Justification as ignorance and epistemic Geach principles

Julien Dutant

Sven Rosenkranz's *Justification as Ignorance* (2021) is built on an admirably detailed and careful examination of structural principles of knowing and being in a position to know.[1] It proposes two systems for the logic of these notions, an 'idealized' system that endorses single-premise closure and a more 'realistic' one that doesn't. Even though the systems differ substantially, are both weaker than standard epistemic logic, and do not endorse luminosity principles for knowledge or being in a position to know, 'something beautiful emerges' (81): they agree on a strong set of principles for states of second-order ignorance, namely *not being in a position to know that one knows* and *not being in a position to know that one is not in a position to know*. The book's central thesis is that we should identify those states with doxastic and propositional justification, respectively. Given that identification, principles we obtain notably include the following:

- Justification (of both kinds) is pairwise consistent ($\mathbf{D_J}$, $\mathbf{D_D}$). If one is justified in believing p, one is not justified in believing not-p.
- Propositional justification is positively introspective ($\mathbf{4_J}$). If one is

---
[1]References are to Rosenkranz (2018) unless otherwise indicated.

justified in believing p, one is justified in believing that one is.

- Propositional justification is negatively introspective ($\mathbf{5_J}$). If one is not justified in believing p, one is justified in believing that one is not.

In short, a strongly internalist conception of justification is built out of a strongly externalist conception of knowledge.

My focus here is on the pairwise consistency principles ($\mathbf{D}$). Given the book's central thesis, pairwise consistency for doxastic justification ($\mathbf{D_D}$) amounts to the principle that one is in a position to know that one doesn't know something, or that one doesn't know its negation. Following Rosenkranz in using k and K as operators for knowing and being in a position to know, respectively, that is the schema: $K\neg kp \vee K\neg k\neg p$. As we'll see the schema is *prima facie* appealing and Rosenkranz shows that it cannot be as easily dismissed as you might intially think. I will nonetheless argue it faces counterexamples, situations in which one is neither in a position to know that one doesn't know p nor that one doesn't know not-p. Given that knowing entails being in a position to know these are also counterexamples to the Geach principle for knowledge, $k\neg kp \vee k\neg k\neg p$.[2] Moreover, I argue that these are also cases in which one is not in a position to know that one is not in a position to know p, nor that one is not in a position to know not-p. Hence the Geach principle for being in a position to know, $K\neg Kp \vee K\neg K\neg p$, fails too. Given the book's central thesis, these should be cases in which the pairwise consistency of

---

[2]In modal logic the Geach schema states that what is possibly necessary is necessarily possible ($\Diamond\Box p \rightarrow \Box\Diamond p$, which is equivalent to $\Box\neg\Box p \vee \Box\neg\Box\neg p$). In epistemic logic this becomes the claim that if one knows that one does not know that not-p then one doesn't know that one knows p, or equivalently, that one doesn't one knows that one doesn't know something or that one doesn't know its negation.

propositional justification ($\mathbf{D_J}$) fails. However, they are not: they aren't situations in which one is justified in believing p and justified in believing not-p. Hence they also challenge the book's central thesis.

Rosenkranz doesn't simply assume the Geach principles but derives them from plausible assumptions together with his principle $\mathbf{L}$ of luminosity of not being in a position to know that one is not in a position to know.[3] Principle $\mathbf{L}$ further underpins the introspection principles $\mathbf{4_J}$ and $\mathbf{5_J}$. Rosenkranz puts forward two powerful arguments for $\mathbf{L}$ (chap. 4 and appendix). My cases suggest that these arguments are not sound. I will indicate which premise they cast doubt on. That in turns threatens the introspection principles.

Motivating epistemic Geach principles

Start with the mixed Geach principle that one is in a position to know that one doesn't know p, or that one doesn't know not-p. Roughly, you are in a position to know p if you would know p upon reasonable and careful reflection (henceforth: reflection) on whether p that doesn't alter your initial epistemic position with respect to p.[4] A simple argument for the principle is this: upon reflection on whether p, you would not both believe p and believe not-p, and you would know that you do not believe p, or that you do not believe not-p (or both). Since knowledge requires belief, you would thereby know that you

---

[3]In the idealized system, the $\mathbf{D}$ principles follow from $\mathbf{L}$ together with single-premise closure and basic assumptions about knowledge (87). In the realistic system, $\mathbf{D_J}$ is primitive but still motivated by the same considerations (91-92), and $\mathbf{D_D}$ derived from $\mathbf{D_J}$ (89).

