Some Thoughts on the JK-Rule[1]

Martin Smith

University of Glasgow

In 'The normative role of knowledge' (2012), Declan Smithies defends a 'JK-rule' for belief: One has justification to believe that P iff one has justification to believe that one is in a position to know that P. Similar claims have been defended by others (Huemer, 2007, Reynolds, 2013). In this paper, I shall argue that the JK-rule is false. The standard and familiar way of arguing against putative rules for belief or assertion is, of course, to describe putative counterexamples. My argument, though, won't be like this – indeed I doubt that there are any intuitively compelling counterexamples to the JK-rule. Nevertheless, the claim that there *are* counterexamples to the JK-rule can, I think, be given something approaching a formal proof. My primary aim here is to sketch this proof. I will briefly consider some broader implications for how we ought to think about the epistemic standards governing belief and assertion.

I. RULES FOR ASSERTION AND BELIEF

According to the 'K-rule' for assertion, one has justification to assert that P iff one is in a position to know that P[2]. The 'justification' here should be read as *epistemic* rather than practical or prudential. The K-rule is intended to articulate the epistemic standards that govern assertion – but meeting the epistemic standards for assertion is obviously compatible with there being overwhelming practical or other considerations compelling one to hold one's tongue.

Defenders of the K-rule for assertion argue that it provides the best overall explanation for our intuitive verdicts about a range of hypothetical cases. While it's undeniable that the K-rule does account for a broad range of intuitive verdicts, certain well known cases – including Gettier cases – pose a prima facie problem. Suppose Sarah and Julia are, unbeknownst to them, driving through barn facade county – the locals have erected a number of convincing paper mache barn facades that can easily be mistaken for the real thing. As they approach the one true barn in the county, Sarah remarks to Julia 'There's a barn up ahead. Let's stop and take a look'.

---

[1] This paper was presented at the Northern Institute of Philosophy, University of Aberdeen in October 2011. Thanks to all of those who participated on that occasion. Particular thanks to Dylan Dodd, Andreas Fjellstad, Michael Lynch, Aidan McGlynn, Grant Reaber and Crispin Wright.

[2] The general idea that knowledge serves as the norm or standard of assertion has been endorsed by Williamson (2000, chap. 11), DeRose (2002), Hawthorne (2003) and Hawthorne and Stanley (2008) amongst others – though more precise formulations of this claim tend to vary from philosopher to philosopher. The formulation used here is that given by Smithies (2012).

The conventional verdict about this case is that Sarah does not know, and is in no position to know, that there is a barn up ahead, in which case the K-rule predicts that her assertion is unjustified. On the face of it, this would appear to be the wrong verdict. Under the circumstances, Sarah *would* seem to be justified in asserting what she does. Surely any of us would be inclined to make the same assertion if placed in Sarah's position. Defenders of the K-rule, however, have a somewhat standard response to these cases – although Sarah's assertion is unjustified, she is nevertheless *excusable* for having made it. Unjustified assertions can, in circumstances like Sarah's, be excused (Williamson, 2000, section 11.4, DeRose, 2002, section 2.1, footnote 23, Hawthorne and Stanley, 2008, introduction).

The distinction between justified assertion and unjustified but excusable assertion is surely a legitimate one, as evidenced by cases like the following: Suppose Malcolm is slipped a drug that undetectably dulls his analytical skills before surveying some reasonably complicated data concerning the efficacy of a recent advertising campaign. While the data strongly suggest that the campaign was ineffective, Malcolm, as a result of the drug's effect, leaps to the very opposite conclusion. If Malcolm were to then go around asserting 'The campaign was a great success!', his assertions would clearly be unjustified. Nevertheless, under the circumstances, he would surely be excusable for having made them. While there are possible cases of unjustified but excusable assertion, it remains dubious, however, whether Sarah's case, and others like it, belong in this category. Indeed, Sarah's case seems intuitively very different to Malcolm's – Sarah, unlike Malcolm, doesn't seem to have done anything for which an excuse is needed.

