The $\neg K\neg K$ rule and the structurally unknowable

Yiwen Zhan¹

Received: 6 October 2021 / Accepted: 9 December 2021 / Published online: 16 February 2022
© The Author(s), under exclusive licence to Springer Nature B.V. 2022

Abstract

Rosenkranz (2021) offered a logic and a detailed account of justification, according to which justification that p can be analyzed as a form of second-level ignorance: $\neg K\neg Kp$. An intuition behind the analysis is that the justified subject has the potential, at least in a nearby world, to either come to know p or come to know $\neg Kp$. However, given Rosenkranz's hyperintensional semantics for modeling knowledge states, we can always construct, out of an $\neg K\neg K$ -agent's knowledge state, epistemic possibilities that prohibit the agent from having that potential. This shows that Rosenkranz's analysis is still insufficient for carrying all the epistemic weight underlying the notion of justification. A better analysis would require a deeper look into an agent's knowledge structure and its involved modalities, which could be a lot more complex than what the $\neg K\neg K$ rule tells us.

Keywords Justification · Ignorance · The $\neg K\neg K$ rule · Rosenkranz · Knowability

1 The $\neg K\neg K$ rule

In *Justification as Ignorance* (Rosenkranz, 2021), Rosenkranz has developed, with rich and inspiring details, a view that justification can be analyzed as a kind of ignorance. As he argues, propositional justification signifies a particular agnostic epistemic status in which the subject is not in a position to know that she is not in a position to know that the proposition at issue is true.¹ It implies that, for a proposition P , if I cannot already tell whether the fact that I am in a position to know P obtains, the second best thing is to stay agnostic about whether the opposite fact, i.e. that I am not in a position to know P , obtains.

¹ In this paper, unless specified, I will be exclusively focusing on propositional justification and propositional knowledge in terms of being-in-a-position-to-know.

✉ Yiwen Zhan
y.zhan@bnu.edu.cn

¹ School of Philosophy, Beijing Normal University, Beijing, China

The Rosenkranzian justification is hence a kind of second-level ignorance: let ' K ' stand for the operator which says that the subject 'is in a position to know', and ' J_R ' be the operator for the (Rosenkranzian) propositional justification; then, for a subject to have justification to believe ϕ is for her to be in a situation where she is in no position to rule out $K\phi$:

$$J_R\phi \models \neg K \neg K\phi \quad (1)$$

Formulated as an $\neg K \neg K$ rule, Eq. 1 tells us that justification is not yet knowledge but a particular state of ignorance. Moreover, such kind of ignorance is not a purely 'negative' one but rather has a much 'positive' aspect; namely, it says that ϕ is propositionally justified for a subject S if and only if, according to S 's information state, $K\phi$ *might* be the case.

To compare, imagine I say that the train *might* arrive at 5pm. In saying so, I am not merely making a negative claim of my shortage of knowledge; i.e., I am not merely claiming that the possibility that p (that the train arrives at 5pm) is not something eliminable given all that I am in a position to know, just like many other possibilities are. Rather, my expression also carries a more 'positive' tone, stressing that p is a kind of '*epistemic opportunity*' worth noting, to the extent that at least at nearby worlds, if not already at the actual world, p obtains as a fact. Switching p to Kp , then we have the Rosenkranzian justification.

Of course, according to the standard interpretation of epistemic modals, expressions like '*might p*' are simply analyzed as ' $\neg K \neg p$ '. If so, there is nothing mysterious about the kind of '*epistemic opportunity*' such an expression implies, since it says nothing over and above the fact that the possibility that p is uneliminated. In recent years, however, many have argued that the standard interpretation of epistemic '*might*' (i.e. qua ' $\neg K \neg$ ') cannot really capture the kind of '*more positive*' modality for the '*epistemic opportunity*' that the speaker intends to convey.² If this is true, then one might raise a similar worry if Rosenkranz's modal analysis of justification based on the $\neg K \neg K$ rule has also failed to fully capture the kind of '*epistemic opportunity*'—and hence the full epistemic juice—that the notion of justification intuitively implies.

