming to the rule and breaking it. All other norms for assertion are the joint outcome of the C(P) rule and considerations not specific to assertion. If an assertion satisfies the rule,

whatever derivative norms an assertion violates, it is correct in a salient sense.<sup>1</sup> Call this account the C(P) account, and any account of this form, *simple*.

More complex accounts of assertion are conceivable. Some rules make some assertions obligatory; silence satisfies the C(P) rule. There might be several rules of assertion. There might be none. Assertion might be wholly or partly constituted by a norm or norms whose normativity is not rule-like. Such a norm might be essentially comparative: mastery of the speech act would involve grasping a scale on which assertions could be assessed as better or worse than each other, but not grasping a threshold for an assertion to be “good enough”—that could be left to the discretion of individual speakers with particular purposes. Alternatively, assertion might be constituted only by non-normative features. Nevertheless, a simple account of assertion would be theoretically satisfying, if it worked. This paper defends a simple account, leaving the examination of more complex accounts to others.

One obvious candidate to play the role of the condition that C(P) is truth—the truth of the content that P. Since truth behaves disquotationally, this is equivalent to the condition that P. Thus, the rule boils down to this:

*The truth rule* One must: assert that P only if P.

The truth rule forbids false assertions. It would be broken often, but so are many rules. The truth account—a simple account of assertion based on the truth rule—explains other norms as the joint outcome of the truth rule and considerations not specific to assertion. In particular, it explains epistemic norms as norms of evidence for truth: satisfying these secondary norms consists in having evidence that one satisfies the primary norm. Grice describes the category of Quality in his account of the rules of conversation in this vein: the supermaxim “Try to make your contribution one that is true” leads to two more specific maxims, “Do not say what you believe to be false” and “Do not say that for which you lack adequate evidence” (Grice 1989, 27). ‘Evidence’ here

---

<sup>1</sup>In contrast, Dummett’s notion of correctness for assertions is undifferentiated: “to say one thing when there is no point in saying it rather than something more usual in those circumstances can be misleading if it prompts the hearers to suppose a point that does not exist; it, too, is therefore incorrect in the general sense” (1991, 168).

does not mean *mere* evidence; when one noninferentially knows that P, one's evidence for the assertion that P may simply be that P ("self-evidence").

Unlike truth, other candidates to play the role of the condition that C(P) are sensitive to the epistemic circumstances of the asserter. A speaker who satisfies such a condition will be described as having warrant to assert that P, in a schematic sense of 'warrant'. On such views, the rule becomes:

*The warrant rule* One must: assert that P only if one has warrant to assert that P.

The warrant rule forbids unwarranted assertions. For any reasonable notion of warrant, a true assertion based only on a lucky guess will satisfy the truth rule without satisfying the warrant rule. Even so, versions of the warrant rule can be embedded in radically different simple accounts of assertion. On one kind of account, the content of an assertion consists in the condition for having warrant to make it; or perhaps it consists in that condition and the conditions for having warrant to make other structurally related assertions—for example, the negation of the assertion. This account somehow reduces truth to an abstraction from warrant, and derives the norm of truth from the warrant rule. Such an account can be called *antirealist*, although the term could equally well be applied to accounts of content in which truth plays no role at all.

The warrant rule can also be embedded in a different kind of account, on which having warrant to assert that P amounts to *knowing* that P. Then the warrant rule takes this form:

*The knowledge rule* One must: assert that P only if one knows that P.

The knowledge rule would be broken even more often than the truth rule, but so are many rules. The knowledge account—a simple account of assertion based on the knowledge rule—explains the norm of truth as a mere corollary of the knowledge rule: satisfying the former is a necessary but not sufficient condition of satisfying the latter. If one knows that P then P. Nevertheless, this account in no way limits the transcendence of truth over warrant; still less does it make the former an abstraction from the latter. "One knows that P" is not conceptually prior to "P."<sup>2</sup> This account

---

<sup>2</sup>The knowledge account does not by itself guarantee realism. The term

can be called *realist*. Given plausible connections between knowledge, belief, evidence, and truth, the knowledge account explains what is right about Grice's two more specific maxims of quality, "Do not say what you believe to be false" and "Do not say that for which you lack adequate evidence," as well as the supermaxim "Try to make your contribution one that is true."<sup>3</sup>

This paper defends the knowledge account. The ramifications of the account for the debate between realism and antirealism will not be pursued here (see Williamson 1995b). The account can be roughly summarized in the slogan "Only knowledge warrants assertion." 'Warrant' is used here as a term of art, for that evidential condition (if any) which plays the role of the condition that C(P) in the correct simple account of assertion. This use need not correspond exactly to the use of 'warrant' in everyday English. It is not denied that false assertions are sometimes warranted in the everyday sense that they are sometimes reasonable; the claim is rather that the reasonableness of such assertions is explicable as the joint outcome of the knowledge rule and cognitive considerations not specific to assertion. Still, if the account is correct, ordinary speakers are implicitly sensitive to the knowledge rule, for they must have implicitly grasped it in mastering assertion. It is just that they need not use the word 'warrant' for that norm. Much of the evidence for the knowledge account comes from the ordinary practice of assertion.

## 1. The truth account

It is somehow good to assert the true and bad to assert the false. Is that idea articulated by the truth account, the simple account of

---

'warranted assertibility' goes back to Dewey (1938), who combines a kind of antirealism with the identification of warranted assertibility and knowledge (7, 143) by means of a pragmatist conception of knowledge.

<sup>3</sup>This point is central to the case for the revised maxim of Quality, "Say only that which you know," in Gazdar 1979 (46–48) ('say' should be read as 'assert'; see n. 4 below). Note that Grice's maxim of sincerity requires someone who asserts that P to lack the belief that not P, not (as might be expected) to have the belief that P. But if one asserts that P while agnostic as to whether P, one is insincere in a way that seems to flout conversational rules. Conversely, knowing that P while unable to rid myself of what I recognize as an irrational belief that not P, I might appropriately assert that P, contrary to Grice's maxim but not to the more obvious sincerity condition.



assertion based on the truth rule? This section argues that such an account is incorrect, and that its defects recommend the knowledge account.

One doubt about the truth account is that assertion is not the only speech act to aim at truth. For many speech acts A, normatively different from assertion and from each other, it is somehow good to perform A with a true content and bad to perform A with a false content. By definition, the truth account entails that the truth rule is individuating, that is, that assertion is the *unique* speech act A whose unique rule is "One must perform A with the content that P only if P." In this sense, the truth account claims that assertion is more intimately associated with the aim of truth than any other speech act is. But no basis is discernible for assigning this privilege to assertion in preference to all those other speech acts.

There is, for example, a speech act of *conjecturing* that P, for which the evidential norms are more relaxed than they are for assertion. Although it is somehow good to conjecture the true and bad to conjecture the false, it is quite acceptable to conjecture that P, but not assert that P, when it is merely more probable than not that P on one's evidence. In English, one can perform this speech act by using the words 'I conjecture' parenthetically, as in "P, I conjecture." Equally, there is a speech act of *swearing* that P, for which the evidential norms are more stringent than they are for assertion. Not only is it somehow good to swear the true and bad to swear the false, it is acceptable to swear that P only if one has grounds for unusual certainty that P, more than is required to assert that P. In English, one can perform this speech act by using the words 'I swear' parenthetically, as in "P, I swear."<sup>4</sup> Indeed, there is a whole range of speech acts, differing in their evidential norms but all in some sense aiming at the truth. Why should assertion be the only one of them to be a speech act A whose unique rule is

---

<sup>4</sup>Compare Slote 1979 (182–87) on parenthetical uses of 'I believe'. Even *saying* may be weaker than asserting (Unger 1975, 267, with credit to Harman). Note that what matters to the argument is not the ordinary use of 'conjecture', 'swear', and 'say', but the possibility of speech acts of the kind described. Note also that attempts to differentiate the speech acts and uphold the truth rule by adding rules about the gravity of breaches of it depart from the structure of a simple account. They also fail to meet a subsequent objection to the truth account.