In 'The normative role of knowledge' Declan Smithies argues that Gettier cases give us prima facie reason to doubt the K-rule for assertion and to experiment with alternatives such as the JK-rule: One has justification to assert that P iff one has justification to believe that one is in a position to know that P (Smithies, 2012, section 2). According to the JK-rule, one can meet the epistemic standards for asserting that P even if one is not in a position to know that P – all that is required is that one have justification to believe that one is in a position to know that P.

The JK-rule offers a different verdict about Sarah's case. Sarah presumably has justification to believe, albeit falsely, that she isn't surrounded by barn facades and, thus, has justification to believe, albeit falsely, that she's in a position to know that there's a barn up ahead. As well as offering an intuitive verdict in Gettier cases, the JK-rule can also, according to Smithies, accommodate the various intuitive verdicts that have been adduced in favour of the K-rule – and, as such, enjoys an overall explanatory advantage.

My interest in this paper is not primarily with the JK-rule for assertion but, rather, with a corresponding rule for *belief* that Smithies also endorses: One has justification to believe that P iff one has justification to believe that one is in a position to know that P.

(Smithies, 2012, section 5). Others have defended claims in the near vicinity, such as Huemer (2007) and Reynolds (2013)[3].

Adopting the JK-rule for assertion puts us under some pressure to adopt the JK-rule for belief. It is natural to think that assertion and belief should be subject to the very same standards of justification. After all, it is plausible that assertion serves as the outward expression of belief and belief as the inward correlate of assertion. Smithies expresses the idea in this way: 'The very nature of assertion is to be understood in terms of its role in the expression of belief' (Smithies, 2012, section 3, see also Williamson, 2000, pp238, pp255-256). In this paper I shall argue that the JK-rule for belief is untenable. This will, in turn, place the JK-rule for assertion under threat, via this plausible belief-assertion link. I'll have a little more to say about assertion in the final section.

One very familiar way of arguing against putative rules for assertion or belief is by proposing putative counterexamples – cases in which our intuitions seem at variance with what the rule predicts. One way to argue against the JK-rule, then, would be by describing a case in which one intuitively has justification to believe a proposition but intuitively lacks justification to believe that one is in a position to know it. Lottery cases are one sort of case that might fit the bill. Suppose I hold a single ticket in a fair lottery with 1000 tickets and a single guaranteed winner. Prior to the draw it seems intuitive that I have justification to believe that my ticket will lose. It's clear, though, that I am in no position to *know* that my ticket will lose and that I lack justification to believe that I am.

Smithies anticipates this kind of objection to the JK-rule and argues at some length against this reaction to lottery cases. He argues, in particular, that one cannot have justification to believe that a ticket will lose a fair lottery, purely on the basis of the odds involved – one can, at best, have justification for investing a high level of confidence in this proposition (Smithies, 2012, section 5). I am broadly sympathetic to what Smithies says here – and I've defended elsewhere the view that one lacks justification to believe such lottery propositions, their high probability notwithstanding (Smith, 2010). My argument against the JK-rule, in any case, won't rely upon any putative counterexamples to the rule – at least, not in a direct way.

---

[3] Huemer defends the principle that, if one believes that P then one is rationally committed to the view that one's belief qualifies as knowledge (Huemer, 2007, section 4). From here it is a relatively small step to the JK-rule – the only supplementary principle that is needed is something along the following lines: If one has justification to believe that P and believing that P would rationally commit one to believing that Q, then one has justification to believe that Q.

Reynolds (2013) defends the view that if one justifiably believes that P then it appears to one that one knows that P. It's plausible that, if it appears to one that one knows that P, then one has justification to believe that one knows that P. This still doesn't get us quite to the JK-rule, as Reynolds' claim pertains to *doxastic*, rather than propositional, justification. A further principle is needed – if it is a requirement upon doxastic justification that one have justification to believe that one knows that P, then it's a requirement on propositional justification that one have justification to believe that one is in a position to know that P. The forthcoming arguments against the JK-rule will also weigh against Huemer's and Reynolds' claims – though they will be cushioned, to an extent, by these bridging principles.