Let us say that a subject S has *generalized propositional justification* to believe ϕ iff S *might* be in a position to know ϕ . Let us then use ' M ' to denote the kind of epistemic opportunity involved in the '*might*' expression. One way to analyze justification, then, is by endorsing the following:

$$J\phi \models MK\phi \quad (2)$$

Obviously, ' $J\phi$ ' will be equivalent to ' $J_R\phi$ ' whenever ' M ' stands for ' $\neg K \neg$ '. Call the account of propositional justification which not only endorses Eq. 2 but also

²It has been recently argued by quite some people that, looking deeper, these expressions reveal some more intricate structure of information-sensitivity that was largely neglected by the standard interpretation: If by asserting that the train might arrive at 5pm, and if my assertion is interpreted simply as an $\neg K \neg$ -statement about my own informational state at the time, it can neither explain the linguistic data that you may still raise a genuine disagreement against my statement (should you be convinced that the train will not arrive at 5pm), nor explain the perceived need for me to retract my previous assertion upon learning that the train is delayed (see, e.g. Beddor and Goldstein (2021)).

maintains the equivalence of ' M ' and ' $\neg K \neg$ ' the *congruent view*. My suspicion is precisely that the congruent view (i.e. the $\neg K \neg K$ rule) has some serious limitation in delivering the right understanding of justification. Perhaps, to generate a proper analysis of justification, at least certain form of *incongruent view* will be needed, such that the kind of modalities involved in ' M '—whatever it might be—may nonetheless differ from the modality involved in ' $\neg K \neg$ '.

The exact right way to interpret ' M ' for justification is hard to come by. But we can rule out some candidate interpretations. Rosenkranz has been clear that it would be unhelpful to understand justification in terms of metaphysical possibilities to know: on the one hand, it is perfectly okay for one to have justification to believe in something metaphysically impossible; on the other hand, there does not seem to be any epistemically innocent bridge principle available, which guarantees that the metaphysical possibility for one to know p amounts to the propositional justification for her to believe p . On this I agree. It is also probably true to say that the $\neg K \neg K$ rule at least provides a necessary condition for justification. But the worry is that it falls short of generating a *sufficient condition* for justification.

2 A toy model

To begin with, let us take a brief look at the semantic model for Rosenkranz's analysis. In his book, Rosenkranz has presented two different systems of logic for K : an idealized one and a realistic one. In his words, the latter is 'realistic' to the extent that it applies to a 'suitably improved version of ourselves'. It not only fails to agglomerate over conjunction and fails to be closed under known implication, but also invalidates the *monotonic rule*. This is motivated by his hyperintensional, ultra fine-grained (i.e. propositional content at the 'mode of presentation'-level of granularity) account of knowledge-attribution, which allows the subject to know a coarse-grained proposition under one mode of presentation without knowing it under another, even though the latter is (logically or metaphysically) necessarily equivalent to the former.

Hyperintensional semantics works well for modeling a subject's realistic *belief states*. But here, the ambition is to analyze propositional justification using models for a subject's hyperintensional, non-monotonic *knowledge state*. Assuming, as mentioned earlier, that a subject S has generalized propositional justification to believe ϕ iff S might be in a position to know ϕ . The hyperintensional modeling would then permit the epistemic 'might'-operator to incorporate some impossible states of affairs into S 's knowledge state in a certain way, preventing S from settling on factive, non-modal knowledge. For instance, let $\|p\|$ be the coarse-grained proposition that Fermat's Last Theorem is true. Since p is a mathematical truth, we have $\|\neg p\| = \emptyset$. But now consider a realistic subject S who is in no position to know that p is true. To this, we can simply conceive of a mathematically impossible state of affairs—such that the non-empty, fine-grained proposition $\|\neg p\| \neq \emptyset$ nonetheless obtains—being incorporated in S 's knowledge state in a certain way, so that S is in no position to rule out that ' $\neg p$ ' might be the case. Moreover, since knowledge is factive, the impossibility of ' $\neg p$ ' implies the impossibility of ' $K \neg p$ '. Hence, the justification for

Fig. 1 The actual world w_1 validates $\neg K\phi$ but also $\neg K\neg K\phi$

$\neg p$, according to the MK -analysis, would similarly require an impossible state of affairs for the obtainment of $|K\neg p|$ being incorporated in the subject's knowledge state in a certain way, so that S is not in a position to rule out that $\neg K\neg p$ might be the case.