“One must perform A with the content that P only if P,” as the truth account requires?

It might be held that although not all assertings are swearings, all swearings are assertings—swearing that P would be a solemn way of asserting that P. This would not upset the argument. Conjecturing that P is no way of asserting that P. The evidential standard required for asserting would still be intermediate between those required for conjecturing and swearing. The question would remain: is that intermediate standard more intimately connected with the aim of truth than all the other standards are?

Simple accounts of assertion based on evidential rules face no such difficulty. They correspond to simple accounts of conjecturing and swearing based on evidential rules which require more and less respectively than does the rule of assertion. The speech acts are thereby differentiated from each other.

Although the preceding doubt about the truth account suggests (without showing) that the rule of assertion is evidential, it fails to indicate an appropriate standard of evidence. A stronger objection to the truth account will now be developed, one that does cast light on the appropriate standard of evidence.

Assertion obviously has some kind of evidential norm. It is somehow better to make an assertion on the basis of adequate evidence than to make it without such a basis. Now assume the truth account, for an eventual *reductio ad absurdum*. Then the evidential norm is derivable from the truth rule. One ought to have evidence for one's assertions *because* they ought to be true.

The proposed derivation is simple. Its core is an inference from the premise that one must not assert something when it is not true to the conclusion that one should not assert it when one lacks evidence that it is true. Since evidence that an assertion is true just is evidence for that assertion, the truth account implies that one should not make an assertion for which one lacks evidence. The underlying principle is quite general and has nothing to do with the fact that it is truth that the rule of assertion is supposed to require. It may be stated as a schema, with parentheses to indicate scope:

- (1) If one must ( $\phi$  only if P), then one should ( $\phi$  only if one has evidence that P).

The transition from ‘must’ to ‘should’ represents the transition

from what a rule forbids to what it provides a reason not to do. For example, if one must not bury people when they are not dead, then one should not bury them when one lacks evidence that they are dead. It is at best negligent to bury someone without evidence that he is dead, even if he is in fact dead. The proposed explanation of the evidential norm substitutes 'assert that P' for ' $\phi$ ' in (1). Clearly, there is much room for variation in the letter of (1) without violation of its spirit.

On a charitable reading of (1), the required weight of evidence that P will vary with the badness of  $\phi$ ing when not P. One should take more care to avoid killing people than to avoid offending them, if the risks are equal in probability. The question is whether (1), so read, can explain the weight of evidence that we require speakers to have for their assertions in terms of the degree of badness that we attribute to making an untrue assertion. Is the former proportionate to the latter?

Consideration of lotteries suggests a negative answer. Suppose that you have bought a ticket in a very large lottery. Only one ticket wins. Although the draw has been held, the result has not yet been announced. In fact, your ticket did not win, but I have no inside information to that effect. On the merely probabilistic grounds that your ticket was only one of very many, I assert to you flat-out, "Your ticket did not win," without telling you my grounds. Intuitively, my grounds are quite inadequate for that outright unqualified assertion, even though one can construct the example to make its probability on my evidence as high as one likes, short of 1, by increasing the number of tickets in the lottery. You will still be entitled to feel some resentment when you later discover the merely probabilistic grounds for my assertion. I was representing myself to you as having a kind of authority to make the flat-out assertion that in reality I lacked. I was cheating.<sup>5</sup>

Can the fault in my assertion be explained by appeal to some version of (1)? The explanation would have to be that making an untrue assertion is so bad that one should not run even a minute

---

<sup>5</sup>Dudman 1992 uses the point to argue against the thesis that the assertibility of a conditional varies with the conditional probability of its consequent on its antecedent; Lowe 1995 (44) follows Dudman; contrast Edgington 1995 (287). Of course, there is a special jocular tone in which it is quite acceptable to say "[Come off it—] Your ticket didn't win," but the tone signals that the speaker intends not to make a serious assertion.

risk of doing so. Is that plausible? We may well regard both honesty and the pursuit of truth as very serious matters, but it does not follow that we must regard every untrue assertion as a serious crime; the pursuit of truth would not get very far if we did. When we discover that we have inadvertently asserted something false on some casual matter, most of us are wracked by no more guilt than we feel when we inadvertently tread on someone's toes. In the present case, let it be common knowledge between us that the result of the lottery will be announced within a few minutes, and that you care little whether your ticket wins. Thus the bad consequences of the falsity of my assertion that I risk inflicting on you—but do not actually inflict, since my assertion is in fact true—amount to your briefly having a false belief (if you believe my assertion) on a matter about which you care little. Ordinarily, we should not regard the fact that an action of mine involved a one-in-a-million risk of inflicting consequences of such limited badness on you as much cause for criticism. Yet you are entitled to insist that I was quite wrong to make the assertion that I did, for I had no authority to do so. That criticism of me does not derive from the kind of consideration embodied in (1). No assessment of the probability or gravity of untruth is even relevant to the criticism. The point is simply that, in making the assertion, I exceeded my evidential authority. In other cases, where untruth is less improbable or worse in its consequences if it does occur, the speaker is no doubt subject to *further* criticism on those grounds, but they should not be allowed to obscure the possibility of criticizing speakers simply for exceeding their evidential authority.

Could a defender of the truth account explain what is wrong with my assertion by appeal to Gricean rules of conversation? The idea would be that my assertion was misleading because you, to whom I was speaking, were entitled to assume that the grounds on which I made it were not obviously already available to you, so were entitled to assume that I had inside information about the result of the lottery. For making the assertion on grounds obviously already available to you might be held to violate one of the maxims of Quantity, "Do not make your contribution more informative than is required" (Grice 1989, 26). However, if that Gricean point were the objection to my assertion, then the objection would extend to case (a), in which I assert, "Your ticket is almost certain not to have won," and the objection would not extend to case (b),

in which I assert, “Your ticket did not win,” but (unlike the previous cases) it is not obvious that you know how many tickets other than your own have been sold. For in case (a), parallel Gricean reasoning would indicate that you are entitled to assume that the grounds on which I made my assertion were not obviously already available to you, and therefore that you are entitled to assume that I had inside information about the result of the lottery—for example, evidence that it was almost certain to have been rigged in favor of someone else. In case (b), my grounds for the assertion—the number of tickets sold—are not obviously already available to you; so the assumption to that effect, which the argument supposes you to be entitled to make, is true, and the Gricean objection lapses. In fact, however, the problem behaves in the opposite way to that predicted by the Gricean explanation.<sup>6</sup> It does not extend to case (a), in which the worst to be said of my assertion is that it is banal and unkind. The problem does extend to case (b), in which you are still entitled to feel resentment at the merely probabilistic grounds for my assertion. Probabilistic evidence warrants only an assertion that something is probable.

A further problem for the Gricean explanation is that I should be able to remove the objection to my assertion by explicitly canceling the supposed conversational implicature. I am not. 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.”

A different defense of the truth account appeals to the fact that, for each ticket, I have a similar basis for asserting that it did not win: if I make all those assertions, I shall have asserted something false. But how does that explain what is wrong with making any one of them, granted the truth rule? Consider an analogy. I am faced with an enormous pile of chocolates. I know that exactly one of them is contaminated and will make me sick; alas, I cannot tell them apart. I have a strong desire to eat a chocolate. I can quite reasonably eat just one, since it is almost certain not to be contaminated, even though, for each chocolate, I have a similar reason for eating it, and if I eat all the chocolates, I shall eat the contaminated one, and my sickness will be overdetermined. No plausible

---

<sup>6</sup>The criticism is not of Grice’s theory of conversational implicature but of an overenthusiastic application of it.

principle of universalizability implies that, in the circumstances, any reason for taking one chocolate is a reason for taking them all; the most to be implied is that, in the circumstances, any reason for taking one chocolate is a reason for taking any other chocolate instead. The truth account does not supply the resources to rule out the possibility that there is adequate evidence for each of the assertions "*t* did not win" but not for their conjunction. If each conjunct is true then the conjunction is also true, of course, but it does not automatically follow that the same goes for adequate evidence of truth. That entitlement to assert each conjunct implies entitlement to assert the conjunction may be an independently plausible principle; but the truth account cannot explain it.

