Asian Journal of Philosophy (2022) 1: 29
https://doi.org/10.1007/s44204-022-00021-7

BOOK SYMPOSIUM

Replies to critics

Sven Rosenkranz[1,2]

Received: 28 February 2022 / Accepted: 31 March 2022 / Published online: 22 June 2022

Abstract
Responding to the critical comments by the symposiasts Julien Dutant, Niccolò Rossi, Martin Smith, Daniel Waxman, and Yiwen Zhan, I further elaborate, refine, and defend the account of epistemic justification I advanced in my book *Justification as Ignorance* (Oxford: Oxford University Press, 2021). Central issues tackled here include logical omniscience and non-normal epistemic logics, the luminosity of epistemic justification, methods for telling what one is in no position to know, and the relation between doxastic and propositional justification.

Keywords Epistemic justification · Non-normal epistemic logics · Luminosity · Safety · Knowledge · Being in a position to know · Epistemic methods · Cognitive limitations · Epistemic possibility · Doxastic versus propositional justification

Abbreviations
JAI Justification as Ignorance

1 Introduction

I am very grateful to Nikolaj Pedersen, editor of the *Asian Journal of Philosophy*, for organizing this symposium, as one of the journal's first, and to Julien Dutant, Niccolò Rossi, Martin Smith, Daniel Waxman, and Yiwen Zhan for their excellent and challenging comments that made me rethink central aspects of the account of justification presented in *Justification as Ignorance* (Rosenkranz, 2021; henceforth JAI) and, at the same time, gave me the opportunity to clarify or amend others. Lucky me. It is very rare to have such a formidable cast of thinkers dedicate so much time and energy in order to digest and criticize one's ideas. In what follows, I will do my best to answer their challenges.

✉ Sven Rosenkranz
rosenkranz.sven@gmail.com

[1] ICREA, Passeig de Lluís Companys, 23,08010 Barcelona, Spain

[2] Department of Philosophy, University of Barcelona, C. Montalegre, 6-8, desp. 4009, 08001 Barcelona, Spain

Although I have conceived of my replies as self-standing pieces, they clearly hang together, which is why they frequently include cross-references. To the interested reader familiar with JAI, a brief roadmap might be helpful. In my reply to Rossi, I further clarify my misgivings about normal epistemic logics underlying much of chapters two, four, and five of JAI. My reply to Smith elaborates, clarifies, and refines the argument for the luminosity of propositional justification, on my construal of it, that I gave in chapter four of JAI. It enables me to stave off Smith's own anti-luminosity challenge, or so I argue. In different ways, Waxman, Dutant, and Zhan all challenge the pair of theses, put forward in chapter six of JAI, that being in no position to rule out that one is in a position to know *p* suffices for having propositional pro toto justification for *p* and that being in no position to rule out that one knows *p* suffices for having doxastic pro toto justification for *p*. Addressing Waxman's alleged counterexample to the first, I further refine and amend my conception of what it is to do one's best in order to decide a given question—a notion already at work in my reply to Smith. Dutant primarily targets the second thesis—more specifically its implication that one is always in a position to know either that one doesn't know *p* or that one doesn't know *p*'s negation. In response, I clarify what I take to be the evidential force of misleading impressions that one knows the answer to a given question, in the process of doing one's best to find out that one actually doesn't. Like Waxman, Zhan presents what he takes to be a counterexample to the first of the aforementioned theses which is meant to improve upon an example I discuss in chapter six of JAI. My reply gives me the opportunity to say more about what I take to be the connection between propositional and doxastic justification.

2 Reply to Niccolò Rossi

In chapter five of JAI, I devised two distinct bimodal systems for both knowledge (k) and being in a position to know (K), one I called 'idealized', the other 'realistic' (JAI: 80-101). The realistic system is contained in the idealized system. Both systems are non-normal logics, rejecting the Kripke axiom for either modality, alongside a number of other principles that I had argued in earlier chapters were objectionable on general epistemological grounds. Unlike the realistic system, the idealized system nonetheless still contains implausibly strong assumptions—which is why I settled for the realistic system in the end. Instead of establishing the soundness of the realistic system directly, in terms of a semantics geared to its specific needs that allows for models that are not, at the same time, models of the idealized system, in JAI I did so indirectly by devising models for the idealized system. In this context, I also provided a class of models for the idealized system invalidating many of the unwanted principles earlier identified. However, as Niccolò Rossi observes, the models in that class still failed to invalidate either of the following:

$$\mathbf{RN}_k \qquad \varphi / k\varphi$$

$$\boldsymbol{K\text{–}k} \qquad K\varphi \supset k\varphi.$$

According to the first, if φ is a theorem, so is ⌜$k\varphi$⌝, implying that it is correct to partly characterize what the targeted epistemic agent knows by using any given logical truth. According to the second, whatever such an agent is in a position to know, they do know.

In his insightful contribution to the symposium, Rossi devises a better class of models for the idealized system that invalidate both $\mathbf{RN}_k$ and ***K-k*** in addition to those invalidated by the ones I devised (Rossi, 2022). This is a welcome improvement. Invalidating ***K-k*** is of particular interest, as this principle would lead to a collapse of the two modalities, given only that its converse is valid in both systems. Failure to provide models that invalidate ***K-k*** is a clear shortcoming of JAI (even if such models can easily be given). According to Rossi's assessment, showing that the idealized system has models invalidating $\mathbf{RN}_k$ is equally important, since $\mathbf{RN}_k$ is one ingredient of the so-called thesis of logical omniscience, involving, as Rossi contends, 'a strong idealization, requiring that one knows any logical truth, even the most convoluted' (Rossi, 2022: 3).

Since the realistic system is strictly weaker than the idealized system, Rossi's result shows that the realistic system, too, has models invalidating all those principles at once. To invalidate those principles all at once has the benefit that, in Rossi's words, 'we don't need to assume one [unwanted principle] to invalidate another'. This is particularly important with respect to the purpose of invalidating ***K-k*** alongside other unwanted principles (and it is not as easy to come up with such models as it is to come up with models that invalidate ***K-k*** in isolation).

However, it transpires that the idealized system still has unwanted features that ultimately make the attempt to show that it has models invalidating $\mathbf{RN}_k$ somewhat moot. To begin with, note that the following monotonicity rule is valid in the idealized system:

$$\mathbf{RM}_k \qquad \varphi \supset \psi / k\varphi \supset k\psi.$$

As Rossi observes, with $\mathbf{RM}_k$ in place, models invalidating $\mathbf{RN}_k$ must involve worlds in which nothing is known (Rossi, 2022: 4). $\mathbf{RM}_k$ straightforwardly entails the equivalence rule

$$\mathbf{RE}_k \qquad \varphi \equiv \psi / k\varphi \equiv k\psi.$$

The logic for 'k' is intended as a logic of knowledge ascriptions. As it turns out, the validity of $\mathbf{RE}_k$ already requires a conception of such ascriptions that, at the same time, assuages any worries about $\mathbf{RN}_k$.

To see this, compare two basic kinds of conceptions of knowledge ascriptions (JAI: 20-27). According to the first, the prejacents of such ascriptions give some indication, however indirectly (e.g. via translation), of the distinctive way the thing known is articulated in the mind of the subject to whom knowledge is being ascribed—a distinctive *guise* of the thing known. I say 'articulated in the subject's mind' rather than 'present to the subject's mind', intending a broad, non-Cartesian conception of the mind that encompasses inferential dispositions, as the knowledge ascribed may be implicit.[1] According to the second kind of conception, the prejacents of knowledge ascriptions give no such indication, even if it is agreed that, whenever such an ascription is true, there must be some way or other in which the thing known is articulated in the ascribee's mind—i.e. some guise of the thing known under which the ascribee knows that thing. In the semantics Rossi and I respectively provide for the idealized system, 'the thing known' is a coarse-grained proposition corresponding to a subset of possible worlds.

Plausibly, the thought that $\mathbf{RN}_k$ involves too strong an idealization—one that requires subjects to know 'any logical truth, even the most convoluted' (Rossi, 2022: 3)—presupposes a conception of the former kind. For, plausibly, we shouldn't be worried by the idea that subjects know the universal proposition *in one of its guises* and, if knowledge ascriptions are silent on what the relevant guise might be, it doesn't follow from the validity of $\mathbf{RN}_k$ that subjects know the universal proposition *in all its guises*, corresponding to all logical truths expressing that proposition.

However, a conception of the first kind casts equal doubt on the validity of $\mathbf{RE}_k$. For, if knowledge ascriptions are understood in accordance with such a conception, $\mathbf{RE}_k$, too, imposes overly strong demands on epistemic subjects, viz. that, whenever they know a given coarse-grained proposition under the guise indicated by one prejacent, they know it under the guise indicated by any other, logically equivalent prejacent. But, then, the idealized system ultimately presupposes a conception of the second kind—in which event the worries about $\mathbf{RN}_k$ should likewise dissipate.

What here holds for knowledge ascriptions holds for ascriptions of epistemic states in general, including those that specify what subjects are in a position to know. The idealized system validates the luminosity thesis

$$\mathbf{L} \qquad \neg K\neg K\varphi \supset K\neg K\neg K\varphi.$$

L implies that, in any world of any of the idealized system's models, the subject is in a position to know at least one proposition. The idealized system likewise validates

$$\mathbf{RM}_K \qquad \varphi \supset \psi / K\varphi \supset K\psi.$$

Taken together, these principles allow us to derive

$$\mathbf{RN}_K \qquad \varphi / K\varphi.$$

If we presupposed a conception of truths of the form ⌜$K\varphi$⌝ according to which the prejacent φ indicated the guise under which the subject in question is in a position to know the coarse-grained proposition assigned to φ as its semantic value, then $\mathbf{RN}_K$ would likewise be problematic and almost as problematic as $\mathbf{RN}_k$. Thus, we may ask: if it is important to have a logical system with models invalidating $\mathbf{RN}_k$,

[1] See, however, principle (1) in my reply to Dutant (2022) below.

because of legitimate worries about logical omniscience, why isn't it just as important to have a logical system with models invalidating $\mathbf{RN}_K$?[2]

All this goes to show that, on the first kind of conception of ascriptions of epistemic states, the idealized system proves inadequate.[3] Once a conception of the second kind is assumed, however, there is no longer any sense in which $\mathbf{RN}_k$, $\mathbf{RE}_k$, $\mathbf{RN}_K$, and $\mathbf{RE}_K$, are expressive of any controversial thesis of logical omniscience. Arguably, having the resources to express such a controversial thesis, if only to reject it, is itself a desideratum which, on a conception of the second kind, we simply cannot fulfil. Although models for the idealized system are models for the realistic system, by design they never invalidate all the principles that legitimate concerns about logical omniscience might identify as objectionable. Rossi is, of course, fully aware of this. However, the desire to find models for *both* systems that invalidate $\mathbf{RN}_k$, if driven by such concerns, is ultimately misplaced, because such models cannot but come at the cost of validating $\mathbf{RN}_K$, $\mathbf{RE}_k$, and $\mathbf{RE}_K$.

Things are in fact worse for the idealized system. For, even on a conception of the second kind, $\mathbf{RM}_k$ and $\mathbf{RM}_K$ are wildly implausible. Trivially, contingent coarse-grained propositions entail other, equally contingent ones. There is no guarantee that, if a subject knows, or is in a position to know, some such proposition under one guise or other, she knows, or is in a position to know, all the contingent propositions it entails under some such guise. For, if the entailing proposition is distinct from the proposition entailed, any guise of the former will differ from any guise of the latter and, while it may be deemed uncontroversial to think that we always know, or are in a position to know, the universal proposition in some of its guises, the same cannot be said about all the contingent propositions entailed by something we know or are in a position to know (JAI: 26, 31-32).

The lesson is that, on either conception of the ascriptions of epistemic states, the idealized system proves inadequate. The difficult task that lies ahead, at least for those sympathetic to the overall project, is to devise models for the realistic system that are not, at the same time, models for the idealized system. The task is difficult inter alia because the realistic system still contains principles that are instances of $\mathbf{RM}_k$ and $\mathbf{RM}_K$ and also contains **L**, which latter, as we've seen, ensures that subjects are always in a position to know something—indeed infinitely many things.

3 Reply to Martin Smith

One of the key theses defended, and exploited, in JAI was

[2] One might still find it desirable to leave room for subjects who simply haven't come around to knowing any logical triviality at all, although they are in a position to do so. But such a desideratum is not the same as the desideratum to leave room for a failure of logical omniscience.

[3] In JAI, the idealized system was merely used to establish the soundness of the realistic system and to show that the latter has models in which certain unwanted principles fail. It served as a ladder to be kicked away, once those results had been established, and wasn't meant to have any further, independent interest.

L If one is in no position to know that one is in no position to know *p*, one is in a position to know that this is so.

If we let '*K*' be short for 'one is in a position to know' and conceive of cases as triples of worlds, subjects, and times, we can rephrase **L** as follows: for any case γ, if ⌜¬*K*¬*Kp*⌝ holds in γ, so does ⌜*K*¬*K*¬*Kp*⌝—where, throughout, 'one' refers to the subject at the centre of γ.

L implies that, for any *p*, ⌜¬*K*¬*Kp*⌝ encodes a *luminous* condition, i.e. a condition such that, if it obtains in one's situation, one is in a position to know that it obtains (Williamson, 2000: 13). As such, it is threatened by Timothy Williamson's argument to the conclusion that no non-trivial condition is ever luminous (Williamson, 2000: 126-29; see also Srinivasan, 2013). A condition will be non-trivial, if we can construct a series of cases running from cases in which the condition is clearly met to cases in which it clearly isn't met. There is no denying that, for any given *p*, we can construct such a series for ⌜¬*K*¬*Kp*⌝ (for a nice example, see Smith, 2022: 5). Accordingly, in JAI, I undertook to show that **L** can after all escape the threat posed by Williamson's argument. In doing so, I relied on a number of premises that I sought to independently justify.

In his exceptionally clear and illuminating comments, Martin Smith reconstructs and critically assesses my attempt to insulate **L** from this anti-luminosity challenge and then proceeds to formulate a challenge of his own (Smith, 2022). While he generously concedes that I succeed in defusing the first challenge, he argues that I didn't provide any resources to defuse the second. Below I rehearse, and further clarify, my defence of **L** and argue that it, or a natural extension of it, after all allows me to answer Smith's challenge.