[4]See (1d) p. 41 for a much sharper formulation and section 3 below. The 'only if' direction is questionable because it may not be impossible for reflection on a matter not to alter your initial epistemic position on that matter (e.g. reflecting on p may be incompatible with p's truth, see Fara 2010, 68–70; Yli-Vakkuri and Hawthorne forthcoming; Rosenkranz 2021, 38). It is not assumed here.

don't know p or that you don't know not-p. Hence you are in a position to know that you don't know p or in a position to know that you don't know not-p (or both).[5]

The principle faces simple counterexamples. Rocks aren't in a position to know anything, including that they don't know p and that they don't know not-p. But we may restrict the principle to subjects capable of possessing knowledge. One may lack the concept of knowledge or concepts involved in p. In reply, we may restrict the principle's instances to subjects who have the required concepts or weaken the requirements for being in a position to know. Rosenkranz opts for the former (7). For the latter, consider whether Julius Caesar was in a position to know that he had no idea what smartphones are. Presumably, knowing that smartphones are F requires having at least some minimal concept of smartphones. But we may think that for most, if not all, non-trivial F, there is a minimal concept of smartphones that Caesar could have acquired without altering his epistemic position with respect to whether smartphones are F. And we may think that, for him to have been in a position to know that he didn't know that smartphones are F it is sufficient that upon reflection and conceptual improvement that would not alter his epistemic position on that matter, he would have known that he did not know it. If so, for most F, Caesar was in a position to know that he didn't know that smartphones were F. Fans of analyticity would argue that for some Fs, knowing that smarphones are F is required to have even a minimal concept of smartphones. If so, our tests for being in a position to know cannot be

---

[5]The argument doesn't cover cases in which reflection that doesn't alter your epistemic position is impossible. But if it holds in all cases where it is possible, this gives a *prima facie* reason to generalize it to cases where it is not.

4

applied to show that Caesar was in position to know that he didn't know those propositions. But they cannot be applied to show that he wasn't either—so we haven't secured a counterexample after all. Alternatively, one may further weaken the requirements for being in a position to know and say that for one to be in a position to know that p, it suffices that upon reflection and conceptual improvements that do not affect one's epistemic position with respect to p *on non-analytic matters*, one would know p. It would follow that Casear is in a position to know that he doesn't know the negation of analytic truths, so no counterexample to the mixed Geach principle would be forthcoming either.

The case against Epistemic Geach principles

A couple of warm-up cases first. Suppose you're asked the first name of the actor playing Truffaut's on-screen alter ego. You feel you have it on the tip of your tongue but can't produce it. In a first case, you do in fact know that (p) it is 'Jean-Pierre'. It escapes you right now but it's in your memory and will later come back to you. Arguably, you are then in a position to know that you don't know that it's not 'Jean-Pierre'. For *ex hypothesi* you know that it's 'Jean-Pierre', and since (by the factivity of knowledge) this obviously entails that you don't know that it isn't, you arguably know in the same implicit manner that you don't know that it isn't. Alternatively, suppose you falsely believe that it's not 'Jean-Pierre'. Say you have another, incorrect, name stored in your memory, though again it escapes you at the moment. Arguably you are then in a position to know that you don't know that it's 'Jean-Pierre'. For you arguably believe in the same implicit matter that you don't know it,

and (let us grant) this belief constitutes knowledge.[6] Now for the main case:

MISLEADING TIP OF THE TONGUE. You feel you have the name on the tip of your tongue. But the feeling is misleading: you have no name stored in memory. And it is persistent: you'll keep having it whenever you examine the question.

In this case, you do not believe that the name is 'Jean-Pierre', nor that it isn't. But, I claim, you are not in a position to know that you lack either one of these beliefs. Upon reasonable and careful reflection on the matter, you would still experience the tip of the tongue feeling. So, for all you know, the name 'Jean-Pierre' is stored in your memory and you do know that it is 'Jean-Pierre', and for all you know, another name is stored in your memory and you do know that it's not.