It is not even essential to the lottery case that each ticket have an equal chance of winning. Consider a variant lottery in which each ticket is assigned a publicly known weight proportional to its probability of winning, and your ticket has a somewhat lower weight than the others. I am still not entitled to assert that your ticket will not win, even though my evidence that it will not win is now better than for any other ticket. Alternatively, the lottery might even be one in which there was probably *no* winning ticket (DeRose 1996).

It might finally be protested that if I lacked warrant to assert, "Your ticket did not win," then we lack warrant to make most of our ordinary assertions, because few of them are quite certain. Of course, it follows that I had warrant only given the anti-skeptical premise that we do have warrant to make most of our ordinary assertions. Moreover, the protest simply assumes that no other account of assertion can discriminate between "Your ticket did not win" and most of our ordinary assertions. That assumption needs testing; it is rejected in the next section. In any case, for whatever reasons, probabilistic bases are ordinarily taken to be inadequate for assertion.

The truth account does not explain something that it is committed to explaining: the evidential norms for assertion. It should therefore be rejected. More positively, assertion seems to be governed by a nonderivative evidential rule, which my assertion in the lottery case broke; I was cheating.

One possible explanation is this: The rule of assertion is the knowledge rule; one must not assert that *P* unless one knows that *P*. In the lottery case, it is intuitively clear, given the nature of my

evidence, that I did not know that your ticket did not win.<sup>7</sup> Thus, my assertion violated the rule of assertion. After all, the natural way for you to articulate the criticism that I lacked evidential authority for my assertion is by saying, “But you didn’t *know* that my ticket hadn’t won!” This argument will be developed in the next section.

## 2. The knowledge account

One may lack the evidential authority to make an assertion about a lottery, even though it is very highly probable on one’s evidence. I will now argue that the underlying phenomenon is general to assertions about any subject matter.

Let “P” express a proposition whose truth value is known to an expert but about which you have no evidence. The expert holds a lottery. There are a million tickets, of which you have one. However, he does not announce the number of the winning ticket; he merely hands each participant a slip of paper. If your ticket won, the true member of the pair {“P”, “Not P”} is written on your slip; if your ticket lost, the false member of the pair is written there. There is no doubt that this is the arrangement. You are not in a position to confer with other participants. “Not P” turns out to be written on your slip. On your evidence, there is a probability of one in a million that your ticket won and “Not P” is true, and a probability of 999,999 in a million that your ticket lost and “Not P” is false. Thus, if you assert “P,” the probability on your evidence that your assertion is true is 999,999 in a million. Intuitively, however, you are not entitled to assert “P” outright. That intuition can be supported. On your evidence, you can certainly assert “P if and only if my ticket did not win”; by hypothesis, that biconditional is not in doubt. If you assert “P,” you are therefore surely in a position to detach, and assert “My ticket did not win.” But you are not entitled to assert that, for your only evidence is that your ticket was one in a million. That “Not P” rather than “P” was written on

---

<sup>7</sup>Arguing for an externalist view of knowledge, Armstrong denies that we have knowledge in the lottery case (1973, 185–88); in defending an internalist view of knowledge from Armstrong’s argument, BonJour concedes the absence of knowledge but proposes an internalist explanation of it (1985, 56). See also Harman 1968, Dretske 1981 (99–102), Craig 1990 (98–103), DeRose 1996, and Lewis 1996.

your slip tells you nothing, for you have no independent evidence for or against those propositions. Thus you are not entitled to assert "P," even though it has a probability on your evidence of 999,999 in a million.

"P" could be *any* assertion about which you happen to have no evidence. Indeed, even if you have probabilistic evidence that tends to support "Not P" against "P," the number of tickets in the lottery can be made so large that your probabilistic evidence from the lottery for "P" will overwhelm your other evidence against "P." Thus, the argument indicates that, for almost any kind of proposition at all, very high probability on one's evidence does not imply assertibility. The propositions not covered by the argument are those for which one is bound to have independent nonprobabilistic evidence, for example, "I exist": but they are not plausible candidates for assertion on a merely probabilistic basis. The obvious moral is that one is *never* warranted in asserting a proposition by its probability (short of 1) alone. What matters in the original lottery case is not the subject matter of the assertion but the probabilistic basis on which it was made.

To say that no probability short of 1 warrants assertion is not yet to say that only knowledge warrants assertion. Some nondeductive forms of inference might be held sometimes to warrant assertion nonprobabilistically without providing knowledge; an example is inference to the best explanation. It is hard to see how inference to the best explanation could ever generate numerical probabilities, but even if it does lead to conclusions of high probability short of 1, it would not warrant assertion *in virtue* of doing so. The implication is that one might have warrant to assert the conclusion of an inference to the best explanation, even though one lacked warrant to assert an equally probable proposition whose high probability had a different basis; no inference to the best explanation provided the probabilistic evidence that your ticket did not win. Such a view is consistent, but is it plausible? If one has warrant to assert a proposition of probability less than 1 on one's evidence, then in some lottery case one lacks warrant to assert a proposition—perhaps the very same proposition—of higher probability on one's evidence.

Assume, plausibly, that if it is less probable that P than that Q on one's evidence, and one has warrant to assert that P, then one has warrant to assert that Q. Given our intuitions about lotteries,



it follows that one never has warrant to make assertions of probability less than 1 on one's evidence. This conclusion might appear to be a skeptical one, even a *reductio ad absurdum*. For it is easy to suppose that almost all our ordinary empirical assertions are of probability less than 1 (Edgington 1995, 287). But what kind of probability is in question? If it is objective probability, then the problem affects only assertions about the future, for only they have an objective probability other than 1 or 0. But objective probability is too objective to warrant assertion: of two past tense assertions whose objective probability is 1, I may have excellent evidence for one and none for the other. Equally, subjective probability (degree of belief) is too subjective to warrant assertion: I do not gain warrant to assert that I am Napoleon merely from my baseless conviction that I am Napoleon, even if my conviction is so dogmatic that the assertion has subjective probability 1. If any probability warrants assertion, it is probability *on one's evidence*.

What is one's evidence? A simple answer is available to the knowledge account: one's evidence is just what one knows. This answer was not assumed in the earlier discussion of evidence; for present purposes, it does not matter whether it elucidates the prior notion or stipulatively sharpens it, for either way the result is a tenable notion of evidence.<sup>8</sup> Without making any substantive assumptions about the conditions for knowledge, this view makes it trivial that if one knows that P, then the probability that P on one's evidence is 1. This does not imply that no discovery could shake one's confidence that P, for discoveries can undermine knowledge. Nor does it imply that one would in practice bet one's life against a penny that P; that test defines no useful notion of probability (substitute a moderately complicated tautology for 'P'). The standard of probability 1 on one's evidence is no more demanding than the standard of knowledge.

The denial of knowledge in the lottery case might also be feared to have skeptical implications, on the grounds that virtually all our

---

<sup>8</sup>See Williamson forthcoming for more on the conception of one's evidence as what one knows. On this view, "That P is part of *a*'s evidence for the proposition that Q" entails "*a* knows that P"; it does not entail "*a* knows that Q," for that Q is not in general part (or all) of the evidence for the proposition that Q. The view makes evidence propositional; non-propositional accounts of evidence find it hard to discount unknown features of the evidence in an appropriate way.

empirical knowledge has a probabilistic basis. For example, our perceptual processes are subject to random error. However, one must distinguish between causal and evidential senses of the word 'basis'. The causal connection between the environment and our perceptual beliefs about it is no doubt probabilistic, but it does not follow that those beliefs rest on probabilistic evidence. On the view above of evidence, when those beliefs constitute knowledge, they are part of our evidence. Moreover, they may constitute knowledge simply because perceiving that P counts as a way of knowing that P; that would fit the role of knowledge as evidence (Williamson 1995a). I certainly did not *perceive* that your ticket did not win. The matter deserves further exploration; it must suffice here to say that there is no obvious valid argument from the denial of knowledge in the lottery case to its denial in perceptual and other cases in which we ordinarily take ourselves to know. The knowledge rule provides a better explanation of the inadequacy of probabilistic grounds for assertion than do accounts on which something less than knowledge warrants assertion.