In his reconstruction of my defence of **L**, Smith identifies the crucial premises on which I relied as

$\mathbf{PK_1}$ If ⌜*Kp*⌝ holds and one has done everything one is in a position to do in order to decide *p*, one knows *p*.

$\mathbf{PK_2}$ If one does everything one is in a position to do in order to decide *p*, one does everything one is in a position to do in order to decide ⌜¬*p*⌝, and conversely.[4]

$\mathbf{PK_3}$ If one does everything one is in a position to do in order to decide ⌜*Kp*⌝, one does everything one is in a position to do in order to decide *p*.

While $\mathbf{PK_2}$ is trivial, at least given a natural reading of 'whether' and classical logic, $\mathbf{PK_3}$ is not. Smith is nonetheless willing to accept $\mathbf{PK_3}$ for the sake of argument. Smith assumes that any characterization of the notion of being in a position to know will sanction $\mathbf{PK_1}$. Although closely related, these are not, in fact, the premises I ultimately used in arguing for **L**. With hindsight, it is best to first explain why.

Smith writes that $\mathbf{PK_1}$ 'already gives us the result that being in a position to know is factive' (Smith, 2022: 2). This is incorrect. $\mathbf{PK_1}$ only implies that being in

[4] To be perfectly accurate, Smith doesn't explicitly mention the converse conditional. But I here presume that, as part of a reconstruction of what I say in JAI, it is likewise intended (cf. JAI: 42-43).

a position to know *and* doing everything that one is in a position to do in order to decide p is factive. But, one may be in a position to know p without doing everything that one is in a position to do in order to decide p. $\mathbf{PK_1}$ doesn't tell us anything about such cases. In this regard, it is better to replace $\mathbf{PK_1}$ by:

$\mathbf{PK_1'}$ If ⌜Kp⌝ holds, then, were one to do everything one is actually in a position to do in order to decide p, one would know p.

Unlike $\mathbf{PK_1}$, $\mathbf{PK_1'}$ has implications for cases in which one is in a position to know p without knowing p, viz. cases in which one does not, but could do everything that one is in a position to do in order to decide p. However, ⌜Kp⌝ may hold, even if it is impossible to do *everything* that one is actually in a position to do in order to decide p: implementation of some of the procedures that one is in a position to successfully implement to this end might preclude implementation of others (JAI: 41). This is why it is better to replace $\mathbf{PK_1'}$ by

$\mathbf{PK_1''}$ If ⌜Kp⌝ holds, then, were one to do the best one is actually in a position to do in order to decide p, one would know p.

With hindsight, it is more convenient to rephrase $\mathbf{PK_1''}$ in terms of cases as follows:

P1 For any $\delta \in \mathrm{B}(\gamma, p)$, if ⌜$Kp$⌝ holds in γ, ⌜kp⌝ holds in δ,

where $\mathrm{B}(\gamma, p)$ is the set of possible cases, close to γ and centred on the same subject as γ, in which the subject of γ has done the best they are in γ in a position to do in order to decide p (and has done nothing they are not in γ in a position to do; see JAI: 40).[5] Correspondingly, $\mathbf{PK_2}$ and $\mathbf{PK_3}$ should be replaced by

P2 $\mathrm{B}(\gamma, p) = \mathrm{B}(\gamma, \ulcorner \neg p \urcorner)$

P3 $\mathrm{B}(\gamma, \ulcorner Kp \urcorner) \subseteq \mathrm{B}(\gamma, p)$.[6]

Plausibly, it's always possible to do the best that one is in a position to do in order to decide a given question—even if that may amount to doing very little—where some such possibility is furthermore close: if any possibility in which one Fs is remote, one isn't in a position to F. Thus, $\mathrm{B}(\gamma, p) \neq \varnothing$. Accordingly, **P1** implies that, if ⌜Kp⌝ holds in γ, it is a close possibility that one knows p after having done the best that one is in γ in a position to do in order to decide p. To the extent that doing something in order to decide a given proposition leaves the truth-value of that

[5] For more on the intended reading of 'the best', see principle **P7** in my reply to Daniel Waxman (2022) below.

[6] In JAI, I also assumed $\mathrm{B}(\gamma, p) \subseteq \mathrm{B}(\gamma, \ulcorner Kp \urcorner)$ (see principle (2c) in JAI: 46-48). But, for reasons that emerge in the context of my reply to Waxman (2022), this was overly simplistic.

proposition unaffected (Williamson, 2009), **P1** will then also ensure the factivity of '*K*'.

These are not the only premises I need, however, and I will invoke two further principles below. But before we get to that, let me make two observations:

- Since knowledge implies belief, **P1** in turn implies that, if ⌜Kp⌝ holds in γ, one wouldn't do the best that one is in γ in a position to do in order to decide *p*, were one not to form a belief in *p* in the process.

Accordingly, if one is in a position to know *p*, one is in a position to believe *p* in the course of doing the best that one is in a position to do in order to decide *p* (JAI: 39).

- **P1** does *not* imply the conditional that, if ⌜Kp⌝ holds in case γ, then, for any case γ′ centred on the same world and subject as γ and succeeding γ by very little, if $B(\gamma, p) \neq \varnothing$, then, for some $\delta \in B(\gamma, p)$, $\delta \in B(\gamma', p)$.

To use an analogy for illustration, it may be a close possibility in γ that I jump through the loop, and a close possibility in γ′ that I jump through the loop, where γ′ succeeds γ by very little, while it is not a close possibility in γ′ that I jump through the loop having jumped through the loop a little earlier. If so, then, even if I do in fact jump through the loop in γ, in γ′ it may still be a close possibility that I jump through the loop, albeit one in which I haven't jumped through the loop a little earlier (something, we assume, I might easily not have done).

Let us now assume a series of cases, centred on the same subject and world, running from cases in which ⌜$K\neg Kp$⌝ clearly holds to cases in which ⌜$\neg K\neg Kp$⌝ clearly holds, where for any pair of adjacent cases in the series, their temporal distance can be arbitrarily small (this is the assumption Smith labels '$\mathbf{AL_1}$'). Suppose that I am that subject, that α is the last case in the series in which ⌜$K\neg Kp$⌝ holds, that α′ is the first case in the series in which ⌜$\neg K\neg Kp$⌝ holds, and that like any other pair of adjacent cases in the series, α and α′ are exactly one nanosecond apart.

In his reconstruction of Williamson's anti-luminosity argument, Smith assumes, as Williamson (2000), Srinivasan (2013), and others before him, that throughout the envisaged series of cases, I may in fact, in *each* case, do the best that I am, in that case, in a position to do in order to decide ⌜$\neg K\neg Kp$⌝. This implies that, in α′, it is an easy possibility that I do the best that I am in α′ in a position to do in order to decide ⌜$\neg K\neg Kp$⌝ after having done the best that I was in α in a position to do in this regard. Given that '*K*' does not agglomerate over conjunction (Heylen, 2016), this assumption, though hardly ever questioned, should immediately raise suspicion.

The assumption is indeed far from innocent. Ex hypothesi, ⌜$K\neg Kp$⌝ holds in α. Suppose that, for purely psychological reasons, whenever one knows a given proposition *q*, one couldn't easily believe, a nanosecond later, that one is in no position to know *q*. Call this thesis '**T**'. Given **P1**, **T** implies that, if, in α, I have done the best that I am in α in a position to do in order to decide ⌜$\neg Kp$⌝, I don't believe ⌜$\neg K\neg Kp$⌝ in α′ and hence do not know ⌜$\neg K\neg Kp$⌝ in α′. This alone doesn't show that ⌜$K\neg K\neg Kp$⌝ doesn't hold in α′. It at most implies that, if I do my best in α and ⌜$K\neg K\neg Kp$⌝ holds in α', then, in α', I don't do the best that I am in α' in a position to do in order to decide $\ulcorner \neg K \neg Kp \urcorner$. For, there may then still be a case β' close to α' such that, in β', I do the best that I am in α' in a position to do in order to decide $\ulcorner \neg K \neg Kp \urcorner$ and come to know $\ulcorner \neg K \neg Kp \urcorner$ as a result—even if, given **T**, any such case would have to be one in which I haven't done the best that I was in a position to do to this end a nanosecond earlier.[7] To demand that, if $\ulcorner K \neg K \neg Kp \urcorner$ holds in α', as **L** predicts, there should be some such β' in which I have also done the best that I was in a position to do to this end a nanosecond earlier is to pitch **T** against **L**. Although I see in the end no compelling reason to accept **T**,[8] I don't see any compelling reason either why I should be committed to Smith's assumption which implies a conflict between **T** and **L**.[9]

If $\ulcorner K \neg K \neg Kp \urcorner$ holds in α', then, as we've observed, given **P1**, for any case γ' close to α' in which I do not believe (and hence do not know) $\ulcorner \neg K \neg Kp \urcorner$, I will not, in γ', have done the best that I am in α' in a position to do in order to decide $\ulcorner \neg K \neg Kp \urcorner$. Forming a belief in $\ulcorner \neg K \neg Kp \urcorner$ will then be a necessary condition on doing the best that I am in α' in a position to do in order to decide $\ulcorner \neg K \neg Kp \urcorner$. The same is true, mutatis mutandis, of any case γ close to α in which I do not believe (and hence do not know) $\ulcorner \neg Kp \urcorner$: in such a case γ, I will likewise not have done the best that I am in α in a position to do in order to decide $\ulcorner \neg Kp \urcorner$, as $\ulcorner K \neg Kp \urcorner$ holds in α. Forming a belief in $\ulcorner \neg Kp \urcorner$ is a necessary condition on doing the best that I am in α in a position to do in order to decide $\ulcorner \neg Kp \urcorner$.

Insofar as I am rational, as suitably idealized subjects are supposed to be, I will not believe $\ulcorner \neg Kp \urcorner$ and believe $\ulcorner \neg K \neg Kp \urcorner$ on the same occasion. Let us follow Smith and define, for any case γ in the particular series of cases under consideration: δ is a *near duplicate* of γ iff (i) $\ulcorner \neg K \neg Kp \urcorner$ holds in γ iff it holds in δ, (ii) $\gamma \in B(\gamma, \ulcorner \neg K \neg Kp \urcorner)$ iff $\delta \in B(\delta, \ulcorner \neg K \neg Kp \urcorner)$, and (iii) γ and δ 'are very similar in all other respects' (Smith, 2022: 3-4).[10] Then, on assumption that, for some $\beta' \in B(\alpha', \ulcorner \neg K \neg Kp \urcorner)$, I believe $\ulcorner \neg K \neg Kp \urcorner$ in β', the foregoing already falsifies the following anti-luminosity premise (corresponding to Smith's $\mathbf{AL_2}$):

7 While the cases in the series immediately preceding or succeeding α' are among the cases close to α', the latter also include cases that are centred on a world distinct from that of α'.

8 Below, I'll argue that, if, in α', I do the best that I am in α' in a position to do in order to decide $\ulcorner \neg K \neg Kp \urcorner$ and believe $\ulcorner \neg K \neg Kp \urcorner$ as a result, α' will be a case in which I believe $\ulcorner \neg K \neg Kp \urcorner$ inter alia because I believe p and, if, in α, I do the best that I am in α in a position to do in order to decide $\ulcorner \neg K \neg Kp \urcorner$ and believe $\ulcorner \neg Kp \urcorner$ as a result, α will be a case in which I believe $\ulcorner \neg Kp \urcorner$ inter alia because I don't believe p. The assumption that I do both doesn't seem to be psychologically any more problematic than the idea that, while I continue to do the best that I am in a position to do in order to decide p, I at some point stop believing p as a result.

9 This may call for revision of some of the things I say about opportunities to know in JAI. It no longer strikes me as helpful to demand that, if I do not know p, but am in a position to know p, I have the opportunity to go on and do what I am so far merely in a position to do. Am I in a position to know that there is, for the flicker of a moment, a gnat on the rim of my glass, if I miss it because (a) I happen to blink, (b) I happen to have yawned a few seconds earlier so as to make my eyes water, or (c) I happen not to have started to pay attention in time? I think we would still want to say that the gnat's sitting on the rim of my glass in plain view is the kind of condition, however fleeting, that I am in a position to ascertain (borderline cases aside). Thanks to Julien Dutant for discussion.

10 All near duplicates of γ will be understood to be cases close to γ, but not vice versa.

A2 For any $\beta' \in B(\alpha', \ulcorner \neg K \neg Kp \urcorner)$ such that I believe $\ulcorner \neg K \neg Kp \urcorner$ in β', for some $\beta \in B(\alpha, \ulcorner \neg K \neg Kp \urcorner)$, there is a near duplicate of β in which I believe $\ulcorner \neg K \neg Kp \urcorner$.

For, given **P2** and **P3**, any near duplicate of any $\beta \in B(\alpha, \ulcorner \neg K \neg Kp \urcorner)$ will be a case in which I believe $\ulcorner \neg Kp \urcorner$ and so, being rational, don't believe $\ulcorner \neg K \neg Kp \urcorner$.

For any q, if γ is a case near the borderline between q's being true and q's being false, then, in γ, one is in a position to believe q—and so there is a case γ' close to γ in which one does believe q. Ex hypothesi, case α' *is* near the borderline between $\ulcorner \neg K \neg Kp \urcorner$'s being true and $\ulcorner \neg K \neg Kp \urcorner$'s being false. Thus, it is here unproblematic to assume that I am, in α', indeed in a position to believe $\ulcorner \neg K \neg Kp \urcorner$. As we've observed, in order for $\ulcorner K \neg K \neg Kp \urcorner$ to hold in α', believing $\ulcorner \neg K \neg Kp \urcorner$ must be a necessary condition on doing the best that one is in α' in a position to do in order to decide $\ulcorner \neg K \neg Kp \urcorner$. The question accordingly is whether, in α', I am, furthermore, in a position to believe $\ulcorner \neg K \neg Kp \urcorner$ *upon doing the best that I am in* α' *in a position to do in order to decide* $\ulcorner \neg K \neg Kp \urcorner$—and so whether there is a $\beta' \in B(\alpha', \ulcorner \neg K \neg Kp \urcorner)$ such that I believe $\ulcorner \neg K \neg Kp \urcorner$ in β'. Insisting that I am in no such position would seem to beg the question against the thesis, licensed by **L**, that $\ulcorner K \neg K \neg Kp \urcorner$ holds in α'.

