

Is Justification Knowledge?

B.J.C. Madison
The University of Notre Dame Australia

School of Philosophy and Theology
The University of Notre Dame Australia
19 Mouat Street (PO Box 1225)
Fremantle, Western Australia, 6959

bmadison@nd.edu.au

Abstract: Analytic epistemologists agree that, whatever else is true of epistemic justification, it is distinct from knowledge. However, if recent work by Jonathan Sutton is correct, this view is deeply mistaken, for according to Sutton justification *is* knowledge. That is, a subject is justified in believing that p iff he knows that p . Sutton further claims that there is no concept of epistemic justification distinct from knowledge. Since knowledge is factive, a consequence of Sutton's view is that there are no false justified beliefs.

Following Sutton, I will begin by outlining kinds of beliefs that do not constitute knowledge but that seem to be justified. I will then be in a position to critically evaluate Sutton's arguments for his position that justification is knowledge, concluding that he fails to establish his bold thesis. In the course of so doing, I will defend the following rule of assertion: (The JBK-rule) One must: assert p only if one has justification to believe that one knows that p .

Is Justification Knowledge?

B.J.C. Madison

(Forthcoming in the *Journal of Philosophical Research*)

Draft Version – Do Not Cite Without Approval

Abstract: Analytic epistemologists agree that, whatever else is true of epistemic justification, it is distinct from knowledge. However, if recent work by Jonathan Sutton is correct, this view is deeply mistaken, for according to Sutton justification *is* knowledge. That is, a subject is justified in believing that p iff he knows that p . Sutton further claims that there is no concept of epistemic justification distinct from knowledge. Since knowledge is factive, a consequence of Sutton's view is that there are no false justified beliefs.

Following Sutton, I will begin by outlining kinds of beliefs that do not constitute knowledge but that seem to be justified. I will then be in a position to critically evaluate Sutton's arguments for his position that justification is knowledge, concluding that he fails to establish his bold thesis. In the course of so doing, I will defend the following rule of assertion: (The JBK-rule) One must: assert p only if one has justification to believe that one knows that p .

I. Introduction

Analytic epistemologists agree that, whatever else is true of epistemic justification, it is distinct from knowledge. However, if recent work by Jonathan Sutton is correct, this view is deeply mistaken, for according to Sutton, justification *is* knowledge.¹ That is, a subject is justified in believing that p iff he knows that p . Sutton further claims that there is no concept of epistemic justification distinct from knowledge. Since knowledge is factive, a consequence of Sutton's view is that there are no false justified beliefs.

There are two putative types of unknown justified beliefs. First, there are those in which the subject knows that he does not know the proposition in question and yet intuitively he is justified nonetheless (known unknowns). Secondly, there are those in

which the subject is ignorant of his ignorance (unknown unknowns). In this paper I shall argue that Sutton fails to establish that unknown unknowns are not justified.

Unknown unknowns fall into two categories: justified false beliefs and the justified true beliefs of Gettier cases. Putative justified false beliefs include my belief that there is a dog on the lawn on the basis of its looking to me like there is a dog on the lawn, its sounding to me like there is a dog on the lawn, Stephen's testimony that there is a dog on the lawn, etc. But in fact there is no dog on the lawn, but rather only a cleverly disguised sheep. Here what I believe is false, but intuitively, I am justified nonetheless.

On other hand, a Gettier case is minimally an instance of *justified*, true belief where the subject fails to know what is believed. Gettier cases arise when there is *too much* epistemic luck around. For instance, suppose Al seems to see a dog on the lawn and on that basis infers that there *is* a dog on the lawn. Suppose further that what he sees is not a dog, but rather only a cleverly disguised sheep. Unbeknownst to Al, however, there is *also* a dog on the lawn (perhaps just out of view behind a tree). Here Al believes that there is a dog on the lawn, it is true that there is a dog on the lawn, and since he has overwhelming evidence that there is a dog on the lawn (it plainly looks like a dog in front of him, after all!), intuitively Al is *justified* in believing that there is a dog on the lawn.

However, does Al *know* that there is a dog on the lawn? Surely not. There is too much luck around. He could have formed the same belief using a very similar method that could have very easily been false. That is, there is a very nearby possible world where Al forms the same belief that there is a dog on the lawn on the basis of there seeming to see one in front of him, but the belief is false, since there is no dog on the lawn just out of view to make his belief true, but just the cleverly disguised sheep that

looks like a dog. In the actual world, it is just sheer luck (in the relevant sense) that Al's belief is true.ⁱⁱ The standard lesson of such cases is that knowledge is incompatible with such luck. A corollary of such cases is that, despite not knowing the proposition in question (e.g. that there is a dog on the lawn), the subject is nevertheless *justified* in believing as he does.

A consequence of Sutton's proposal is that Gettier cases are impossible. There will still be cases where a subject fails to know a proposition because there is too much luck around, but these will not be cases of *justified*, true beliefs that fail to be knowledge; hence, there will be no Gettier cases since these just are cases of justified, true belief that are not knowledge.

Sutton agrees that it is the presence of too much of the relevant kind of luck that prevents subjects like Al from having knowledge (Sutton 2007: 9). He goes on to say that the reason we are inclined to call such cases instances of justified belief is because the subject *would* have knowledge if it had not been for this luck. Sutton thinks that this shows that justification is parasitic on the concept of knowledge. As he says, "Justification in the relevant sense is perhaps a disjunctive concept – it is knowledge or would-be knowledge" (Ibid.: 9-10). On the face of it, this is a somewhat odd claim for Sutton to make, given his official stance – and especially since he thinks that the subjects of Gettier cases are not justified. This is because if we understand justification as a disjunctive concept – i.e. as knowledge or would-be-knowledge – it seems that we ought to say that beliefs held by the protagonists of Gettier cases *are* justified. After all, those beliefs count as would-be knowledge, and so satisfy the disjunctive concept. Perhaps a charitable reading of Sutton here construes him as saying that what is merely *called*

justification in Gettier cases is a disjunctive concept, although it is not *really* epistemic justification proper, since it is not knowledge.