Things get even more straightforward once we adopt Rosenkranz's $\neg K\neg K$ -analysis: to say that ϕ might be the case is simply to say that ϕ is true at a (possible or impossible) world epistemically indistinguishable to the subject from the actual world; and to say that $K\phi$ might be the case is to say that there is a (possible or impossible) world epistemically indistinguishable to the subject at which $K\phi$ is true. Based on this, we can devise a toy model for a non-trivial case of justification (case in which the subject has justification but is short of knowledge) as shown in Fig. 1.³

In a sense, our toy model reflects both the 'negative' and the 'positive' aspects of justification. By the negative aspect, the subject is incapable of distinguishing the actual world w_1 from w_2 , hence cannot know whether it is ϕ or $\neg\phi$ that is actually the case. By the positive aspect, the subject also 'sees' a world w_3 and hence the possibility for $K\phi$ to hold. Since the subject is not in a position to know $\neg K\phi$, we have $\neg K\neg K\phi$ validated by w_1 , too. (In fact, $\neg K\neg K\phi$ is validated by all the three worlds in the model, making it a luminous condition as Rosenkranz desires.)

The idea can also be put in the way of dynamic epistemic logic: the model above is an indication of how the subject's knowledge state *could* potentially evolve by upcoming information: Given new information, the aim is that the subject would be able to either eliminate w_2 or w_3 . Eliminating w_2 , she would be in a position to know the truth of ϕ . Eliminating w_3 , she would be in a position to know $\neg K\phi$. Thus, under the $\neg K\neg K$ rule's construal, a non-trivial case of justification (case when $K\phi$ does not already hold) corresponds to an epistemic state with two kinds of potentials: the subject could either get to know the truth of ϕ (or $\neg\phi$), or get to know that it is actually something one is in no position to know. Either way, the intuition behind the analysis is that it suffices to indicate that the subject in such a state has already done her due diligence in trying to settle the truth for ϕ , and is thus qualified to count as being justified.

³In his book, Rosenkranz offered a semantics for his logics using neighborhood frames (cf. Rosenkranz 2021, §5.2). But in this paper we will be using standard Kripke frames instead. These can be turned into neighborhood frames if one wishes. It can also be verified that our toy model is indeed a model for Rosenkranz's logic of K , satisfying all its related axioms but also invalidating some undesired formulas like (4_K) and (5_K) , i.e. the ones for positive and negative introspection.

3 The structurally unknowable

The worry, however, is that the model above may nonetheless fall short of showing that the subject has indeed done her due diligence for justification: There could be handicaps underlying the subject's actual epistemic state, preventing the subject from achieving either knowledge of ϕ or knowledge of $\neg K\phi$. Rosenkranz has in fact discussed such a case in his book, which goes as follows:

Yonsei Library I want to find out whether ϕ . I know that if the answer to the question is to be found anywhere, it is to be found in the library of Yonsei University. I already have a tentative list of books that I need to consult, but for all I know there may be more. Unbeknownst to me, I am suspected to be a spy and the Korean authorities have instructed the library staff to secretly remove any relevant book from the shelves whenever I approach the aisle. Thus, there are factors that jointly determine that, as I browse the library, I am in no position to know that ϕ . But since the library staff are very skilled, I am, in that situation, in no position to know about their actions, and hence in no position to know that I am in no position to know that ϕ . Yet, clearly in that situation, I have no justification for $|\phi|$. Rosenkranz (2021, 125)

In response, Rosenkranz points out that:

The structure of the example is this: there is a ' ψ ' such that ' ψ ' is true, its being true determines the truth of ' $\neg K\phi$ ', and ' $\neg K\psi$ ' holds. From this characterization of ' ψ ' it is then inferred that ' $\neg K\neg K\phi$ ' holds. However, even if ' ψ ' is true and its being true determines the truth of ' $\neg K\phi$ ', it by no means follows that the only way for one to be in a position to know that $\neg K\phi$ is by being in a position to know that ψ . [...] By doing what I am in a position to do to find out whether ϕ , I am in a position to realize that my attempts are frustrated, without coming to know why. [...] Soon enough I realize that I have exhausted the options available to me in my present situation. (Ibd.)