II. AGAINST THE KK-THESIS

One thing that we can immediately observe is that the JK-rule for belief, unlike the JK-rule for assertion, is *recursive*, in that it can be reapplied to the results of previous applications. If one has justification to believe that P (JP) then, by the JK-rule, one must have justification to believe that one is in a position to know that P (JKP). But if one has justification to believe that one is in a position to know that P (JKP) then, by the JK-rule, one must have justification to believe that one is in a position to know that one is in a position to know that P (JKKP) and so on... In general, we have it that $JP \supset JK^xP$ for any positive integer x.

Here is an immediate worry that we might have about this: At some point in the infinite sequence P, KP, KKP, KKKP...we will, presumably, reach propositions that can no longer be entertained, let alone believed, by any normal human subject. And yet, according to the JK-rule, one must have justification to believe such propositions whenever one has justification to believe P. There is, undoubtedly, something unsettling about this – but rejecting the JK-rule on these grounds alone would, I think, be hasty.

As Smithies himself emphasises, the idea that one can have justification to believe a proposition that one is psychologically incapable of believing is not obviously incoherent (Smithies, 2012, section 2). On one sort of view, what one has justification to believe is purely a matter of what one's evidence supports – and one's evidence surely could stand in the support relation to propositions that one is, for one reason or another, incapable of actually believing. If the JK-rule is correct, then this kind of situation will be far more widespread than we might have expected (at any given moment, one will have justification to believe an infinite number of propositions that one is incapable of believing) – but this is not obviously incoherent either. In any case, I won't pursue this discussion further here. Even if this immediate worry can be answered, a further and, I think, more troubling worry looms.

The JK-rule effectively generates an infinite sequence of necessary conditions on having justification to believe a proposition. If these conditions all turned out to be logically equivalent – if JKP, JKKP, JKKKP etc. all turned out to be logically equivalent – then perhaps this need not be any cause for concern. I think that there are good reasons, however, for supposing that each of these generated conditions is, in fact, logically distinct. Most contemporary epistemologists reject the so-called 'KK'-thesis according to which, if one is in a position to know that P, then one is in a position to know that one is in a position to know that P. For most epistemologists, second order knowledge is, in some way, more *epistemically demanding* than first order knowledge – and, in general, n+1th order knowledge is more epistemically demanding than nth order knowledge. As Williamson puts it '...iterating knowledge is hard, and each iteration adds a layer of difficulty' (Williamson, 2000, pp122, see also Nozick, 1981, Dretske, 2004). But if each proposition in the sequence KP, KKP, KKKP... is logically stronger than its predecessor, then it's natural to think that every proposition in the sequence JKP, JKKP, JKKKP... is also logically stronger than its predecessor. If this is right, though, then the JK-rule not only generates an infinite sequence of necessary conditions on having justification to believe a proposition – it generates an infinite sequence of necessary conditions *of ever increasing logical strength*. And this, surely, is something that should give us pause.

One way of capturing the idea that each iteration of K places one under greater epistemic strain is by exploiting a *safety* condition upon knowledge. A number of epistemologists have been attracted to the idea that knowledge requires a certain safety from error – if a belief is to qualify as knowledge then it must be the case that it could not *easily* have been false, given the way that it was formed and the evidence upon which it was based. There are various different ways of making this kind of requirement more precise, most of which will serve for present purposes. For ease, I will work here with the following: Let Δ be the total body of evidence that one possesses. Δ can be thought of as a set of propositions or as a set of experiences or other mental states or, indeed, in any of a range of further ways. The safety condition can be formulated as follows: In order for one to be in a position to know a proposition P, it must be the case that in all very similar possible worlds in which Δ is one's total evidence, P is true. According to this condition, the possession of Δ must ensure the truth of P throughout the local modal neighbourhood – otherwise it would leave one too exposed to error to qualify as an adequate basis for knowledge.