But we can here say something more powerful in support of both this idea and **L**. Thus, consider the following condition:

C For any γ and any $\delta \in B(\gamma, p)$, if p is false in γ, then, in δ, one doesn't believe p.

It's easy to check that, given **P1** to **P3** and suitable rationality assumptions, **C** is satisfied for the following choices of p: $\ulcorner \neg Kp \urcorner$, $\ulcorner \neg K \neg Kp \urcorner$, etc. This already explains why I have no trouble accepting the following anti-luminosity premise (corresponding to Smith's $\mathbf{AL_3}$):

A3 For any $\beta' \in B(\alpha', \ulcorner \neg K \neg Kp \urcorner)$, if I know $\ulcorner \neg K \neg Kp \urcorner$ in β', then, for any $\beta \in B(\alpha, \ulcorner \neg K \neg Kp \urcorner)$, there is no near duplicate of β in which I falsely believe $\ulcorner \neg K \neg Kp \urcorner$.

Against this backdrop, I take the following principle to be defensible:

P4 If, for some $\delta \in B(\gamma, \ulcorner \neg Kp \urcorner)$, one believes $\ulcorner \neg Kp \urcorner$ in δ, then $\ulcorner K \neg Kp \urcorner$ holds in γ.[11]

P4 implies that, if one is in α' in a position to believe $\ulcorner \neg K \neg Kp \urcorner$ upon doing the best that one is in α' in a position to do in order to decide $\ulcorner \neg K \neg Kp \urcorner$, then $\ulcorner K \neg K \neg Kp \urcorner$ holds in α'. Similarly, **P4** implies that, if $\ulcorner \neg K \neg Kp \urcorner$ holds in α', as ex hypothesi it does, one is *not* in α' in a position to believe $\ulcorner \neg Kp \urcorner$ upon doing

[11] Below, I will address one worry one might still have about **P4**, viz. that it doesn't take proper account of a dimension of safety that concerns beliefs formed on the basis of *relevantly similar methods* (see footnote 14). For another worry, see my reply to Waxman (2022).

the best that one is in α′ in a position to do in order to decide ⌜¬*Kp*⌝. Given **P2** and **P3**, *this should render it sufficiently plausible that, if* ⌜¬*K*¬*Kp*⌝ *holds in α′, one is in α′ in a position to believe* ⌜¬*K*¬*Kp*⌝ *upon doing the best that one is in α′ in a position to do in order to decide* ⌜¬*K*¬*Kp*⌝ (JAI: 72-74). By **P4**, it then follows that ⌜*K*¬*K*¬*Kp*⌝ holds in α′—just as **L** predicts.

Note that the same cannot be said about ⌜¬*Kp*⌝. Even with **P1** to **P4** in place, we cannot derive that, if one is in γ in a position to believe *p* upon doing the best that one is in γ in a position to do in order to decide *p*, then ⌜*Kp*⌝ holds in γ. For, **C** only holds for certain choices of *p* and, if, throughout, we replaced '¬*Kp*' by '*p*' in **P4**, the result wouldn't be defensible. For the same reason, we cannot derive the implausible conclusion that, if ⌜¬*Kp*⌝ holds in γ, one is *not* in γ in a position to believe *p* upon doing the best that one is in γ in a position to do in order to decide *p* (JAI: 73).[12]

Note also that nothing yet implies that, for any case γ, if ⌜¬*K*¬*Kp*⌝ holds in γ and one believes ⌜¬*K*¬*Kp*⌝ upon doing the best that one is in γ in a position to do in order to decide ⌜¬*K*¬*Kp*⌝, one is *rationally compelled* to believe ⌜¬*K*¬*Kp*⌝. For all that has so far been said, such beliefs aren't rationally compelling in borderline cases.

Smith suggests that this type of defence of **L** assumes 'that there is a kind of constitutive connection between the obtaining of [⌜*K*¬*Kp*⌝] and our beliefs about it' in the following sense: '[p]rovided that I'm doing everything I am in a position to do to determine whether [⌜¬*K*¬*Kp*⌝ holds], my believing [that ⌜¬*K*¬*Kp*⌝ holds] would, given [**P1** to **P3**], prevent [⌜*K*¬*Kp*⌝] from obtaining' (Smith, 2022: 6-7). I find this talk slightly misleading (even if we here replace 'everything' by 'the best'). For, what one ends up believing, upon doing one's best, depends on what one's epistemic situation affords, rather than the other way around. Thus, what I'd rather want to say is that, if ⌜*K*¬*Kp*⌝ holds and I do the best that I am thereby in a position to do in order to decide ⌜¬*K*¬*Kp*⌝, then these facts, in conjunction with facts about my rationality, prevent that I believe ⌜¬*K*¬*Kp*⌝ as a result. Accordingly, if I rationally believe ⌜¬*K*¬*Kp*⌝, upon doing the best that I am in a position to do in order to decide ⌜¬*K*¬*Kp*⌝, that will be so partly because ⌜¬*K*¬*Kp*⌝ holds.

In formulating his own challenge to **L**, Smith invokes a principle—which he labels '$\mathbf{AL_4}$'—whose presently relevant implication is this:

A4′ If I do not believe ⌜¬*K*¬*Kp*⌝ in α, there is a near duplicate of α′ in which I do not believe ⌜¬*K*¬*Kp*⌝.

As noted earlier, Smith assumes that both α ∈ B(α, ⌜¬*K*¬*Kp*⌝) and α′ ∈ B(α′, ⌜¬*K*¬*Kp*⌝). In order to free ourselves from this unnecessarily controversial assumption, let us replace **A4′** by

A4 For any β ∈ B(α, ⌜¬*K*¬*Kp*⌝), if I do not believe ⌜¬*K*¬*Kp*⌝ in β, then, for some β′ ∈ B(α′, ⌜¬*K*¬*Kp*⌝), there is a near duplicate of β′ in which I do not believe ⌜¬*K*¬*Kp*⌝ either.

[12] This is one way in which it transpires that my defence of **L** is not, by extension, a defence of the **S5** axiom for '*K*'. See also footnote 13.

If acceptable, this principle would indeed spell trouble for **L**. The question is why we should accept it. Let's assume, as we may, that I do not believe ⌜¬*K*¬*Kp*⌝ in some β ∈ B(α, ⌜¬*K*¬*Kp*⌝). In the light of this and our earlier observations, **A4** would simply beg the question against **L**. For, according to **L**, ⌜*K*¬*K*¬*Kp*⌝ holds in α′. From this, it straightforwardly follows, by the factivity of knowledge and **P1**, that ⌜¬*K*¬*Kp*⌝ holds in any β′ ∈ B(α′, ⌜¬*K*¬*Kp*⌝). From this, by Smith's definition of 'near duplicate', it in turn follows that, for any near duplicate γ′ of such a β′, both ⌜¬*K*¬*Kp*⌝ holds in γ′ and γ′ ∈ B(γ′, ⌜¬*K*¬*Kp*⌝). But, according to **L**, in any such near duplicate γ′, I do know ⌜¬*K*¬*Kp*⌝ and hence believe it.

To put it more succinctly, **A4** alleges that, if I don't believe ⌜¬*K*¬*Kp*⌝ upon doing the best that I am, in α, in a position to do in order to decide ⌜¬*K*¬*Kp*⌝, which we assume, there should be some case in which (i) ⌜¬*K*¬*Kp*⌝ holds, (ii) I do the best that I am in a position to do in order to decide ⌜¬*K*¬*Kp*⌝, yet in which (iii) I don't believe ⌜¬*K*¬*Kp*⌝ as a result. But, according to **L**, *no* case in which one fails to believe ⌜¬*K*¬*Kp*⌝, yet in which ⌜¬*K*¬*Kp*⌝ is true, will be a case in which one does *the best* that one is in a position to do in order to decide ⌜¬*K*¬*Kp*⌝. This is not to deny that, even if ⌜¬*K*¬*Kp*⌝ holds in a given case, one may fail to believe ⌜¬*K*¬*Kp*⌝ upon doing *slightly less* than the best that one is, in that case, in a position to do in order to decide ⌜¬*K*¬*Kp*⌝. Thus, the proponent of **L** has no problem accommodating

A4* For any β ∈ B(α, ⌜¬*K*¬*Kp*⌝), if I do not believe ⌜¬*K*¬*Kp*⌝ in β, then, for some β′ ∈ A(α′, ⌜¬*K*¬*Kp*⌝), there is a near duplicate of β′ in which I do not believe ⌜¬*K*¬*Kp*⌝,

where, for any case γ and any proposition *q*, we let A(γ, *q*) be defined as the set of possible cases, close to γ, in which one does almost, but not quite the best that one is in γ in a position to do in order to decide *q* (and does nothing that one is not in γ in a position to do). As far as I can see, Smith provides no argument for **A4** *as opposed to merely* **A4***. If **L** holds, type II errors may still occur, but only if one behaves sub-optimally (which, for all that has been said, one might easily do).

Even if Smith's challenge might thus be defused in the end, the manoeuvre invites the friends of anti-luminosity to have another comeback. Thus, suppose they suggest combining **A3** with

A3* For any β′ ∈ B(α′, ⌜¬*K*¬*Kp*⌝), if I know ⌜¬*K*¬*Kp*⌝ in β′, then, for any β ∈ A(α, ⌜¬*K*¬*Kp*⌝), there is no near duplicate of β in which I falsely believe ⌜¬*K*¬*Kp*⌝.

Suppose that β′ ∈ B(α′, ⌜¬*K*¬*Kp*⌝) and, in β′, I believe ⌜¬*K*¬*Kp*⌝. Since ⌜*K*¬*Kp*⌝ holds in α, by **P1** to **P3**, it follows that, for any β ∈ B(α, ⌜¬*K*¬*Kp*⌝), I know ⌜¬*Kp*⌝ in β—and hence, ⌜*K*¬*Kp*⌝ also holds in β. From this, it in turn follows that, for any near duplicate γ of such a case β, ⌜*K*¬*Kp*⌝ also holds in γ, γ ∈ B(γ, ⌜¬*K*¬*Kp*⌝), and I know ⌜¬*Kp*⌝, and hence do not falsely believe, ⌜¬*K*¬*Kp*⌝ in γ. But, for all that, it may then still be that, for some β ∈ A(α, ⌜¬*K*¬*Kp*⌝) and some near duplicate γ of β, I falsely believe ⌜¬*K*¬*Kp*⌝ in γ. Any case is a near duplicate of itself. For simplicity's sake, assume that α itself is such that, while $\alpha \notin B(\alpha, \ulcorner\neg K\neg Kp\urcorner)$, nonetheless $\alpha \in A(\alpha, \ulcorner\neg K\neg Kp\urcorner)$ and that I falsely believe $\ulcorner\neg K\neg Kp\urcorner$ in α. It would then follow from **A3*** that I do not know $\ulcorner\neg K\neg Kp\urcorner$ in β′—and so, since ex hypothesi, in β′, I've done the best that I am in α′ in position to do in order to decide $\ulcorner\neg K\neg Kp\urcorner$, I am in no position to know $\ulcorner\neg K\neg Kp\urcorner$ in α′—in which event **L** fails.

A3* is highly controversial, however, and I think there are good reasons to reject it. To see why, first note that **A3*** is meant to be a kind of safety principle and recall that such principles ultimately derive from a general idea about knowledge-conducive methods, viz. that, if one knows *p* on the basis of applying some method *m*, there is no close possibility in which one falsely believes *p* or a closely related proposition on the basis of applying any method relevantly similar to *m*. Here, closely related propositions are best conceived as propositions within the range of a method relevantly similar to *m*, where, trivially, *m* is relevantly similar to itself (cf. JAI: 27-28, 269-70).

Whether **A3*** is acceptable accordingly depends on whether doing the best that one is in a position to do in order to decide a given question, and doing almost, but not quite, the best in this sense, are guaranteed to involve use of relevantly similar methods. Methods that belong to different types aren't relevantly similar. (This coheres with the observation that, unlike identity of type, relevant similarity isn't transitive—to the extent, namely, that there may then still be three methods of the same type such that the first is relevantly similar to the second, and the second to the third, although the first isn't relevantly similar to the third.) If a method *n* is in some way constitutively geared to other methods *M*, then, plausibly, *n′* will be of the same type as *n* only if there are methods *M′* of the same type(s) as *M* such that *n′* is constitutively geared to *M′* in that same way. In line with **P1** to **P4**, assume

P5 The best that one is in a position to do in order to decide $\ulcorner Kp\urcorner$ is such that, upon doing it, one believes $\ulcorner\neg Kp\urcorner$ iff one doesn't believe *p*,

where this is so because, if one believes $\ulcorner\neg Kp\urcorner$ upon doing the best that one is in a position to do in order to decide $\ulcorner Kp\urcorner$, one applies a method *n* for telling that $\ulcorner\neg Kp\urcorner$ holds such that, among one's methods there are methods *M* for telling that *p* holds, where

i. always, if $\ulcorner Kp\urcorner$ holds and one does the best that one is in a position to do in order to decide *p*, there is some *m* among *M* such that one applies *m* to *p*,
ii. always and constitutively, for any *m* among *M*, if *m* is applied to *p* and out puts *p*, *n* doesn't output $\ulcorner\neg Kp\urcorner$, and
iii. always and constitutively, if *n* is applied to $\ulcorner\neg Kp\urcorner$ and outputs $\ulcorner\neg Kp\urcorner$, then, for any *m* among *M*, *m* doesn't output *p*,

where, if a conditional holds constitutively, it sustains corresponding counterfactuals.

P5 is a mouthful. Yet, complex as it is, upon reflection, **P5** strikes me as rather plausible: the outputs of one's best method *n* for telling that one is in no position to know *p* shouldn't get ahead of, but rather be appropriately geared to, the outputs of one's best method *m* for telling that *p* holds. In the appendix to JAI, I gave some reasons in **P5**'s favour that were of the same kind as those I gave for **P3** in chapter 3 (see JAI: 43-48, 271-72). But, for lack of space, I cannot here provide a more thorough argument for it (however, see my reply to Waxman, 2022, where I say more about methods).

