Justification and Being in a Position to Know

Daniel Waxman
National University of Singapore

1. Introduction

The notion of *being in a position to know* has increasingly drawn the attention of epistemologists.[1] One of the boldest recent deployments of the notion is due to Sven Rosenkranz (2018, 2021), who has argued that it can be used to give a theoretically appealing account of propositional justification. Let K and J be propositional operators expressing 'the agent is in a position to know that' and 'the agent is propositionally justified in believing that'. In these terms, Rosenkranz's account of propositional justification is

$$J\phi \leftrightarrow \neg K \neg K\phi$$

where the biconditional is to be read as implying that its two sides express the same state of affairs.[2] Let us call this thesis $J = \neg K \neg K$.[3]

$J = \neg K \neg K$ claims a promising explication of justification in terms of knowledge. One source of appeal is that $\neg K \neg K$ has certain structural features in common with justification as traditionally conceived. Just as knowledge is traditionally said to entail

[1] Its contemporary prominence arguably stems from Williamson (2002), where it is used to formulate many fundamental concepts and theses including luminosity and the KK principle. Hawthorne (2003) and Schaffer (2007) use it to formulate closure principles. It has also received attention in its own right, for instance in Heylen (2016), Rosenkranz (2016), and Hawthorne and Yli-Vakkuri (forthcoming).

[2] Rosenkranz (2021: 109). The objections presented below do not turn on this 'metaphysical' reading of the biconditional; they would apply equally well on a material reading.

[3] In addition to $J = \neg K \neg K$, Rosenkranz (2021: 108) defends the analogous claim that *doxastic* justification D is equivalent to $\neg K \neg k$, where k expresses 'the agent knows'. Nothing I say below directly bears on this thesis, except insofar as it derives some plausibility from $J = \neg K \neg K$.

justification, K entails $\neg K \neg K$ (by the factivity of K); and just like traditional conceptions of justification, $\neg K \neg K$ is *not* itself factive. Indeed, $J = \neg K \neg K$ may be especially attractive to those sympathetic to the 'knowledge first' project in epistemology.[4] Some hard-line knowledge-firsters argue that justification is simply knowledge.[5] But the resulting factive and externalist notion of justification is arguably a radical departure from the ordinary notion. $J = \neg K \neg K$ thus holds out the intriguing prospect of allowing knowledge-firsters to accommodate many justification-theoretic and internalist intuitions, all while taking the fundamental notion in epistemology to be resolutely knowledge-theoretic.

Furthermore, $J = \neg K \neg K$ is well-suited to theorizing. It allows for simple model building: instead of two epistemic accessibility relations – one expressing what is consistent with what the agent is in a position to know and one expressing what is consistent with what the agent is justified in believing – only a single accessibility relation is needed. Perhaps for this reason, $J = \neg K \neg K$ has gained some currency in recent work in epistemology and epistemic logic; see for instance Goldstein (forthcoming), Carter and Goldstein (2021), Williamson (2013a), and Williamson (2013c).[6]

Unfortunately, as I will now argue, $J = \neg K \neg K$ faces clear counterexamples.

[4]Most famously pursued in Williamson (2002).

[5]See for instance Sutton (2005), Littlejohn (2012), and Williamson (2013b).

[6]Some of these papers explicitly appeal to closely related but distinct theses: for instance (Carter and Goldstein (2021: 2510)) that S is justified in believing p iff it is not the case that S *knows* that it is not the case that S *knows* that p, and (Williamson (2013a); cf. Lenzen (1978), Stalnaker (2006), Halpern et al. (2009)) that S *believes* p iff it is not the case that S knows that it is not the case that S knows that p. But actually, since these authors make modelling assumptions that collapse various distinctions, the differences between their theses and Rosenkranz's $J = \neg K \neg K$ (stated in terms of *being in a position to know*) is not clear. For instance, Carter and Goldstein (2021: 2512, fn 5) make the 'simplifying assumption that the agent knows everything that they are in a position to know'; and Williamson (2013a: 11) notes that in his model, 'all [the agent's] beliefs are justified in an internalist sense', i.e. has no unjustified beliefs.