What then, is required in order for one to be in a position to know that one is in a position to know that P? By the safety condition, it must be the case that in all very similar worlds in which Δ is one's total evidence, one is in a position to know that P. Then, by another application of the safety condition, in each of *these* worlds, it must be true that, in all similar worlds in which Δ is one's total evidence, P is true. Let a Δ-world be a world in which Δ is one's total evidence. In order for one to be in a position to know that P, P must be true in all Δ-worlds that are very similar. In order for one to be in a position to know that one is in a position to know that P, P must be true in all Δ-worlds that are very similar to Δ-worlds that are very similar. This will, in general, be an *expanded* set of worlds. Close similarity is not a transitive relation. Just because world w_1 is very similar to the actual world and world w_2 is very similar to world w_1 it does not follow that world w_2 is very similar to the actual world. In order for one to be in a position to know that one is in a position to know that P, Δ would have to be *stronger* evidence for P, at least along one dimension – it would have to guarantee the truth of P across an expanded set of worlds.

Say that Δ is safe evidence for P just in case P is true in all Δ-worlds that are very similar to the actual world. Say that Δ is *safely* safe evidence for P just in case P is true in all Δ-worlds that are very similar to Δ-worlds that are very similar to the actual world. If being in a position to know something requires the possession of safe evidence but not, in general, the possession of safely safe evidence, then we will have the required counterexamples to the KK-thesis. We need only imagine a situation in which P is true in all very similar Δ-worlds and false in some Δ-worlds that can be reached by two close similarity steps. In this case Δ will be safe evidence for P but not safely safe evidence for P, and a subject with evidence Δ may be in a position to know that P but will not be in a position to know that he's in a position to know that P.

There is, of course, much in this argument that is discussable. One might dispute, for instance, whether there really is a safety condition on knowledge. And even one broadly sympathetic to safety could certainly take issue with the specific formulation that I've used. I think that there are very good reasons for positing a safety condition on knowledge, and one that can be formulated in such a way as to sustain this kind of argument – but I won't pursue these points here. For when it comes to rejecting the KK-thesis, there is another, stronger, argument available – and this is the argument outlined by Timothy Williamson in chapter 5 of *Knowledge and Its Limits*.

The following is a variation on Williamson's argument: Suppose Mr. Magoo is staring out of his window at a tree on the other side of his garden that is exactly 200cm tall. Suppose that, when estimating heights at this distance, Magoo is accurate to within a margin of about 10cm, but can't discriminate more finely than that – his eyesight is too poor. Presumably Magoo is in a position to know a range of propositions about the height of this tree – he's in a position to know, for instance, that the tree is taller than 1cm, that it's taller than 100cm etc. Plausibly, though, he's in no position to know that the tree is taller than *199cm* – his eyesight just isn't good enough for that. Even though this proposition happens to be true, it would have been false if the tree were just 1cm shorter and this is too small a difference for Magoo to judge.

More generally, if the tree is exactly n+1cm tall then Magoo, given his discriminative limitations, is in no position to know that the tree is taller than n cm. Contraposing this, if Magoo *is* in a position to know that the tree is taller than n cm, then the tree must be taller than n+1cm. In order for Magoo to be in a position to know that the tree is taller than n cm, we might say, n must be cushioned by a certain *margin for error* – and 1cm is too small a margin. If we let L_n be the proposition that the tree is over n cm tall, then we have the following 'margin for error' principle: For any positive integer x, $KL_x \supset L_{x+1}$. If Magoo is sufficiently aware of his own discriminative limitations, and has some appreciation of what knowledge requires, then presumably he could be in a position to know each instance of this margin for error principle: $K(KL_x \supset L_{x+1})$ for any positive integer x. Suppose that this is so. Given these assumptions – all of which seem perfectly reasonable in the case described – it is possible to prove, using the KK-thesis, the absurd conclusion that Magoo is in a position to know that the tree is over 200cm tall or over 1000cm tall or, indeed, over x cm tall for *any* positive integer x.