With **P2**, **P3**, and **P5** in place, it follows that, for any *p*, the best that one is in a position to do in order to decide ⌜¬*K*¬*Kp*⌝ is such that, upon doing it, one believes ⌜¬*K*¬*Kp*⌝ iff one doesn't believe ⌜¬*Kp*⌝ iff one believes *p*. Given that knowledge implies belief and given **P1**, it then likewise follows that, if one believes ⌜¬*Kp*⌝ upon doing the best that one is in a position to do in order to decide ⌜¬*Kp*⌝—and so, by **P2** and **P3**, upon doing the best that one is in a position to do in order to decide *p*—one uses a method for telling that ⌜¬*Kp*⌝ holds such that one couldn't easily falsely believe ⌜¬*Kp*⌝ on its basis—a method, namely, that tracks those failures to believe *p* that wouldn't come to pass if ⌜*Kp*⌝ was true.[13] Similarly, if one believes ⌜¬*K*¬*Kp*⌝ upon doing the best that one is in a position to do in order to decide ⌜¬*K*¬*Kp*⌝—and so, by **P2** and **P3**, upon doing the best that one is in a position to do in order to decide ⌜¬*Kp*⌝—one uses a method for telling that ⌜¬*K*¬*Kp*⌝ holds such that one couldn't easily falsely believe ⌜¬*K*¬*Kp*⌝ on its basis—a method, namely, that tracks those failures to believe ⌜¬*Kp*⌝ that wouldn't come to pass if ⌜*K*¬*Kp*⌝ was true (JAI: 271-72).

Doing almost, but not quite the best that one is in a position to do in order to decide ⌜¬*K*¬*Kp*⌝ may break these constitutive connections—with the effect that the method for telling that ⌜¬*K*¬*Kp*⌝ holds that one applied in the course of doing the suboptimal in this regard would no longer be of the same type.[14] This is the reason why I reject **A3***. It thus remains the case that, even if **L** should ultimately be shown to fail, its failure anyway isn't obvious.

4 Reply to Daniel Waxman

One of the key theses of the account of justification given in JAI is the following:

H1 One has propositional justification for believing *p* iff one is in no position to know that one is in no position to know *p*.

If we let '*J*' be short for 'one has propositional justification for believing', and '*K*' for 'one is in a position to know', and conceive of cases as triples of worlds, subjects, and times, we can rephrase **H1** as follows: for any case γ, ⌜*Jp*⌝ holds in

[13] Note that this does not imply that, whenever ⌜¬*Kp*⌝ holds, ⌜*K*¬*Kp*⌝ holds: whatever method *n* for telling that ⌜¬*Kp*⌝ holds might be, in the way described, constitutively geared to whatever methods *M* for telling that *p* holds, one might believe *p* on the basis of one of the methods *M* in circumstances in which ⌜¬*Kp*⌝ holds. This is another way in which it transpires that my defence of **L** is not, by extension, a defence of the **S5** axiom for '*K*'.

[14] The same reasons will also defuse a remaining worry about **P4**, viz. that it unduly ignores beliefs in ⌜¬*Kp*⌝, formed upon doing the best that one is in a position to do in order to decide ⌜¬*Kp*⌝, that are unsafe because methods relevantly similar to those used may churn out propositions that do not satisfy **C**. If **P5** is correct, there are no such relevantly similar methods.

γ iff ⌜¬*K*¬*Kp*⌝ holds in γ—where, throughout, 'one' refers to the subject at the centre of γ. (Note that, by contrast, in JAI '*J*' was just a metalinguistic abbreviation of '¬*K*¬*K*'.)

H1 is a conjunction of two conditionals. Of these, the right-to-left conditional is the one most vulnerable to the threat of counterexample. The most obvious candidates for such counterexamples can be dismissed from the start because, as I made clear, **H1** is only intended to hold for suitably idealized subjects (JAI: 108). Idealizations can protect theses from easy refutation. But they must not be carried too far, lest those theses lose much of their interest. For instance, crediting the subjects in question with unlimited epistemic powers would go too far. Theses exclusively concerned to capture general features of the epistemic lives of subjects with such unlimited resources won't tell us anything about what, in our best moments, we can aspire *our* epistemic lives to be. Thus, **H1** is only meant to hold for subjects whose epistemic powers, while they outstrip our own, are still bounded—what we may call 'finitely improved versions of ourselves'.

Concerns about too strong idealizations were also on my mind when, in JAI, I was searching for a suitable logic for '*K*'. For this reason, I refused to buy into the principle, characteristic of normal epistemic logics, according to which one is in a position to know whatever is logically implied by what one is in a position to know. Improved versions of ourselves, at least of the kind I envisaged, still are subject to computational and other limitations and so fail to live up to the standard set by that principle (JAI: 58-59; see also my reply to Rossi, 2022).

In his challenging comments, Daniel Waxman applauds the lessons I draw for epistemic logic, but goes on to devise what he takes to be a counterexample to the right-to-left direction across **H1**, in the sense in which it is intended, assuming that, throughout, 'one' denotes a single, suitably idealized epistemic agent of the requisite kind that still is subject to computational limitations (Waxman, 2022). Below, I clarify the structure of the challenge and, in addressing it, argue that the threat to **H1** can after all be averted.

If one's computational powers are bounded, there will be mathematical proofs that are too computationally complex for one to construct or even go through. To fix ideas, and for the sake of simplicity alone, Waxman conceives of this computational limitation in terms of the length of proofs (Waxman, 2022: 3). There will accordingly be a maximal length a proof might have, if one is to be able to construct or go through it.

Waxman then envisages a case α in which a given mathematical claim *t* only admits of proofs, or disproofs, whose length exceeds that maximal length, in which one does not know, and has no independent justification to believe, *t*, its negation, or ⌜¬*Kt*⌝, and in which any candidate proofs that one is in a position to construct, or go through, will in fact leave *t* undecided—in the sense that one will then recognize that one neither knows *t* nor its negation (Waxman, 2022: 3-5).

To the extent that one has no propositional justification for *t* in α, as Waxman assumes (2022: 4), **H1** forces us to conclude that ⌜*K*¬*Kt*⌝ holds at α. If ⌜*K*¬*Kt*⌝ holds at α, however, then there must be a method that one is in a position to apply in order to find out that one is in no position to know *t*. But, according to Waxman, no such method is forthcoming, even if it should be conceded that, in α, one has knowledge of the extent of one's computational limitations. For, according to Waxman, finding out that one is in no position to know *t* would require constructing, or going through, all candidate proofs for deciding *t* that one is in a position to construct or go through. This, however, is known to be too demanding a task for an agent with limited computational powers (Waxman, 2022: 4-5).

To bring Waxman's challenge into sharper relief, we can simplify his example by assuming that, while in α one neither knows nor has independent justification for *t*, $\ulcorner \neg t \urcorner$, or $\ulcorner \neg Kt \urcorner$ (e.g. by testimony), one knows in α that there are *only two* alternative proof strategies one is in a position to individually pursue—but which, given one's computational limitations, *one cannot both pursue*, simultaneously or successively—in order to decide *t*, and whose pursuit eventuates in either case β_1 or case β_2. If we let '◯' stand for a situation in which *t* remains undecided and '●' for a situation in which *t* is proved, then we can distinguish between (at least) three conceivable scenarios,[15] of which the first, ex hypothesi, is actual:

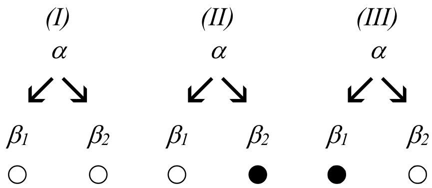

Let '*k*' be short for 'one knows'. We assume that, for any $\beta \in \{\beta_1, \beta_2\}$, if *t* remains undecided in β, then $\ulcorner k\neg kt \urcorner$ holds at β and hence so does $\ulcorner K\neg kt \urcorner$. Similarly, we assume that, for any $\beta \in \{\beta_1, \beta_2\}$, if *t* is proved in β, then $\ulcorner kt \urcorner$ holds at β. Given one's computational limitations, one is in fact in scenario (I). However, whichever proof strategy one pursues, in scenario (I), one will end up in a situation in which one cannot exclude *both* (II) and (III). More specifically, given that one is in fact in scenario (I), in β_1, one neither knows nor is in a position to know that $\ulcorner \neg kt \urcorner$ holds at β_2—and so cannot exclude (II)—and, in β_2, one neither knows nor is in a position to know that $\ulcorner \neg kt \urcorner$ holds at β_1—and so cannot exclude (III). Consequently, given that one is in scenario (I), one is not, in α, in a position to exclude that one is in either (II) or (III). But, importantly, according to Waxman, in scenarios (II) and (III), one *is* in α in a position to know *t*. From this, he concludes that, in scenario (I), one is not, in α, in a position to know that one is in no position to know *t*—i.e. that, in (I), $\ulcorner \neg K\neg Kt \urcorner$ holds at α. By **H1**, it would then follow that, in α, one has propositional justification for *t*—which seems intuitively unacceptable.

Indeed, the latter conclusion is also unacceptable on the account of justification given in JAI. For, that account implies

[15] For reasons of simplicity, I deliberately ignore conceivable scenarios in which either β_1 or β_2 is a case in which *t* is *dis*proved.

H2 One has doxastic justification for *p* iff one is in no position to know that one doesn't know *p*.

Ex hypothesi, in scenario (I), ⌜*K*¬*kt*⌝ holds at both β_1 and β_2. Consequently, given **H2**, in (I), one lacks doxastic justification for *t* in both β_1 and β_2—and this is hard to reconcile with the assumption that, in (I), one has propositional justification for *t* in α (cf. my reply to Zhan, 2022).

There is something initially puzzling about Waxman's confidence that, in (I), ⌜¬*K*¬*Kt*⌝ holds at α. To the extent, namely, that it is anyway known in α that after pursuing one of the two proof strategies one cannot go on to pursue the respective other, in scenario (I), it should hold at both β_1 and β_2, not only that one knows one doesn't know *t* (i.e. ⌜*k*¬*kt*⌝), but also that one knows one is in no position to know *t* (i.e. ⌜*k*¬*Kt*⌝).[16] This implies that, in the present context, we anyway face no violation of the letter of one of the key principles of JAI that I already alluded to in my reply to Smith (2022), viz.

P1 For any δ ∈ B(γ, *p*), if ⌜*Kp*⌝ holds in γ, ⌜*kp*⌝ holds in δ,

where, for any γ and any *p*, B(γ, *p*) is the set of possible cases, close to γ and centred on the same subject as γ, in which the subject of γ has done the best they are in γ in a position to do in order to decide *p* (and has done nothing they are not in γ in a position to do). In the present example, in scenario (I), B(α, *t*) = $\{\beta_1, \beta_2\}$. Given

P2 B(γ, ⌜¬*p*⌝) = B(γ, *p*)

P3 B(γ, ⌜*Kp*⌝) ⊆ B(γ, *p*),

it follows that, in scenario (I), B(α, ⌜¬*Kt*⌝) ⊆ $\{\beta_1, \beta_2\}$. Accordingly, if he wants to spell trouble for the claim that, in (I), ⌜*K*¬*Kt*⌝ holds at α, Waxman would seem to need something stronger than **P1**, viz.

W1 For any β ∈ B(α, ⌜¬*Kt*⌝), if ⌜*K*¬*Kt*⌝ holds in α, then there is a term '*x*' referring to α such that ⌜*k*¬$K_x t$⌝ holds in β,

where ⌜K_x⌝ abbreviates 'one is in *x* in a position to know⌝. Note that **W1** cannot be derived from **P1** to **P3** in conjunction with the factivity of '*K*' alone. For, even if ⌜*K*¬*Kt*⌝ holds in α only provided that ⌜¬$K_\alpha t$⌝ holds, this alone does not entail that, for any β ∈ B(α, ⌜¬*Kt*⌝), ⌜*k*¬*Kt*⌝ holds in β only if ⌜*k*¬$K_x t$⌝ holds in β, for some '*x*' referring to α.

[16] In Waxman's original example, such knowledge will only be reached if one has exhausted one's resources to conduct any further proof searches and knows that one has done so. Waxman wouldn't seem opposed to the idea that, insofar as one knows the extent of one's computational limitations and is a suitably idealized agent, such a point can indeed be reached (see Waxman 2022: 4; cf. also Waxman forthcoming).

Accordingly, a first line of response would be to reject **W1**. However, even if **W1** may in the end be motivated on intuitive grounds, there is another assumption in play that we aren't forced to accept—viz. Waxman's assumption that, in scenarios (II) and (III), ⌜Kt⌝ holds at α. Given **P1**, this assumption implies that, in scenario (II), B(α, t) = {β_2} and, in scenario (III), B(α, t) = {β_1}. If, on the contrary, in all scenarios, including (II) and (III), B(α, t) = {β_1, β_2}, then, to the extent that the latter is also known in α, in scenario (I), one's knowledge of **P1** will allow one to know, in both β_1 and β_2, that ⌜$\neg K_\alpha t$⌝ holds—in which event, even with **W1** in place, there is no reason to deny that, in scenario (I), ⌜$K\neg Kt$⌝ holds at α and Waxman's challenge can successfully be defused.

This is the second line of response I wish to pursue. But, before motivating its key assumption—i.e. that, in all scenarios, B(α, t) = {β_1, β_2} –, it is best first to tease out the implications this line of response has beyond the confines of Waxman's particular example. In this context, it is important to realize that Waxman's challenge is of a far more general kind than his focus on computational limitations and the failure of logical omniscience might lead one to expect.[17] Below I will therefore devise another example that exhibits a more general structure than Waxman's own and that, at the same time, will provide for a good test of the viability of the line of response just sketched (see also Waxman forthcoming).

