Assertion, Knowledge, and Lotteries

One of the central claims of Williamson's ground-breaking epistemology is the claim that knowledge is the norm of assertion. I believe this viewpoint is mistaken, and will argue that here. I will first explain Williamson's path to the conclusion he holds, identifying the two major arguments that he uses to support his claim that knowledge is the norm of assertion. I will then summarize the prima facie case for an alternative view,[1] following which I will address the tension between this prima facie case and Willamson's arguments. I will argue that a proper resolution of the conflict results in a denial of the idea that knowledge is the norm of assertion. Instead, I will maintain that to the extent that appropriate assertion is subject to epistemic constraints, those constraints have to do with justification rather than knowledge.

One note before beginning. Williamson's discussion of the norm or norms of assertion assumes that such norms are indefeasible–that one always does something wrong when one asserts something one does not know. I see no reason for such a strong reading of the norm, though I do understand how one might get to this idea. If we think of sports games, such as baseball or basketball, we find rules or norms governing the game that must be interpreted as indefeasible. If the rules were defeasible, then we'd expect that, e.g., if one's life depended on getting away with a foul in the last seconds of a close game, the referee should overlook the infraction. In one sense, this can happen, for in an emergency, the rules are violated by abandoning the game. But if we are playing the game, the rules cannot be overridden by factors outside the rule book. Rules of a game of this sort are quite different from ordinary norms, such as norms of reasoning, social norms, norms of etiquette, aesthetic norms, or norms of polite discourse. Treating the norms of assertion as indefeasible requires an argument that, as

Williamson puts it, there are constitutive norms for assertion akin to rules of a game. Williamson gives no argument that there are any such norms for assertion, but only suggests that we proceed on the assumption that there are such norms to see where it leads.[2]

This assumption has two negative consequences for Williamson's view. The first consequence is that it makes the norm or norms of assertion quite different from the norms just mentioned that govern politeness, etiquette, art, reasoning, and the like. In general, the norms that govern human behavior are defeasible, so it would be surprising to find the norms governing assertion to be indefeasible.

The second consequence of Williamson's assumption is that it makes it more difficult to defend the view that knowledge is the norm of assertion. There is a variety of cases in which the norm is violated and yet we do not view the assertion as deserving of criticism. Skeptics claim not to know anything, and if we consider the possibility that they are correct, we need not demand that they speak no more. Eliminative materialists say that there are no beliefs, and they don't have to flout the norms of assertion to discuss this philosophical possibility, even if it should turn out that they are right. William James counsels nontheists who are convinced by his pragmatic argument to go to Mass and hope for the best. When they do, and engage in the practice of espousing theological claims that they don't (yet) believe, we may find humor in the situation, but there is nothing wrong with the attempt (barring some unforeseen argument that theistic belief and practice is always and everywhere irrational). Philosophers espouse positive views in all of its subfields that, on reflection, they agree that they don't know to be true; but they continue the practice unrepentant. Visualization techniques by sports psychologists often involve rehearsing a positive outcome regarding what is going to happen next in a competitive situation,

and such rehearsal often involves thoroughly unwarranted assertions, such as "He's going to hang a curveball and I'm going to turn on it and drive it right down the left field line." Such techniques are often employed by aspiring academicians: "I'm going to send this paper to the best journal in my field and they'll accept it and hundreds, no, *thousands*, will read it; it will change the landscape of my discipline and I'll be famous!" (For those in this situation, don't be overly concerned about the example: just apply the other part of the technique, which is to ignore information that contributes to a defeatist mentality...)

In other cases, we assert things we don't believe because of a social role we inhabit. For example, a judge may give instructions to the jury that he doesn't believe are correct instructions. But his role requires that he give these instructions, and in order to comply, he asserts them without believing them. In other cases, a teacher is required by law to teach certain material. Suppose, for example, that a teacher is required to teach certain sex education material to high school students. One might do so to comply with the law even when one doesn't believe that what one is teaching is true. Note that in some such cases, the material will be false, and yet teaching it will not be criticizable. In other cases, the material will be true, but will be disbelieved by a teacher with, let us say, "outlier" beliefs. Perhaps the teacher is a thorough-going fatalist who thinks that precautions are a waste of time, or perhaps has an unusual theory about reproduction and sexually transmitted diseases.

I have belabored this point in order to prevent the impression that stopgap measures might be useful in sustaining the view that the norms of assertion are indefeasible. For example, one might attempt to treat assertion not backed by belief as a kind of theater. In a play, actors vocalize sentences and do not believe what the sentences express, and perhaps apparent assertion

not backed by belief could be modeled on this practice. I'm not impressed by this treatment of the case. In the case of theatrics, we can pull the actor out of the context and ask whether they are asserting what their sentences express, and they'll say "no". That's an easy test to apply when a question arises about whether a person is asserting anything, and applied to the above cases, the test won't turn out well for the view that assertions are not being made. Perhaps the speakers will also express some reservations if we remove them from the scripted situation, but the reservations will be about the truth of what they are asserting, not about whether they are asserting the claims in question. Regardless of the plausibility of specific stopgap measures, however, the examples above display how widespread are the examples in which assertions not backed by knowledge occur without deserving criticism. The simplest explanation of this data is that the norms of assertion are defeasible. This explanation has a further advantage as well over any collection of stopgap measures aimed at avoiding this explanation. This further advantage is that it construes the norms of assertion the same way that is appropriate for treating the other kinds of norms listed earlier: social norms, rules of etiquette, norms of politeness, aesthetics norms, and norms of reasoning, to name a few.

I raise this issue to put it aside in what follows. I will assume that the norms of assertion are defeasible in what follows, and making this assumption has the advantage that it makes it easier to defend Williamson's view from the above examples. If the norms of assertion are defeasible, then we have the option of explaining all of the above cases in terms of factors that defeat the claim that you shouldn't say what you don't know to be true. The fundamental issues that separate Williamson and me on the relationship between assertion and epistemology are not here, however, so I will not require of the view that knowledge is the norm of assertion that it be

an indefeasible norm.