Before sketching the proof, though, one further principle is needed – this is a closure principle to the effect that, if one is in a position to know φ and is in a position to know that ψ follows from φ then one is in a position to know that ψ. For present purposes this can be formulated as the following distribution schema: $K(\varphi \supset \psi) \supset (K\varphi \supset K\psi)$. With this in mind, the proof proceeds as follows: As a base case we have it that KL_1. Assume, for induction, that KL_n for some integer n. By the KK-thesis we have it that KKL_n. Since Magoo is in a position to know the relevant instance of the margin for error principle, we have it that $K(KL_n \supset L_{n+1})$. Presumably Magoo could deduce, from KL_n and $KL_n \supset L_{n+1}$ that L_{n+1}. Since Magoo is in a position to know both of these premises it follows, by closure, that he is also in a position to know the conclusion – we have KL_{n+1} as required. Given the base case and the induction step, we have proved that KL_x for any integer x.

More formally:

(1) KL_1	Premise
(2) KL_n	Induction Hypothesis
(3) KKL_n	2, KK-thesis
(4) $K(KL_n \supset L_{n+1})$	Premise
(5) $KKL_n \supset KL_{n+1}$	4, Closure
(6) KL_{n+1}	3, 5, Modus Ponens
(7) $\forall x \in \mathbf{Z}^+, KL_x$	1, 2, 6, Proof by Induction

If we hold on to closure, and to our original assumptions, then the only way to avoid this absurd conclusion is to concede that there is some positive integer y, such that KL_y and $\sim KKL_y$. We have a counterexample to the KK-thesis.

One way to resist this argument, of course, is to challenge the closure principle upon which it relies. It's important to note that the closure principle used effectively guarantees that the K operator is closed under relations of *multiple premise* deductive consequence – and one significant concern about closure principles of this sort is that multiple premise deductive inferences can aggregate the risk of error. Put simply, the conclusion of a multiple premise deductive inference may be less probable than any of the premises, taken individually. If one thinks that being in a position to know can tolerate a small error risk, but not a large one, then one will have reason to reject this closure principle, at least in full generality. Williamson anticipates this objection to the argument and responds by outlining a revised argument that appears to dispense with any problematic multiple premise closure principle. In the appendix, I shall attempt something similar, though my argument differs, at least in some details, from Williamson's.

The margin for error principle exploited in this argument might be thought to reflect a general safety condition upon knowledge – but it certainly doesn't presuppose any such condition. Indeed, the Magoo argument doesn't rely upon *any* general presumptions about the requirements for knowledge – and herein lies its strength. Aside from the closure principle, the argument rests only upon a few compelling intuitions about one very specific case.

III. AGAINST THE JK-RULE

If there is a safety condition on knowledge, of the kind formulated in the last section, then the JK-rule has consequences that seem clearly unacceptable. Suppose my evidence Δ provides modest support for a proposition P – enough support for me to have justification to believe that P, but no more. By the JK-rule, I must have justification to believe that KP which entails that Δ is safe evidence for P and to believe that KKP which entails that Δ is safely safe evidence for P and, indeed, to believe that $K^{100}P$ which entails that Δ is safely, safely, safely, ... safe evidence for P. Clearly, though, this would be an outrageous overestimation of the strength of my evidence – so how could I possibly have justification to believe something that entails it?