Throughout, I shall continue to assume **P1** to **P3** and likewise that the inferences from ⌜$K\neg Kp$⌝ to ⌜$K\neg kp$⌝ and from ⌜$k\neg Kp$⌝ to ⌜$K\neg Kp$⌝ are in good standing (for the latter two inferences, see JAI: 37-38, 86-89). Since we are here concerned with challenges to the right-to-left conditional across **H1**, we also need a dialectically neutral criterion that helps to determine if subjects lack propositional justification. Reliable pre-theoretic intuitions about the lack of propositional justification are hard to come by, if only because propositional justification is a mere potential for doxastic justification and so need not consist in any piece of evidence currently in one's possession (see my reply to Zhan, 2022). Dialectic neutrality doesn't mean theoretical neutrality. Thus, for the purposes of Waxman's challenge and in line with the result he intends, I suggest that the following criterion will do:

CRIT For all β ∈ B(α, ⌜Kp⌝), if ⌜$K\neg kp$⌝ holds in β (and so, according to **H2**, one lacks doxastic justification for p in β),[18] then one lacks propositional justification for p in α.[19]

Let me now turn to the advertised example (cf. also Waxman forthcoming). Suppose that, in case α, one knows—and so it is true—that there are two distinct boxes, box_1 and box_2, such that, in α, one has no indication or clue what either

[17] This point was made during the symposium, at least in basic outline, by Julien Dutant.

[18] Note that this only exploits the left-to-right conditional across **H2**, which is far less controversial (cf. JAI: 118).

[19] For more on the link between propositional and doxastic justification, see my reply to Zhan (2022).

box contains, such that one is in a position to open either one of the boxes, but not both, thereby advancing to either β_1 (in which case one opens box_1) or β_2 (in which case one opens box_2), and such that, given what one presently knows or is presently justified in believing, all of the following conceivable scenarios are equally likely so that, upon opening one of the boxes and finding it either empty or non-empty, one has no indication or clue as to whether the other box is empty or not (where '$\square_i$' indicates that box_i is empty and '$\boxtimes_i$' indicates that box_i is non-empty):

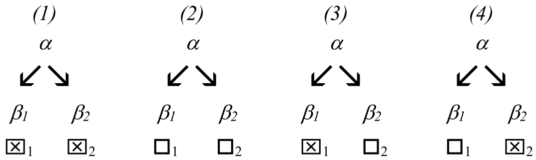

Unlike in Waxman's original example, we make no assumption about which of these scenarios is actual. Let q be the proposition *that at least one of the two boxes is non-empty*. We assume that, for any $\beta \in \{\beta_1, \beta_2\}$, $\ulcorner kq \urcorner$ holds at β iff the box one opens in β is non-empty. In scenario (2), q is false. A fortiori, in (2), $\ulcorner \neg Kq \urcorner$ holds at α. In fact, (2) is the only scenario in which q is false. Given the setup, we may assume that one already knows in α that, no matter which box one opens and no matter what's inside, one will at most be able to exclude two of the four scenarios. Consequently, in all four scenarios, $\ulcorner K\neg K\neg q \urcorner$ holds at α.

For the time being, let us suppose, with Waxman, that, in all the other scenarios, (1), (3), and (4), $\ulcorner Kq \urcorner$ holds at α. (This corresponds to Waxman's earlier assumption that, in scenarios (II) and (III), $\ulcorner Kt \urcorner$ holds at α.) Then, in scenarios (1), (3), and (4), $\ulcorner \neg K\neg Kq \urcorner$ holds at α, and by **P1** to **P3**, in (3), both $\beta_2 \notin B(\alpha, q)$ and $\beta_2 \notin B(\alpha, \ulcorner \neg Kq \urcorner)$, while in (4), both $\beta_1 \notin B(\alpha, q)$ and $\beta_1 \notin B(\alpha, \ulcorner \neg Kq \urcorner)$. Given this supposition, it is then no problem to maintain that, in those three scenarios, one has propositional justification for q in α. For, then, in those scenarios, if one does the best that one is in α in a position to do in order to decide $\ulcorner \neg Kq \urcorner$, one will know q and so be in no position to know that one doesn't know q (and hence, according to **H2**, will have doxastic justification for q).

Given Waxman's supposition, the only potentially problematic scenario will then be (2) (which corresponds to scenario (I) above). In (2), $\ulcorner K\neg kq \urcorner$ holds at both β_1 and β_2. Therefore, by CRIT, **P2** and **P3**, one has, at α, no propositional justification for q. Applied to the present example, Waxman's diagnosis then is that, in (2), $\ulcorner \neg K\neg Kq \urcorner$ nonetheless holds at α—in which event **H1** will fail.

We have in effect already seen that, for this diagnosis to be forced upon us, we would in addition need a principle of the kind already exemplified by **W1**, viz.

W2 For any $\beta \in B(\alpha, \ulcorner\neg Kp\urcorner)$, if $\ulcorner K\neg Kp\urcorner$ holds in α, then there is a term 'x' referring to α such that $\ulcorner k\neg K_x p\urcorner$ holds in β.

Let us not here take issue with **W2**. Instead, we can respond by insisting that, *pace* Waxman, in each of the four scenarios, $B(\alpha, q) = \{\beta_1, \beta_2\}$—thereby rejecting the aforementioned supposition that $\ulcorner Kq\urcorner$ holds at α in scenarios (3) and (4).[20]

We can now also see more clearly that rejection of this supposition is independently called for, as long as we wish to maintain that, in (2), $\ulcorner K\neg Kq\urcorner$ holds at α. For, *if* (2) is indeed the only scenario in which $\ulcorner\neg Kq\urcorner$ holds at α, and one knows this much about the setup, then, since (2) is also known to be the only scenario in which q is false, $\ulcorner K(\neg Kq \supset \neg q)\urcorner$ should hold at α in all four scenarios. Accordingly, if $\ulcorner K\neg Kq\urcorner$ holds at α in (2), then, given the availability of a harmless closure step, in scenario (2), one would after all in α be in a position to know that one is in that scenario—which, given the setup, is impossible.

In α, one already knows that, no matter which box one opens and no matter what's inside, one will at most be able to exclude two of the four scenarios. If, in all four scenarios, $B(\alpha, q) = \{\beta_1, \beta_2\}$, then, given **P1**, (1) will be the only scenario in which $\ulcorner Kq\urcorner$ holds at α. To the extent that this is a further thing one knows in α, consequently, in all four scenarios, not only does $\ulcorner K\neg K\neg q\urcorner$ hold at α, but so does $\ulcorner K\neg KKq\urcorner$. But, if $\ulcorner K\neg KKq\urcorner$ holds at α, we cannot similarly exploit $\ulcorner K(Kq \supset r)\urcorner$ in order to arrive at $\ulcorner Kr\urcorner$ by a single closure step, where r is the proposition *that none of the boxes is empty*. Equally, then, we may assume, in line with **W2**, that, for any $\beta \in \{\beta_1, \beta_2\}$, $\ulcorner k\neg Kq\urcorner$ will hold at β iff $\ulcorner k\neg K_\alpha q\urcorner$ holds at β (i.e. iff the box one opens in β is empty). Accordingly, given the inference from $\ulcorner k\neg Kq\urcorner$ to $\ulcorner K\neg kq\urcorner$, scenario (2) no longer presents a problem for **H1**, all the while (1) continues to present no such problem either.

Given our rejection of Waxman's supposition, the only potentially troublesome scenarios now are (3) and (4). For, *if*, in both (3) and (4), $B(\alpha, \ulcorner Kq\urcorner) = \{\beta_1, \beta_2\}$, **P1** and **P2** will imply that, in those scenarios, $\ulcorner\neg K\neg Kq\urcorner$ holds at α, because, upon opening a non-empty box, one will *not* know that one is in no position to know p. Given **H1**, CRIT is equivalent to

P6 For any $\delta \in B(\gamma, \ulcorner Kp\urcorner)$, if $\ulcorner K\neg kp\urcorner$ holds in δ, then $\ulcorner K\neg Kp\urcorner$ holds in γ.

[20] This implies rejection of the converse of **P1**, i.e.

CP1 For any $\delta \in B(\gamma, p)$, if $\ulcorner kp\urcorner$ holds in δ, then $\ulcorner Kp\urcorner$ holds in γ,

which I had relied on in my argument for **L** (see JAI: 41-42, 72). However, the instance of **CP1** that I needed for the latter purpose was in fact

For any $\delta \in B(\gamma, \ulcorner\neg Kp\urcorner)$, if $\ulcorner k\neg Kp\urcorner$ holds in δ, then $\ulcorner K\neg Kp\urcorner$ holds in γ,

which can be derived from **P2**, $\ulcorner k\neg Kp \supset K\neg kp\urcorner$, and principle **P6** below, which latter I argue remains acceptable (cf. principle (8) in JAI: 72).

Given **P6**, that ⌜¬*K*¬*Kq*⌝ holds at α will then imply that ⌜*K*¬*kq*⌝ holds at neither β_1 nor β_2—which is patently false, because, upon opening an empty box, one *will* know that one is in no position to know *q*.

However, as long as it can in addition be maintained, in line with **P1** to **P3**, that, in (3), $\beta_1 \notin$ B(α, ⌜¬*Kq*⌝), while in (4), $\beta_2 \notin$ B(α, ⌜¬*Kq*⌝), any looming trouble can be averted.[21] We can derive this additional assumption as follows. Recall another principle that I invoked in the context of my reply to Smith (2022):

P4 If, for some δ ∈ B(γ, ⌜¬*Kp*⌝), one believes ⌜¬*Kp*⌝ in δ, then ⌜*K*¬*Kp*⌝ holds in γ.

One might think that the combination of **W2** and **P1** puts pressure on **P4**. For, given **W2** and **P1**, ⌜*K*¬*Kq*⌝ will hold in α, only if, for all β ∈ B(α, ⌜¬*Kq*⌝), one knows, and hence believes, ⌜¬$K_\alpha q$⌝ in β, for some '*x*' referring to α. Yet, given our assumption that it is already known in α that ⌜*Kq*⌝ holds only if ⌜*kq*⌝, and hence ⌜*Kq*⌝, holds at both β_1 and β_2, this alone is insufficient to refute **P4**. For, given this assumption, plausibly,

for any β ∈ B(α, *q*), if one believes ⌜¬*Kq*⌝ in β, one also believes ⌜¬$K_\alpha q$⌝ in β.

Accordingly, unless it is faulted on other grounds, **P4** can be upheld even in the light of **W2**. Ex hypothesi, B(α, *q*) = {β_1, β_2} and, in scenarios (3) and (4), ⌜¬*Kq*⌝ holds at α, while, in (3), ⌜*k*¬*Kq*⌝ holds at β_2, and in (4), ⌜*k*¬*Kq*⌝ holds at β_1. Given **P2** and **P3**, plausibly and quite generally,

for any δ ∈ B(γ, *p*), if one believes ⌜¬*Kp*⌝ in δ, then δ ∈ B(γ, ⌜¬*Kp*⌝).

Consequently, in (3), β_2 ∈ B(α, ⌜¬*Kq*⌝) and, in (4), β_1 ∈ B(α, ⌜¬*Kq*⌝). So, by **P4**, ⌜*K*¬*Kq*⌝ holds in α, and so, by **P1** to **P3**, in (3), $\beta_1 \notin$ B(α, ⌜¬*Kq*⌝), while in scenario (4), $\beta_2 \notin$ B(α, ⌜¬*Kq*⌝).

Given this combination of assumptions, the observation that, according to CRIT, in (3) and (4), one lacks propositional justification for *q* in α doesn't cause trouble for **H1**. For, then, in those two scenarios, if one does the best that one is in α in a position to do in order to decide ⌜¬*Kq*⌝, one will know that ⌜¬*Kq*⌝ holds and hence be in a position to know that one does not know *q*. In the same vein, **P6** can then be maintained.

What, on this line of response, still stands in need of an explanation is the key assumption that, in some scenarios, viz. those in which box_1 is empty iff box_2 isn't, it is the case that B(α, *q*) = {β_1, β_2}—although ⌜¬*kq*⌝ holds in β_1 iff ⌜*kq*⌝ holds in β_2. Any such explanation must cohere with the finding that, nevertheless, in those very scenarios, for any β ∈ {β_1, β_2}, β ∈ B(α, ⌜¬*Kq*⌝) iff ⌜¬*kq*⌝ holds in β (i.e. iff the box one opens in β is empty)—with the consequence that {β_1, β_2} ⊈ B(α, ⌜¬*Kq*⌝). The

[21] This requires a correction of what I said in JAI. For, there, I assumed—not only **P2** and **P3**—but that, for any proposition *p* and any case γ, B(γ, *p*) ⊆ B(γ, ⌜*Kp*⌝) (cf. principle (2c) in JAI: 47). At most, it seems, we get the weaker principle that, if ⌜*Kp*⌝ holds in γ, then B(γ, *p*) ⊆ B(γ, ⌜*Kp*⌝) (which will in fact do for the work assigned to (2c) in JAI: 75).

sought-after explanation will have to say something suitably instructive about what it is to do the best that one is in a position to do in order to decide a given question.

Let us say that a given method *m* is among one's ex ante best methods for telling whether *p* holds iff before, and independently from the results furnished by, the use of any of one's methods for telling whether *p* holds, one has propositional justification for thinking that *m* is no less promising a method to use to this end than any other of one's methods.[22] Let a method's conditions of application to a given proposition *p* within its range be those conditions under which alone the method yields any output with regard to *p*—including, as the case may be, YES, NO, or NOUGHT. Now consider:

P7 If, among one's ex ante best methods for telling whether *p* holds, there are some whose conditions of application to *p* one is in a position to bring about, then one does the best that one is in a position to do in order to decide *p* iff for some such method, one brings about its conditions of application to *p* and applies it to *p*.

It follows that, if, among one's ex ante best methods for telling whether *p* holds, there is one method such that one brings about its conditions of application to *p* and applies it to *p*, then one does the best that one is initially in a position to do in order to decide *p*. It likewise follows that, if one does the best that one is in a position to do in order to decide *p*, then, if, among one's ex ante best methods for telling whether *p* holds, there is indeed one whose conditions of application to *p* one is in a position to bring about, one brings about the conditions of application to *p* of one such method and applies it to *p*.

Given a suitably realist conception of the individuation of methods and their application conditions—a conception that allows us to clearly distinguish between methods and their individual applications, as well as between circumstances in which a given method yields the output NOUGHT and circumstances in which it isn't applied at all—we can now respond to the explanatory challenge as follows. (What I am about to say chimes with what I say towards the end of my reply to Smith, 2022.)