Conversational patterns confirm the knowledge account.<sup>9</sup> Consider a standard response to an assertion, the question "How do you know?" The question presupposes that it has an answer, that somehow you do know. If not only knowledge warrants assertion, what makes that presupposition legitimate? The question "Where did you read that?" is not normally appropriate in response to an assertion, because someone who asserts that P is not usually com-

---

<sup>9</sup>Much of the relevant evidence is marshaled in Unger 1975 (250–65) and Slote 1979. The thesis defended by Unger and Slote is that, in asserting that P, one *represents* oneself as knowing that P; see also DeRose 1991 (598–605). These authors say little about the general notion of representation, which this paper scarcely employs. The knowledge account subsumes the Unger-Slote thesis under more general principles. In doing anything for which authority is required (for example, issuing orders), one represents oneself as having the authority to do it. To have the (epistemic) authority to assert that P is to know that P. The Unger-Slote thesis follows. Lloyd Humberstone suggests that Max Black may have originated this talk of representing oneself: "In order to use the English language correctly, one has to learn that to pronounce the sentence 'Oysters are edible' in a certain tone of voice is to *represent oneself* as knowing, or believing, or at least not disbelieving what is being said. (To write a cheque is to represent oneself as having money in the bank to honour the cheque)" (Black 1952, 31). G. E. Moore makes a related claim: "by asserting *p* positively you *imply*, though you don't assert, that you know that *p*" (1962, 277; see also 1912, 125).

mitted to having read somewhere that P. But “How do you know?” is normally appropriate.<sup>10</sup> Of course, the questioner does not always *believe* the presupposition of the question, for it is sometimes (not always) intended as a challenge to the assertion. Nevertheless, it is an *implicit* challenge: the questioner politely grants that the asserter does know that P, and merely asks how, perhaps suspecting that there is no answer to the question. If not only knowledge warrants assertion, the absence of an answer would not imply the absence of a warrant; why should the question constitute even an implicit challenge? The hypothesis that only knowledge warrants assertion makes good sense of the phenomenon.

A less standard and more aggressive response to an assertion is the question “Do you know that?” Its aggressiveness is easy to understand on the hypothesis that only knowledge warrants assertion, for then what it calls into question is the asserter’s warrant for the assertion. On the hypothesis that not only knowledge warrants assertion, the aggressiveness of the question is hard to understand, for the asserter might truthfully answer “No” and still have warrant for the assertion.

A related argument starts from a version of Moore’s paradox, with ‘know’ in place of ‘believe’ (Moore 1962, 277; Unger 1975, 256–60; Jones 1991). Something is wrong with any assertion of the form “P and I do not know that P,” even though such assertions would often be true if made. What is wrong can easily be understood on the hypothesis that only knowledge warrants assertion. For then to have warrant to assert the conjunction “P and I do not know that P” is to know “P and I do not know that P.” But one cannot know “P and I do not know that P.” One knows the conjunction only if one knows each conjunct, and therefore knows that P (the first conjunct); yet one knows the conjunction only if it is true, so only if each conjunct is true, so only if one does not know that P (the second conjunct); thus the assumption that one knows “P and I do not know that P” yields a contradiction. Given that only knowledge warrants assertion, one therefore cannot have warrant to assert “P and I do not know that P.”<sup>11</sup> In contrast, the

---

<sup>10</sup>It is silly to ask “How do you know?” when the questioner obviously knows as well as the asserter how the latter knows—for example, when someone has said “I want to go home.”

<sup>11</sup>If knowing entails believing, a similar explanation reveals what is wrong with any assertion of the more familiar Moorean form “P and I do not believe that P.” See Sorensen 1988 (15–56) for a more general account

hypothesis that not only knowledge warrants assertion makes it hard to understand what is wrong with an assertion of that form. One often has good evidence that P while knowing that one does not know that P; in such cases one has good evidence short of knowledge for the conjunction "P and I do not know that P." If good evidence short of knowledge warranted assertion, one would have warrant to assert "P and I do not know that P"—but one has not.<sup>12</sup>

Knowledge is not even a cancelable implication of assertion (Slo-  
te 1979, 179). For if the implication could be canceled, the second  
conjunct, "I do not know that P," would cancel it, and it would be  
acceptable to assert the conjunction; but it is not acceptable.

One might fear that such arguments would prove too much. After all, something is wrong even with the assertion "P and I cannot be certain that P." Does that not suggest that only something *more* than knowledge warrants assertion? What seems to be at work here is a reluctance to allow the contextually set standards for knowledge and certainty to diverge. Many people are not very happy to say things like "She knew that P, but she could not be certain that P." However, we can to some extent effect such a separation, and then assertibility goes with knowledge, not with the highest possible standards of certainty. For example, one may have warrant to assert "P and by Descartes's standards I cannot be absolutely certain that P," where the reference to Descartes holds those standards apart from the present context. Again, it would often be inappropriate to respond to the assertion "P" by asking "How can you be so certain that P?" The word 'so' flags the invocation of unusually high standards of certainty. By ordinary standards you may have had warrant to assert that P even if you could not be *so* certain that P.<sup>13</sup>

---

of Moorean paradoxes along similar lines. The claims in the text apply only to utterances of the conjunction within a single context. If the contextual standards for knowledge are raised between the utterance of the first conjunct and that of the second, the assertion might be acceptable. Its unacceptability within a single context must still be explained.

<sup>12</sup>Hintikka's explanation of the defectiveness of such assertions by appeal to their epistemic indefensibility assumes that assertions ought to be epistemically defensible in his sense (1962, 78–102). That assumption is unmotivated unless only knowledge warrants assertion.

<sup>13</sup>These issues obviously impinge on the skeptical arguments that precede the account of assertion in Unger 1975; see Lewis 1996. However, the present account does not concede the truth of the skeptic's claims in the

Considerable evidence has emerged that our ordinary linguistic practice acknowledges the knowledge rule. Certain phenomena are nevertheless likely to be adduced as counterevidence. The next section considers such objections.

### 3. Objections to the knowledge account, and replies

That false beliefs are often reasonable is a commonplace. Evidence can be misleading. When one reasonably but falsely believes that P, is it not reasonable to *assert* that P, even though one does not know that P? If so, what becomes of the claim that only knowledge warrants assertion?

It may sometimes be reasonable to believe that P, even though one knows that one does not know that P. Perhaps it is reasonable for me to believe that I shall not be run over by a bus tomorrow, even though I know that I do not know that I shall not be run over by a bus tomorrow (Slote 1979, 180; I am not confined in a bed, lost in a jungle, etc.). Such cases do not threaten the hypothesis that only knowledge warrants assertion, for they are ones in which, intuitively, assertion is not warranted. It would be foolish of me baldly to assert that I shall not be knocked down by a bus tomorrow; it would invite the objection "You don't know that." As in the lottery case, I should assert no more than that it is very unlikely that I shall be knocked down by a bus tomorrow. Such cases support the hypothesis that only knowledge warrants assertion.

Sometimes one knows that one does not know that P, but the urgency of the situation requires one to assert that P anyway. I shout, "That is your train," knowing that I do not know that it is, because it probably is and you have only moments to catch it. Such cases do not show that the knowledge rule is not the rule of assertion; they merely show that it can be overridden by other norms not specific to assertion. The other norms do not give me warrant to assert that P, for to have such warrant is to satisfy the rule of assertion. Similarly, when I am speaking a foreign language, the urgency of the situation may require me to speak ungrammatically, because it would take me too long to work out the correct gram-

---

context in which they are made. The picture of knowledge in Williamson 1995a is very different from Lewis's.

mathematical form for what I want to say; it does not follow that my utterance satisfied the rules of grammar in that context.