Once again, one could attempt to resist this argument by challenging the safety condition on knowledge. Once again, I won't attempt to defend the condition here because an alternative argument – one that dispenses with it – is available. This is, in fact, a straightforward adaptation of the Magoo argument above. In the case originally described, Mr. Magoo is not only in a position to know that the tree is taller than 1cm, he also, quite clearly, has justification to believe that the tree is taller than 1cm. Further, if Magoo is in a position to know each instance of the margin for error principle by reflection, then he is also in a position to justifiably believe each instance of the margin for error principle by reflection. Finally, it's plausible that justification satisfies an analogous closure principle to knowledge: If one has justification to believe that φ and has justification to believe that ψ follows from φ then one has justification to believe that ψ. Once again, this can be formulated as a distribution schema: $J(\varphi \supset \psi) \supset (J\varphi \supset J\psi)$. Given these assumptions, it is possible to prove, using the JK-rule, the absurd conclusion that Magoo has justification to believe that the tree is over 200cm tall or over 1000cm tall or, indeed, over x cm tall for any positive integer x.

Consider the following: As a base case we have it that JL_1. Assume, for induction, that JL_n for some integer n. By the JK-rule we have it that JKL_n. Since Magoo has justification to believe the relevant instance of the margin for error principle, we have it that $J(KL_n \supset L_{n+1})$. Presumably, Magoo could deduce, from KL_n and $KL_n \supset L_{n+1}$ that L_{n+1}. Since Magoo has justification to believe both of these premises it follows, by J-closure, that he also has justification to believe the conclusion – we have JL_{n+1} as required. Given the base case and the induction step, we have proved that JL_x for any integer x.

More formally:

(1) JL_1	Premise
(2) JL_n	Induction Hypothesis
(3) JKL_n	2, JK-rule

(4) $J(KL_n \supset L_{n+1})$ Premise

(5) $JKL_n \supset JL_{n+1}$ 4, J-Closure

(6) JL_{n+1} 3, 5, Modus Ponens

(7) $\forall x \in \mathbf{Z}^+, JL_x$ 1, 2, 6, Proof by Induction

If we hold on to J-closure and to our original assumptions, then the only way to avoid this absurd conclusion is to concede that there is some positive integer y, such that JL_y and $\sim JKL_y$. We have a counterexample to the JK-rule. Once again, one could attempt to resist the argument by taking issue with the principle of J-Closure. But, once again, I think that an alternative argument, dispensing with the principle, is available. Details are provided in the appendix.

IV. BACK TO RULES FOR ASSERTION AND BELIEF

While the difficulties explored in the previous section pertain directly to the JK-rule for belief, they can, as hinted in the first section, be used to apply substantial pressure to the JK-rule for assertion as well. Indeed, the only way to insulate the JK-rule for assertion from these difficulties is to insist that the epistemic standards for assertion are more stringent than the epistemic standards for belief. Suppose one has justification to assert that P. By the JK-rule for assertion, one has justification to believe that one is in a position to know that P. If one has justification to assert any proposition that one has justification to believe, it follows that one has justification to assert that one is in a position to know that P at which point an analogue of the Magoo argument will engage. If we are to avoid this consequence, it must be that justification to believe does not suffice for justification to assert.

Do the arguments of the previous section push us back, then, in the direction of the K-rule for assertion? Perhaps this is the right lesson to draw – but it's worth noting that any defender of the K-rule for assertion may also have to concede, albeit for slightly different reasons, that assertion is subject to harsher epistemic standards than belief. Consider again the Gettier case described in the first section. According to the K-rule, as we've seen, Sarah's assertion that there's a barn up ahead is unjustified. But, so long the defender of the K-rule endorses the conventional verdict about Gettier cases such as this one, he will have to admit that Sarah's *belief* that there is a barn up ahead *is* a justified one[4]. Once again we have

[4] There is, I think, a certain tension in *Knowledge and Its Limits* concerning this very point. In chapter 11 Williamson famously defends a knowledge rule for both assertion and belief. But a rule of this kind, as noted, is in some tension with the standard verdict about Gettier cases – and, in particular, with the verdict that the beliefs held by Gettiered subjects are justified. And yet, in chapter 1, Williamson explicitly *endorses* the standard verdict about Gettier cases, and this endorsement forms a part of his case for the claim that knowledge is unanalysable. There may well be a range of ways to resolve this conflict – but I won't explore this here. Comesaña and Kantin (2010) also suggest, though for slightly different reasons, that Williamson's overall

the result that justification to believe does not suffice for justification to assert. If this is a point against the JK-rule for assertion, then it is a point against the K-rule for assertion as well.