I think Rosenkranz is right in pointing out that, in the *Yonsei Library* case, even if the truth of ψ is supposed to determine the truth of $\neg K\phi$, ignorance of the former does not necessarily imply ignorance of the latter. As he rightly maintains, there could be some other way for the subject to be in a position to eventually come to a conclusion as to whether $\neg K\phi$ de facto holds. Let us call a fact ϕ *first-order structurally unknowable* with respect to ψ if, and only if, $\neg K\phi$ is true whenever $\neg K\psi$ is. But indeed, our counterstrategy should *not* be about finding such a ψ , such that one's ignorance of it makes $\neg K\phi$ true: If ψ represents the fact that I am an envatted brain, being ignorant of this surely guarantees that I am in no position to know mostly anything. But that is nothing problematic after all, since the concern here is only about my justification in believing in certain things.

Nevertheless, I do think there are certain kinds of structural unknowability that can threaten one's justification. For this, we just need to give the above analysis a further twist: instead of looking for a ψ such that the ignorance of ψ makes $\neg K\phi$ true, our counterstrategy should be to find a ψ such that the ignorance of ψ makes $\neg K\neg K\phi$

true, and at the same time, makes it bad enough for the subject to count as justified. Let us say that ϕ is *second-order structurally unknowable* with respect to ψ if, and only if, being ignorant of ψ , the subject is in no position to know her ignorance of ϕ . We only need show that, at least for some such second-order cases of structural unknowability, the subject's ignorance of ψ undermines her status of justification. Consider the following case:

Chemistry Professor Suppose Professor Moore is a chemistry professor who wants to hire a lab assistant. There are two candidates, Amy and Ben. Professor Moore can see that Ben has a slightly higher average grade in chemistry courses. He takes it to be important and, after careful deliberation, decides to hire Ben. But actually, Professor Moore is also a sexist and is biased against female candidates, which he himself is not aware of. Let us even imagine that he has hired quite a few lab assistants before, all of whom are male. He hired each of them after similar careful deliberations and they all later turned out to be good choices—Those assistants all have had very good performances at the lab indeed. Finally, suppose that Professor Moore has since long been detached from day-to-day lab work. So, unbeknownst to him, it is actually the candidate's programming and data analysis skills that is by far the more important quality for being a good lab assistant. But Professor Moore fails to take that into consideration completely. It is just mere coincidence that all those male assistants he had hired before happen to have good programming and data analysis skills. And once again by coincidence, Ben happens to have better programming and data analysis skills than Amy does.

In the above example, Professor Moore is confident of his hiring strategy, which is based on careful deliberation of each candidate's scores in chemistry courses. Also relying on that false belief, he is confident that Ben is the better choice, without ever being able to find out that his confidence is actually grounded in his sexist biases. Let $|\phi| = |\ulcorner \text{Ben is the better candidate to hire} \urcorner|$. Professor Moore cannot form a safe belief of ϕ , thus he is in no position to know ϕ , even though it is true at the actual world. Furthermore, let $|\psi| = |\ulcorner \text{The candidate's average grade in chemistry courses is not the most important quality to consider} \urcorner|$, which is true at the actual world. Professor Moore's predicament is precisely due to his ignorance of ψ in his knowledge structure. This makes him unable to form a safe belief of $\neg K\phi$. Hence, he is in no position to know that he is in no position to know ϕ . But intuitively, neither does he have justification to believe ϕ .

More formally, we can consider the toy model presented in Section 2 again. In that model, w_3 is the only world that makes $K\phi$ true. Hence, w_3 's indistinguishability from the actual world w_1 makes $\neg K\neg K\phi$ true. Now, if we add a proposition $|\psi|$ to the model, such that ψ is true at w_1 and w_2 but false at w_3 , the subject's ignorance of ψ will then ensure that she is in no position to know $\neg K\phi$. In other words, ϕ is second-order structurally unknowable with respect to ψ .