In short, for Sutton the subjects of Gettier cases have unjustified beliefs. Their beliefs are unjustified precisely because they are not knowledge. How then does Sutton think of justification? Sutton considers five candidates: justification as Warrantⁱⁱⁱ; justification as blamelessness^{iv}; justification construed deontologically^v; justification construed evaluatively; and justification as reasonableness.^{vi} While Sutton goes on to say interesting things about each, I will only consider here his treatment of justification as reasonableness, since it is the clearest case of epistemic justification that is not materially equivalent to knowledge.

II. Justification as Reasonableness

The final concept of justification that Sutton considers is spelled out in terms of reasonableness. Sutton asserts that ordinary usage counts in favour of applying the word “justified” primarily to actions, and applies the word “reasonable” to beliefs whenever and wherever philosophers talk of epistemic justification. We commonly talk of reasonable but false belief, or the subjects of Gettier cases believing reasonably. So if justification is a matter of reasonableness, it is false that justification is knowledge (since one cannot know falsehoods, and one fails to know in a Gettier case).

Why then, one might ask, is justification not to be understood in terms of reasonableness? Sutton’s argument is shaky here. He claims that justification is primarily meant to be a property of *beliefs*, whereas reasonableness is supposed to primarily apply to *people*, and only derivatively apply to beliefs. Sutton writes,

A belief is reasonable in the circumstances in which it is held if a reasonable person would or could hold it in those circumstances. [...] The notion of a reasonable *person* is understood in terms of knowledge; he is one whose belief-forming faculties and habits (e.g. inferential habits) are such as to deliver knowledge when conditions are right. (Sutton 2007: 35-6)

Given this account, people can have reasonable but unjustified beliefs when they lack knowledge (say, when they are victims of a Gettier case, or when they believe something false on the basis of good evidence), and reasonableness is not equated with mere blamelessness. Sutton intends these concessions to make his view more palatable.

But in making these concessions, one cannot help but wonder whether Sutton concedes too much. Has he in effect conceded that there is a concept of epistemic justification distinct from knowledge? He is right that there is a property that is distinct from mere blamelessness that people's beliefs can enjoy even when those beliefs are false. This property, even if not identical with reasonableness, is certainly one enjoyed when the subject is being reasonable. Why is not *this* property, which is not identical with knowledge, *justification*? Sutton's primary reason for denying this is that justification is "supposed to" be primarily a property of beliefs, and reasonableness, while it surely applies to beliefs, is supposed to primarily be a property of persons. But then again, it seems that 'reasonable people' could be defined as people whose beliefs are by and large reasonable, and so reasonableness applies primarily to *beliefs* and only derivatively to persons.

In short, why think that reasonableness is primarily a property of persons rather than of beliefs? Ordinary usage allows predication of reasonableness to persons, propositions, and attitudes, such as beliefs; there seems to be no *prima facie* reason to assume the priority of one over the others. Especially given the importance of this claim

for Sutton's case, what is needed is an argument for the view.

One might expect that issues concerning fundamental and derivative sorts of reasonableness will parallel issues concerning fundamental and derivative sorts of epistemic justification, e.g. when comparing the relations between personal, doxastic and propositional justification. For example, Kvanvig and Menzel (1990) argue that propositional justification is the basic sort, that personal justification can be defined in terms of doxastic justification, and that doxastic justification is in turn understood as propositional justification plus proper basing.

On the other hand, virtue epistemologists, and those inspired by them, argue that doxastic justification is the fundamental property, and that propositional justification is to be understood in terms of it. For example, John Turri has recently proposed the following account of the relation between propositional and doxastic justification:

(PJ) Necessarily, for all S , p , and t , if p is propositionally justified for S at t , then p is propositionally justified for S at t *because* S currently possesses at least one means of coming to believe that p such that, were S to believe p in one of those ways, S 's belief would thereby be doxastically justified. (original emphasis) (Turri forthcoming: 13-14)

These reliable 'means' of coming to believe something are thought of as intellectual virtues.^{vii}

Adjudicating which kind of justification/reasonableness is more fundamental than other sorts, if any, is beyond the scope of this paper. But it should be clear that the answer to this question is far from self-evident, and accordingly, Sutton is not entitled to assume without argument that reasonableness is necessarily primarily a property of persons rather than of beliefs.

Given this, what if some philosophers have simply misunderstood the sense in which justification attaches to beliefs? That is, it would surely be objectionable if a theory allowed that justification were not a property of beliefs in *any* sense, but is it really that objectionable if it applies to beliefs in a derivative rather than a primary sense?

In a broadly Moorean spirit, what are we more sure of: i) that justification applies to beliefs in a primary, rather than derivative way; or ii) that there are no justified false beliefs -- that there is no concept of epistemic justification distinct from knowledge? The truth of i) seems far less obvious than the falsity of ii). Therefore, if the truth of ii) is implied by the truth of i), the first claim should be rejected, or at least not presupposed as Sutton does.

On the model of justification I am considering, a belief is justified in its particular circumstances if a reasonable person would or could hold it in those circumstances.^{viii} I take it that how the notion of a reasonable person is understood is an open question. Perhaps Sutton is right: a reasonable person is to be understood in terms of knowledge; he is one whose belief-forming faculties and habits are such as to deliver knowledge when conditions are right. But even if this is right and justification must be understood in terms of knowledge, it does not follow that justification just *is* knowledge, which is his primary thesis. Sutton is arguing for that bold thesis, not the weaker one that Timothy Williamson defends that it is 'knowledge first', the primitive in terms of which all other epistemic notions are understood.^{ix} That is, the concept of justification being parasitic upon the concept of knowledge is one thing; the material equivalence of justification and knowledge is quite another.

On the other hand, perhaps the notion of a reasonable person is to be understood independently of knowledge, e.g. a reasonable person is one who proportions her beliefs to the evidence (whether that delivers knowledge or not). In that case, reasonableness would be an independent epistemic goal from knowledge and so be a distinct form of evaluative justification. And if there is an obligation to believe reasonably distinct from the obligation to believe only if one knows (e.g. one might have the obligation to believe only if one's belief fits the available evidence), reasonableness would be a distinct form of deontological justification. Conceding a notion of epistemic reasonableness distinct from knowledge (at least in the sense of not equating reasonableness with knowledge, even if it is not more conceptually fundamental), Sutton undermines the thesis that epistemic justification is materially equivalent to knowledge even before giving his arguments for it.