The Knowledge Account

Williamson's discussion proceeds as follows. He begins with the working hypothesis that assertion has a constitutive rule. He grants that this hypothesis is by no means obvious, but suggests that pursuing the idea to see where it leads may prove fruitful. He then proposes two main accounts of a constitutive rule for assertion: the truth account and the knowledge account. According to the truth account, one has a warrant for asserting *p* only if *p* is true; according to the knowledge account, one has warrant for asserting *p* only if one knows *p*. The notion of warrant here is stipulatively related to that of the idea of a constitutive rule: to have warrant for asserting a claim is simply to have satisfied the constitutive rule for assertion.

After proposing these two accounts, Williamson first argues against the truth account and in favor of the knowledge account. The remainder of his defense of the knowledge account consists in rebutting attempts to weaken the knowledge account so as allow there to be warrant for asserting something false. Perhaps, for example, one only needs to believe, or rationally believe, that one knows in order to have such warrant.

The primary argument Williamson employs against the truth account involves lottery claims, such as the claim that a given ticket will lose. According to Williamson, there is something wrong with the assertion that one's ticket will lose, even given the truth of the claim. Williamson notes a natural response we have to such an assertion, which is to chastise the assertor by saying, "But you don't know that!"

This chastisement leads to the primary arguments Williamson offers on behalf of the knowledge account. Williamson first notes the conversational propriety of questions such as "How do you know that?" and "Do you know that?" Such questions presuppose the knowledge account, according to Williamson, for if the knowledge account were false, the questions would appear to be irrelevant.

This line of argument may prove too much, however. It is also appropriate to ask, "Are you certain?", "Are you absolutely sure?" If the conversational propriety of various questions is an argument in favor of the knowledge account, the propriety of these questions is an argument in favor of a stronger account: that one must be absolutely certain in order for a claim to be assertible. Moreover, it is also appropriate to ask of an assertor, "Do you have any good reason to think that?" In fact, when asking "How do you know?", the answer we are looking for is one that cites the reasons that person has for thinking that the claim in question is true, and when we get such an answer, we are satisfied. If conversational propriety of questions were our only data, we'd have no good reason to single out the knowledge account as opposed to some weaker or stronger position.

It may be that once we have a good argument for a position about the norm or norms of assertion, we will be in a position to return to this issue of appropriate questions and explain why one of the ways of questioning is primary and the others derivative. My point here is not to show that there is no way to explain away data not in accord with one's preferred theory. My point is only that the data about conversationally appropriate questions doesn't settle the matter as to the precise nature of the norm or norms of assertion.

Williamson's other major argument on behalf of the knowledge account is Moorean in

character, appealing to the defective character of the assertion "*p* but I don't know that *p*." Williamson notes the viewpoint widely shared in recent epistemology that to assert *p* is to represent oneself as knowing *p*. Williamson cites as primary exponents of the view G.E. Moore, Michael Slote, Max Black, and Peter Unger; put in different order, we can honor these names by referring to the representational claim they all endorse as the BUMS view.[3] Williamson claims that a virtue of the knowledge account of assertion is that it offers a general explanation of the truth of the BUMS view: since the norms of a practice are presupposed by those involved in the practice, hearers expect speakers to know and speakers expect to be taken to know, thereby creating a situation in which an assertion represents the speaker as having knowledge. .

Both the BUMS view and the knowledge account of assertion provide a basis for explaining the impropriety of asserting "*p* but I don't know *p*." On the BUMS view, the second conjunct contradicts the representation made by asserting the first conjunct, making incoherent the entire assertion. The knowledge account implies something stronger. It implies that there can be no propriety for the assertion, since the second conjunct directly contradicts the condition required for the propriety of asserting the first conjunct.

This argument relies on the intuitive impropriety of asserting "*p* and I don't know *p*," and this reliance raises a version of the problem noted earlier, for it is equally inappropriate to say "*p* but I'm not certain whether *p*." Williamson's answer is to insist that the standards for knowledge and certainty in a context typically are the same, and when they diverge, the knowledge account predicts the correct response. Thus, he holds, if one asserts "*p* and I lack Cartesian metaphysical certainty that *p*," no impropriety results.

The Justification Account

There are two other prominent alternatives to the knowledge account, one of which I will defend here. The two accounts are the belief account and the justification account. On the first approach, the only requirement on assertion is sincerity: say what you think.[4] This view will strike most everyone as excessively weak, but there is a way to make it a bit more plausible. We might say that belief is the fundamental norm of assertion, and that among the norms for belief are the other things that lead one to think that the sincerity requirement alone is too weak. So, for example, if one thinks that good reasons are required for assertion, that would be because it is a norm of belief that you shouldn't believe things without good reasons. Assertion thereby inherits this derivative norm because of the sincerity requirement.

A full defense of the belief account of assertion would thus involve an account of the types of activity that are norm-governed and a relegation of each norm to its natural home or homes. I doubt that any such regimentation can be accomplished, however. More carefully, I'm not sure that any one regimentation will be defensible as the best one, nor that it will be better than a view that refuses the project. Perhaps it can be done, but it is hard to see why it needs to be done, or what the value of it would be, except in service of the sincerity requirement itself. In what follows, then, I propose to ignore the issue of which norms are fundamental to a practice and which are derivative, and focus instead only on the norms of whatever level that govern the practice of assertion.