Without these modelling assumptions, the claim that justification (either propositional or doxastic) is equivalent to *S not knowing that S does not know that p* has little to recommend it. Take propositional justification first. Consider some claim p such that S has neither entertained p nor $\neg p$. S will then not know that she doesn't know p nor that she doesn't know $\neg p$. But, on this view, S has propositional justification to believe both p and $\neg p$. This is not only implausible on its own terms, but complicates the logic of justification (amounting to a rejection of the widely-accepted D axiom; cf. Rosenkranz (2021: Ch 5)). Next take doxastic justification. Consider a case where S has good reasons available for believing p but instead believes p on the basis of other, bad, reasons. S may well believe that she knows p on the basis of the bad reasons, and so (if consistent) will fail to believe and therefore fail to know that she doesn't know p. Such examples are paradigm cases of propositional but not doxastic justification; but the present view diagnoses S as doxastically justified.

2. A counterexample to $J = \neg K \neg K$

An initial source of resistance to $J = \neg K \neg K$ comes from the idea that justification cannot consist in the *absence* of potential knowledge. After all, 'small infants, dogs, trees, madmen, drunkards and dead people' (Rosenkranz (2021: 108)) may be in no position to know anything about what they are in a position to know, due to lack of cognitive and epistemic capacities; but this hardly means they have justification to believe everything. Defenders of $J = \neg K \neg K$ are well aware of objections of this kind. To deal with them, Rosenkranz (2021: 108) advances $J = \neg K \neg K$ as applying only to what we might call *realistically idealized agents*:

> suitably improved versions of ourselves whose epistemic powers finitely extend our own, who can grasp every thought expressible in the language, and who have other epistemic virtues such as freedom of irrationality, bias and compulsion, freedom of attention deficiencies, and freedom of other ills that affect the epistemic lives of ordinary subjects.

Even supposing that this response succeeds, I will argue that there are counterexamples to $J = \neg K \neg K$ involving realistically idealized agents. Consider:

Missing Prize: Based on a fair die roll, a prize has been placed either behind door A, behind door B, or removed altogether. S is in a position to open door A or to open door B, but in no position to open both doors. S is in no position to do anything else to figure out the location of the prize, nor does S have any other knowledge or information relevant to its location. S knows all of the above. In fact, unbeknownst to S, the prize has been removed altogether.

Let p be the proposition *the prize is behind one of the doors*. It is, I think, as intuitively clear as anything in philosophy that S is not propositionally justified in believing p.[7] I

[7] A probabilistic view of justification might give a contrary verdict if justification is identified with rational credence of 2/3 or higher. But, aside from the implausibility of such a view, the case can easily be tweaked to make one's initial rational credence in p arbitrarily low.

will now argue that S is in no position to know that she is in no position to know p, so that $\neg K \neg K p$.

Suppose that S is in a position to know that she is in no position to know p. Then S is in a position to know that neither opening door A nor opening door B would result in her knowing p. But, if so, S must be in a position to know that the prize is behind neither door A nor B. But this is simply false, given the setup of the scenario: S is in a position to know at most what is behind one of the doors. So the supposition that S is in a position to know that she is in no position to know p is false.

Two brief points are worth mentioning. First, note that it is essential to the case that it involves a failure of agglomeration of K over conjunction:

$$K\phi \wedge K\psi \rightarrow K(\phi \wedge \psi)$$

because S is in a position to know that the prize is not behind door A, and S is in a position to know that the prize is not behind door B, but S is in no position to know that the prize is behind neither door. This, however, should not be disturbing: following Heylen (2016), it is widely accepted (see also Rosenkranz (2016) and Hawthorne and Yli-Vakkuri (forthcoming) for agreement) that agglomeration is false, for precisely the reason that, as Rosenkranz (2021: 13) puts it, 'one may be in a position to know each of two propositions, and yet, doing what one is thereby in a position to do to come to know one of them precludes coming to know the respective other.'

Second, note that the case is fully consistent with S being realistically idealized: nothing in it appeals to bias, or a failure of rationality or reflectiveness or attention, and so on. S is in no position to know that she is in no position to know p, not because of any reflective deficiency on her part, but simply because of the predicament she finds herself in (in particular her inability to open both doors).

So, we have a case where $\neg Jp$ and $\neg K \neg K p$, and thus a clear counterexample to $J = \neg K \neg K$.

3. Generalizing the counterexample

Reflection on the form of Missing Prize points the way to a more general recipe:

Counterexample Recipe: For all S knows, p is true. There are two procedures S is in a position to carry out, either of which might (for all S knows) give a positive result or a negative result. A positive result would verify p (thereby giving S knowledge that p), but a negative result would not refute p. S is in no position to carry out both procedures. S knows all this. As it happens, unbeknownst to S, p is false (and so neither procedure would in fact give a positive result).