In other cases, one reasonably believes that P and is in no position to know that one does not know that P. One cannot discriminate between one's actual circumstances and circumstances in which one would know that P. For example, it is winter, and it looks exactly as it would if there were snow outside, but in fact that white stuff is not snow but foam put there by a film crew of whose existence I have no idea. I do not know that there is snow outside, because there is no snow outside, but it is quite reasonable for me to believe not just that there is snow outside but that I know that there is; for me, it is to all appearances a banal case of perceptual knowledge. Surely it is then reasonable for me to *assert* that there is snow outside.

The case is quite consistent with the knowledge account. Indeed, if I am entitled to assume that knowledge warrants assertion, then, since it is reasonable for me to believe that I know that there is snow outside, it is reasonable for me to believe that I have warrant to assert that there is snow outside. If it is reasonable for me to believe that I have warrant to assert that there is snow outside, then, other things being equal, it is reasonable for me to assert that there is snow outside. Thus, the knowledge account can explain the reasonableness of the assertion. However, granted that it is reasonable for me to believe that I have warrant to assert that P, it does not follow that I do have warrant to assert that P. The term 'warrant' has been reserved for the condition that C(P) in the rule of assertion. There may be other evidential norms for assertion, if they can be derived from the knowledge rule and considerations not specific to assertion. The reasonableness of asserting that P when one reasonably believes that one knows that P has just been derived in exactly that way.

One can think of the knowledge rule as giving the condition on which a speaker has the *authority* to make an assertion. Thus, asserting that P without knowing that P is doing something without having the authority to do it, like giving someone a command without having the authority to do so. Characteristic standards of authority thus play a constitutive role in the speech act of assertion, as they do in other institutions. The distinction between having warrant to assert that P and reasonably believing oneself to have such warrant becomes a special case of the distinction between

having the authority to do something and reasonably believing oneself to have that authority. Someone who does not know that *P* lacks the authority to assert that *P*, and therefore cannot pass that authority on to me by asserting that *P*, no matter how plausibly he gives me the impression that he has done so.<sup>14</sup>

We are not always in a position to know whether we know that *P*. The knowledge account therefore implies that we are not always in a position to know whether we have warrant to assert that *P*. We are liable to error and ignorance about warrant, just as we are about almost everything else. This view of warranted assertibility is in sharp contrast with its treatment in antirealist theories of meaning to which the notion of the assertibility conditions of sentences is crucial. Such theories characteristically assume that one has no difficulty in knowing whether one has warrant to assert that *P*. Independently of the knowledge account, there is reason to doubt that there could be a norm of the kind postulated by antirealist theories.<sup>15</sup>

The knowledge account may seem to imply that speakers should

---

<sup>14</sup>The assimilation of warrant to authority is misleading in one respect. Authority, even intellectual authority, usually extends over an area; it is not confined to a single proposition. In the present sense, testimony can give one the authority to assert that *P*, even if one is pitifully ignorant about neighboring questions, and no extent of knowledge about neighboring questions can give one the authority to assert that *P* if one happens to be mistaken on that single point.

<sup>15</sup>Williamson 1995b argues independently of the knowledge account that Dummett's objections to truth-conditional theories of meaning also apply to assertibility-conditional theories unless the schema "If it is assertible that *P*, then it is assertible that it is assertible that *P*" holds, and that the schema fails for any reasonable notion of assertibility. Given the primacy of knowledge, the schema becomes the notorious KK principle, "If one knows that *P*, then one knows that one knows that *P*" (criticized in Williamson 1992). There are related difficulties for the account of assertion in Brandom 1983, where a central normative role is played by a notion of justification for which it is claimed that "a justification is whatever the community treats as one—whatever its members will let assertors get away with" (644) and knowledge is a matter of "an appropriately justified true belief" (647). The community is neither omniscient nor infallible in judging when true belief amounts to knowledge. The usual problems would attend an appeal to epistemically ideal counterfactual conditions. Brandom's emphasis on the ability to articulate inferential justifications of one's assertions in response to challenges (641–42) also seems to overestimate the link between warranted assertion and a smooth tongue. See Brandom 1994 for a much fuller account of his views.

always be at great pains to verify a proposition before asserting it. The wide variety of situations in which speakers go to no such pains may therefore seem to threaten the account: consider a lively seminar discussion, or gossip. To rule that speakers are not making genuine assertions in such situations would be to trivialize the account. In natural languages, the default use of declarative sentences is to make assertions, and the situations at issue are not special enough to cancel the default. Rather, the point is that the knowledge account does not imply that asserting that *P* without knowing that *P* is a terrible crime. We are often quite relaxed about breaches of the rules of the game we are playing. If the most flagrant and the most serious breaches are penalized, the rest may do little harm. It is said that the offside rule in rugby union is being breached most of the time. Similarly, many of the utterances in an ordinary conversation are syntactically ill-formed even by the standards of the speaker's own idiolect, for example, as a result of intentional or unintentional changes of direction in mid-sentence. Breaches of the rules are more serious in writing than in speech; that applies to the rule of assertion too. When assertions come cheap, it is not because the knowledge rule is no longer in force, but because violations of the rule have ceased to matter so much.

To be relaxed in applying a rule is not to replace it by a different rule. Even in a lively seminar discussion, or gossip, the knowledge rule does not give way to a rule of reasonable belief. For example, even in gossip, it would be cheating to assert "Mr. Jones won nothing in the lottery again last week" merely on the basis of its high probability. Similarly, if I overhear someone in a room full of logicians say, "A flaw has just been found in the proof of the main theorem in his last paper," when it is 99% probable that the person he is looking at is Professor X, I may form a reasonable belief that a flaw has just been found in the proof of the main theorem of Professor X's last paper. But even in a lively seminar discussion, it would be cheating for me to answer someone who bases an objection to my views on that theorem by asserting, without qualification, "A flaw has been found in the proof of that theorem." Such assertions are unacceptable because the speaker knows that he lacks the requisite knowledge, even though he has a reasonable belief. When we are relaxed in applying the rule, we feel entitled to assert that *P* whenever we are not confident that we do not know



that P. We still try to obey the knowledge rule, but we do not try very hard.

In debate, we are often willing to assert that P when we do not expect to persuade our interlocutors that P. However, knowing that P is quite consistent with being unable to persuade other people that P. Knowledge often depends on good judgment, the speaker may have better judgment than the hearer, and most speakers value their own judgment more highly than they know their hearers do.

Some people use the locution “I assert that . . .” only when they cannot supply compelling grounds; the implied contrast is with “I can prove that . . .” or the like. For the reason just given, they are not conceding that they do not know. The simplest analysis of what one does in uttering the syntactically declarative sentence “I assert that P” is that one asserts that P by asserting that one asserts that P—just as, in uttering “I promise to  $\phi$ ,” one promises to  $\phi$  by asserting that one promises to  $\phi$  (Lewis 1983, 224, from Lewis 1970; Ginet 1979). On that view, one obviously knows that one asserts that P, and therefore is warranted in asserting that one asserts that P. This may help to distract attention from the more problematic question, Is one warranted in asserting that P? One dodges that question by focusing one’s hearers’ attention on the less contentious assertion.

#### 4. The BK and RBK accounts

I reasonably believe that your lottery ticket did not win; I am not warranted in asserting that it did not win. I reasonably believe that I shall not be knocked down by a bus tomorrow; I am not warranted in asserting that I shall not be knocked down by a bus tomorrow. Neither belief nor reasonable belief warrants assertion. Nevertheless, it might still be thought that some false assertions are warranted in the technical sense that they obey the rule of assertion. One proposal along such lines is that the rule for assertion is this:

*The BK rule* One must: assert that P only if one believes that one knows that P.