I will not provide a further taxonomy of concepts of epistemic justification. This is because my aim is merely to defend the platitude that justification and knowledge are distinct epistemic statuses from Sutton's claim that justification = knowledge, which threatens to undermine it. A *single* class of justified beliefs, therefore, that falls short of knowledge will falsify Sutton's thesis. With a discussion of epistemic justification now in place, we can turn to Sutton's arguments that identify it with knowledge.

III. The Arguments

A. Argument i) The Assertion Argument

Following Williamson, Sutton accepts the Knowledge Rule of Assertion: one must assert that p only if one knows that p . Sutton argues that the connection between

warranted assertability and knowledge is mysterious unless there is the same connection between justification and knowledge. That is, “one must: believe p only if one knows p , [is] a norm that is as constitutive of belief as Williamson’s knowledge rule is for assertion” (Sutton 2007: 44).

However, if sound, this first argument shows that the standards governing assertion need to be the same as those governing belief. This in itself is not surprising if the rules governing assertion are simply the ‘external’ counterparts of rules governing belief. It is a further question what norm(s) do govern assertion. In other words, the cogency of this argument depends in turn on the cogency of arguments for the Knowledge Rule of Assertion.

Sutton embraces Williamson’s three arguments for the knowledge rule of assertion. Williamson’s arguments are as follows:

- 1) The Knowledge Rule is needed to explain what goes wrong in quasi-Mooreian paradox. The hypothesis is that the reason it is improper to assert “ p and I don’t know that p ” is because by asserting p one implies that one knows it, which is then denied in the second conjunct, which is contradictory (Williamson 2000: 253-5).
- 2) The Knowledge Rule is needed to explain the unacceptability of assertions that one will lose in Lottery Paradox cases. It is allegedly improper to assert that one’s ticket will lose the lottery, no matter how likely that fact is. The best explanation of the source of this impropriety, it is suggested, is that one does not *know* that one will lose the lottery solely on the basis of its being highly improbable (Williamson 2000: 249-252).

3) The final argument for the Knowledge Rule is, as Sutton says, “simply the appropriateness of challenges and rebukes to assertions such as ‘How do you know?’ and ‘You don’t know that!’” (Sutton 2005: 375) The fact that appropriate challenges to assertions involve questioning the asserter’s knowledge provides some evidence of the correctness of the Knowledge Rule (Williamson 2000: 249-252).

As I said, this argument for the thesis that justification is knowledge presupposes the correctness of the Knowledge Rule of Assertion. If some other principle better explains the phenomena of assertion than the Knowledge Rule, this argument fails to go through. I will consider such a principle below. Sutton goes on to offer an argument based on the Lottery Paradox as well as one he calls the Modesty Argument. However, I will not lay them out here since both of these arguments also presuppose the correctness of Williamson’s Knowledge Rule of Assertion, as Sutton himself concedes (Sutton 2007: 59-60).^x

Instead, I will turn to Sutton’s fourth argument for the conclusion that justification is knowledge, since it is the only one he explicitly presents whose soundness is independent of the Knowledge Rule of Assertion. I will then consider Sutton’s claims about loose versus strict ways of speaking of justified belief since, despite Sutton’s not identifying it as an argument, one might maintain that these claims also work to support his case.^{xi}

B. Argument ii) The Posterior Evaluation Argument

If a belief that p is one that would be justified were one to form it, and it is in one's interest to have a belief in whether or not p , and one is capable of forming such a belief, then, in some intuitive sense, one should believe that p , although that 'should' is not purely a 'should' of epistemic obligation. It is rather a 'should' generated by prudential considerations of self-interest interacting with the epistemic goals that determine an evaluative notion of justification. (Sutton 2007: 57)

Sutton claims that if we substitute 'would be justified' with 'would constitute knowledge' in the above passage, then the claim as stated is clearly true. His thesis is that this claim is not true if any other epistemic concept other than knowledge is inserted in its place. I will show that this is false.

This thesis is argued for by means of a variation on the original Gettier case. In Sutton's case one has the desire to drive a Ford and will receive an incentive for doing so. One has lots of evidence that one's colleague owns a Ford, but this evidence is misleading. However, one does have another colleague who owns a Ford, so if one formed the belief that a colleague owns a Ford, one would have a justified true belief that would fail to be knowledge. However, one fails to take notice of one's evidence, and so fails to form any beliefs about who might own what car.

Suppose also, Sutton continues, that one later comes to appreciate the implications of one's earlier evidence, but also learns that the evidence was misleading, so that if one had formed a belief on its basis, one would have formed a true belief that would have been 'Gettiered', so to speak. Sutton now asks, "Is there an intuitive sense in which one will *now* judge that one should *then* have formed the belief that one had a colleague who owned a Ford – an allegedly justified true belief?" (Sutton 2007: 58) Sutton says it is clear that there is, and he is right about one particular sense in which one

should have formed the belief. One should have then formed the belief because one would have gotten the incentive. And it is true that that same non-epistemic sense in which one should have formed the belief is present in a version of the case where one forms the belief about the colleague's Ford in the absence of any evidence whatsoever. A true belief not based on any evidence at all would have equally gotten the incentive. However, only in the misleading evidence case, and not the absence of evidence case, is there an *additional* sense to the one that Sutton acknowledges in which one should have formed the belief, viz. that was the well-evidenced belief to form.