Figure 2 gives it a further generalization: we add another world w_4 to the model, which makes $K\psi$ true. The addition of w_4 only serves the purpose of avoiding $\neg K\psi$ from being a necessary truth. (It also ensures that both ψ and $K\psi$ (and their negations) have different extensions as ϕ and $K\phi$ (and their negations) do, such that they

Fig. 2 A case of second-order structural ignorance of ϕ with respect to ψ

are all mutually independent facts.) It is easy to check that, according to the new model, while ϕ is true at w_1, w_3, w_4 , and ψ is true at w_1, w_2, w_4 , we have $K\phi$ true at w_3, w_4 , and $K\psi$ true only at w_4 ; furthermore, both $\neg K\neg K\phi$ and $\neg K\neg K\psi$ are true at all worlds.

Again, we can view the model from the perspective of dynamic epistemic logic: Should Professor Moore acquire some new evidence, he might be able to eliminate worlds from the model. Of course, if he gets direct evidence for ϕ (e.g. by a reliable testimony), he would be able to eliminate w_2 and know ϕ . But it could also be the case that he eventually comes to know that he is in no position to know ϕ . According to the $\neg K\neg K$ rule, that is the (only) case where the subject loses justification for ϕ . Yet by our model, coming to know $\neg K\phi$ requires the elimination of w_3 , i.e. it requires the subject to be in a position to know ψ first; thus, whenever the subject is in no position to know ψ , he must have justification to believe ϕ . However, at least from *Chemistry Professor*, we discover that being unaware of the truth of ψ actually makes Professor Moore fall short of having justification for ϕ .

At this point, one might raise the objection that the assumption that Professor Moore is in no position to know ψ is in itself either questionable or ad hoc: by doing what he is in a position to do to find out whether Ben should be hired, there is no good reason to deny that Professor Moore *is* in a position to know that programming and data analysis skills are more important for the hiring process; thus, there is no good reason to maintain that $\neg\psi$. In response, notice that, rather than just providing a single counterexample, the *Chemistry Professor* case points towards a pattern for generating more counterexamples, based on the following, general formal feature in Rosenkranz's analysis: According to his hyperintensional semantics for knowledge, given any non-trivial case of justification—any case in which $K\phi$ does not already hold, so that there are at least a ϕ -world and an $\neg\phi$ -world epistemically indistinguishable to the subject (just like w_1 and w_2 in our model)—we can *always freely* construct an epistemically 'possible world' w_{3x} , and accordingly, some unknown fine-grained proposition $|\psi_x|$, which prohibits the subject from forming a safe belief of $\neg K\phi$. Therefore, our model in Fig. 2 is rather an exemplification of a general

pattern for constructing such $|\psi_x|$.⁴ And since a proposition like $|\psi_x|$ is generally free for construction in one's knowledge structure, it suffices for us to claim that there is *some* $|\psi_i|$ such that (i) the subject is ignorant of ψ_i and (ii) while the ignorance of ψ_i implies $\neg K \neg K \phi$, it is bad enough for the subject to count as justified.⁵

This is not to deny that, for some propositions $|\psi_x|$, the scenarios generated may be too improbable to affect our intuitive assessment of the subject as being justified. For instance, consider a different scenario in which Professor Moore has been right in holding that one's grades in chemistry courses are indeed the most important quality to consider, but let $|\psi_1| = |\ulcorner \text{Ben is providing a forged transcript of his courses} \urcorner|$. Suppose that, unbeknownst to Professor Moore, ψ_1 is actually true. In such case, again, Moore's ignorance of ψ_1 makes him unable to form a safe belief of $\neg K \phi$. (Had the transcript that Ben provided *not* been forged, Moore would be in a position to know that Ben is the right one to hire.) But perhaps this will not deprive him of the justification in believing ϕ anyway. Nevertheless, since Rosenkranz's semantics puts virtually no formal constraint on which kind of $|\psi_x|$ is admissible and which is not, it is just implausible to maintain that the subject must be justified *for any* proposition $|\psi_x|$. In fact, as I will discuss later, the subject's justification seems to be sensitive to what particular kind of content ψ_x represents.