(Thijssse forthcoming, citing similar views from Lenzen 1980) What one believes oneself to know need not be true. Can the BK account explain the phenomena?

The BK account can explain many of the conversational phenomena that were used as evidence for the knowledge account by adapting the latter's explanations to its own use. For example, I can follow the proof which shows that I cannot know "P and I do not know that P" and should therefore refrain from believing that I know "P and I do not know that P": if I do so refrain, then the assertion "P and I do not know that P" would violate the BK rule. Similarly, if I am committed to believing that I know by my assertion, then the challenge "How do you know?" has an obvious relevance.

One problem for the BK account is that my belief that I know that P may be as irrational as any other belief. The BK account's analysis of the modified Moorean sentence depends on the assumption that if "Q" is inconsistent then "I believe that Q" is inconsistent, which is invalid for subjects who are logically capable of irrationality. Suppose that I have an irrational belief that I know that G. E. Moore was a serial killer. On the BK account, my assertion "G. E. Moore was a serial killer" satisfies the rule of assertion. Neither its falsity nor its irrational basis constitutes a breach of the BK rule. So far, nothing is wrong with the assertion itself. Plenty is wrong with the asserter, for I have a completely irrational belief, but that is another matter. Although I have obeyed the BK rule only by expressing an irrational belief, the BK account lacks the resources to explain why that is a fault in the assertion itself. Defenders of the BK account cannot deny that we distinguish faults in the assertion from faults in the asserter. If I am asked, "Do you really have the belief that G. E. Moore was a serial killer?" (a question about me, not about G. E. Moore) then, in the circumstances, I ought to answer "yes," which is to assert that I have the belief that G. E. Moore was a serial killer; the assertion itself is quite in order, even though its being so depends on my irrational belief that G. E. Moore was a serial killer. The fault there is clearly in the asserter, not in the assertion. Since the BK account cannot explain why we regard the assertion "G. E. Moore was a serial killer," not just its asserter, as faulty, the account should be rejected. In contrast, the knowledge account has no difficulty in explaining what is wrong with the assertion, for the assertion breaks the knowledge rule.

On an obvious revision of the BK account, the rule for assertion is this:

*The RBK rule* One must: assert that P only if one rationally believes that one knows that P.

The added condition of rationality both improves the analysis of modified Moorean sentences and eliminates the counterintuitive consequences above. Nevertheless, all is not well with the RBK account. One problem concerns conjunctive assertions. Consider a complicated paradox, in which a contradiction is deduced from a very large number of premises  $P_1, \dots, P_n$ . For each  $i$ , it seems intuitively obvious that  $P_i$ ; indeed, it seems intuitively obvious that we know that  $P_i$ . Even on reflection, we are quite unsure which premise to blame for the contradiction. Suppose also that it is unlikely that more than one of the premises is false; each premise seems to have a quite different basis from the others, so that its falsity would be unlikely to infect them. Then we might easily, for each  $i$ , rationally believe ourselves to know that  $P_i$ . For each  $i$ , on the RBK account, we therefore have warrant to assert that  $P_i$ . Nevertheless, we know that the conjunction that  $P_1$  and  $\dots$  and  $P_n$  is false, because it entails a contradiction; thus, it is not rational to believe ourselves to know the conjunction, so, on the RBK account, we lack warrant to assert it. Warrant to assert would not be closed under conjunction. This consequence of the RBK account is disturbing, but not clearly absurd.<sup>16</sup>

The RBK account shares a simpler problem with the BK account: the analogue for falsity of the latter's problem about irrationality. Suppose that I rationally believe myself to know that there is snow outside; in fact, there is no snow outside. On the BK and RBK accounts, my assertion "There is snow outside" satisfies the rule of assertion. Yet something is wrong with my assertion; neither the BK nor the RBK account implies that there is. They can allow that something is wrong with my belief that I know that there is snow outside, for it is false, but that is another matter. The BK and RBK accounts lack the resources to explain why we regard the false assertion itself, not just the asserter, as faulty.

A further objection to the BK and RBK accounts, an obvious

---

<sup>16</sup>Arguably, the knowledge rule does not preserve warrant to assert under conjunction either, since one might know some propositions without having entertained their conjunction. However, this would be an insignificant failure of closure compared with the one in the text, which remains even when the subject has carefully considered the conjunction.

methodological one, is that they are less simple than the account based on the knowledge rule. Their adoption might be reasonable if the latter were refuted, but it has not been. The onus of proof is on the RK and RBK accounts. After all, when I assert that P, why should it matter whether I rationally believe myself to know that P if I am not required to know that P? Of course, the truth account is even simpler than the knowledge account, so the onus of proof is on the latter against the former; but it has already been discharged.

One possible motivation for belief-based accounts of assertion (hinted at by Thijsse) is the idea that what warrants assertion should be a mental state of the asserter. On a common view, believing and reasonably believing oneself to know that P are mental states, while knowing that P is not. However, it is arguable that knowing that P is a mental state, of an externalist kind (Williamson 1995a). Indeed, the combination of that idea with the idea that falsity is a fault in the assertion itself, so that what warrants asserting that P entails that P, implies that what warrants asserting that P is a mental state which entails that P. Knowing that P is the best candidate for such a state. Those issues cannot be pursued here. If a more internalist conception of mental states is adopted, the question becomes more pressing: why should what warrants assertion be a mental state of the asserter? One bad answer would be that one can always tell whether one is in a given mental state. One cannot. It may be hard to tell whether one's confidence that one knows that P is high enough for one to count as believing oneself to know that P, and even harder to tell whether it is rational enough for one to count as rationally believing oneself to know that P. There is no good reason to accept a belief-based account of assertion. Indeed, our attitude to false assertions is misrepresented by any simple account on which what warrants assertion does not entail truth.

## 5. Mathematical assertions

The rule of assertion is easier to identify in more formal situations, of which mathematics provides some of the best examples. Assertibility in mathematics has the additional interest that attempts to construct assertibility-conditional theories of meaning have taken

the intuitionistic proof-conditional account of mathematics as a paradigm. Assertion in mathematics will therefore be considered. The mathematical case is, it will be argued, more representative than has often been supposed.

In mathematics, the distinction between warranted and unwarranted assertions is striking. Count the propositions that are axiomatic for working mathematicians as having one-line proofs. Then, to a first approximation, in mathematics one has warrant to assert that  $P$  if and only if one has a proof that  $P$ . On the knowledge account, that is so because, to a first approximation, in mathematics one knows that  $P$  if and only if one has a proof that  $P$ . One has a proof that  $P$  when one has followed such a proof and retains some memory of it, in particular of its conclusion. Those are just first approximations, but where having warrant to assert that  $P$  diverges from having a proof that  $P$ , so does knowing that  $P$ . Conversely, where knowing that  $P$  diverges from having a proof that  $P$ , so does having warrant to assert that  $P$ . Having warrant to assert that  $P$  and knowing that  $P$  do not diverge from each other; the knowledge account is confirmed.

The word 'proof' has just been used in the informal sense common in ordinary mathematics, in which only truths have proofs; a working mathematician who says that it has been proved that  $P$  does not leave it open whether  $P$ . This notion is not relativized to an arbitrary formal system; if it were, the connection with (unrelativized) assertibility would be lost. The axioms that have one-line proofs do so in virtue of their status in the practice of mathematics, not in virtue of their place in a particular formal system. 'Proof' will be used in this informal sense below.