The other alternative to the knowledge and truth accounts is the justification account. According to this view, the propriety of an assertion is a function of one's justification for the

content of the assertion rather than a function of whether one knows the content to be true. The primary argument for the justification view arises from two factors. The first factor is that there are four distinctive necessary conditions for knowledge, conditions described by the defeasibility theory of knowledge. According to this theory, where the notion of defeat is properly clarified, knowledge is (ultimately) undefeated justified true belief. This assumption may seem to conflict with Williamson's aversion to the idea that knowledge can be analyzed,[5] but in fact it does not. My assumption is only that the four distinctive conditions of the defeasibility theory are necessary conditions for knowledge, not that they constitute an analysis of knowledge. It may be that these four conditions are only some of the necessary conditions for knowledge, or it may be that these four conditions specify fully the nature of knowledge. Either view is compatible with a denial that knowledge can be analyzed. Like Williamson, I have little interest in the project of analysis, though I do have an interest in the nature of knowledge. If we give a suitable construal of each claim, it is correct to say that knowledge requires belief, truth, and a justification that is ultimately undefeated. The most controversial of these assumptions is the assumption that justification is required for knowledge, but the objections to this view generally begin by assuming that justification has some special character–that it is a deontological notion, that it requires voluntary control over belief, that it requires that one be able to construct an argument for each justified belief, etc. In my view, justification is a function of the evidence one possesses, including mental states such as beliefs and experiences (though, of course, the assumption is not that all mental states are part of one's evidence). No account of knowledge can succeed without making justification, construed in this broad and general way, a necessary condition for knowledge.

The second central feature of the prima facie case for the justification view, beside the claim about these four necessary conditions for knowledge, concerns our attitudes when our assertions are corrected. In some cases of correction, we *take back the content of our speech act*, and in other cases *we apologize for, and regret, the very act itself*. For example, if we assert a claim and then are shown that the claim is false, we take back the content of our speech act, but we needn't apologize for or regret the very act itself. Randy says, "I've studied music all my life; there's no piece of group music even moderately well-known in the U.S. where part of the group is playing in 15/16 time and another part in 17/16 time," to which Michael responds, "That's certainly a reasonable judgment, except that you don't know enough about King Crimson. They are moderately well-known, and they have just such a piece." Michael then shows Randy the piece (so, I'm assuming that Michael is correct), to which Randy says, "I was wrong, I take it back." Now Randy may regret his assertion if he is the sort of person who strongly dislikes confronting his own fallibility. He may even vow to be much more careful not to say anything at all when he risks being wrong in order not to repeat this embarrassing moment, though such a response is surely overblown. Chagrin is normal, even mild embarrassment, but apologizing would be unctuous and overwrought. As I told the story, Randy responds appropriately. He doesn't apologize for making the assertion, but what he does instead is take back the content of the assertion. In fact, were he to apologize, the natural response would be dismissive: "Give it a rest, nobody's always right..."

In other cases, such remonstrance for apology would not be forthcoming. If we assert something that we do not in fact believe, we naturally apologize for the speech act itself when our assertion is challenged, and the apology is fully appropriate. Joe says, "I'm terrible at

baseball, I can't hit a curve ball to save my life!", to which Jared responds, "Don't say that! You're just feeling bad because you're in a hitting slump." Joe's too down on himself to respond at the time, but by the next game, he's feeling better, and says, "Jared, you're right, I shouldn't have said that–I was just discouraged after striking out again." Here Joe is apologizing for and regretting the speech act itself, not just taking back its content. The apology is perfectly in order, unlike the case above where the only fault concerns the content of the assertion.

This same distinction plays out with the other two conditions as well. Suppose you've been gettiered, and you assert, "Someone in this room owns a Ford." When it is pointed out that you've been duped by Nogot, you will take back the content of your assertion. You will also experience chagrin at being duped–it is not high on the usual lists of enjoyable experiences–but no apology for the speech act is in order. Practical jokes play off of this very fact. For those who find humor in such, the humor is in getting a person to respond in a fully appropriate way, even though the circumstances are not what the person reasonably judges them to be. In such a case, only a person suffering from a "private, preponderant horror at being duped"[6] would apologize for the action, and anyone who would so apologize deserves remonstrance for that action. The appropriate response is to take back the content of what is said, not apologize for the saying of it.

Things are different when you don't have justification for what you say. Suppose Billy Bob sees a headline on a tabloid in the checkout line at Walmart that reads "George Bush–A closet communist!" Billy Bob is a Texas Democrat and will believe almost anything negative about Bush, he despises him so much. So he believes that Bush is a communist, and upon seeing his best friend Bobbie Sue, says, "Hey, ja'hear Dubbya's a commie!" Bobbie Sue doesn't like Bush either, but she's skeptical: "Naw, yer jokin'," she says. "Naw I ain't–I saw't in da paper

over at the Wawl-Mort!" "Billy Bob, them's *tabloids*! You don't trust *tabloids*!" Billy Bob, donning the accent of what those outside of Texas take as a mark of intelligence, reflects and says, "Yes, you're right, I shouldn't have believed that paper and I shouldn't have said what I did. I take it back." Here he does the right thing; it is appropriate to regret and apologize for the speech act in question, even if by some bizarre twist in reality he turns out to be right.

The feature I am pointing out plays a central role in the prima facie case for the justification account of the norm of assertion. Only when the speech act itself is at fault, do we have reason to think that some norm of assertion is at work; when only the content of the assertion needs to be taken back, the assertion itself is not at fault. This point should be self-evident, but I emphasize it so that it is not overlooked: norms of assertion are norms governing a certain type of human activity, and thus relate to the speech act itself rather than the content of such an act. Notice that when we look at the four conditions for knowledge above, the only ones regarding which apology or regret for the speech act itself is appropriate are the belief and justification conditions. There is, therefore, a prima facie case that knowledge is not the norm of assertion, but rather justified belief is.

It is important as well to notice that the distinction here should not be accounted for in terms of a distinction between responsibility and blameworthiness. This distinction comes into play when our behavior is irresponsible though not blameworthy. For example, if Jim offends Billy Bob by making fun of Republican policies, not realizing that Billy Bob is a fan of these policies, Jim may be morally responsible for offending Billy Bob, but not blameworthy. After all, he might have had good reason to think that Billy Bob shared his disparaging attitudes toward Republican policies (I'm assuming, of course, that such attitudes can be justified). Not knowing

better, and not having been negligent in any way for not knowing better, Jim shouldn't be blamed for offending Billy Bob, even though he is morally responsible for doing so. (Note that I am not implying here that he has done something that is, all things considered, morally wrong. The only point needed here is that he is morally responsible for offending Billy Bob, and that fact is a prima facie indicator of having done something morally wrong.)