4 Justification and the structural incongruence in epistemic opportunities

There are two possible responses from $\neg K \neg K$ -theorists. The first response is to label Professor Moore as defective agents—who is disqualified from being 'suitably improved version of ourselves', so that the $\neg K \neg K$ rule does not apply. In contrast, if there was a 'suitably improved version' of Professor Moore, he would be free of any biases, sexist or any sort; and we would have to grant that he *is* capable of forming a safe belief either of ϕ or of $\neg K \phi$. However, we have seen that in the toy model,

⁴This hence opposes what Rosenkranz claims in his argument for Thesis (7) (Rosenkranz 2021, 70). My tentative diagnosis is that Rosenkranz's Thesis (5) for safe belief may not hold in cases where structural unknowability is at stake, either because in such cases what is 'the best one is in a position to do to decide' could differ substantially, or because what counts as the 'closest' cases could become significantly vague. Also note that, in our toy model, we have only considered two propositions, $|\phi|$ and $|\psi|$. In reality, much more background information can be at stake. But this does not have to affect how w_3 and $|\psi|$ can be constructed, because in simple cases, the worlds in the model have no need to differ from each other in the truths of other propositions apart from those of $|\phi|$ and $|\psi|$. (Of course, the worlds also differ on what epistemic accessibility relations each of them instantiates. But these would be meta-modal information about the subject's knowledge structure that the subject herself should have no access to.)

⁵It might be tempting to also consider the following proposition $|\chi| = |\ulcorner \text{Moore fails to carry out his decision procedure in a fair and unbiased way} \urcorner|$. Suppose that χ is actually true. Then, if Moore is in a position to know χ , that seems to be *sufficient condition* for him to achieve negative introspection, i.e. for $K \neg K \phi$ to hold. But this also means that, unlike ψ , ignorance of χ may not be a sufficient condition for $\neg K \neg K \phi$. Despite this, the behavior of χ is an interesting point worth further exploration. E.g., it appears to be representing a higher-order proposition implicitly quantifying over various propositions $|\neg K \psi_x|$. Thanks to an anonymous referee for drawing my attention to this point.

w_3 and $|\psi|$ is free for construction. Thus, $\neg K \neg K$ -theorists who adopt this strategy would have to maintain that, for *any* proposition like $|\psi|$, it cannot be the case that the subject is in no position to know it. This would essentially mean a rejection of Fig. 2 as a valid model. But if so, the intended construal of the ‘suitably improved version of ourselves’ would actually be requiring a stricter logical account of knowledge beyond the realistic system that Rosenkranz has offered.

The second response, on the contrary, is to admit that Professor Moore in our story can indeed be seen as a member of the ‘suitably improved version of ourselves’. This will lead to the acceptance, contra some of our intuition, that Professor Moore *does* have justification to believe that Ben is the better candidate to hire. But again, as we have seen, given the ultra fine-grained modeling of knowledge states, one will also have to claim that the subject’s status of justification will not be undermined by *any* of the indefinitely many propositions that could be assigned to ψ_x . It is hence a burden on the defenders of the $\neg K \neg K$ rule to show that, for any such $\neg K \neg K$ -agent and any $|\psi_x|$, the agent indeed has justification for ϕ . Therefore, while the first type of response is too demanding, this second response risks making it too permissive for one to count as having justification. This is no surprise, since the formal structure of the $\neg K \neg K$ analysis is unable to differentiate between the various contents of ψ_x .

It is also worth noting that the problem of the structurally unknowable is different from the problem of the ‘pathological’ kind of ignorance that ‘trees, madmen, drunkards and dead people’ are subject to, which Rosenkranz has reasonably dismissed (Rosenkranz 2021, 119). There are of course pre-semantic, pre-modeling, objective conditions for someone to be eligible as a subject of the analysis. But the barriers for a dead person to count as a ‘suitably improved versions of ourselves’ is clearly not to be modeled epistemically. In contrast, the structurally unknowable seems to be a purely epistemic barrier. It is part of the formal feature derived from the hyperintensional structure itself, which is essentially freed from the need to take any other non-epistemic modalities—neither the logical nor the ‘objective’ ones—into consideration.