While it would be question-begging to build in $\neg Jp$ to the description of Counterexample Recipe, it should be clear enough by now that instances can be formulated in which S is not propositionally justified in believing p. Furthermore, on basically any (independent) theoretical view of justification I can think of, instances can be concocted so that S has no evidence or epistemic grounds or reasons or sufficiently high rational credence or reliable method or safe method or... [insert your favourite view of justification here] for believing p. So it is easy to elaborate Counterexample Recipe so that $\neg Jp$.

But, I claim, $\neg K \neg Kp$. The reasoning mirrors that above in Missing Prize; nothing there appealed to special features of the case that are not more generally shared with instances of Counterexample Recipe. Briefly: although it is true that S is in no position to know p, S herself cannot realize that: for to do so would require her to know that *neither* procedure would return a positive result, which she is in no position to do. So, any case with the same structure will also constitute a counterexample to $J = \neg K \neg K$.

In thinking through why Counterexample Recipe works, it might be helpful to see why certain of its features are essential. First, why is it important that there are at least *two* procedures S is in a position to carry out? Because, if the setup is changed so that just one procedure is available, S *is* in a position to know that she is in no position to know p. For, once the procedure yielded a negative result, S could then conclude that she was not initially in a position to know p after all. Second, why is it important that a negative result does not refute p (in the sense of providing knowledge that $\neg p$)? Because, if the setup is changed so that a negative result refutes p, then S *is* in a position to know that she is in no position to know p. For once the procedure yielded a negative result, S would know $\neg p$, and so would arguably be in a position to know that she is in no position to know p (since knowledge is factive).

In the remainder of the paper, I will consider and reject two responses that a defender of $J = \neg K \neg K$ might make.

4. Another way of knowing $\neg Kp$?

In Missing Prize, I argued that S is in no position to know that she is in no position to know that the prize is behind one of the doors. But here is a potential response. Suppose that S picks a door at random and (as per the setup of the case) finds no prize. Then consider the following argument:

(1) The best S was in a position to do to know p is to pick one of the doors at random.

(2) S did in fact pick one of the doors at random.

So: (3) S did the best that S was in a position to do to know p, and it didn't result in S knowing p.

(4) If S did the best that S was in a position to do to know p, and it didn't result in S knowing p, then S was in no position to know p.

So: (5) S was in no position to know p.

What is more, it might be argued, the reasoning in the above argument can be carried out by S herself. If so, then S is in a position to know the conclusion of the argument – that is, she is in a position to know that she was not initially in a position to know p, contrary to the analysis of the case given above.[8]

[8] Sven Rosenkranz raised an argument of this form in a response to comments I gave at the *Asian*

Although this line of thought is ingenious, I do not believe it works. A first point to make is that it is subject to a dilemma. The question to ask is this: by running through the argument above, can S thereby come to know that the prize is behind *neither* door?

Answering 'yes' is tantamount to rejecting the case: for it was stipulated that S is initially in a position to open only one of the doors (and has no other way of figuring out the location of the prize). So if S can in fact come to know that the prize is behind neither door, there must be some incoherence in the setup of the case. But there is none: once we recognize that agglomeration can fail, it should be clear that it can fail in cases with the structure of Missing Prize.

Answering 'no' is similarly unpalatable. Recall that p is the claim that the prize is behind one of the doors: so this horn amounts to saying that $K\neg Kp$ while $\neg K\neg p$. But there is a straightforward argument that these claims are incompatible. For in Missing Prize, $p \leftrightarrow Kp$. (The right-to-left direction follows from the factivity of knowledge. For the left-to-right direction, suppose the prize is behind one of the doors. Whichever door it is, S is in a position to open that door and thereby come to know p.) What's more, S is in a position to know this by reflecting on her knowledge of the setup; so $K(p \leftrightarrow Kp)$. But then if, as supposed, S is in a position to know $\neg Kp$, S is also in a position to carry out the foregoing reasoning and thereby come to know $\neg p$ – precisely what this horn denies.

I think this dilemma causes trouble for the thought that S can come to know $\neg Kp$ using the above argument. Even so, it would be more satisfying to explain where the reasoning has gone wrong. Fortunately, a plausible diagnosis is available.

Notice that there is an ambiguity between two notions of 'the best S was in a position to do' (Cf. Rosenkranz (2021: 41).. There is an *ex ante* notion, i.e. the best thing to do according to the information S had at the time, and an *ex post* notion, i.e. the thing to do which would have in fact delivered knowledge, or been most objectively likely to do so. With this distinction in hand, it should be clear that (4) is plausible only when

Journal of Philosophy symposium on his book (the exchange is forthcoming as Waxman (forthcoming) and Rosenkranz (forthcoming)).