Sutton thinks that there is a 'partially epistemic' sense of a 'should' that arises only in a case where the original evidence was *not* misleading and so the subject would have gained knowledge. In such a case, Sutton comments as follows:

Evaluating one's former cognitive situation after the fact, is there an intuitive sense in which one should have formed the belief apart from the consideration that it would have been true? I contend there is: one should have known that one has a colleague who owned a Ford, and this is to say, that in the case at hand, that one let oneself down doxastically speaking in a way that goes beyond having merely failed to form a true belief. (Sutton 2007: 59)

I agree that in that case there is a sense that one should have formed the true belief, but it is the *same* partially epistemic sense present in the misleading evidence case. One failed to appreciate one's evidence and proportion one's belief to it, which is plausibly an epistemic obligation regardless of whether, as it happens, the evidence is misleading. In the misleading evidence case, one should have believed that one's colleague has a Ford, and one lets oneself down doxastically by failing to proportion one's belief to the evidence. Even if in hindsight the subject regrets having formed the

belief he did, given that his evidence was misleading, this does not undermine the claim that he should have formed it nevertheless – after all, he was diligently following his evidence, which is the most that can be reasonably asked of him. So contrary to what Sutton says, there is an intuitive sense in which one should have formed the belief that one’s colleague has a Ford in the misleading case apart from the non-epistemic sense in which believing truly would have been beneficial. The intermediate case between knowledge and the merely pragmatic case is the one where one should have formed the belief because that is what the evidence supported (even though the evidence was misleading). The existence of the intermediate case shows that justification is not knowledge.

On the face of it, these considerations also count against knowledge norms of belief and hence the Knowledge Rule of Assertion, since this is a case where one should believe that *p* even though one does not know that *p* -- in which case it is false that one should believe (assert) only what one knows.

It might be argued that Sutton would object to the way I have been using the term evidence throughout this discussion. This might be so, given that according to Sutton, *all evidence is knowledge*, whereas I have been holding that misleading evidence is still evidence.^{xii}

However, Sutton holds that,

...it is a very short path to my strict conception of evidence [on which evidence is knowledge] once one abandons the idea that justification, in any important sense, differs from knowledge. In any case, my argument for my strict conception is not that it is intuitive; it is that it gives satisfying and consistent diagnoses of the problem cases, *together with the fact that it follows from the identification of justification and knowledge for which I have already argued.* (emphasis added) (Sutton 2007: 136-137)

Therefore, a knowledge-based conception of evidence does not *support* justification being knowledge, but rather, in Sutton's view, it is a *consequence* of it. Given that I will provide sufficient argument to undermine that initial identification, it is at this point an open question what constitutes the correct account of evidence.

A possible reply here is that the knowledge-based view of evidence is only partly based on the fact that it allegedly follows from the identification of justification with knowledge; it is also supported by the fact that it casts light on some puzzles concerning evidence.^{xiii} Perhaps such an account can do this philosophical work; perhaps not. But to explore this topic here would be to digress from the central thesis that justification is knowledge, which is my concern. Not evaluating the claim that all evidence is knowledge is permissible given that as Sutton notes, his view of evidence is consistent with, and perhaps even follows from his view of knowledge and justification, and therefore cannot be used as an argument *for* that position.^{xiv}

C. Argument iii) Justified Belief: Loose v. Strict Talk

Sutton argues that just as we can confuse believing that *p* with believing that *probably p*, we use the phrase 'justified belief that *p*' *loosely* to refer to both justified belief that *p* AND to what is, *strictly speaking* mere justified belief that *probably p* (Sutton 2007: 66). In an aim to undermine further intuitive support for the distinction between justification and knowledge, Sutton argues, for example, that only *loosely speaking* do subjects in Gettier cases have justified beliefs, whereas *strictly speaking* their beliefs are not justified since they fail to constitute knowledge.

For example, we might say a Gettiered subject's belief is justified when speaking loosely, Sutton argues, since there are known, and hence justified beliefs 'in the vicinity': e.g. a Gettiered subject knows, and is therefore justified in believing, that there is *probably* a barn in the field when looking at a real barn deep in fake barn country; or she knows that it is *probably* 12 o'clock when looking at a clock that has unbeknownst to her broken twenty four hours ago, but has by luck stopped on what is now the correct time. (Ibid.) But *strictly speaking* the Gettiered beliefs are not justified, and are not justified precisely because they are not known.

It is true that we can and do sometimes speak loosely about justified belief, but that is not to say that we *must*. We judge that *strictly speaking* we can have justified beliefs that fail to be knowledge. To see this, consider an example Sutton provides:

Consider a stranger who asks where the post office is. One does not know, but has a vague idea that it is a mile to the right. Consequently, one says 'I believe that it is a mile down that way'. This is a perfectly proper utterance; I suggest that it is also not literally true. One does not believe categorically that the post office is a mile to the right; one believes that it is more likely than not. Ironically, a sentence of the form 'I believe that *p*' is sometimes used to convey precisely that one does *not* believe that *p*, strictly speaking. (Sutton 2007: 64)

In this case it is plausible that the subject does not believe categorically where the post office is, but believes where it likely is. Here the subject could have just as easily asserted the hedged, probabilistic claim, and would confirm that this is exactly what was meant if asked to clarify his statement. But often we assert the stronger, unhedged belief, and we judge that we do so properly, and with justification.

Consider Postie the Postmaster on the day of retirement from his forty-year career at the post office. Suppose one were to ask him where the post office is, to which he replies, “I believe that it is a mile down that way”. While this might be somewhat odd to assert rather than merely asserting that, “It is a mile that way” (Postie being an eccentric fellow), it could nevertheless be true and justified. If one were to ask Postie to clarify his assertion by asking further, “Do you think it is merely *likely* that is where the office is, or do you believe it categorically?” he would likely assume that one was joking. “No, I am certain that is where it is. I have worked there for forty years!” he might reply, somewhat annoyed. Postie’s belief about the location of the office is surely justified. This remains the case even if one minute before being asked where the post office is, it is burned to the ground, thus making Postie’s belief about where the office is false, as the office would have been destroyed. Since in that case his belief would be false, it could not be known; it would be justified nevertheless. Justification is not knowledge.

IV. The Main Objection

Towards the end of his defence Sutton discusses an objection that, if sound, devastates the first three of his arguments (cf. Sutton 2007: 59-60), of which I have discussed the first. And given that I have provided sufficient materials above to undercut his fourth argument (the Posterior Evaluation argument), as well as an argument from strict versus loose ways of speaking, the objection becomes even more pressing. The objection is that a weaker norm-governing assertion will do all the work that the Knowledge Rule (K-rule) is supposed to. That is, if a weaker norm can handle the quasi-Moorean Paradox cases, the Lottery Cases, and the cases that involve challenging or

rebuking assertions, then the K-rule is unmotivated. This would be devastating for Sutton as the K-rule is an explicit premise of his assertion argument. Sutton considers that the best and simplest candidate to possibly replace the K-rule is:

(The J-rule): One must: assert p only if one has a justified belief that p (Sutton 2007: 61).