Consider, first, putative cases in which one has warrant to make a mathematical assertion but lacks a proof. In the simplest cases, one knows by testimony that there is a proof that  $P$ : but then one knows by testimony that  $P$ , and thereby satisfies the knowledge rule. In rarer cases, the nondeductive evidence for a mathematical proposition may be strong enough to warrant its assertion (Steiner 1975, 93–108). Nevertheless, this biconditional remains plausible: the evidence is strong enough to warrant the assertion that  $P$  if and only if it is strong enough to provide knowledge that  $P$ . What the knowledge account will not grant is that one can have warrant to assert that  $P$  without a proof that  $P$  by having grounds for a mistaken belief that one has a proof that  $P$  (or that there is such

a proof). When one has such a belief, on the knowledge account, one at best mistakenly believes that one has warrant to assert that P. Even if expert mathematicians play a practical joke and inform you falsely that it has been proved that P, you do not really acquire warrant to assert that P (if you did, the joke would be still less funny).<sup>17</sup> You acquire only misleading evidence that you have such warrant. Although your belief that you have it is reasonable, that does not make it true. The reasonableness in question can be explained as derivative from the knowledge rule; you reasonably believe yourself to know that P, so you have reason to believe that you have warrant to assert that P. This view of the matter is independently defensible. Testimony is a special source of warrant because one speaker can *pass on* a warrant to another. Since the expert mathematicians have no warrant to assert that P themselves, they have none to pass on to you.

Now consider putative cases in which one has a proof of a mathematical proposition but lacks warrant to assert the proposition. The possibility of such cases is sometimes denied, on the Cartesian grounds that genuine proofs are transparent to the subject. That denial does little justice to the complexity of many actual proofs. It can take months of effort by the mathematical community to decide whether a purported proof is genuine. When I have a genuine proof, expert mathematicians may tell me falsely that it contains a fallacy. They may give me a complicated explanation of the supposed fallacy, blinding me with science. I may recall other occasions on which what I believed for broadly similar reasons to be a proof really did turn out to be fallacious. In such cases, it would be unreasonable for me to assert that P, for it is unreasonable for me to believe that I have warrant to assert that P. It does not immediately follow that I have no warrant to assert that P. One may have the authority to do something even when it is unreasonable for one to believe that one has that authority. What the knowledge account implies is just that I cease to have warrant to assert if and only if I cease to know. That biconditional remains plausible. If I know that P, I thereby have warrant to assert that P. Conversely, there is no reason to expect my possession of a proof that P to give

---

<sup>17</sup>See Kitcher 1983, especially 55–56 and 89–91, for discussion of the social character of mathematical knowledge. Whether this character permits it to be *a priori* is another matter.

me warrant to assert that P independently of giving me knowledge that P. The more plausible of the two ways for the biconditional to hold is for me to lose both knowledge and warrant to assert: the appearance of ignorance undermines knowledge in a way in which the appearance of knowledge does not undermine ignorance. But even if I retain both knowledge and warrant to assert, the knowledge account stands.

The remaining cases are those in which one has a proof that P if and only if one has warrant to assert that P. They are still less likely to threaten the proposed equivalence between knowing that P and having warrant to assert that P.

How untypical are mathematical assertions? Proofs are often supposed to warrant them in a way inapplicable to most or all empirical assertions. Proofs, it is said, are conclusive, while empirical warrants are not. However, the nature of the contrast is unclear. No doubt new information cannot make a proof into a non-proof. But the issue is not whether proofs continue to be proofs; it is whether they continue to warrant assertion. Define a way of having warrant to assert that P to be *defeasible* just in case one can have warrant to assert that P in that way and then cease to have warrant to assert that P merely in virtue of gaining new evidence. A way of having warrant to assert that P is *indefeasible* just in case it is not defeasible. Most ways of having warrant to make empirical assertions are defeasible, but the considerations above about the social character of mathematical knowledge suggest that even grasping a proof of a mathematical proposition is a defeasible way of having warrant to assert it. One can have warrant to assert a mathematical proposition by grasping a proof of it, and then cease to have warrant to assert it merely in virtue of gaining new evidence about expert mathematicians' utterances, without forgetting anything. If so, mathematical propositions do not differ from empirical ones in point of defeasibility.

The notion of indefeasibility should not be confused with that of factiveness. A way of having warrant to assert that P is *factive* just in case a necessary condition of having warrant to assert that P in that way is that P. Grasping a proof of a mathematical proposition is a *factive* way of having warrant to assert it: a necessary condition of grasping a proof that P is that P. Factiveness does not entail indefeasibility. Knowing that P is always a *factive* way of having warrant to assert that P; it is almost never an indefeasible way.

New evidence can almost always undermine old knowledge; perhaps there is an indefeasible way of having warrant to assert that one exists. On the knowledge account, any way of having warrant to assert something is factive.<sup>18</sup> Thus, mathematical propositions do not differ from empirical ones in point of factiveness either.<sup>19</sup>

On the showing of this section, mathematical practice is consonant with the knowledge account in a way that generalizes smoothly to practice outside mathematics.

## 6. Why assert?

It is pointless to ask why the knowledge rule is the rule of assertion. It could not have been otherwise; a constitutive rule is not a convention.<sup>20</sup> It is, however, pointful to ask why we have such a speech

---

<sup>18</sup>By itself, indefeasibility does not entail factiveness. If warrant to assert that P consisted merely in good reason to believe that P, then the inhabitants of a universe created six thousand years ago with every appearance of having existed for millions of years might have an indefeasible nonfactive warrant to assert that they are not inhabitants of a universe created six thousand years ago with every appearance of having existed for millions of years. The account defended in this paper guarantees factiveness independently of indefeasibility.

<sup>19</sup>The distinction between indefeasibility and factiveness is blurred in this passage from Dummett 1973:

When we first learn language, we are taught to make assertions only in the most favoured case, namely in that situation in which the speaker can recognize the statement as being true. One way to observe the convention to try to utter only true assertoric sentences would of course be to utter them only in this most favoured situation; but for the great majority of forms of statement that is not at all what we do. Some forms of statement—those in the future tense, for example—are never uttered in the situation which conclusively establishes their truth; and others, which we originally learned to utter only in such situations, we later learn to utter in circumstances in which we may turn out to have been mistaken. (355)

It seems to be assumed that (a) recognizing a truth requires conclusively establishing it and (b) ordinary standards permit us to assert propositions whose truth has not been conclusively established. If 'conclusively' means indefeasibly, (a) is false. If 'conclusively' means factively, we have no reason to accept (b). See also McDowell 1982 and Wright 1993 on criteria and defeasibility.

<sup>20</sup>The knowledge rule may be linked to conventions in a different way. Suppose that a language  $\mathcal{L}$  assigns to each sentence type  $s$  in some domain a proposition  $\mathcal{L}(s)$ . Then it might be a convention in a particular community that in normal contexts one should utter  $s$  only if one knows  $\mathcal{L}(s)$ . Such a convention of *knowledgeableness* in  $\mathcal{L}$  might even be part of what it



act as assertion in our repertoire. Could we not have done otherwise? No doubt we need a speech act something like assertion, to communicate beliefs, but could we not have done so just as well by using a speech act whose rule demanded less than knowledge? It would have to permit testimony and inference to enable us to utter new instances on the basis of old ones, just as they do for assertion. But the knowledge rule is not the only rule to underwrite that possibility; the truth rule is another.<sup>21</sup>

One obvious answer is that we need assertion to transmit knowledge.<sup>22</sup> In normal circumstances, a speaker who asserts that P thereby puts a hearer in a position to know that P if (and only if) the speaker knows that P.<sup>23</sup> That answer is probably right, as far as it goes, but leaves at least two points to be explained. First, why could we not transmit knowledge by means of a speech act whose rule required only (for example) truth? The idea might be that when successful communication occurs, what is transmitted is what is *overt* in the assertion that P, and what is overt in the assertion is satisfaction of the rule, so what is transmitted is satisfaction of the

---

is for £ to be the language of that community. This convention in £ is an obvious variant on the convention of truthfulness in £ used by David Lewis to define what it is for £ to be the language of a community: "To be truthful in £ is to act in a certain way: to try never to utter any sentences of £ that are not true in £. Thus it is to avoid uttering any sentence of £ unless one believes it to be true in £" (1983, 167, from Lewis 1970). Of course, the account of what it is for £ to be the language of a community must somehow take into account the probability that speech acts other than assertion can be performed in £ (Lewis 1983, 172). The shift from conventions of truthfulness to conventions of knowledgeableness also has repercussions in the methodology of interpretation. The appropriate principle of charity will give high marks to interpretations on which speakers tend to assert what they know, rather than to those on which they tend to assert what is true, or even what is reasonable for them to believe.