Note, however, that in such cases, upon finding out the true nature of the situation, the proper thing for Jim to do is to apologize (assuming that he grants that he shouldn't have offended Billy Bob). He cannot appeal to the inappropriateness of blame as a justification for not apologizing. The key here is that a correct understanding makes the apology appropriate. What we are called on to apologize for goes with responsibility, even when one is blameless for the action in question.

Applied to the topic at hand, this result shows that one can't avoid the above prima facie case for the justification account by claiming that what one apologizes for tracks some secondary normative notion such as blameworthiness, and that a correct account of the norms of assertion should track a primary normative notion such as responsibility. The point of focusing on which speech acts one ought to apologize for, after learning the true nature of one's action, is to show that in the primary sense of responsibility, rather than the secondary sense of blameworthiness, there is strong evidence on behalf of the justification condition. Apologies are appropriate for being insincere and for speaking without adequate evidence. For the other necessary conditions for knowledge, the appropriate response is to take back the content of the assertion, but not to apologize for the speech act itself.

It is also important to note that one may need to apologize for the *consequences* of a

speech act, even if the speech act itself doesn't call for one. If I tell you something false, but have very good reasons for what I think, and you are harmed in the future because you take my word on the matter, it would be appropriate and necessary for me to apologize for causing you harm. But it is one thing to apologize for the harm one's actions cause, and quite another to apologize for the speech act itself. In such cases, one feels in a special way the weight of responsibility, and one agonizes over the potential one has for causing unintended harm. There is a clear difference, however, between appropriate attitudes when one has caused harm by, say, driving while intoxicated and when one has caused harm when driving with an ordinary modicum of care. In both cases, one may regret driving because of the harm caused. In the first case, an apology for getting behind the wheel is certainly in order. In the second case, one may regret having gotten behind the wheel, thinking to oneself "if only...". One might even vow never to drive again, and though such a decision may count as a rational one, it is certainly not morally required. Apologizing for harm caused is one thing, and apologizing for the actions that caused the harm is another.

I'm assuming, in this example of driving with normal care and in the previous one of saying something false, that one has not done something negligent or careless in the course of performing the action in question that would make our assessment of the case more like the case of driving while intoxicated. That is precisely the right assumption to make, however, to compare with cases in which nature has conspired so as to deceive one regarding the truth. You have taken the ordinary modicum of care in the opinion in question, and such care is sufficient for justifying you in believing what you sincerely report. In such a case, just as in the ordinary driving case, there are grounds for regret and sorrow when one's actions cause harm. One may

even apologize for ever getting in the car or opening one's mouth. To the extent that such an apology is fitting, however, to that extent we have reasons to doubt that there was enough evidence to justify the assertion in question. In the normal case, such an apology would count as an overreaction albeit an understandable one. It is one thing to need to apologize for the harm caused, and it is quite another thing for the action itself to have been morally irresponsible.

So, two points are in order about this prima facie argument for the justification view. The proper approach to the topic of norms of assertion should focus on speech acts themselves. Hence, which acts of assertion demand apology is a good guide for determining where the norms of assertion lie. Moreover, by focusing on what demands apology, we have a criterion for distinguishing between blameworthiness and irresponsibility. Norms of assertion should concern what we are responsible for rather than what we are blameless for, and focusing on what demands apology places the focus on the concept of responsibility where it belongs.

This case for the justification account is prima facie only, and its implications cannot be endorsed without showing how the justification view fares once Williamson's arguments on behalf of the knowledge view are taken into account. As we have seen, the two primary arguments for that view involve the explanation of Moorean sentences and the related BUMS view, and considerations related to lottery sentences. I turn then to the question of how powerful these arguments are for the knowledge account, and whether they are powerful enough to undermine the prima facie case for the justification account.

Moorean Sentences and the BUMS View

Williamson's Moorean argument arises from the claim that it is incoherent to say that p is true but that one does not know that p is true. This argument counts in favor the knowledge account and against the justification account only if the justification account cannot explain the data in question. It can, if we employ the right kind of justification in the explanation. The kind of justification in question is the kind that puts one in a position to know, i.e., that kind which, if ungettiered and combined with true belief, yields knowledge. I will term this kind of justification "epistemic justification." The crucial feature of such justification is that one cannot have it for believing *p*, but not enough for believing *p* to count as knowledge. That is not to say, of course, that such justification is sufficient for knowledge. Justified false beliefs are possible, as are justified true beliefs that are gettiered. What can't happen though is that a true belief has an epistemic justification that is ungettiered, but where the level of justification falls below the threshold needed for knowledge.

In such a case, if one is justified in believing *p*, and one knows that one believes *p*, then one is justified in believing that one knows *p*. Showing that this claim is true will take some work, but before taking on that task, it is worth seeing what the payoff will be. This fact, I will claim, is sufficient to explain the untoward character of a Moorean assertion. We ordinarily assume that speakers know whether they are being sincere, and if the justification account is correct, they should not say anything without being justified in believing what they say. But if they are justified in believing *p*, then they are justified in believing that they know *p* as well (given that they know whether they are being sincere). So, by whatever basis they have for

asserting *p*, they also have a basis for asserting that they know *p*. From this point and a mundane principle about justification, it follows that if they have a basis for asserting *p*, they lack a basis for asserting that they don't know *p*. So the Moorean sentence is out-of-order, and the justification account can explain this impropriety.