Sutton thinks that the J-rule would handle the relevant phenomena that the K-rule is supposed to, plus (and I think this is a virtue of the J-rule) it implies that certain false assertions are warranted. However, Sutton rejects the J-rule because he thinks it falls prey to an anti-complexity argument, as I will explain below. However, despite Sutton's claims, it is not clear how the J-rule accounts for the data that it is supposed to.

One might think that a J-Rule handles all the cases it is meant to in the following ways: first, if someone asserts ' p and I don't know that p ', one might say that that is inappropriate since by asserting p the subject implies he is justified in believing it, but then implicitly denies this. However, if justification is distinct from knowledge, then one can have justified beliefs that do not amount to knowledge. Therefore, to admit one's ignorance is not to say one is not justified. There is no contradiction.

So, although the J-rule makes sense of the incongruence of asserting ' p and I am not justified in believing that p ', it does not account for the tension in asserting ' p and I do not know that p '. So it seems that the stronger K-rule has an advantage over the J-rule.

It is important to distinguish two kinds of Moorean data:

- 1) Assert: ' p but I don't know that p ' (because the subject has no justification to believe that p)

- 2) Assert: '*p* but I don't know that *p*' (even though the subject has a highly justified belief that *p*)

It is the latter datum that the advocate of the J-rule must explain, since this is where the J-rule and the K-rule diverge in what they are able to account for. The J-rule can accommodate only the former datum.

A similar point holds in the challenges cases, i.e. 'How do you know?' and 'You don't know that!' (cf. Williamson 2000: 249-252) In making the challenge, e.g., in a Lottery Paradox case, the challenger may recognize that there is a huge amount of justification for *S* to believe that *p* but still feel able to properly challenge *S*'s assertion that *p* by questioning whether *S* knows that which they are asserting. Lottery Paradox cases are supposed to be paradigm cases whereby a subject's belief that his ticket will lose is extremely probable on the evidence, and yet he should not assert that his ticket will lose. Here we have justification without knowledge, and indeed a case where one cannot explain the data with the J-rule since the J-rule predicts that one *can* assert that one will lose the lottery. But intuitively this is wrong, since such assertions, while justified, are inappropriate nevertheless.

In light of this, it seems that the best and simplest candidate to replace the K-rule (and the J-rule) is:

(The JBK-rule) One must: assert *p* only if one has justification to believe that one knows that *p*.^{xvixvi}

Such a principle explains what is wrong in the quasi-Moorean and Lottery cases, as well as why assertions can be challenged by challenging the subject's knowledge. In

each case, the subject lacks justification to believe that he knows what he is asserting, e.g. in quasi-Moorean paradox cases, it is inappropriate to assert ‘ p and I don’t know that p ’ because by asserting p one is representing oneself as having justification to believe that one knows that p , which is then denied; in Lottery Paradox cases, it is improper to assert that one knows one’s ticket is a loser, no matter how likely, because one lacks justification to believe that one knows one has a losing ticket; finally, challenging assertions by saying ‘How do you know?’ and ‘You don’t know that!’ are appropriate because one is challenging the asserter’s justification to believe that she knows what she is asserting.

Therefore, unlike the J-rule, the JBK-rule can handle *all* the cases that the K-rule can; unlike the K-rule, however, it allows that subjects are sometimes warranted in asserting what they do not know, e.g. in Gettier cases, as well as cases of justified but false belief. The JBK-rule, therefore, shares all the explanatory virtues of both the K-rule and the J-rule, without sharing their vices.

Why then does Sutton accept the K-rule over rival rules of assertion, such as the JBK-rule? It is because he accepts a version of Williamson’s complexity argument that he thinks that the K-rule has a presumption in its favour over other rules since it provides a simpler account of warranted assertion.

The J-rule is allegedly more complex than the K-Rule because justification is allegedly parasitic on knowledge. Sutton would therefore think that the JBK-rule is more complex than the K-rule, since it relies on the concepts of both epistemic justification and knowledge. Recall from the beginning of this discussion where I noted that Sutton says that the reason we are inclined to call Gettier cases instances of *justified* belief is because

the subject *would* have knowledge if it had not been for his bad luck (e.g. if the subject had not been in fake barn country, he would have known that he was looking at a barn). As Sutton says, “Justification in the relevant sense is perhaps a disjunctive concept – it is knowledge or would-be knowledge.” I suggested above that perhaps justification can be spelled out independently of knowledge, e.g. in terms essentially tied to evidence and not to knowledge. If such an account can be given, then the J-rule is just as simple as the K-rule. However, now we see that the best rule of assertion, the JBK-rule, does invoke the notion of justification, as well as knowledge. Perhaps also justification is parasitic on knowledge. What then follows for the thesis that justification *just is* knowledge?

If knowledge is the primitive in terms of which all other epistemic concepts must be understood, does that count decisively in favour of the K-rule? Of course not. A complexity argument is not a definitive argument for anything. Rather, simplicity provides a (defeasible) *presumption* in favour of one theory over the other. A theory needs to have more virtues than mere simplicity, and sometimes therefore, the more complex theory is to be preferred over the simpler one. At least one other essential criterion a philosophical principle must meet is to deal adequately with the relevant cases.

As Williamson himself notes, “Much of the evidence for the knowledge account [or any account of assertion, for that matter] comes from the ordinary practice of assertion” (Williamson 2000: 243). It is for this reason that Williamson endorses the K-rule over the admittedly more simple truth rule of assertion (one must: assert *p* only if *p* is true) (See Williamson 2000: 242; 262). The Truth Rule (T-rule) does not explain what is wrong with the Lottery cases, or cases where a subject asserts that *p* truly but has absolutely no evidence that *p*.