<sup>21</sup>Brandom 1983 gives an interesting account of the role of testimony and inference in the social practice of assertion; however, it is based on a notion of justification too generic to distinguish assertion from the other speech acts envisaged in the text. Craig 1990 attempts to explain the concept of knowledge in terms of the need for such a practice, but the knowledge account is not an immediate consequence.

<sup>22</sup>"[T]he essential character of the assertoric use of language lies in its availability for communicating, in the sense of transmitting knowledge about the subject matter of assertions" (McDowell 1980, 128; compare Evans 1982, 310).

<sup>23</sup>'Normal circumstances' adverts to the hearer's knowledge that the speaker asserted that P, the speaker's lack of a reputation for unreliability, and so on.

rule; thus, knowledge is transmitted if and only if it is what the rule requires. However, the relevant notion of overttness is hard to pin down. If the overttness of the satisfaction of the rule put the hearer in a position to know that the rule was satisfied, then if the rule required truth, the hearer would be in a position to know that the assertion was true, and the truth rule would suffice for the transmission of knowledge after all. Second, it will be asked why the transmission of knowledge is what matters, rather than the transmission of true belief, or reasonable belief, or some other cognitive attitude.

A comparison between knowing and doing reveals a clue to a further line of thought. One may think of knowing that *P* as standing to believing that *P* as (intentionally) bringing it about that *P* stands to desiring that *P*. If one knows that *P*, then *P*; likewise, if one brings it about that *P*, then *P*. But even if *P* and one believed that *P*, it does not follow that one knew that *P*; likewise, even if *P* and one desired that *P*, it does not follow that one brought it about that *P*. In each case, the fit between content and world is insufficient because it may have been "accidental." Both knowing that *P* and bringing it about that *P* are ways of ensuring that *P*; what differs is the direction of fit. If one brings it about that *P*, one's actions (characterized in environment-dependent ways) ensure that *P*; likewise, if one knows that *P*, one's mental states (characterized in environment-dependent ways) ensure that *P* (Williamson 1995a).

Obedience to a command, as ordinarily understood, involves bringing something about; what matters is not simply the fit between content and world, but someone's responsibility for that fit. To issue a command with appropriate authority is to confer a responsibility; to obey a command is to discharge that responsibility.<sup>24</sup> The point emerges more distinctly for negative commands, where what is commanded is not itself an intentional action. You shout, "Don't move!" I try to move, but find myself stricken by paralysis.

---

<sup>24</sup>Obedience to a command may also require an appropriate causal connection between the command and the action; if I do not hear the command to halt, but halt anyway, there is a sense in which I have not obeyed the command. Similarly, if I know that *P* and assert that *P*, but the asserting is causally independent of the knowing, then something is wrong with the asserting. Perhaps what is wrong can be derived from the knowledge account.

In one sense I did not obey your command. Although its content was fulfilled, I did not ensure that it was; I did not bring it about. The knowledge account extends the analogy between commanding and asserting. To 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. Our possession of such speech acts is no more surprising than the fact that we have a use for relations of responsibility.

*University of Edinburgh*

### References

- Armstrong, D.M. 1973. *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.
- Black, M. 1952. "Saying and Disbelieving." *Analysis* 13:25–33.
- BonJour, L. 1985. *The Structure of Empirical Knowledge*. Cambridge: Harvard University Press.
- Brandom, R. B. 1983. "Asserting." *Noûs* 17:637–50.
- . 1994. *Making It Explicit*. Cambridge: Harvard University Press.
- Craig, E. J. 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Oxford University Press, Clarendon Press.
- DeRose, K. 1991. "Epistemic Possibilities." *Philosophical Review* 100:581–605.
- . 1996. "Knowledge, Assertion and Lotteries." *Australasian Journal of Philosophy* 74:568–80.
- Dewey, J. 1938. *Logic: The Theory of Inquiry*. New York: Henry Holt.
- Dretske, F. I. 1981. *Knowledge and the Flow of Information*. Oxford: Blackwell.
- Dudman, V. 1992. "Probability and Assertion." *Analysis* 52:204–11.
- Dummett, M. A. E. 1973. *Frege: Philosophy of Language*. London: Duckworth.
- . 1991. *The Logical Basis of Metaphysics*. London: Duckworth.
- Edgington, D. 1995. "On Conditionals." *Mind* 104:235–329.
- Evans, M. G. J. 1982. *The Varieties of Reference*. Oxford: Oxford University Press, Clarendon Press.
- Gazdar, G. 1979. *Pragmatics: Implicature, Presupposition, and Logical Form*. New York: Academic Press.
- Ginet, Carl. 1979. "Performativity." *Linguistics and Philosophy* 3:245–65.
- Grice, H. P. 1989. *Studies in the Way of Words*. Cambridge: Harvard University Press.
- Harman, G. 1968. "Knowledge, Inference and Explanation." *American Philosophical Quarterly* 5:164–73.

- Hintikka, J. 1962. *Knowledge and Belief*. Ithaca: Cornell University Press.
- Jones, O. R. 1991. "Moore's Paradox, Assertion and Knowledge." *Analysis* 51:183–86.
- Kitcher, P. 1983. *The Nature of Mathematical Knowledge*. Oxford: Oxford University Press.
- Lenzen, W. 1980. *Glauben, Wissen und Wahrscheinlichkeit*. Vienna: Springer-Verlag.
- Lewis, D. K. 1970. "General Semantics." *Synthese* 22:18–67. Reprinted in Lewis 1983.
- . 1975. "Languages and Language." In *Minnesota Studies in the Philosophy of Science*, vol. 7, ed. K. Gunderson. Minneapolis: University of Minnesota Press. Reprinted in Lewis 1983.
- . 1983. *Philosophical Papers*, vol. 1. Oxford: Oxford University Press.
- . 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74: 549–67.
- Lowe, E. J. 1995. "The Truth about Counterfactuals." *Philosophical Quarterly* 45:41–59.
- McDowell, J. H. 1980. "Meaning, Communication and Knowledge." In *Philosophical Subjects: Essays Presented to P. F. Strawson*, ed. Z. van Straaten. Oxford: Oxford University Press, Clarendon Press.
- . 1982. "Criteria, Defeasibility, and Knowledge." *Proceedings of the British Academy* 68:455–79. Partly reprinted with revisions in *Perceptual Knowledge*, ed. J. Dancy. Oxford: Oxford University Press, 1988.
- Moore, G. E. 1912. *Ethics*. London: Thornton Butterworth.
- . 1962. *Commonplace Book: 1919–1953*. London: Allen and Unwin.
- Slote, M. A. 1979. "Assertion and Belief." In *Papers on Language and Logic*, ed. J. Dancy. Keele: Keele University Library.
- Sorensen, R. A. 1988. *Blindspots*. Oxford: Oxford University Press, Clarendon Press.
- Steiner, M. 1975. *Mathematical Knowledge*. Ithaca: Cornell University Press.
- Thijsse, E. G. C. Forthcoming. "The doxastic-epistemic force of declarative utterances." In *Abduction, Beliefs and Context: Studies in Computational Pragmatics*, ed. W. J. Black and H. C. Bunt. London: University College Press.
- Unger, P. 1975. *Ignorance: A Case for Scepticism*. Oxford: Oxford University Press.
- Williamson, T. 1992. "Inexact Knowledge." *Mind* 101:217–42.
- . 1995a. "Is Knowing a State of Mind?" *Mind* 104:533–65.
- . 1995b. "Does Assertibility Satisfy the S4 Axiom?" *Critica* 27:3–22.
- . Forthcoming. "Knowledge as Evidence." *Mind*.
- Wright, C. J. G. 1993. *Realism, Meaning and Truth*. 2d ed. Oxford: Blackwell.