It can explain this impropriety, assuming that we can defend the principle used in this explanation, the principle that if one is justified in believing *p* and knows that one believes *p*, then one is justified in believing that one knows *p*. The argument for this claim proceeds via the standard conditions for knowledge, showing that if a person is justified in believing *p*, then they are justified in believing that each of the conditions for knowledge obtains. First, the principle stipulates knowledge that the belief condition holds, and the truth condition itself is trivial: if one is justified in believing *p*, one is justified in believing that *p* is true. The truth condition would not be trivial if the consequent of this conditional required actual belief, but it does not. The phrase 'justified in believing' can be interpreted in two ways, one of which requires belief and the other one not. Since the present topic concerns what the quality of evidence needed to put one in a position to know, we should not interpret the phrase as requiring actual belief. If we do not, the principle only requires that one's total evidence confirm the claim that *p* is true when that evidence confirms *p*, and if anything in epistemology is an unremarkable truth, that is. Thus, one is justified in believing that the belief condition and the truth condition holds if one is justified.

Assume that we hold an internalist conception of justification. In one sense, a defender of the justification account is free to hold whatever conception of justification he wishes, so this assumption is not especially troubling. One implication of such a conception is usually put as

follows: one's justification will be accessible to one on reflection. I will assume this point as well for now, though later I will show that it was an unnecessary assumption. An implication of the claim that one's justification is accessible by reflection is that one will be able to ascertain on reflection whether one's belief is justified, and if one can so ascertain, then one's total evidence is sufficient for showing not only that *p* is true but also that one is justified in believing *p*.

One way to resist this argument is to question whether, in reflecting, one would be creating additional evidence so that the body of evidence apart from reflection need not itself be sufficient to justify believing that one has a justification for the claim in question. Here is the proper response, however. Take a person who is justified in the internalist sense, and then imagine an evil angel implanting the belief in this person that this person is justified in their belief. The proper thing to say about such a case is that the belief is one whose content is warranted, but where the believing itself is not warranted. It is not properly based, but it is simply false that the content of the belief has little going for it epistemically, as would be the case if the evil angel implanted the belief that there are an even number of grains of sand on a certain beach. No, the dark angel has done something to prevent doxastic justification (the kind of justification that attaches to a belief when its content is (propositionally) justified and the belief is properly based), but has also prompted a belief that could have been justified had it only been properly based. That is what suitable reflection provides. The role of the reflection is to bring it about that the belief formed is well founded or properly based, that it is suitably related to the evidence that is already present, so that the meta-belief that is formed is doxastically justified and not just propositionally justified. It is different if the metabelief is planted by an evil angel. In that case, the proper thing to say would be that the belief is propositionally justified though not

doxastically justified because improperly based.

This argument assumes internalism about justification as well as a particular version of internalism, usually called access internalism. There are, of course, alternative conceptions of justification on which one can be justified and yet be in no advantaged position to tell that this is so. That need not bother us here, for in the present context, we are free to specify whatever kind of justification we wish in giving an account of the norms of assertion. So the fact that there are other conceptions of justification or other kinds of justification is no objection to the claim that internalist justification has the property noted.

I do wish to point out, however, that the access aspect of the internalism used here is not necessary to the account. What really matters to the above argument is that the logic of justification honor the following principle:

(J) J$sp \vdash$ JJsp.

Talk of access to justificatory status is just an imaginative way of putting this point. Strictly speaking, however, (J) only claims that if a body of information stands in the confirmational relationship to p, then that same body of information stands in that same relationship to the claim that p is justified. Talk of access will ordinarily mean something like principle (J), but it can also fail to do so in unusual cases. For example, suppose there is a demon who is devoted to keeping you from reaching any reflective conclusions. The demon is powerful and able to make good on his intentions. Then if you were to reflect about anything, you'd come to no conclusions at all. Such a result is a problem for access internalism, but not for principle (J). So the appeal to access internalism here should be viewed as simply an imaginative way of endorsing principle (J) in the sort of internalism being assumed. That principle is true in virtue of some capacity or

disposition of a person who has evidence confirming *p*, but that ability or capacity or disposition need not be cashed out in terms of some subjunctive conditional concerning what would happen upon reflection. What matters for the internalism being assumed here is that evidence one possesses is a matter of which internal states one is in, and that the subject in question have the ability to detect these states in such as way that it can guide behavior, both in terms of belief and action. Given this limited internalism, we can still secure the result that if a body of information justifies a certain claim for a person, it also confirms for that person that the claim in question is justified by the relevant body of information.

One might wonder here whether it is possible to construct a sorites-like objection to this view, by constructing a sequence of bodies of evidence which are pairwise indistinguishable to an agent, but which begins with a body of evidence that definitely fails to confirm a given claim and ends with a body of evidence that definitely confirms that same claim. I don't think such an argument can be successful here, but a full defense of this claim depends on the argument I will give next concerning the gettier condition. For the present, I'll merely assume the argument below in order to disarm this objection. The assumption is that, when you are epistemically justified in believing *p*, you are epistemically justified in believing that further inquiry into the truth of *p* is not needed. Given this principle, we can disarm this sorites-like counterargument. The central claim I wish to make in disarming this argument is this: if you can't tell whether a given body of information confirms that further inquiry is not needed, it doesn't confirm that claim. So if two bodies of evidence differ with respect to the degree of justification they provide for a given claim, they either both yield epistemic justification or neither yields it; there can't be two bodies of evidence that as far as one can tell provide the same degree of confirmation for

closing off inquiry regarding a claim and yet where one of the bodies of evidence legitimates closing off inquiry and the other one doesn't. That's not to say, of course, that one will be infallible in detecting the proper course of action to take given a particular body of evidence. It is only to say that there is a detectable difference between bodies of information that warrant abandoning further inquiry and those that do not.