'best' is read as 'best-ex-post'. Here is a case that demonstrates this:

Hidden Prize: Based on a fair die roll, a prize has been placed either behind door A, door B, door C, or removed altogether. *S* is in a position to do one of two things: (i) open doors A and B; (ii) open door C. Unbeknownst to *S*, the prize is in fact behind door C.

Let *p* again be the proposition *the prize is behind one of the doors*. The best-ex-ante that *S* is in a position to do to come to know *p* is to pick option (i): after all, that is what is most likely (by her lights) to reveal the location of the prize. And suppose *S* in fact picks option (i) and comes away empty handed. Although *S* did the best-ex-ante she was in a position to do to know *p*, and doing so didn't result in knowledge, it surely doesn't follow that *S* was in no position to know *p* – actually she was, because she was in a position to pick option (ii), and doing so would have resulted in her knowing *p*.

So if (4) is to come out as true, it must be formulated in terms of best-ex-post. And so for the above argument to be sound, it must be formulated in terms of best-ex-post throughout. But even though the argument, so reformulated, is sound, *S* is in no position to reason through it. In particular, *S* is in no position to know (1), once that premiss is explicitly formulated in terms of what is best-ex-post. Although *S* picked randomly, and thus did the best-ex-ante that she was in a position to do, for all *S* is in a position to know it would have been better-ex-post to pick the *other* door, since for all *S* is in a position to know, that is where the prize lies. That is why *S* cannot, in fact, use the above argument to come to know (5).

5. A more austere conception of *being in a position to know*?

A different reaction one might have is to question whether we have been misapplying the notion of *being in a position to know*. In our cases, it might be objected, we have been considering procedures that, in some sense, give the agent new evidence. In contrast, there is a more austere notion of *being in a position to know* on which it involves merely processing the evidence one already has.[9] Perhaps there is a hint of this in the gloss given by Williamson (2002: 95), according to which, when one is in a position to know, 'The fact is open to one's view, unhidden, even if one does not yet see it.'

Consider how this might handle Missing Prize. If what S is in a position to know is strictly a function of the evidence she already has, then perhaps (since she is by hypothesis realistically idealized) she is in a position to reflect on the fact that she initially has no evidence that bears on p, and to thereby come to know that she is in no position to know p in the sense under consideration.

A preliminary point in response is that this kind of austere conception is by no means universal. For instance, Rosenkranz is explicit that he has the more liberal notion in mind, and indeed, this appears essential to his view in many places.[10] So adopting an austere conception of K would require a wholesale reassessment of the case for $J = \neg K \neg K$.

But, putting this dialectical point aside, a considerably more powerful response is available. In Counterexample Recipe, we said that S is in a position to carry out procedures which might (for all S knows) lead to knowledge in p. But notice that it is by no means essential that these procedures *yield new evidence* in the relevant sense. Consider:

Geometric Figures: S undergoes a visual experience as of two many-sided geometric figures, one on the left side of her phenomenal visual field and one on the right. While S knows that there are two geometric figures in her phenomenal visual field, she does not know exactly how many sides each has. S is in a position to either (a) attend to the figure on the left and count how many sides it has,

[9] I've gestured at the austere conception in terms of the potentially-theoretically-loaded notion of evidence, but I believe the points made below are robust to other ways in which it might be fleshed out.

[10] Thus Rosenkranz (2021: viii): "But recently, the notion of being in a position to know has begun to receive treatments that allow for the possibility that more may have to happen, in order for one to come to know what one is so far merely in a position to know, than merely registering what it needs no further epistemic agency on one's part to unearth. To be in a position to know is to have the opportunity to know, and seizing such an opportunity may require implementing an investigative procedure. The notion will here be understood in such a more liberal sense [. . .]". This liberal conception is implicit in, for instance, the responses to potential counterexamples discussed in Rosenkranz (2021: Ch 6.3). For other recent work that rejects the austere conception, see Heylen (2016) and Hawthorne and Yli-Vakkuri (forthcoming).

or (b) do the same for the figure on the right. But *S* is in no position to count both; the visual experience will be too fleeting for that. *S* knows all this. In fact, unbeknownst to *S*, the figure on the left has 49 sides, while the figure on the right has 51.

Let p be the proposition *at least one of the geometric figures has exactly 50 sides*. Then *S* is in no position to know that she is in no position to know p, for exactly the same reasons as before. But here the procedure in question involves merely attending to one's visual experience.