Given our choice for *q*, in α, one already knows that one will be unable to tell that ⌜¬*q*⌝ holds, never mind what one does, and knows that one will be unable to tell that *q* holds upon opening an empty box, but able to tell that *q* holds upon opening a non-empty box. So, there is, in α, no method for telling whether *q* holds that one is in a position to apply and that would ever output NO. There is, in α, only one ex ante best method *m* for telling whether *q* holds: looking inside one of the boxes and telling whether the box in question is empty. If the box in question is non-empty, *m* will output YES, and if the box in question is empty, *m* will output NOUGHT. To pretend to be able to identify two methods here—looking inside an empty box and looking inside a non-empty box—would be to individuate methods too finely and to mistake methods for their individual applications. Accordingly, *m*'s conditions of application to *q* are the conditions under which one looks into one of the boxes. To bring about those conditions is to open *any* one of the boxes (to the exclusion of the respective

[22] Propositional justification here is again to be understood in terms of **H1** (and in accordance with **L**, i.e. with ⌜$\neg K\neg Kp \supset K\neg K\neg Kp$⌝).

other). Ex hypothesi, in α, one is in a position to do so. So, by **P7**, in α, the best one is in a position to do in order to decide *q* is to open one of the boxes, look inside, and then tell whether the box is empty. Accordingly, in scenarios (3) and (4), B(α, *q*) = {β_1, β_2} and so, by **P1**, one is in a position to know *q* in α only if one knows *q* in both β_1 and β_2 (i.e. only in scenario (1)).

Given our choice for *q*, in α, one already knows that one will be unable to tell that ⌜*Kq*⌝ holds, never mind what one does, and knows that one will be unable to tell that ⌜¬*Kq*⌝ holds upon opening a non-empty box, but able to tell that ⌜¬*Kq*⌝ holds upon opening an empty box. So, there is, in α, no method for telling whether ⌜*Kq*⌝ holds that one is in a position to use and that would ever output YES. There is, in α, only one ex ante best method *n* for telling whether ⌜*Kq*⌝ holds: concluding from *m*'s output of NOUGHT (in combination with what else one knows about one's situation) that one is in no position to know *q* (see principle **P5** in my reply to Smith, 2022). If *m* outputs NOUGHT, *n* outputs YES. But, if *m* doesn't output NOUGHT (and so, assuming *m* is applied, outputs YES instead), then *n* isn't even applied and so doesn't output anything, not even NOUGHT: that *m* outputs NOUGHT is part of *n*'s conditions of application to ⌜¬*Kq*⌝. To bring about these conditions is to open an *empty* box and apply *m*. Ex hypothesi, in α, one is in a position to do so. So, by **P7**, in α, the best one can do in order to decide ⌜*Kq*⌝ is to open an empty box, apply *m*, and then conclude that one is in no position to know *q*. Accordingly, in scenario (3), B(α, ⌜¬*Kq*⌝) = {β_2} and so, by **P1**, one is in a position to know ⌜¬*Kq*⌝ in α only if one knows ⌜¬*Kq*⌝ in β_2, even if one doesn't know ⌜¬*Kq*⌝ in β_1. Similarly, in scenario (4), B(α, ⌜¬*Kq*⌝) = {β_1} and so, by **P1**, one is in a position to know ⌜¬*Kq*⌝ in α only if one knows ⌜¬*Kq*⌝ in β_1, even if one doesn't know ⌜¬*Kq*⌝ in β_2.

In a nutshell, then, my response to Waxman is this. Waxman contends that (i) when none of the routes one is in a position to take in order to decide a given *p* leads to a decision, while one is in no position to take them all, one is initially in no position to know that one is in no position to know *p*, and that (ii) when one of those routes does lead to one's knowledge of *p*, one is initially in a position to know *p*. I reject both these assumptions. Instead, I contend that (iii) when some but not all those routes lead to one's knowledge of *p*, one is initially in no position to know *p* (that is, even if there then is some route such that, were one to take it, one would end up knowing *p*), and that (iv), because of that, one is in a position to recognize that one was not initially in a position to know *p* by taking one such route that leaves *p* undecided.

The smaller B(γ, *p*) is, the weaker are the demands that **P1** imposes on being in a position to know *p* in γ. The foregoing thus meshes well with the contention, implicit in much of what I say in JAI, that it is never harder to be in a position know one's ignorance than it is to be in a position to know the condition one is ignorant of. Even so, I expect that, unless more is said about the nature of methods—and about the nature of methods for telling if one is in no position to know, in particular—few will be swayed by these considerations. So, more work needs to be done.

5 Reply to Julien Dutant

The two key theses of the account of justification I gave in JAI are

H1 One has propositional justification for believing *p* iff one is in no position to know that one is in no position to know *p*.

H2 One has doxastic justification for believing *p* iff one is in no position to know that one is in no position to know *p*.

(See JAI: 108.) If we let '*J*' be short for 'one has propositional justification for believing', '*D*' for 'one has doxastic justification for believing', '*k*' for 'one knows', and '*K*' for 'one is in a position to know', and conceive of cases as triples of worlds, subjects, and times, we can rephrase **H1** and **H2** as follows: for any case α, ⌜Jp⌝ holds in α iff ⌜$\neg K \neg Kp$⌝ holds in α, and for any case α, ⌜Dp⌝ holds in α iff ⌜$\neg K \neg kp$⌝ holds in α—where, throughout, 'one' refers to the subject at the centre of α. (Note that, by contrast, in JAI, '*J*' and '*D*' were just metalinguistic abbreviations of '$\neg K \neg K$' and '$\neg K \neg k$', respectively; see JAI: 83.)

Both types of justification were conceived of as being pro toto (all things considered). For this reason, it was essential to establish both of the following

$\mathbf{D}_J$ $\neg K \neg Kp \supset K \neg K \neg p$

$\mathbf{D}_D$ $\neg K \neg kp \supset K \neg k \neg p$.

For, plausibly, a proposition and its negation are never both pro toto justified. $\mathbf{D}_J$ is the so-called Geach axiom for '*K*'. In chapter 5 of JAI, I derived $\mathbf{D}_J$ from the following two principles which I had argued for in chapter 4:

L $\neg K \neg Kp \supset K \neg K \neg Kp$

A6 $K \neg K \neg Kp \supset K \neg K \neg p$

(see JAI: 62, 68-76). **L** is a controversial luminosity principle for which, however, I gave arguments (JAI: 68-79, 269-74; see also my reply to Smith, 2022). **A6** was motivated by the factivity of '*K*' (i.e. the validity of ⌜$Kp \supset p$⌝) and instances of normal modal reasoning of the kind that, despite my rejection of the idea that '*K*' behaves normally throughout, I found unobjectionable, given the idealizations in play (JAI: 62, 86). $\mathbf{D}_D$ was then derived from $\mathbf{D}_J$ and

D-J $K \neg Kp \supset K \neg kp$

which latter can be motivated by the validity of ⌜$kp \supset Kp$⌝ and similarly unobjectionable normal modal reasoning (JAI: 86-87, 89).

In his challenging review, Julien Dutant (2022) argues that $\mathbf{D}_D$ has counterexamples which, given ***D-J***, are also counterexamples to $\mathbf{D}_J$ and, to the extent that **A6** remains acceptable, in turn counterexamples to **L**. Worse still, the alleged counterexamples are cases in which, for some proposition, one has neither doxastic nor propositional justification for that proposition or its negation—so that **H1** and **H2**, too, would be refuted. Accordingly, Dutant can be seen as mounting a broadside attack on the account developed in JAI.

To this end, he invites us to consider cases in which one is in no position to know, and a fortiori does not know, the answer to a given *yes*-or-*no* question, and yet, once the question is raised and understood, one nonetheless has the impression that one knows the answer—an impression that, misleading though it is, persists throughout one's best efforts to decide the question, without ever giving way to any belief in a positive or negative answer to that question. Dutant's preferred examples involve rather resilient tip-of-the-tongue phenomena. He contends that one's persistent impression that one has the answer 'on the tip of one's tongue', and thus that one knows it, debars one both from being in a position to know that one doesn't know *p* and from being in a position to know that one doesn't know the negation of *p*, where the question at hand is whether *p* holds (Dutant, 2022: 4-5). If so, we would have

(#) $\neg K\neg kp \wedge \neg K\neg k\neg p$

and thus a counterexample to $\mathbf{D}_{\boldsymbol{D}}$ (and hence, *modulo* ***D-J*** and **A6**, to $\mathbf{D}_{\boldsymbol{J}}$ and **L**). Plausibly, in the cases as described, one has no doxastic or propositional justification for *p* or its negation and, accordingly, to the extent that (#) holds, **H1** and **H2** would founder, too.

In what follows, I wish to cast doubt on Dutant's contention that the persistent tip-of-the-tongue impression suffices for (#). Note that, in order to do so, I am under no obligation to postulate any epistemic asymmetry between *p* and its negation (which, given that one is ex hypothesi neither in a position to know *p* nor in a position to know its negation, would anyway be a tall order). For, the account proposed in JAI clearly allows for cases in which both ⌜$K\neg kp$⌝ and ⌜$K\neg k\neg p$⌝ hold (i.e. is consistent with the failure of the converse of $\mathbf{D}_{\boldsymbol{D}}$).

Knowledge implies belief where such a belief need not be occurrent. Still, in line with the idealizations I put in place (JAI: 108), I shall assume that, just as one will know *p* if one is in a position to know *p* and does the best that one is in a position to do in order to decide *p* (see **P1** in my reply to Smith, 2022), the following holds:

(1) If one knows the answer to a *yes*-or-*no* question that has been raised and understood, and one has made one's best efforts to come up with that answer in response to that question, then one will come up with that answer,

where coming up with an answer here implies *making* that answer *explicit* (if only in one's mind). (1) rejects the suggestion that one's knowledgeable beliefs may remain implicit forever despite one's best efforts to bring them to mind, although one has all the conceptual capacities necessary to do so. For all that (1) implies, there may be a notion of implicit belief that has this feature, but (1) anyway rejects the thought that knowledge implies a belief of this sort. Should Dutant insist that some of our knowledge implies a belief of that sort, and be able to make a successful case for this assumption, then, as a last resort, I still have the option to concede the point, restrict the scope of '*k*', and ask readers to understand the account of justification devised in JAI accordingly (and as it was intended; see JAI: 22).

Let me now make a number of observations on which I will rely in my response to Dutant's challenge:

(i) Impressions, or seemings, may persist although one knows better or knows something that is evidence for their misleading nature.

Textbook examples are provided by optical illusions, e.g. the Müller-Lyer illusion. Such persistent impressions, or seemings, are typically called '*phenomenal*' and contrast with so-called *epistemic* impressions, or seemings, that may be defeated by what one knows (Brogaard, 2013: 274-76; JAI: 51n2).

(ii) One may have evidence of a given kind even if one doesn't know that one has evidence of that kind.

This is plausible on a conception on which one's evidence consists in what one knows: $\ulcorner kq \supset kkq \urcorner$ is arguably invalid and, similarly, what one's evidence is evidence for, and how strong such evidence it is, need not always be known (even, that is, if one knows that one has that evidence) (cf. Williamson, 2000). This is also plausible on a conception on which one's evidence is what one is in a position to know (JAI: 204-205): $\ulcorner Kq \supset KKq \urcorner$ is arguably invalid and, similarly, what one's evidence is evidence for, and how strong such evidence it is, isn't luminous (even, that is, if one is in a position to know that one has that evidence) (JAI: 66, 227).

(iii) To heed evidence of a given kind, one needn't have evidence that one has evidence of that kind.

This is plausible on a conception on which one's evidence consists in what one knows: to heed one's evidence *p* is to act as it would be rational to act were one to know *p*, yet to act as it would be rational to act were one to know *p* isn't the same as to act as it would be rational to act were one to know that one knows *p*. On a conception on which one's evidence consists in what one is in a position to know, plausibly, in order to heed one's evidence *p* one must know *p*, so that, again, to heed one's evidence *p* is to act as it would be rational to act were one to know *p*. Yet, to act as it would be rational to act were one to know *p* isn't the same as to act as it would be rational to act were one to know, or at least to be in a position to know, that one knows *p*.

(iv) We are familiar with there being propositions *p*, features *F*, and potentially co-referring concepts *a* and *b* such that, first, *a* refers iff *a* and *b* co-refer iff some favourable condition obtains, and secondly, if knowing <*a* has *F*> would put one in a position to know *p*, then, all things being equal, knowing <*b* has *F*> would put one in that same position.

For instance, let *a* be *my visual experience generated by exercise of my senses under normal conditions* and let *b* be *my visual experience*. If knowledge of <*a* is as of a cat on the mat> would put me in a position to know that some beast is on some mat, then, to the extent that my visual experience does in fact result from the exercise of my senses under normal conditions, knowledge of <*b* is as of a cat on the mat> would, all things being equal, put me in that same position. To give another example, let *a* be *the result of my error-free calculation* and let *b* be *the result of my calculation*. If knowledge of <*a* implies that the pressure is above threshold *t*> would put me in a position to know that the pressure is too high, then, to the extent that my calculation is in fact error-free, knowledge of <*b* implies that the pressure is above *t*> would, all things being equal, put me in that same position. Examples can be multiplied.

Against the backdrop of these observations and in the light of (1), I take the following to be eminently plausible:

(2) If a *yes*-or-*no* question has been raised and understood and if, even after prolonged efforts to come up with an answer in response to that question, one doesn't come up with any such answer, then the observation that one doesn't come up with any such answer even after those efforts are made will eventually provide one with inductive evidence for the thought that one doesn't know the answer to that question.

(2) doesn't say how long one must have unsuccessfully tried to come up with an answer, before observation that one fails to do so in spite of such attempts turns into inductive evidence for the negative conclusion that one doesn't know the answer to the question at hand. It merely claims that some such point will eventually be reached—a point at which whatever evidential weight one's tip-of-the-tongue impression had will be outweighed (cf. JAI: 203). Beyond that point, the longer the observed unsuccessful efforts last, the stronger the inductive evidence for the negative conclusion. Note that, given (ii), having inductive evidence of this kind doesn't require having evidence that one has evidence of this kind. Given (iii), nevertheless, one may then still heed that evidence.