Given this result, the only condition still to be considered is the gettier condition. What we have shown to this point is that if you are justified in believing *p*, and know that you believe *p*, then you are justified in believing that your belief is true and justified. The consequent of the principle we are attempting to establish claims that you are justified in believing that you know *p*, and to finish the argument for this principle we only need to show that if you are justified, your total body of evidence confirms that you've satisfied the gettier condition as well. In order to finish the argument, we need to take advantage of a special feature of knowledge. Knowing involves a justified attitude of closure to further inquiry, confirmed by the fact that it makes little sense to say, "I know that it is raining outside, but I think I should go check to make sure," or to say, "I know today is Thursday, but further inquiry regarding today's date is probably appropriate." Just imagine a news conference in which a scientist says, "We've investigated thoroughly the claim that there is some connection between electrical power lines and certain forms of cancer, and as a result, we now know the answer: living close to electrical power lines simply does not cause these kinds of cancer. But we still need to investigate further..." Such a remark would be utterly bizarre, and the reason it would be bizarre is because knowledge involves a legitimate closure of investigation. In particular, it involves an inquiry that is of sufficient quality that it licenses the conclusion that any further learning could undermine one's

present opinion only by presenting one with misleading information.

The point I am making should not be confused with a stronger point about the relationship between knowledge and further learning. One may be tempted here to say that inquiry adequate for knowledge licenses one to conclude that further learning could only confirm one's present opinion, but that claim is too strong. It is too strong because of the possibility of misleading pockets of information. A simple statistical case will suffice as an example of such. Suppose statistical knowledge is possible, and that one's investigation has given one knowledge that most tosses of a given die will not be sixes. It is consistent with such knowledge that were one to investigate further, any finite string of sixes could occur on future tosses of the die. In some such cases, one's opinion would have to change in order to be rational (because the string that occurs is so long that it swamps the evidence obtained prior to that string). So the proper claim is not that no further learning could rationally undermine present opinion, but rather the weaker claim that any further learning that would rationally undermine present opinion would involve misleading information.

We need only tie this notion of misleading information to the defeasibility theory in order to complete our case that justification implies justification that one has not been gettiered. It is well-known that knowledge is different from undefeated justified true belief, since some defeaters are misleading. I will not here present an account of the difference between misleading and non-misleading defeaters, but even without a precise characterization of the difference, the following can be noted. If one's total evidence confirms that further investigation could undermine present opinion only by uncovering misleading information (as it must in order for the justification in question to be epistemic justification, the kind of justification that puts one in a

position to know), then one's total evidence confirms that further learning could at worst reveal only misleading defeaters. Hence, if one is epistemically justified in believing *p*, one is justified in concluding that one's justification is ungettiered, i.e., one's justification is defeated at most by misleading defeaters.

It is worth repeating here that these claims do not require any belief on the part of the epistemic agent regarding further learning. Nor do these claims require that the individual in question has the concept of justification or of further learning. The claims in question are claims about what total bodies of evidence confirm, not about what an individual believes or is capable of believing.

All the above to say this: the justification account of assertion can explain just as well as the knowledge account why the Moorean sentences are improper, and can explain the BUMS view as well. The reason the justification account can succeed here is that to have the kind of justification in question, one must have a justification that also justifies the claim that one has knowledge, given that one knows what one believes. So when a person makes an assertion, followed by a reflective judgment regarding a lack of knowledge, we justifiably suspect incoherence. The justification requirement on assertion implies that the person is in a position to assert that they know the claim in question, so long as they reflect and notice that they believe the claim; and any reasonable account of first-person epistemic authority regarding our beliefs warrants the assumption that the speaker knows whether he or she is sincere in asserting the first conjunct of the Moorean sentence. So the justification view can explain the incoherence of Moorean assertions precisely because such remarks run counter to the kind of evidence needed to justify assertion and the reasonable assumptions we all make about the kind of access we have to

our own beliefs.

The justification account can also explain the BUMS view, according to which we represent ourselves as knowing a claim when we assert it. The grounds needed for legitimate assertion constitute reasons accessible to the speaker for thinking that the speaker has met the non-psychological conditions for knowing, and no one is entitled to assert a claim insincerely, i.e., without believing it. So in asserting, one is committed to the claim that the assertion is legitimate and sincere, and thus committed to the claim that one knows what one asserts (because that which legitimates the assertion also confirms that one has met the non-psychological conditions for knowledge). I take it that such commitments are sufficient for endorsing the claim that in asserting one *represents* oneself as knowing, but if the notion of representation involves some further aspect, then I think the BUMS view is incautious and should be replaced by the commitment claim, that in asserting, one commits oneself to the claim that one knows.

So the knowledge account cannot be judged superior to the justification account on the basis of Moorean considerations or the attitudes of BUMS. That leaves the primary stumblingblock the issue of lottery propositions.

Lottery Propositions

Regarding lottery propositions, Williamson's view seems to be straightforward. His view seems to be that the belief that my ticket will lose cannot be known to be true, but can be justified. It would appear, then that on the justification account it is permissible to assert that my ticket will lose, but it is not permissible to so assert on the knowledge account. Since it isn't

permissible to assert that, the knowledge account wins.

There are several points at which this argument is questionable. The first point concerns the claim that we cannot know that a given ticket will lose. Though I share Williamson's viewpoint here, there is reason to question whether our opinions here are a holdover from the infallibilist epistemology we've all inherited from our ancestors. For much of the history of epistemology, it was assumed that one could not have knowledge if one's evidence failed to rule out every possible alternative to one's present viewpoint, and the continuing temptation of certain forms of skepticism witnesses to the present effects of this heritage. So perhaps Williamson and I are just wrong: we can know that a given ticket will lose.

This point becomes more plausible when one acknowledges the possibility of statistical knowledge. If we acknowledge this possibility and yet deny that we can know that a given ticket will lose, it is very difficult to explain the difference between statistical knowledge and lottery cases. We can, of course, find differences, but the question is whether the differences constitute an adequate explanation of the difference in epistemic status. The longer we search for such an explanation and fail to find one, the more plausible becomes the position that our refusal to grant the existence of knowledge in lottery cases is just a holdover from a bad epistemological heritage.

If we acquiesce and relinquish the view that knowledge is not possible in lottery situations, lottery cases no longer provide any sort of argument for the knowledge view over the justification view. That result I find comforting, but unfortunately, I agree with Williamson that we don't know that a given ticket will lose. So I will proceed on the assumption that Williamson is correct in this matter, since I find that comforting too!