One may object that careful reflection should involve explicitly considering the name 'Jean-Pierre'. And in ordinary tip of the tongue cases, we tend to reach a verdict when a candidate is presented to us. But we may suppose you are not like that. If presented with 'Jean-Pierre', you would have a persistent tip of the tongue-like feeling that you know whether it is so or not, without being able to judge either way.

Such counterexamples are easily multiplied: one only needs a persistent, misleading impression that one knows (or has a belief as to) whether p. So the

---

[6]The argument is less straightforward here. You may believe that you do not know it's 'Jean-Pierre' in virtue of (mistakenly) believing that it's not 'Jean-Pierre'; or you may believe that you do not know it because you believe that you don't believe it; or both. Does the first constitute knowledge, even though it is based on a false belief? Is the latter, introspective, belief a belief that you do genuinely have in such a case? I'll leave those worries aside. Even if they can be answered the principle would face the counterexample introduced below.

mixed Geach principle fails, and with it, the Geach principle for knowledge. These are also counterexamples to the Geach principle of being in a position to know. In those cases, for all you are in a position to know, you are in a position to know p, and for all you are in a position to know, you are in a position to know not-p.[7]

Moreover, you are clearly not both (propositionally or doxastically) justified in believing that the name is 'Jean-Pierre' and justified in believing that it isn't. So contrary to the book's central thesis, second-order ignorance is not sufficient for justification.

The case involves a failure to introspect one's lack of belief. Can we rescue the Geach principles for subjects who are always in a position to know whether they lack belief? The mixed Geach principle, yes.[8] But not the Geach principle for being in a position to know. For even if one would, upon reflection, realize that one doesn't currently have a belief as to whether p, one could still be under the misleading, persistent impression that an answer is within reach. Hence for all one would come to know upon reflection, one would be in a position to know whether p. Hence for all one is in a position to know, one is in a position to know whether p.

---

[7]The latter is independently plausible, but also follows from the failure of the mixed Geach principle together with the fact that knowledge obviously entails being in a position to know and one step of closure for being in a position to know: if $k\phi$ entails $K\phi$ then (by a closure step for K) $K\neg K\phi$ entails $K\neg k\phi$ so by contraposition $\neg K\neg k\phi$ entails $\neg K\neg K\phi$; if so a case in which $\neg K\neg kp \wedge \neg K\neg k\neg p$ holds is one in which $\neg K\neg Kp \wedge \neg K\neg K\neg p$ holds.

[8]We may also rescue the Geach principle for knowledge for the even more restricted class of subjects who are not merely in a position to know, but do know, whether they believe p, for any p (Stalnaker 2006).

Against the luminosity of second-order ignorance

As Rosenkranz shows, the Geach principle for being in a position to know follows from two claims. They'll be easier to discuss by introducing a dual operator of K, M, which we read as 'not being in a position to know not-p' or 'for all one is in a position to know, p'. The claims are (Rosenkranz's labels):

- **L**. $MKp \to KMKp$.
- **A6**. $KMKp \to KMp$.

Put together they entail the Geach schema for K, $MKp \to KMp$, i.e. $K\neg Kp \lor K\neg K\neg p$. If I'm right, then, one of these claims fails. The second is plausible and it isn't put in doubt by misleading tip of the tongue cases. It follows from the factivity of being in a position to know (Williamson 2000, 95; Rosenkranz 2021, 37) and a couple of closure steps that hold even in Rosenkranz's realistic system (Rosenkranz 2021, 85–86, 91–92). Since being in a position to know p entails p ($\mathbf{T}_K$) then (by one instance of closure) if for all one is in a position to know, one is in a position to know p, then for all one is in a position to know, p (**A5**, contraposed). Given the latter, (by a second instance of closure) if one is in a position to know that for all one is in a position to know, one is in a position to know p, then one is in a position to know that for all one is in a position to know, p (**A6**).