My aim here is not to adjudicate between alternate rules for assertion. But it is worth noting, at least, that there is a rule for assertion that avoids the difficulties of the previous section, gives the intuitive verdicts in Gettier cases and preserves perfect symmetry between the epistemic standards for assertion and belief. This is a simple J-rule: One has justification to assert that P iff one has justification to believe that P. Smithies does in fact endorse the J-rule for assertion, but appears to regard it as being close to trivial – almost like a placeholder for something more substantial (namely the JK-rule). And the corresponding J-rule for *belief* is, of course, a genuine triviality: One has justification to believe that P iff one has justification to believe that P.

When it comes to rules articulating the epistemic standards that govern assertion, however, I'm unsure whether we really need, or should expect, anything more substantial than the simple J-rule. In a similar vein, I'm unsure whether we need, or should expect, *any* non-trivial rule articulating the epistemic standards that govern belief – after all, it is here that epistemic standards non-derivatively apply. Rules of this kind are, in any case, only part of the final story that we might hope to tell. Understanding the epistemic standards that govern a practice is one thing – understanding the conditions under which those standards are satisfied is quite another. By combining analyses of the nature of epistemic justification with the J-rules, further and more informative rules could be derived.

REFERENCES

Comesaña, J. and Kantin, H. (2010) 'Is evidence knowledge?' *Philosophy and Phenomenological Research* v80(2), pp447-454

DeRose, K. (2002) 'Assertion, knowledge and context' *Philosophical Review* v111(2), pp167-203

Dretske, F. (2004) 'Externalism and modest contextualism' *Erkenntnis* v61, pp173-186

Hawthorne, J. (2003) *Knowledge and Lotteries* (Oxford: Oxford University Press)

Hawthorne, J. and Stanley, J. (2008) 'Knowledge and action' *Journal of Philosophy* v105(10), pp571-590

position in *Knowledge and Its Limits* is in tension with the standard verdict about Gettier cases (see Comesaña and Kantin, 2010, partic. section 3).

Huemer, M. (2007) 'Moore's paradox and the norm of belief' in Nuccetelli, S. and Seay, G. eds *Themes from G.E. Moore* (Oxford: Oxford University Press)

Nozick, R. (1981) *Philosophical Explanations* (Cambridge, MA: Harvard University Press)

Reynolds, S. (2013) 'Justification as the appearance of knowledge' *Philosophical Studies* v162(2), pp367-383

Smith, M. (2010) 'What else justification could be' *Noûs* v44(1), pp10-31

Smithies, D. (2012) 'The normative role of knowledge' *Noûs* v46(2), pp265-288

Williamson, T. (2000) *Knowledge and its Limits* (Oxford: Oxford University Press)

APPENDIX: DISPENSING WITH MULTIPLE PREMISE CLOSURE

In the appendix I shall outline revised versions of the Magoo arguments against the KK-thesis and JK-rule that dispense, respectively, with the closure and J-closure principles and use instead principles that are immune to risk aggregation worries. I shall begin with the original argument against the KK-thesis. As noted in section II, the risk aggregation worries about the closure principle exploited in this argument arise against the background of a view on which being in a position to know can tolerate small error risks but not large ones. To make this a little more precise, let Pr be a probability function representing one's epistemic or evidential probabilities and suppose that, in order for one to be in a position to know that φ, the epistemic probability of φ must exceed a threshold t that is close to, but less than, 1: $K\varphi \supset Pr(\varphi) > t$. The closure principle allows us to infer $K\psi$ from $K(\varphi \supset \psi)$ and $K\varphi$. This clashes with the preceding constraint precisely because $Pr(\varphi \supset \psi) > t$ and $Pr(\varphi) > t$ are jointly compatible with $Pr(\psi) \leq t$.