If it still seems to one that one knows the answer to the question at hand, then inductive evidence for thinking that one doesn't know the answer to that question is also evidence for thinking that those seemings are misleading. Accordingly, given (2), to the extent that one should heed one's evidence, in the case described by Dutant, as she continues to rake her brain, the subject should give less and less weight to her tip-of-the-tongue impression and become more and more inclined to conclude that it is misleading, even if, phenomenally, that impression might persist. Subjects who observe that they don't come up with an answer to the question at hand even after prolonged efforts to do so—and continue making observations of the same kind as they continue making such unsuccessful efforts—will fail to behave in the rational ways in which idealized agents were stipulated to behave, if they nonetheless continue to give as much weight to their tip-of-the-tongue impression as before. For, this would be an evidential bias of the kind the idealization was designed to exclude (see, e.g. JAI: 108).

I think it is telling that Dutant nowhere assumes that the subject in question *believes* that they know the answer to the question; for, their rational failing would then be more evident. But even the refusal to acknowledge, after so many observed unsuccessful attempts to come up with an answer, that the tip-of-the-tongue impression ceases to be a reason to continue trying is a rational failing. So-called epistemic seemings, unlike phenomenal seemings, 'go away in the presence of a rebutting defeater, if the agent is rational' (Brogaard, 2013: 274). So, if the tip-of-the-tongue impression persists with equal force, as

Dutant assumes it does and as (i) acknowledges it may, then, if the subject is fully rational, the impression in question is phenomenal rather than epistemic. As such, it does not have the force Dutant assumes it to have, viz. to debar the subject from acknowledging that they do not know the answer to the question. This also casts doubt on Dutant's diagnosis that his case refutes what I say about the availability of responsive beliefs in the context of my argument for **L** (Dutant, 2022: XX; JAI: 73-74). The notion of responsive belief was explicated in terms of epistemic seemings rather than merely phenomenal ones (JAI: 51-53).

Plausibly, if one has made one's *best* efforts to come up with an answer to the question at hand, and those efforts prove unsuccessful, then one will have tried long enough in order for one's observation that one's efforts remain unsuccessful to constitute pretty strong inductive evidence for the negative conclusion that one doesn't know the answer after all. In the light of this and (ii), I suggest the following:

(3) If a question has been raised and understood and, even after one's *best* efforts to come up with an answer in response to that question, one doesn't come up with any such answer, then the observation that one doesn't come up with any such answer even after those efforts are made puts one in a position to know that one doesn't know the answer to that question.

(3) doesn't say how long one must have tried to come up with an answer to the question, in order for one to have made one's best efforts to come up with such an answer. It merely presupposes that there is a point at which one has done one's best in this regard.

To the extent that one knows (1) and one is sufficiently logically astute, it isn't hard to see that, if one has made one's best efforts to come up with an answer but has failed, observing *that this is so* would put one in a position to know that one doesn't know the answer. But, given (iv), this needn't be the only observation that may put one in a position to know this: if the efforts one has made are de facto one's best, then, plausibly, the observation that one has made the efforts one has made yet without success will likewise put one in that position. (Notice that the formulation of (3) deliberately remains neutral on which of these two kinds of observations is made.) By (ii), in order for one's observation to put one in a position to know that one doesn't know the answer, one needn't have evidence for thinking that it does so.[23]

The following principle seems rather plausible:

(4) If, for any case $\beta \in B(\alpha, p)$, one is in a position to know p in β, then one is already in a position to know p in α,

where, as before, $B(\alpha, p)$ is the set of possible cases, close to α and centred on the same subject as α, in which the subject of α has done the best they are in α in a

[23] Note that, on the account devised in JAI, $\ulcorner K\neg kp \urcorner$ doesn't encode a luminous condition while $\ulcorner K\neg Kp \urcorner$ does. Futhermore note that, on that account, the converse of ***D-J*** fails. Accordingly, could Dutant now concede that $\ulcorner K\neg kp \urcorner$ and $\ulcorner K\neg k\neg p \urcorner$ both hold, but nonetheless try to make a case for $\ulcorner \neg K\neg Kp \wedge \neg K\neg K\neg p \urcorner$ (and hence directly against $\mathbf{D}_J$) by exploiting the fact that, if $\ulcorner K\neg Kp \urcorner$ held, then, by my own lights, so would $\ulcorner KK\neg Kp \urcorner$? He might, although the devil is, as always, in the detail. But, in any event, in order for $\ulcorner KK\neg Kp \urcorner$ to hold in either scenario—as I am anyway committed to saying since one lacks propositional justification for p in either of them—it still needn't be the case that one be in a position to know that one's obervation puts one in a position to know that one is in no position to know p.

position to do in order to decide *p*, *and has done nothing they are not in α in a position to do*. For, if, for all β ∈ B(α, *p*), whatever one ends up knowing in β was something that one was already in a position to know in α, if there is any such knowledge that puts one in a position to know *p* in β, plausibly, one was already in that position in α.[24]

Unless Dutant can make a case that there are, besides raking one's brain and continuing to observe that one doesn't come up with any answer, other equally good or even better things one is in a position to do in order to decide whether one knows the answer and that don't put one in a position to know one doesn't know the answer, we may now use (3) and (4) to conclude that, in the kind of case Dutant invites us to consider, both ⌜$K\neg kp$⌝ and ⌜$K\neg k\neg p$⌝ hold and hence (#) fails.

6 Reply to Yiwen Zhan

In his thought-provoking comments (Zhan, 2022), Yiwen Zhan, like Daniel Waxman (2022), targets one of the key theses I advanced in JAI, viz.

H1 One has propositional justification for believing *p* iff one is in no position to know that one is in no position to know *p*.[25]

Zhan contrasts **H1** with the following thesis:

H1* One has propositional justification for believing *p* iff one *might* be in a position to know *p*,

where, according to Zhan, 'might' encodes some kind of 'epistemic opportunity' (Zhan, 2022: 2).[26]

[24] In JAI, I had argued for something stronger than (4) which, in light of the strategy I pursue in my reply to Waxman (2022), now seems too strong (cf. principle (1c) in JAI: 42).

[25] Zhan contends that my account of justification implies that, if I am in no position to know 'whether the fact that I am in a position to know *p* obtains, the second best thing is to remain agnostic about whether the opposite fact, i.e. that I am not in a position to know *p*, obtains' (Zhan 2022: 1). On a natural reading of 'whether', this is to say that the account implies that, if ⌜$\neg KKp \wedge \neg K\neg Kp$⌝ holds, one should remain agnostic about whether ⌜Kp⌝ holds. This implication is plausible quite independently from my account. If Zhan instead intends ⌜$\neg KKp \supset \neg K\neg Kp$⌝, then his contention is false: by the factivity of 'K', ⌜$K\neg Kp$⌝ implies ⌜$\neg KKp$⌝, and so, by the transitivity of the conditional, ⌜$\neg KKp \supset \neg K\neg Kp$⌝ would imply that ⌜$\neg K\neg Kp$⌝ holds, for arbitrary propositions *p*. This is clearly not implied by the account I gave in JAI.

[26] One of Zhan's reasons for claiming that 'it might be that' should not be construed as ⌜$\neg K\neg$⌝ is this: 'If [I assert] 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 [disagree with] 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' (Zhan 2022: 2n2). In Rosenkranz (2007), I tried to show that, on the contrary, such an explanation is after all to be had, arguing that, in reasoned debates, participants who make a disputed claim are committed to denying that any party to the debate is in a position to know that they are in no position to know that claim—where, plausibly, any such party will be in a position to learn by testimony what any other party is in a position to know (cf. also JAI: 127n5).

Given what he says about his first toy model discussed below, Zhan would seem to think of such epistemic opportunities in terms of epistemic states to which one is in a position to advance by acquiring more evidence—evidence that would eliminate some of the actually accessible epistemic possibilities and thereby put one in a position to know something that one wasn't in a position to know before. Since ⌜¬*K*¬*Kp*⌝ and ⌜¬*Kp*⌝ may hold while knowing *p* is impossible, it should initially be clear that, on the account of justification **H1** provides, propositional justification does not entail such epistemic opportunities. Surprisingly, however, Zhan writes that even on 'the ¬*K*¬*K* Rule's construal' of justification'—i.e. even given **H1**— 'a non-trivial case of justification [for φ, i.e. a case in which one is in no position to know the truth of φ] corresponds to an epistemic state with two kinds of potentials: the subject could either get to know the truth of φ (or ⌜¬φ⌝) or get to know that it is actually something [they are] in no position to know' (Zhan, 2022: 4).

This doesn't seem to be correct—at least if 'could get to know' here implies 'is in a position to know'. For, if ⌜¬*K*¬*Kp* ∧ ¬*Kp*⌝ implied ⌜*Kp* ∨ *K*¬*Kp*⌝, there would simply be no 'non-trivial cases of justification', in which event we would be saddled with the hopeless **S5** principle for '*K*', i.e. ⌜¬*Kp* ⊃ *K*¬*Kp*⌝. Given ⌜*K*¬*p* ⊃ *K*¬*Kp*⌝ (see JAI: 62), we would get the same unpalatable result, if ⌜¬*K*¬*Kp* ∧ ¬*Kp*⌝ implied ⌜(*Kp* ∨ *K*¬*p*) ∨ *K*¬*Kp*⌝.[27] Clearly, **H1** has no such unpleasant consequences.[28] If, instead, 'could get to know' here just means 'for all that one is in a position to know, one is in a position to know', then, since the disjunction ⌜(¬*K*¬*Kp* ∨ ¬*K*¬*K*¬*p*) ∨ ¬*K*¬*K*¬*Kp*⌝ trivially follows from ⌜¬*K*¬*Kp*⌝ alone, we don't learn anything illuminating about non-trivial cases of justification. It is unclear what other pertinent readings of 'could get to know' there are.

In fact, what one would have expected Zhan to say instead is that, according to **H1**, cases in which one *lacks* justification for *p* are cases in which one could get to know, if not ⌜¬*p*⌝, then at least ⌜¬*Kp*⌝. That's of course true on the first reading of 'could get to know'.[29] But such cases are not in any way illustrated by his first toy model (or his second toy model, for that matter; see below). This first model consists of three worlds, w_1 to w_3—with w_1 serving as the actual world—where

- *p* is true at w_1 and w_3, but false at w_2
- $\{w_1, w_2, w_3\}$ is the set of worlds accessible from w_1 and w_2, and $\{w_1, w_3\}$ is the set of worlds accessible from w_3

27 If ⌜¬*K*¬*Kp* ∧ ¬*Kp*⌝ implied ⌜*K*¬*Kp*⌝, it would imply a contradiction, and so ⌜¬*Kp* ⊃ *K*¬*Kp*⌝ would follow.

28 To the extent that 'one might be in a position to know *p*' entails ⌜¬*K*¬*Kp*⌝, but not the other way round, if **H1*** should have these consequences, this would be a good reason to reject it. Although I think of **H1*** as being equivalent to **H1**, if a case can be made that it isn't, I'm happy to discard **H1*** and stick to **H1**.

29 Note that it is a matter of controversy whether, on the first reading, ⌜one could get to know *p*⌝ implies ⌜one has the opportunity to come to know *p*⌝. For my more recent worries about the alleged link between being in a position to know and *opportunities* to come to know, see footnote 9 of my reply to Smith (2022).

and where

- ⌜$K\varphi$⌝ is true at a world w of the model iff φ holds in all worlds accessible from w.[30]

It follows that ⌜Kp⌝ is true at w_3 but false at both w_1 and w_2 and that ⌜$\neg K\neg Kp$⌝ is true at all worlds of the model and, hence, so is ⌜$K\neg K\neg Kp$⌝ (just as the luminosity of ⌜$\neg K\neg Kp$⌝ requires) (Zhan, 2022: 4). According to **H1**, then, this model is of cases in which the subject *has* propositional justification for p.

Zhan is, of course, quite right to point out that the model itself doesn't show that it models cases in which the subject has such justification. But, equally clearly, such formal models do not aspire to show this. To cause trouble for **H1**, Zhan needs to argue that there are cases, adequately captured by this model (or a related model preserving all its relevant features), in which one has no justification for p. Devising models alone doesn't do the job.

Setting out to construct such a case, Zhan nonetheless sees the need to devise a more complex model, even if it should be clear by now that even such a more complex model, taken on its own, will remain silent on what subjects at the centre of the cases that it illustrates have, or do not have, justification for. The second model is a model of four worlds, w_1 to w_4—with w_1 again serving as the actual world—where

- p is true at w_1, w_3, and w_4, but false at w_2
- q is true at w_1, w_2, and w_4, but false at w_3
- $\{w_1, w_2, w_3, w_4\}$ is the set of worlds accessible from w_1 and w_2, $\{w_1, w_3, w_4\}$ is the set of worlds accessible from w_3, and $\{w_1, w_4\}$ is the set of worlds accessible from w_4

and where, again,

- ⌜$K\varphi$⌝ is true at a world w of the model iff φ holds in all worlds accessible from w.

It follows that ⌜Kp⌝ is true at both w_3 and w_4, but false at both w_1 and w_2, that ⌜Kq⌝ is only true at w_4, and that all of ⌜$\neg KKp$⌝, ⌜$\neg K\neg Kp$⌝, and ⌜$\neg K\neg Kq$⌝ are true at all worlds of the model and, hence, so are both ⌜$K\neg K\neg Kp$⌝ and ⌜$K\neg K\neg Kq$⌝ (Zhan, 2022: 6-7). According to **H1**, then, all the cases for which this model serves as a model are cases in which one has justification for both p and q.

Zhan contends that this model is a model of cases in which one's failure to be in a position to know q *explains* why one is in no position to know ⌜$\neg Kp$⌝—i.e. cases that exhibit what he calls p's '*second-order structural unknowability*'—where, he claims, these cases importantly include some in which one lacks justification for p (Zhan, 2022: 6).

[30] To do justice to the assumption, explicitly made in JAI, that 'K' creates a hyperintensional context, Zhan assumes impossible worlds.