Advocates of the T-rule, such as Matthew Weiner (2005) concede as much, which is why he is forced to explain the inappropriateness of such assertions by adding to the T-rule general Gricean norms of conversation.^{xvii} For example, if a subject asserts that *p* truly but has absolutely no evidence that *p*, the assertion may violate a Cooperative Principle (that one's utterances have a point) or a maxim of Quantity (to not make a contribution more informative than required), or some other pragmatic norm, by creating an implicature that he has reason for what he asserts. Even if it is true that the T-rule, together with the Gricean norms can explain the data that motivated the knowledge account, it is surely preferable, all things being equal, if we can explain the data using a single principle. It is here that the JBK has the advantage over both the knowledge and the truth accounts of assertion by explaining all the data itself, without having to appeal to Gricean norms. Plus, by not being factive, the JBK account can also allow for the warranted assertions of falsehoods, as well as warranted assertions in Gettier situations.

But Weiner argues that there is additional important data that only his T-rule can handle. He argues that there are cases where the speaker's justification is too low for knowledge (and presumably, therefore, the subject is not justified in believing that he knows, which the JBK-rule requires), but still the assertion is permissible. Here are Weiner's examples:

Take the following case of prediction: Captain Jack Aubrey has had long experience of naval combat against the French Navy. He and young Lieutenant Pullings have been watching French ships maneuver off Mauritius all day. At two p.m., Aubrey says to Pullings,

- (3) The French will wait until nightfall to attack.

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

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

Holmes, at this point, has not found any evidence (in the criminal rather than epistemological sense) incriminating Professor Moriarty, but he is sticking his neck out based on a sense of what Moriarty's crimes are like[...]These cases pose a problem for the knowledge account [and the JBK account, if cogent]. They seem to be proper assertions in the absence of knowledge; the speakers' grounds for belief fall short of those necessary for knowledge, yet the assertions do not seem wrong or odd. (Weiner 2005: 230-1)

However, *pace* Weiner, it is unclear why (3) and (4) are not ascriptions of knowledge, and so are perfectly acceptable according to the K-rule and the JBK-rule. After all, in the first case Aubrey is described as a senior officer with a wealth of naval experience, including combat against the French. Also, he has been watching the ships all day, and is making a prediction about what they will do in only a couple hours time. Therefore, Aubrey has lots of inductive evidence that he can know such things in such circumstances, and that he is in such circumstances. Short of an unreasonable inductive scepticism, it is unclear why one would not say that Aubrey knows what the French will soon do. If, as it happens, his assertion turns out to be false, note that it is still perfectly proper on the JBK account, given that Aubrey is justified in believing that he knows what the French will soon do, based on his day of careful observation, his expertise, and his years of Naval combat experience, an intuition to which the truth and knowledge accounts cannot do justice.

As for Holmes, on at least one reading of the case, it also seems unproblematic that he knows, or is at least justified in believing that he knows, what he asserts. Holmes has lots of experience dealing with Moriarty and his crimes. He therefore has inductive evidence of what Moriarty's crimes are like. Professor Moriarty is Holmes' nemesis after

all; Holmes is the world's greatest detective, an insightful master of abduction. Given a reasonable Fallibilism, it is unclear why Holmes is not at least justified in believing that he knows what he is asserting. If, on the other hand, Holmes is merely making his assertion on the basis of groundless hunch, or for no reason at all, then it is difficult to see that the assertion is warranted in any reasonable sense, simply in virtue of being true. Therefore, insofar as assertions (3) and (4) are intuitively warranted, the JBK-rule is able to explain why this is so. The T-rule is of no advantage.^{xviii}

So, given that the T-rule does not adequately explain what is wrong with the Lottery cases, or cases where a subject asserts that p truly but has absolutely no evidence that p , the more complex K-rule is motivated via these considerations. However, as I have been stressing, our ordinary practice of assertion includes not only the Mooreian and Lottery cases, but also Gettier cases and cases of justified false belief. If possible, therefore, we need a rule of assertion that gives a unified treatment of all these cases. The JBK-rule provides such a treatment.

It is worth considering one last rival rule of assertion that has recently been offered as providing a unified treatment of all the relevant cases. This is Jennifer Lackey's *Reasonable to Believe Norm of Assertion*:

(RTBNA): "One should assert that p only if (i) it is reasonable to believe that p , and (ii) if one asserted that p , one would assert that p at least in part because it is reasonable for one to believe that p ." (Lackey 2007: 608)

The first thing to note is that while this rule explains the permissibility of some false well-evidenced assertions and Gettier assertions, on its own it does not prohibit the Moorean, Lottery or challenge to knowledge assertions that we began with. Lackey

acknowledges as much, which is why, like Weiner, she pairs her rule of assertion up with a pragmatic Gricean rule of conversation. Her preferred maxim is:

Not Misleading Norm of Assertion (NMNA**): “S should assert that *p* in context C only if it is not reasonable for S to believe that the assertion that *p* will be misleading in C relative to the purposes of the exchange in question.” (Lackey 2007: 617)

Even if this pair of rules does explain the relevant data, again it is surely preferable, all things being equal, if we can explain the data using a single principle. It is here that the JBK-rule has the advantage over the RTBNA + NMNA** account of assertion by explaining all the data itself, without having to appeal to Gricean norms.

But like Weiner, Lackey too argues that there is an additional important phenomenon that only her pair of rules can handle. She argues that there are cases that thwart knowledge based accounts of assertion (and therefore principles like the JBK-rule that make essential reference to knowledge). She calls these instances of Selfless Assertion (see Lackey 2007, especially section 2). Summarizing what her putative counterexamples have in common, Lackey writes:

There are three central components to this phenomenon: first, a subject, for purely non-epistemic reasons, does not believe (and hence does not know) that *p* [they also therefore lacks justification to believe that they know that *p*]; second, despite this lack of belief, the subject is aware that *p* is very well supported by all the available evidence; and, third, because if this, the subject asserts that *p* without believing and, hence, without knowing that *p*. The combination of these three features has the following result: a subject offers an assertion in the absence of knowledge [and lack of justification for believing that they know that *p*] and is not subject to criticism in any relevant sense. Indeed, in all of the above cases [the three Lackey introduces], the subject in question, *qua* asserter, is *appropriately subject to praise*. (original emphasis) (Lackey 2007: 599)

While the phenomenon that Lackey discusses is an interesting one, it is not clear that it is a counterexample to knowledge-related norms of assertion. This is because it is not clear that what such speakers would be uttering in the absence of belief is a genuine assertion. Lackey briefly considers and rejects such an objection on the grounds that such a position is committed to holding that a necessary condition of S's asserting that *p* is that S believes that *p*, which she regards as absurd (Lackey 2007: 600-1). The alleged absurdity stems from the fact that this position entails that it is not possible to assert a lie, which Lackey regards as a paradigmatic instance of an (unwarranted) assertion (Ibid.).