The structurally unknowable is somewhat like what Roy Sorenson calls ‘blindspots’ (Sorensen, 1988). They are constraints or handicaps on one’s epistemic opportunities that have nothing to do with one’s cognitive incompetency. A way to understand such constraints might be to regard them as consisting in the fact that knowledge can be information- or topic-sensitive.⁶ The truth of ϕ may be structurally unknowable to a subject if it is not part of the subject’s research questions whether ψ , and this could be a sign of the subject’s epistemic limitations—not just in acquiring specific information at issue, but also, and more importantly, in having the right kind of research questions and knowledge-seeking strategies that are relevant for the subject’s decision procedure for ϕ . For instance, the lesson from *Chemistry Professor* is precisely that, even if all information about each candidates are available for Moore’s consideration, for him to make a justified hiring decision, some extra epistemic ingredients—i.e. the right strategies for finding out the truth of ϕ —are needed

⁶Cf. typically Hintikka (2007) and Berto and Hawke (2021), etc., but also Rosenkranz (2007, 70) and Rosenkranz (2021, VIII), *pace* (Rosenkranz 2021, 159).

to ensure that Moore is able to carry out his decision procedure based on a proper interpretation of the information.

Justification as Ignorance is a formidable book. As a knowledge-first-style analysis of justification, it not only generates a great deal of new understanding about justification, but also invites us to reflect on the role knowledge plays in our cognitive life and decision making. The discussion of the structurally unknowable in this paper is by no means to dismiss the effort to analyze justification. The point I wish to make is only this: while justification is a particular kind of ‘opportunity’ to know, the $\neg K \neg K$ rule might nonetheless be falling short of revealing its full structure; and it remains to be investigated what kind of ‘epistemic opportunity’ justification exactly is. Perhaps, using a single epistemic modality is just inadequate, and a more complicated combination of modalities will be needed. An ‘incongruent view’ in form of $J\phi = MK\phi$, for instance, will need to adopt a more restrictive approach (assuming the fine-grained modeling of knowledge states) on which kind of epistemic possibilities for knowledge that a justified agent has access to. Therefore, for propositions $|\psi_x|$ relative to which ϕ is structurally unknowable, the new modality ‘ M ’ will have to sort them out in such a way that, for any ψ_x that undermines the agent’s status of justification, $MK\phi$ does not hold. These will be an interesting question to further explore but exceed the scope of the current paper.⁷ Nevertheless, we have at least seen that whether an $\neg K \neg K$ -agent really has justification stays an open question, and the analysis of justification has revealed itself to be a harder task yet to be fulfilled.

Acknowledgements An earlier draft of this paper was presented at a book symposium hosted by the Asian Epistemology Network. I would like to thank all the organizers and participants of that event. Special thanks to Sven Rosenkranz for his very helpful comments and feedbacks.

Funding This work is supported by the Fundamental Research Funds for the Central Universities in China.

Availability of data and material The author confirms that all data generated or analysed during this study are included in this published article.

References

- Beddor, B., & Goldstein, S. (2021). Mighty knowledge. *Journal of Philosophy*, 118(5), 229–269.
 Berto, F., & Hawke, P. (2021). Knowability relative to information. *Mind*, 130(517), 1–33.
 Dasgupta, S. (2015). Inexpressible ignorance. *Philosophical Review*, 124(4), 441–480.

⁷Here is a quick sketch of some ideas: The topic-sensitivity of knowledge indicates that there are actually two different kinds of epistemic aims: one about leaning more truths, the other about asking the right research questions and adopting the right knowledge-seeking strategies. If so, the interaction between ‘ M ’ and ‘ K ’ must be constrained by both aims. Moreover, perhaps one can go even further and claim that pure epistemic modalities do not suffice for generating the right kind of analysis of justification, since in many cases the right kind of questions we should ask have to be questions rightly about things in our world (cf. e.g. Dasgupta 2015, 475f.). If so, we may nonetheless need to take certain kind ‘objective modality’ into consideration.

- Hintikka, J. (2007). *Socratic Epistemology*. Cambridge: Cambridge University Press.
- Rosenkranz, S. (2007). Agnosticism as a third stance. *Mind*, 116(461), 55–104.
- Rosenkranz, S. (2021). *Justification as Ignorance*. Oxford: Oxford University Press.
- Sorensen, R. (1988). *Blindspots*. Oxford: Clarendon Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.