The first thing to note is that not only does ⌜$K\neg K\neg Kq$⌝ hold in the model, but so does ⌜$K(Kq \supset Kp)$⌝ and, hence, ⌜$K(Kq \supset \neg K\neg Kp)$⌝. Accordingly, one is in a position to recognize that one cannot rule out a situation in which ⌜Kq⌝ holds—a situation which one is in a position to know would suffice for the truth of ⌜$\neg K\neg Kp$⌝. Given factivity, such a situation would indeed so suffice. This puts some pressure on the idea that the model adequately captures *p*'s second-order structural unknowability, in the way the latter was defined. At least, for all that the model tells us, a situation in which both ⌜Kq⌝ and ⌜$\neg K\neg Kp$⌝ hold (i.e. one corresponding to w_4) is not merely an epistemic but a real possibility. If so, to what extent is the truth of ⌜$\neg K\neg Kp$⌝ *explained* by the truth of ⌜$\neg Kq$⌝, if it then doesn't even require it? In fact, the model validates ⌜$K((Kq \vee \neg Kq) \equiv \neg K\neg Kp)$⌝ and ⌜$K\neg K\neg Kp$⌝. So, one is in a position to know that, whether or not one is in a position to know *q*, one is anyway in no position to rule out that one is in a position to know *p*. So, one is in a position to recognize—and so it is the case—that the truth of ⌜$\neg Kq$⌝ doesn't make any difference on this score. Accordingly, at best *some* of the cases the model serves to model will be cases in which *p* is second-order structurally unknowable—viz. cases in which, while one is in a position to know that ⌜$\neg K\neg Kp$⌝ would hold even if ⌜$\neg Kq$⌝ did not, it is then the truth of ⌜$\neg Kq$⌝, together with some hitherto unspecified features of the actual world, that explains why ⌜$\neg K\neg Kp$⌝ holds. Consequently, the model is not itself a model *of* second-order structural unknowability.

Given that ⌜$K\neg K\neg Kp$⌝ holds in the model, Zhan will now also have to argue that there are counterexamples to

(∗) If one is in a position to know that one is in no position to know that one is in no position to know *p*, then one has propositional justification for believing *p*.

On the account developed in JAI, the right-to-left conditional across **H1** is equivalent to (∗), since, on that account, we have

L $\neg K\neg Kp \supset K\neg K\neg Kp.$

But, **L** remains controversial and so (∗) is arguably weaker than the right-to-left conditional across **H1**. Zhan's focus on cases for which his second model serves as a model therefore inadvertently makes for a heavier dialectical burden on his part. Moreover, as noted, ⌜$K\neg K\neg Kq$⌝, too, holds in the model. Accordingly, lest he be charged with begging the question, he must either allow for the assumption that *q* is justified or else give arguments for thinking that, in the cases he envisages, not only does *p* present a counterexample to (**H1** and) (∗), but so does *q*. The model itself is silent on the matter.

Let us now at last turn to Zhan's example of Professor Moore who selects one of two candidates for a vacancy according to criteria whose application she knows has led to successful choices in the past, all the while the success of those choices was in fact, but unbeknownst to her, pure chance—insofar, namely, as past candidates, like the one she presently chooses, merely happened to meet other criteria which she is in no position to recognize are in fact more important than the ones she applies.[31] Let the chosen candidate be Ben, let C_M be the criteria used by Professor Moore, let *p* be <Ben is the better candidate>, and let *q* be <C_M are not the most important criteria to use>. According to Zhan, while Moore is in no position to know *p*, she is in no position to know that she is in no such position *because* she is in no position to know *q*. Let that indeed be so.

Zhan contends that his second toy model adequately reflects the structure of this case. However, how plausible is it to assume that Professor Moore is in a position to know that ⌜$Kq \supset Kp$⌝ holds, all the while she is in a position to know that she is in no position to rule out that ⌜Kq⌝, and hence *q*, holds? It is hard to see how she might be in the former position as long as she is in the latter. For, what reason does she possibly have for assuming that, if she was in a position to know (and so it was true) that her criteria are not the most important ones—as she is in a position to know she might—she would then still be in a position to know that Ben is the better candidate? This problem is in no way defused by observing that ⌜$Kq \supset Kp$⌝ is a material conditional, and so equivalent to ⌜$\neg Kq \vee Kp$⌝. For, ex hypothesi, she is in no position to know ⌜Kp⌝ either. And, if she really was in a position to know that she is in no position to rule out that she is in a position to know *q*, why wouldn't she thereby have justification for *q*, just as **H1** then predicts? We may ask a similar question about *p*.

Zhan takes great care to describe the case of Professor Moore in such a way as to make intelligible why the professor selects Ben. The professor's selection criteria may not be best, but she has evidence that their past application has always had the desired results. She is, of course, mistaken about the reason why this has been the case, but she possesses that evidence nonetheless. Its possession—combined with her lack of any evidence for thinking that there is after all no real connection between meeting C_M and being a good candidate—already in part explains why she is in no position to recognize that she is in no position to know *p*. The recognition that Ben meets C_M also plays its part, and so does the fact that the professor is in no position to know *q*. For, plausibly, if she were in a position to know *q*, she would also be in a position to recognize that she is in no position to know *p* on the basis of her own criteria. But, or so I contend, this combination of factors likewise explains why the professor has justification for believing *p*. I see nothing in the description of the case that would refute this claim.[32]

[31] I here deliberately ignore Zhan's further assumption that the professor's selection of candidates is (at least in part) the result of her sexist bias, as this would disqualify her from the range of subjects for whom **H1** was supposed to hold (see JAI: 108). See the next footnote.

[32] Perhaps Zhan thinks that, by assuming that the professor's selection decision is partly driven by her unconscious sexist bias, he can ensure that the professor a fortiori has no such justification. While I don't think that such a bias alone prevents one from being justified, more importantly and as highlighted in the previous footnote, **H1** was never meant to apply to subjects suffering from such a condition. Contrary to what Zhan alleges (Zhan 2022: 9), to say this much is *not* to commit to the claim that for no choice of *q* will there be cases whose structure is adequately captured by his second toy model. (As noted, this model on its own both is silent on what justification is, or isn't, available and doesn't capture the kind of explanatory connections second-order structural unknowability requires.)

In his closing section, Zhan anticipates this line of response, but argues that it falls short of guaranteeing that no other cases can be constructed to which his toy model applies but in which the subject lacks justification for the relevant *p*. While I don't think that this observation alone, though accurate, is good enough as a challenge to **H1**, I agree that more needs to be done to defend **H1**. In the remainder, I will accordingly try to say something more principled in response to the allegation that the right-to-left direction of **H1** is threatened by this type of counterexample.

My opponent considers a case α in which ⌜¬*K*¬*Kp*⌝ holds, yet in which the subject at the centre of α has no propositional justification for *p*. It might seem easy enough to find a *p* and a subject who has no propositional justification for *p*—quite independently from whether or not she is in a position to know ⌜¬*Kp*⌝. But one shouldn't underestimate what this involves. It may often be straightforward to point to a piece of information, or evidence, present in α, in order to identify what the subject's propositional justification for a given *p* is in α. It doesn't follow, conversely, that the absence of such a piece of information, or evidence, in α implies that the subject has no such propositional justification for *p* in α. Here is one reason.

Commonly, propositional justification for a given *p* is understood as a potential for doxastic justification, viz. as a potential of one's epistemic situation that one is in a position to realize so as to convert it into the basis for a justified belief in *p*. To the extent that such a basis for belief is taken to consist in a piece of information or evidence—and there are legitimate doubts that this must always be the case—what anyway ultimately matters is its presence *once α's potential has been realized*. The question then is what it takes for this potential to be realized. Here is one plausible answer:

(LINK) For any β ∈ B(α, ⌜*Kp*⌝), one has propositional justification for *p* in α iff one has doxastic justification for *p* in β,

where, as before, for any *q* and γ, we let B(γ, *q*) be the set of possible cases close to γ, and centred on the same subject as γ, in which that subject has done the best they are in γ in a position to do in order to decide *q* (and has done nothing they are not in γ in a position to do; see JAI: 40). Since one is always in a position to do the best that one is in a position to do to decide a given question, even if this may amount to doing very little, we can safely assume that, for any *q* and γ, B(γ, *q*) ≠ ∅.

According to (LINK), one has propositional justification for *p* just in case one would end up having doxastic justification for *p*, were one to do one's best to find out whether one is in a position to know *p*. With (LINK) in place, my opponent will accordingly be committed to claiming that, with respect to the *p* and α in question, for some β ∈ B(α, ⌜*Kp*⌝), one has no doxastic justification for *p* in β. The account of doxastic justification proposed in my book (JAI: 108) implies

H2 One has doxastic justification for believing *p* iff one is in no position to know that one doesn't know *p*.

Given (LINK), **H1** and **H2** accordingly commit me to the following principle:

($LINK_H$) For any $\beta \in B(\alpha, \ulcorner Kp \urcorner)$, $\ulcorner \neg K \neg Kp \urcorner$ holds in α iff $\ulcorner \neg K \neg kp \urcorner$ holds in β.

We can derive ($LINK_H$) as follows. From the fact that knowing implies being in a position to know, we get

(1) If $\ulcorner k \neg Kp \urcorner$ holds in α, $\ulcorner K \neg Kp \urcorner$ holds in α.

As I argued in JAI, we can and should assume that, at least for reasonably idealized subjects, the following likewise holds (JAI: 89):

(2) If $\ulcorner K \neg Kp \urcorner$ holds in α, $\ulcorner K \neg kp \urcorner$ holds in α.

Recall the following two principles that I already invoked in the context of my reply to Martin Smith (2022):

P1 For any $\beta \in B(\alpha, p)$, if $\ulcorner Kp \urcorner$ holds in α, $\ulcorner kp \urcorner$ holds in β

P2 $B(\alpha, p) = B(\alpha, \ulcorner \neg p \urcorner)$.

From **P1** and **P2**, we immediately get

(3) For any $\beta \in B(\alpha, \ulcorner Kp \urcorner)$, if $\ulcorner K \neg Kp \urcorner$ holds in α, $\ulcorner k \neg Kp \urcorner$ holds in β.

From (1) to (3), we then obtain

(4) For any $\beta \in B(\alpha, \ulcorner Kp \urcorner)$, if $\ulcorner \neg K \neg kp \urcorner$ holds in β, $\ulcorner \neg K \neg Kp \urcorner$ holds in α.

This establishes one half of ($LINK_H$). The other half of ($LINK_H$) is equivalent to the following principle that I invoked in my reply to Daniel Waxman (2022):

P6 For any $\beta \in B(\alpha, \ulcorner \neg Kp \urcorner)$, if $\ulcorner K \neg kp \urcorner$ holds in β, then $\ulcorner K \neg Kp \urcorner$ holds in α.

There, **P6** was derived, from a criterion for the lack of propositional justification, by means of **H1** (cf. CRIT in my reply to Waxman, 2022). But, unless my opponent has an *independent* argument against **H1**, I see no reason why I cannot here appeal to **P6**. In fact, **P6** may be defensible irrespective of whether or not **H1** is assumed. So, pending any independent argument against **P6**, with (LINK) and ($LINK_H$) in place, my opponent is now forced to say that, with respect to the p and α in question, for some $\beta \in B(\alpha, p)$, in β, the subject at the centre of β has no doxastic justification for believing p although she is in no position to rule out that she in fact knows p. Yet, what might prevent her from ruling out that she in fact knows p that does not, at the same time, make her belief in p (if any) justified? For starters, any such belief would be epistemically blameless. As I argued in JAI, there is also a clear sense in which such a belief would be epistemically permissible (JAI: 126-37). There are, of course, plenty of competing accounts of doxastic justification that might allow prising these factors apart. But *that* would be an altogether different discussion to have.

My opponent might nonetheless want to insist that structurally the same argument can now simply be rerun with doxastic justification instead of propositional justification. What if there is a fact whose obtaining the subject is in no position to recognize and that, at the same time, debars her from ever advancing to an epistemic state in which she recognizes that she doesn't currently know *p*, all the while she has no doxastic justification for *p*? We are going in circles. For, why shouldn't this fact itself serve as a ground for the subject's doxastic justification for *p*? Evidently, this fact itself will not be a piece of information, or evidence, supporting *p* that is in the subject's possession. For, ex hypothesi, the subject is in no position to know that this fact obtains. But, as I argued in JAI, grounds for doxastic justification needn't be luminous, even if doxastic justification is. Given the connection between the truth of $\ulcorner \neg K \neg kp \urcorner$ and one's epistemic blamelessness in, and epistemic permissibility of, believing *p* (if one does), this just goes to show that doxastic justification for *p* doesn't require possession of any such distinct piece of information or evidence from which one could infer that one is in no position to rule out that one doesn't know *p*.

Acknowledgements The author wishes to thank Nikolaj Pedersen, Editor in Chief of the *Asian Journal of Philosophy*, for organizing the symposium and Julien Dutant, Niccolò Rossi, Martin Smith, Daniel Waxman, and Yiwen Zhan for their invaluable contributions.

Funding Open Access funding provided thanks to the CRUE-CSIC agreement with Springer Nature. Work on this paper received support from the project PGC2018-099889-B-I00, financed by the Spanish Ministry of Science, Innovation and Universities (MICINN).

Declarations

Competing interests The authors declare no competing interests.

References

Brogaard, B. (2013). Phenomenal seemings and sensible dogmatism. In C. Tucker (Ed.), *Seemings and justification: New essays on dogmatism and phenomenal conservatism* (pp. 270–289). Oxford University Press.

Dutant, J. (2022). Justification as ignorance and epistemic Geach principles. *Asian Journal of Philosophy, 1*(14), 1–7.

Heylen, J. (2016). Being in a position to know and closure. *Thought, 5*, 63–67.

Rosenkranz, S. (2007). Agnosticism as a third stance. *Mind, 116*, 55–104.
Rosenkranz, S. (2021). *Justification as ignorance*. Oxford University Press.
Rossi, N. (2022). An enhanced model for Rosenkranz's logic of justification. *Asian Journal of Philosophy, 1*(12), 1–9.
Smith, M. (2022). Is $\neg K \neg Kp$ a luminous condition? *Asian Journal of Philosophy, 1*(6), 1–10.
Srinivasan, A. (2013). Are we luminous? *Philosophy and Phenomenological Research, 90*, 294–319.
Waxman, D. forthcoming: 'Justification and being in a position to know', *Analysis*, https://doi-org.sire.ub.edu/10.1093/analys/anab086
Waxman, D. (2022). Justification as ignorance and logical omniscience. *Asian Journal of Philosophy, 1*(7), 1–8.
Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.
Williamson, T. (2009). Tennant's troubles. In J. Salerno (Ed.), *New Essays on the Knowability Paradox* (pp. 183–204). Oxford University Press.
Zhan, Y. (2022). The $\neg K \neg K$ rule and the structurally unknowable. *Asian Journal of Philosophy, 1*(8), 1–11.

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