Nevertheless, it is unclear why a lie should be thought of as an assertion at all. Rather, it seems a lie should be best thought of as a kind of speech act that is *passed off* as an assertion with the intent to deceive. Part of the reason a person is misled by a lie is that they mistake what is said for an assertion; the liar merely purports to assert that *p* without actually asserting that *p*. Assertion, like a belief, involves commitment in a way that lying does not. Assertion, like belief, involves a kind of sincerity in a way that lying does not. Genuine declarative mood is present when an assertion is made, not mimicked declarative tone like when a lie is told. Given these parallels, it is most natural to think of an assertion as the outward manifestation or the making public of a belief. Just as news readers and actors do not necessarily believe the lines they read simply in virtue of reading them, similarly they do not necessarily make assertions simply in virtue of uttering declarative sentences in what seems to be a declarative tone. Such utterances can be reporting what *others* assert, without the speakers committing themselves to what is said.

A key feature of making a genuine, rather than a merely apparent assertion, therefore, is that the subject sincerely takes things to be as they describe them as being, i.e. they believe what they assert. Since an utterance cannot not be an assertion without belief, Lackey's counterexamples are off point. Accordingly, since they fail to be cases of assertion, they are not cases a theory of assertion needs to be able to explain. The JBK-rule remains unscathed.

Having considered the relevant data, it is clear that the JBK-rule gains superiority over the K-rule and its rivals by handling all the relevant cases the K-rule does, *plus*, it allows for warranted false assertions and warranted assertions in Gettier cases, very intuitive cases that the K-rule prohibits. In short, since the JBK-rule does what the K-rule is supposed to do, and more, by handling a wider range of cases in the simplest possible way without having to appeal to Gricean pragmatic principles, it should be preferred over the K-rule, even if it turns out to be true that knowledge is more conceptually fundamental than justification.

Without the K-rule, by Sutton's own admission, three out of four of his arguments for the thesis that justification is knowledge fail. By Sutton's own admission, the best (and single?) reason to accept the K-rule over justification-based norms of assertion is that it is more simple than its rivals. It is not obvious that it is more simple than principles that invoke notions of justification that are not parasitic on knowledge, e.g. justification supervenes on evidence, or justification is a matter of reasonableness (understood independently of knowledge). But even if the K-rule is simpler than its rivals, there is a principled reason to accept the JBK-rule over the K-rule, viz. that it is intuitive that there is nothing wrong with a false assertion over and above the false belief

that it expresses. We judge that we can be warranted in making false assertions. Also, we judge that our Gettiered assertions are warranted. In short, the JBK-rule of assertion explains the phenomena that motivate the K-rule and its rivals, while better handling a wider range of cases.

V. Conclusion

Justification is not Warrant in Plantinga's sense, nor is it mere blamelessness, but there do seem to be deontological and evaluative forms of justification that appear to amount to a concept of epistemic justification that might be elucidated in term of reasonableness and/or evidence.

As for the Posterior Evaluation argument, I hope to have shown that there is an intermediate case between knowledge and the merely pragmatic case. That is the case wherein one should have formed a belief because that is what the evidence supported (even though the evidence was misleading). The existence of the intermediate case, coupled with the arguments I gave above in support of the JBK-rule over the K-rule, shows that justification is not knowledge.^{xix}

Bibliography:

- Alston, William. 1989a. "Concepts of Epistemic Justification." Reprinted in *Epistemic Justification*. Ithaca: Cornell University Press, 81-114.
- Alston, William. 1989b. "The Deontological Conception of Justification." Reprinted in *Epistemic Justification*. Ithaca: Cornell University Press, 115-152.
- Dodd, Dylan. 2007. "Why Williamson Should Be A Sceptic." *The Philosophical Quarterly* 57: 635-649.
- Goldman, Alvin. 1992a. "What Is Justified Belief?" Reprinted in *Liaisons: Philosophy Meets The Cognitive and Social Sciences*. Cambridge, Mass.: The MIT Press, 105-126.
- Goldman, Alvin. 1992b. "Strong and Weak Justification." Reprinted in *Liaisons: Philosophy Meets The Cognitive and Social Sciences*. Cambridge, Mass.: The MIT Press, 127-141.
- Greco, John. 1999. "Agent Reliabilism." *Philosophical Perspectives* 13: 273-296.
- Grice, Paul. 1989. "Logic and Conversation." Reprinted in *Studies in the Way of Words*. Cambridge, MA: Harvard University Press, 22-40.
- Kvanvig, Jonathan L. and Menzel, Christopher. 1990. "The Basic Notion of Justification." *Philosophical Studies* 59: 235-261.
- Kvanvig, Jonathan. Forthcoming. "Assertion, Knowledge, and Lotteries." In Patrick Greenough & Duncan Pritchard (eds.), *Williamson on Knowledge*. Oxford: Oxford University Press.
- Lackey, Jennifer. 2007. "Norms of Assertion." *Nous* 41: 594-626.
- Plantinga, Alvin. 1993. *Warrant: The Current Debate*. Oxford: Oxford University Press.
- Pritchard, Duncan. 2005. *Epistemic Luck*. Oxford: Clarendon Press.
- Pritchard, Duncan. 2007. "Anti-Luck Epistemology." *Synthese* 158: 277-297.
- Sosa, Ernest. 2007. *A Virtue Epistemology*. Oxford: Oxford University Press.
- Sutton, Jonathan. 2005. "Stick To What You Know." *Nous* 39: 359-396.
- Sutton, Jonathan. 2007. *Without Justification*. Cambridge, Mass.: The MIT Press.