Misleading tip of the tongue cases instead threaten **L**, the luminosity of not being in a position to know that one is not in a position to know. Because the tip of the tongue feeling is persistent, upon reflection on what the actor's name is, it will still seem to you that you know what it is. Hence, upon reflection, you will not rule out the possibility that you know that it's not

Jean-Pierre (not-p). Since knowing that it's not Jean-Pierre would entail being in a position to know that you are not in a position to know that it *is* Jean-Pierre, you would not rule out the possibility that you are in a position to know that you are not in a position to know that it is. Since you would not rule out the possibility that you are in such a position, you would not believe upon reflection that you are not in such a position. Since knowledge requires belief, you would not come to know that you are not in such a position. Since you would not come to know it upon reflection, you are not in a position to know it.

This challenges a crucial step in Rosenkranz's carefully constructed arguments for **L** (chap. 4, appendix). Rosenkranz claims that if, in cases where you implement your best procedure to decide whether ψ, you fail to believe that ψ, then in some such cases you will believe (i) that you do not believe it and (ii) that you do not believe it after having done your best to answer it:

> "As long as one's occurrent (full) beliefs are, like judgements, an all or nothing affair, then amongst cases [in which one does one's best to answer whether ψ and] one does not come to believe that ψ, there will be cases in which one responsively believes not to believe that ψ, and so, given one's awareness that knowledge implies belief, responsively believes not to know that ψ. Even if having done the best that one is in a position to do to decide a given matter, and nothing else, is not a luminous condition, one will often, when one has done so, responsively believe that one has done so." (73-74)

(Here "responsively believing" means believing in response to, and in accordance with, how things seem to one to be (51).) Misleading tip of the tongue cases put pressure on both claims. First, even if, upon doing your best, you do not judge that the actor's name is Jean-Pierre nor that it isn't, it still seems to you that you have you have a belief one way or the other, hence you would not responsively believe that you do not have one. Second, even though you have in fact done your best to decide what the name is (and to decide whether you are in a position to know that you are in a position to know whether it's Jean Pierre) it will still seem to you that you haven't done your best, because it still seems to you that the name is on the tip of your tongue.

Rosenkranz makes clear that whether a procedure to decide whether ϕ is "the best available to you" is not a matter of it being "objectively best" but rather "best in the light of your background beliefs and dispositions" (41). Still, because a given procedure can be presented under various guises, a procedure might be such that you believe it best to decide whether ψ, you execute it, and still you do not believe having done the best to decide whether ψ. In the tip of the tongue case, you may for instance believe that *racking your brains about it as hard as you can* is the best you can do. But even if you *in fact* racked your brains as hard as you can, you would not realize that you have done so; you would still believe that you could have done better, because it will still seem to you have you have a belief one way or another.

Thus tip of the tongue cases provide counterexamples to a central premise of Rosenkranz's argument for **L** (L2, p.72). Furthermore, the internalist principles of positive and negative introspection for justification ($4_\mathbf{J}$, $5_\mathbf{J}$) are

derived from **L** (83 and 89, respectively). $\mathbf{5_J}$ is directly threatened: arguably, in those cases you are not justified in believing that the name is Jean-Pierre, but you are not justified in believing that you aren't so justified. The cases don't directly threaten $\mathbf{4_J}$, but rather two claims central to Rosenkranz's derivation: that (propositional) justification is second-order ignorance and that the latter is luminous.

References

Fara, Michael. 2010. "Knowability and the Capacity to Know." *Synthese* 173 (1): 53–73. https://doi.org/10.1007/s11229-009-9676-8.

Rosenkranz, Sven. 2018. "The Structure of Justification." *Mind* 127 (506): 309–38. https://doi.org/10.1093/mind/fzw057.

———. 2021. *Justification as Ignorance.* Oxford: Oxford University Press.

Stalnaker, Robert. 2006. "On Logics of Knowledge and Belief." *Philosophical Studies* 128 (1): 169–99. https://doi.org/10.1007/s11098-005-4062-y.

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

Yli-Vakkuri, Juhani, and John Hawthorne. forthcoming. "Being in a Position to Know." *Philosophical Studies*, forthcoming, 1–17. https://doi.org/10.1007/s11098-021-01709-x.