Similarly, consider:

Fond Memories: *S* attended two parties, one on Saturday and one on Sunday. Fortunately, her episodic memories are sufficiently intact that, for each party, she is in a position to access her memory of it and thereby come to know exactly which friends attended. Unfortunately, *S* has taken a memory-destroying drug, so she has time to access her memory of at most one of the parties before the drug kicks in. *S* knows all this. As it happens, unbeknownst to *S*, her friend Jack was not at either party.

Let p be the proposition *Jack attended at least one of the parties on the weekend*. As before, *S* is in no position to know that she is in no position to know p. But here the procedure in question involves merely accessing one's episodic memories.

The point of these examples is that, if anything counts as merely processing one's current evidence (as opposed to gaining new evidence), it is surely a procedure of this kind. It is extremely natural to view the contents of one's phenomenal visual field or one's episodic memory as epistemically 'in reach', if anything is, even if one's attention is temporarily directed elsewhere.

Of course, there is room to quibble: these cases draw on deniable assumptions in the philosophy of perception and memory; nor is the gestured distinction between 'adding new' versus 'processing already existing' evidence very precise. But I think

I've said enough to make a more general conclusion plausible. However exactly an austere notion of *being in a position* to know is cashed out, in order to avoid triviality, it must allow for some gap between what one actually knows and what one is in a position to know. That is, it must countenance *some* notion of propositions which are not known but which are 'in reach' of being known upon some suitable procedure being carried out. If that is so, then it is hard to see why Counterexample Recipe cannot then be implemented, using the procedure in question, to once again cause trouble for $J = \neg K \neg K$.[11]

References

Sam Carter and Simon Goldstein. The normality of error. *Philosophical Studies*, 178(8): 2509–2533, 2021.

Simon Goldstein. Fragile knowledge. *Mind*, pages 1–25, forthcoming.

Joseph Y. Halpern, Dov Samet, and Ella Segev. Defining knowledge in terms of belief: The modal logic perspective. *Review of Symbolic Logic*, 2(3):469–487, 2009.

John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, 2003.

John Hawthorne and Juhani Yli-Vakkuri. Being in a position to know. *Philosophical Studies*, 5(1), forthcoming.

Jan Heylen. Being in a position to know and closure. *Thought: A Journal of Philosophy*, 5 (1):63–67, 2016.

Wolfgang Lenzen. Recent work in epistemic logic. *Acta Philosophica Fennica*, 30:1–219, 1978.

Clayton Littlejohn. *Justification and the Truth-Connection*. Cambridge University Press, 2012.

Sven Rosenkranz. Being in a position to know and closure: reply to Heylen. *Thought: A Journal of Philosophy*, 5(1):68–72, 2016.

[11] Thanks to Martin Abreu Zavaleta, Zach Barnett, Sharon Berry, Harjit Bhogal, Tom Donaldson, Julien Dutant, Simon Goldstein, Sebastian Liu, Lavinia Picollo, Sven Rosenkranz, Mattias Skipper, Zeynep Soysal, Jared Warren, and Mike Zhao for helpful discussion and comments.

Sven Rosenkranz. The structure of justification. *Mind*, 127(506):309–338, 2018.

Sven Rosenkranz. *Justification as Ignorance*. Oxford University Press, 2021.

Sven Rosenkranz. Replies? *Asian Journal of Philosophy*, forthcoming.

Jonathan Schaffer. Closure, contrast, and answer. *Philosophical Studies*, 133(2):233–255, 2007.

Robert Stalnaker. On logics of knowledge and belief. *Philosophical Studies*, 128(1):169–199, 2006.

Jonathan Sutton. Stick to what you know. *Noûs*, 39(3):359–396, 2005.

Daniel Waxman. Justification as Ignorance and logical omniscience. *Asian Journal of Philosophy*, forthcoming.

Timothy Williamson. *Knowledge and its Limits*. Oxford University Press on Demand, 2002.

Timothy Williamson. Gettier cases in epistemic logic. *Inquiry: An Interdisciplinary Journal of Philosophy*, 56(1):1–14, 2013a.

Timothy Williamson. Knowledge first. In Matthias Steup John Turri, editor, *Contemporary Debates in Epistemology*, pages 1–10. Blackwell, 2013b.

Timothy Williamson. Response to Cohen, Comesaña, Goodman, Nagel, and Weatherson on Gettier cases in epistemic logic. *Inquiry: An Interdisciplinary Journal of Philosophy*, 56(1):77–96, 2013c.