Turri, John. Forthcoming. "On the Relationship Between Propositional and Doxastic Justification." *Philosophy and Phenomenological Research*.

Weiner, Matthew. 2005. "Must We Know What We Say?" *The Philosophical Review* 114: 227-251.

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

ⁱ See Jonathan Sutton 2005 and 2007.

ⁱⁱ See Pritchard 2005 for a discussion of the kinds of luck (in)compatible with knowledge.

ⁱⁱⁱ See Plantinga 1993.

^{iv} For example, see Goldman 1992b for a notion he calls 'weak justification' which he understands in terms of blamelessness.

^v See Alston 1989b.

^{vi} Why does Sutton consider *these* five conceptions of justification, rather than Reliabilist or other paradigmatically externalist conceptions of epistemic justification? He does not say. In fairness to Sutton, the conceptions of justification he *does* discuss have all been influential and central ways of thinking of epistemic justification in contemporary analytic epistemology; for example, see Alston 1989a. In any case, I will show that his argument fails.

^{vii} For similar accounts that understand propositional justification in terms of doxastic justification, see for example Goldman 1992a; Greco 1999 and Sosa 2007.

^{viii} Of course this is too broad to be terribly informative as it stands. A detailed account is needed; virtue epistemologists attempt to provide such details. (see note 7) Nevertheless, it does seem an open possibility that whether or not a belief is justified depends upon the person who holds it and facts about them, such as their cognitive character.

^{ix} For a defence of this thesis, see Williamson 2000.

^x As Sutton rightly notes, "One objection is so potentially devastating that it deserves separate treatment. If successful, the objection demolishes three of my four arguments; all but the posterior evaluation argument, which is arguably the least intuitively compelling of my arguments in any case. [...] The objection is that a weaker norm governing assertion will account for the phenomena that Williamson cites in support of his knowledge rule" Sutton 2007: 59-60. After critiquing the Posterior Evaluation argument below, I go on to press just such an objection.

^{xi} See Sutton 2007: 63-67 for an elaboration of this distinction.

^{xii} See Sutton 2007 ch. 4, especially pp.128-147. This is weaker than Williamson's well-known view of evidence, since according to Williamson, not only is all evidence knowledge, but also, all knowledge is evidence; see Williamson 2000, ch. 9.

^{xiii} See Sutton 2007: 134-143 where he attempts to make this case.

^{xiv} Knowledge-based theories of evidence are not without their own problems, of course. For example, for an argument that such accounts are committed to scepticism, see Dodd 2007.

^{xv} Williamson himself considers and rejects a similar principle, except casts it in terms of 'rational' rather than justified belief, as well as in terms of doxastic justification (rationality) rather than in terms of propositional justification (rationality) as I do here. See Williamson 2000: 260-63 for a discussion of this similar principle as well as other belief-based accounts of assertion.

^{xvi} Jonathan Kvanvig *forthcoming* defends a norm of assertion he calls 'the justification account of assertion' which is similar in many ways to the principle I am advancing here. There are salient differences, however. The notion of justification that Kvanvig is using is more demanding than the more ordinary sort of justification that Sutton, Williamson and I have in mind. On Kvanvig's account, justification is such that it is sufficient for knowledge in the presence of ungettiered true belief. Further, as Kvanvig writes, "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" (Kvanvig *forthcoming*: 23). In addition to employing this more demanding notion of justification, Kvanvig's

arguments for the correctness of his norm of assertion rely on the correctness of the defeasibility theory of knowledge, i.e. that knowledge is undefeated (epistemically internalist) justified true belief.

Given these features of Kvanvig's account, the principle I offer is superior in that:

- i) it remains neutral on the correct account of knowledge
- ii) it remains neutral on the connection between justification and knowledge, e.g. if justification is necessary for knowledge, and if so, exactly how much is required
- iii) it employs a more natural way of thinking about justification that is not committed to holding that it is a *necessary condition* of being justified in believing that *p simpliciter*, given that one knows what one believes, that one thereby is justified in believing that one knows that *p*. For example, on the account of justification I am advocating, one can be epistemically justified in believing that one's single ticket in a fair lottery is loser on the basis of the probabilistic evidence available to one. Kvanvig is committed to denying this, since a belief justified in this sense is not sufficient for knowledge, even if true and ungettiered.

In short, as I will argue, the principle I propose delivers the same explanatory results that Kvanvig's does, but with fewer contentious theoretical commitments. It should therefore be preferred on that basis.

^{xvii} These are pragmatic rules first outlined by Paul Grice, which he argues are necessary for successfully communication. See Grice 1989.

^{xviii} If Aubrey and Holmes know, or are at least justified in believing that they know, what they assert on the basis of fallible evidence, why then does one fail to know in the Lottery case, given that such subject's have overwhelming fallible evidence? One possibility is to appeal to an anti-luck modal principle as necessary for knowledge, such as:

Safety: it is a necessary condition of knowing that *p* that in most nearby possible worlds in which *S* continues to form her belief about the target proposition in the same way as in the actual world the belief continues to be true. (Pritchard 2007: 281)

So understood, it is clear that lottery beliefs are not Safe: if one believes that a ticket has lost, there is a *very* nearby possible world where one has won but believes falsely that one has lost on the basis of the same probabilistic evidence (of course this world is modally close, not statistically likely – very little would have to be different than it actually is for that lotto winning world to be actual, e.g. the coloured lottery balls would only have to be arranged slightly differently...)

Aubrey and Holmes, by contrast, given their expertise and inductive evidence, do hold beliefs that are Safe, and thus are known if true: there is no nearby possible world where they believe what they do, on the basis they do, but their beliefs are false. Unlike the lottery case, such beliefs are non-luckily true.

^{xix} Many thanks to an Auckland audience at the Australasian Association of Philosophy Annual Conference 2007 (New Zealand Division). Thanks especially to Patrick Greenough, Stephen Hetherington, Vickie Madison, Declan Smithies, Paul Snowdon, Lee Walters, and two anonymous referees for helpful written comments on early drafts of this paper.