In my revised Magoo argument, the multiple premise closure principle $K(\varphi \supset \psi) \supset (K\varphi \supset K\psi)$ is replaced by a single premise closure principle that is not beset by risk aggregation worries: (SPClosure) If $K\varphi$ and $\varphi \supset \psi$ is a logical truth, then $K\psi$. If $Pr(\varphi) > t$ and $\varphi \supset \psi$ is a logical truth, it follows that $Pr(\psi) > t$. The revised argument also makes use of the principle that $K(\varphi \wedge \psi) \supset (\varphi \wedge K\psi)$ is a logical truth. This is perhaps disputable and, thus, worth flagging – but it is a theorem schema of all but the very weakest epistemic logics[5].

In addition to these principles, the revised argument also requires some further assumptions about Magoo – albeit assumptions that seem motivated in the case described. Rather than merely assuming that Magoo is in a position to know every substitution instance

[5] The principle will be a theorem schema of the epistemic logic **MT** – itself a sublogic of **KT**. It can be proved from (M) $K(\varphi \wedge \psi) \supset (K\varphi \wedge K\psi)$ and (T) $K\varphi \supset \varphi$, using the resources of propositional logic.

of the margin for error principle $\forall x \in \mathbf{Z}^+ (KL_x \supset L_{x+1})$, we must now assume that Magoo is in a position to know the principle itself – $K(\forall x \in \mathbf{Z}^+ (KL_x \supset L_{x+1}))$. Let 'E' abbreviate '$\forall x \in \mathbf{Z}^+ (KL_x \supset L_{x+1})$'. We assume that Magoo is in a position to know that E holds. We also assume that Magoo is in a position to know that E holds and the tree is taller than 1cm. These assumptions, as suggested, seem motivated in the case described. Notice that $E \wedge L_{n+1}$ is a logical consequence of $E \wedge KL_n$.

With these principles and assumptions in place, the revised argument proceeds as follows:

(1) $K(E \wedge L_1)$	Premise
(2) $K(E \wedge L_n)$	Induction Hypothesis
(3) $KK(E \wedge L_n)$	2, KK-thesis
(4) $\vdash K(E \wedge L_n) \supset (E \wedge KL_n)$	Epistemic Logic
(5) $K(E \wedge KL_n)$	3, 4, SPClosure
(6) $\vdash (E \wedge KL_n) \supset (E \wedge L_{n+1})$	First Order Logic
(7) $K(E \wedge L_{n+1})$	5, 6, SPClosure
(8) $\forall x \in \mathbf{Z}^+, K(E \wedge L_x)$	1, 2, 7, Proof by Induction

The argument against the JK-rule can be revised in essentially the same way. In this case, the J-closure principle will be replaced by the SPJ-Closure principle: If $J\varphi$ and $\varphi \supset \psi$ is a logical truth then $J\psi$. We also assume that Magoo has justification for believing that E and the tree is taller than 1cm. The argument is as follows:

(1) $J(E \wedge L_1)$	Premise
(2) $J(E \wedge L_n)$	Induction Hypothesis
(3) $JK(E \wedge L_n)$	2, JK-rule
(4) $\vdash K(E \wedge L_n) \supset (E \wedge KL_n)$	Epistemic Logic
(5) $J(E \wedge KL_n)$	3, 4, SPJ-Closure
(6) $\vdash (E \wedge KL_n) \supset (E \wedge L_{n+1})$	First Order Logic
(7) $J(E \wedge L_{n+1})$	5, 6, SPJ-Closure
(8) $\forall x \in \mathbf{Z}^+, J(E \wedge L_x)$	1, 2, 7, Proof by Induction