The second point at which the above argument is questionable, however, is more telling. One of the premises of that argument is that the justification account licenses the assertion of the claim that a given ticket will lose. Williamson holds a view of justification and rationality that supports such a claim. He writes,

> It is plausible, nevertheless, that occurrently believing *p* stands to asserting *p* as the inner stands to the outer. If so, the knowledge rule for assertion corresponds to the norm that one should believe *p* only if one knows *p*. . . . Given that norm, it is not reasonable to believe *p* when one knows that one does not know that *p*. If one knows that what one knows is only that *p* is very probable, then what is reasonable for one to believe is only that *p* is very probable. For example, I should not believe that my ticket will not win the lottery. . . . On this analogy between assertion and belief, the knowledge rule for assertion does not correspond to an identification of reasonable belief with knowledge. . . . The rule makes knowledge the condition for permissible assertion, not for reasonable assertion. One may reasonably do something impermissible because one reasonably but falsely believes it to be permissible. In particular, one may reasonabley assert *p*, even though one does not know *p*, because it is very probable on one's evidence that one knows *p*. In the same circumstances, one may reasonably but impermissibly believe *p* without knowing *p*.[7]

In this passage, Williamson employs a notion of rationality on which it can be rational to assert that one's ticket will lose. Given the close connections between the concepts of justification and rationality, it is fair to conclude that Williamson could also have identified the concept in question as one of justification.

This notion to which Williamson appeals in this quote is not, however, the notion of justification in the view of the norm of assertion under discussion here. Recall the feature of internalistic epistemic justification noted above: such a justification for *p* is a justification for the claim that one knows *p* in the presence of knowledge of what one believes. Given that feature of justification, Williamson's concept of reasonability cannot be identified with the concept of epistemic justification.

The argument for this last claim is as follows. On the notion of justification being employed here, it is false that one can be epistemically justified in believing that a given ticket will lose (if it is false that one cannot know this). The kind of justification in question is that kind which puts one in a position to know, i.e., that is sufficient for knowledge in the presence of ungettiered true belief. One has knowledge only when closure to further inquiry on the issue is legitimated, and such closure is a function of the quality of one's evidence. So on the present view, one's total evidence must confirm that further inquiry would reveal only misleading information about the truth of what is believed. In ordinary lottery circumstances, however, one's evidence does not confirm that further inquiry would be a waste of time. In ordinary circumstances, those circumstances in which we are most inclined to deny knowledge that a given ticket is a loser, we take it to be worth checking the newspaper the next day to see if the improbable has occurred, and this attitude is itself justified. If so, however, our evidence confirms that the newspaper's word will be non-misleading evidence about whether the ticket won. Because of this fact, whatever justification we have for holding that our ticket will lose, that justification is not the epistemic justification used here in the justification account of the norm of assertion.

It is worth putting the argument that there is such a kind of justification in a different way. Imagine two cases of justified true belief, one of which counts as knowledge and the other of which involves a justification that is gettiered. In the case of knowledge, the closure aspect of knowledge is clearly present. Note, however, that it is not in virtue of the external conditions for knowledge that such closure is present. It is rather something about the character or quality of one's justification that licenses such closure. That is to say, the totality of one's evidence not only confirms the claim in question, it also confirms that further inquiry is not necessary in this case. By hypothesis, however, the totality of one's evidence is the same in the case that is not a case of knowledge: the relevant difference between the two cases is not a difference in the totality of one's evidence, but rather in the fact that one's justification is gettiered in one case but not in the other. Hence, given that the totality of evidence is the same in both cases, one is equally justified in holding that further inquiry is not necessary in the non-knowledge case as well. Put more simply, the closure aspect of knowledge is a feature of the kind of justification for knowledge, and is thus a function of the quality and character of one's total evidence. So, as advertised, there is a notion of justification–epistemic justification–on which closure to further inquiry is legitimated by the same evidence that confirms the truth of the claim in question.

This response to the lottery argument–that, in fact, the justification account has precisely the same implications about lottery sentences as does the knowledge account–may be harder to accept if one adopts Williamson's account of evidence, according to which one's body of evidence is simply one's knowledge. On this view of evidence, the account above claims that one's total body of knowledge can confirm that further investigation of a certain claim is not needed. That is obviously true when the claim in question is known, but it is less obvious when

knowledge is not present. In such a case, we would need to ask what must be true of a body of knowledge in order for it to put one in a position to know that further investigation is not needed. If evidence is, and is only, what one knows, it might be plausible to answer that the only cases where one is in a position to know that further investigation is not needed will be cases in which one knows the truth value of the claim in question.

If one refuses to identify evidence with knowledge, I think the temptation to this view is not nearly so strong. As argued above, whether one knows is a function both of the quality of one's total evidence and also the cooperativeness of the environment in which one finds oneself. Given this point, one should expect that the environment can function so as to undermine knowledge of the claim in question without impugning the quality of the evidence in question. If one believes that this claim about epistemic justification doesn't fit well with the view that evidence is knowledge, that strikes me as a reason to reject the knowledge view of evidence rather than reject the characterization of epistemic justification just given.

The reader might be left wondering, however, whether more can be done beyond merely giving my word on the matter. The answer, I think, is "yes," for there are independent reasons to be suspicious of the view that evidence is knowledge.

The implausibility arises when we consider Williamson's account of perceptual evidence in the case of illusions. In such a case, one clearly has perceptual evidence, but that evidence can't be identified with a claim to the effect that, e.g., there is water up ahead, since in fact there isn't any water up ahead. Williamson's response is:

> If perceptual evidence in the case of illusions consists of true propositions, what are they? The obvious answer is: the proposition that things appear to be that way.[8]

But why think that a person possesses that information or believes such a proposition? It is obviously true that one can suffer an illusion and never consider the question of how things appear to be. To this point, Williamson replies,

> ...[O]ne may not consider the cautious proposition that things appear to be that way; one may consider only the unqualified proposition that they really are that way. But it does not follow that one does not know that things appear to be that way, for one knows many propositions without considering them.[9]

This response is true but insufficient to answer the objection. Even if there are cases in which we know propositions without ever having considered them, it would be a monstrous coincidence if in every case of unreflective illusion, the knowledge attribution were true. It is true that I knew that

11,111,111+1=11,111,112

prior to my just considering it now for the first time, but that gives us no reason as far as I can tell to think that in every case of unreflective illusion, I know that I am being appeared to in a certain way. In the simple mathematical case above, it is plausible to hold that I dispositionally believed the claim in question prior to explicitly considering it. In other cases of simple addition, however, the proper thing to say is that I was disposed to believe, rather than dispositionally believed, the claim in question prior to explicitly considering it. Perhaps in some cases of illusion, I dispositionally believe the claim about how I am appeared to. But it is equally plausible to hold that in other cases, I only have a disposition to believe the appearance claim, and thus do not know the way in which I am appeared to. One would expect in this regard that the more salient the possibility of illusion, the better a case there is for holding that a

dispositional belief is present. Correlatively, the more common and ordinary the case, when the possibility of illusion is the furthest thing from one's mind, the better the case for holding that there is only a disposition to believe that things appear to be a certain way.

There is a deeper problem as well with Williamson's treatment of perceptual evidence. He views veridical experience as unproblematic, since, he holds, one can identify the propositional content of the experience and note that the person in question knows this content to be true. This knowledge, Williamson thinks, allows him to explain the evidence provided by veridical experience as a kind of knowledge.

This view strikes me as getting things backwards. When asked to defend the belief that there is a tree with yellowing leaves in the yard south of mine, I will appeal to experience. I don't appeal to my knowledge or to my beliefs. When I say, "I can see it," I'm reporting an experience with propositional content that, in successful cases, is the same content as something I know, but my evidence is that content under a certain modality that is different from the modality of belief or knowledge. It is the modality of experience; my evidence is the content-as-experienced, not that content-as-believed nor that content-as-known. It may be that in reporting one's evidence or in reflecting on one's evidence, the content in question gets transplanted into one's belief box or knowledge box; but that point should not obscure the fact that the evidence existed in the experience itself before the transplant. In reporting the experience, one is simultaneously reporting the existence of something believed and perhaps known to be true, but it is not the belief or the knowledge that is the evidence. The evidence is the experience itself.

The point of this digression was to block any argument against the account of epistemic justification that I am employing here that might arise from the knowledge account of evidence.

If such an argument is forthcoming, it should be rejected, since there are good reasons to reject that account of evidence.

We can therefore return to the main thread of the discussion, which concerns the argument from lottery propositions against the justification account of assertion and in favor of the knowledge account of assertion. I argued above that the central claim of this argument–the claim that the justification account of assertion licenses the assertion that one's ticket will lose–is false when the kind of justification in question is epistemic justification. That is, for internalistic epistemic justification, it is false that one is justified in thinking that one's ticket will lose, and hence if such justification is the norm of assertion, then the assertion that one's ticket will lose is inappropriate. Therefore, Williamson's lottery argument is ineffective at showing the at the justification account of assertion is incorrect.

Conclusion

The proper conclusion to draw, then, is that the prima facie case for the justification view is an ultima facie case as well. The primary arguments that Williamson gives that can be used to try to undermine the justification view fail to do so. Their failure traces to a common flaw, the flaw of failing to notice that the ordinary notion of justification, on which one is justified in thinking that one's lottery ticket is a loser, is not suitable for use in an account of knowledge. The notion that is suitable has the following feature: when it obtains, one has reason for believing that all epistemic requirements for knowledge have been met. As a result, the justification account of assertion can explain the incoherence of Moorean assertions that a claim is true but unknown by

the speaker as well as the BUMS view that in asserting we represent ourselves as knowing. In addition, this version of the justification account of assertion can explain why one shouldn't say that one's lottery ticket is a loser. Even though there is an ordinary notion of justification on which one is justified in believing that one's ticket is a loser, in the appropriate epistemic sense of justification, one does not have such justification. One lacks such justification precisely because the character of one's evidence does not warrant the kind of closure experience with respect to further learning that is characteristic of knowledge. In short, one's justification for thinking that one's ticket is a loser is insufficient justification for knowledge. The failure of these objections to the justification account, combined with the positive defense of that account in terms of the distinction between taking back the content of an assertion versus taking back, regretting, and apologizing for the speech act itself, confirms that the justification account is the better account of the norm of assertion. Put succinctly, the proper counsel to give is not to quit saying what you don't know to be true, but to quit say what you don't believe and what you don't have good epistemic reasons to believe.[10]

Notes

1. The view I defended in *The Value of Knowledge and the Pursuit of Understanding*, (New York: Cambridge University Press, 2003).

2. See, especially, pp. 238-40.

3. Please do not interpret this lighthearted levity as a sign of disrespect. To get the right tone, one needs to imagine an accent from the Bronx and the phrase, "How aah ya, you's bums!"

4. Kent Bach and Mike Harnish defend a version of this view in *Linguistic Communication and Speech Acts*, (Boston: MIT Press, 1979).

5. See page 130 of *Knowledge and Its Limits*, for example.

6.This is the language of William James in describing those who say it is better to go without belief forever than risk being proven wrong.

7.Timothy Williamson, *Knowledge and Its Limits*, (Oxford: Oxford University Press, 2000), pp. 255-6.

8. *Knowledge and Its Limits*, p. 198.

9. *Knowledge and Its Limits*, pp. 198-9.

10. I put the justification qualifier here in a weaker way than requiring that one justifiably believe the claim in question. The reason is that to say that one justifiably believes, in my terminology, implies that one's belief is properly based. I do not believe that proper basing is required for assertion, but arguing for this point goes beyond the scope of this paper, so I merely note the point here without argument.