Theoria, 2026, 41(3), 434-444
https://doi.org/10.1387/theoria.28294

THEORIA
ISSN 0495-4548 – eISSN 2171-679X

The opacity of knowledge: An externalist's case for widespread KK-failure

(La opacidad del conocimiento: un argumento externista a favor del fracaso generalizado del principio KK)

Murali Ramachandran*
Thapar University
https://orcid.org/0000-0001-5611-518X

Keywords

Epistemic externalism
KK
Epistemic closure
Basic knowledge
Easy knowledge
Bootstrapping
Epistemic entitlement

ABSTRACT: An epistemic externalist should allow that whatever proposition about the world, *P*, one knows to be true, there will be some other fact one is not in a position to know, but which is required for the belief in *P* to count as knowledge. We contend that ordinary knowers are cognizant of such 'unknown knowledge-makers'. On this basis, we present a novel case —we distinguish it from rivals— for *widespread* KK-failure: that, quite generally, if one knows *P*, one will not be in a position to know that one knows it.

Palabras clave

Externismo epistémico
KK
Clausura epistémica
Conocimiento básico
Conocimiento fácil
Bootstrapping
Derecho epistémico

RESUMEN: Un externista epistémico debería admitir que, para cualquier proposición sobre el mundo, *P*, que se sepa verdadera, habrá algún otro hecho que uno no está en condiciones de conocer, pero que resulta necesario para que la creencia en *P* constituya conocimiento. Sostenemos que los sujetos cognoscentes ordinarios tienen noticia de tales 'hacedores de conocimiento desconocidos'. Sobre esta base, presentamos un argumento novedoso —que distinguimos de sus rivales— a favor del fracaso generalizado del principio KK: a saber, que, en términos muy generales, si alguien sabe que *P*, no estará en condiciones de saber que lo sabe.

Gako-hitzak

Esternalismo epistemikoa
KK
Itxiera epistemikoa
Oinarrizko ezagutza
Ezagutza erraza
Bootstrapping
Eskubide epistemikoa

LABURPENA: Esternalista epistemiko batek onartu beharko luke munduari buruzko edozein proposizio (*P*) egiazkoa dela jakinda ere beste egitateren bat egongo dela, norberak ezagutzerik ez daukana baina *P* proposizioarekiko sinesmena ezagutzatzat jotzeko ezinbestekoa dena. Baieztatzen dugu ezagutzaile arruntak «ezagutza-sortzaile ezezagun» horien kontziente direla. Hori oinarri hartuta, argudio berritzaile bat aurkezten dugu —beste argudio batzuetatik bereiziaz batera— KK printzipioaren porrot orokortuaren alde: alegia, eskuarki, batek *P* ezagutzen badu, bat horrek ez du jakiterik izango *P* ezagutzen duela.

* ***Correspondence to:*** Murali Ramachandran. Thapar School of Liberal Arts and Science, Thapar University, Bhadson Road, Patiala, Punjab 147004, India. – murali.ramachandran@thapar.edu – https://orcid.org/0000-0001-5611-518X

How to cite: Ramachandran, Murali (2026). «The opacity of knowledge: An externalist's case for widespread KK-failure»; *Theoria. An International Journal for Theory, History and Foundations of Science,* 41(3), 434-444. (https://doi.org/10.1387/theoria.28294).

Received: 18 February, 2026; Final version: 06 September, 2026.
ISSN 0495-4548 - eISSN 2171-679X / © 2026 UPV/EHU Press

This work is licensed under a *Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International* License

1. Preliminaries

The key thought behind *epistemic externalism* is that knowledge must meet some kind of 'reliability' condition, but also that one may not always be in a position to know that this condition is met:

> Knowledge has a *'reliability'* requirement of sorts, e.g. that one's knowing a proposition *P* to be true requires one's belief in *P* to be *reliably-acquired* (*or safe from error, indefeasibly justified, non-accidentally true,* etc.) [**REL**]
>
> One may know a proposition *P* to be true and yet *not* be in a position to know whether one's belief in *P is* reliably-acquired (or safe from error, indefeasibly justified, non-accidentally true, etc.)—Following Cohen (2002), we'll call such knowledge *basic knowledge*[1] [**BK**]

We wish to draw attention to an overlooked argument showing that these conditions lead quite straightforwardly to widespread violations of the KK-principle:

> If one knows a fact *P*, then one knows —or one is at least in a position to know— *that one knows P*. [$K(P) \rightarrow K\,K\,(P)$] [**KK**]

This view is certainly not new; but, we contend, earlier proponents held it on the basis of faulty reasoning. Consider e.g. the following passages from Williams (1991) and Dretske (2004) ('modest contextualism' in the Dretske quote amounts to epistemic externalism as we have characterised it):

> I agree that no item of knowledge guarantees full knowledge of its own reliability conditions: indeed, *this is just another way of stating* the externalist point that knowing does not guarantee knowing that one knows. (Williams, 1991, p. 348, our emphasis)
>
> It is interesting to note, however, that modest contextualism (and, hence, externalism) provides an illuminating explanation of why KK fails. It fails because factual knowledge, according to modest contextualism, depends for its existence on circumstances of which the knower may be entirely ignorant. So the knower can know that *P* without knowing (as required by KK) that he knows that. (Dretske, 2004, p. 176)

What Williams and Dretske are suggesting is that [KK] will fail when, and *because,* the knower is ignorant of pertinent factors —the satisfaction of certain reliability (or safety) conditions— *in virtue of which* she has that knowledge. The underlying thought is presumably this: such a knower, being ignorant of these pertinent factors, *ipso facto* cannot *verify* that these obtain; and, hence, she cannot know *that she knows.*

However, externalism does not demand as much. Consider e.g. a crude causal theory of knowledge: one knows *P* iff one believes *P* and this belief is caused by the fact that *P*.[2] What this account requires for someone to know that they know *P* is only that they believe that they know *P* and this belief is caused by their knowing *P*; they are *not* required to know or be in a position to establish that this causal relation obtains; moreover, they are *not* required to know or be in a position to establish that their belief in *P* itself is caused by the fact that *P*. More generally, Williams's and Dretske's reasoning fall short because the reliability (safety, indefeasibility, etc.) of one's *belief that one knows P* does not require that one is aware of what makes one's *belief that P* reliable (safe, etc.).

---

[1] Actually, Cohen (2002) introduces the term 'basic knowledge' for knowledge one purportedly has without being in a position to know that *its source is reliable* (Cohen, 2002, p. 310). But his central concern—the so-called *problem of easy knowledge* —arises on our broader conception of basic knowledge too. We make some brief remarks about that problem in § 3.4.

[2] Obviously, a feasible causal theory would need to rule out 'deviant' causal chains—but our crude account serves for the point we wish to make.

Of course, Williamson (2000, § 5.1) has famously argued against [KK], making appeal to epistemic *closure*[3]—-i.e. for our purposes:

> If one knows that P_1, ..., and that P_n, and at the same time also knows that these propositions jointly entail Q, then one knows, or is at least in a position to know, that Q [**EC**]

But his argument does not purport to *identify* any proposition which is supposedly known but not known to be known; rather, it aims to establish merely that such a proposition must exist.[4] To this extent, his refutation has little impact on everyday verdicts on second-order knowledge: we can allow that knowers *on the whole* know that they know. By contrast, if the rationale we present in this paper is accepted —which also assumes [EC]— pretty much *whatever* we know about the world, *we do not know that we know it!*

We shall lead up to the key strategy by way of considering two sorts of cases an externalist might categorize as cases of 'unacknowledged knowledge', as it were, and, so, count as examples of KK-failure. (But, just to be clear, these are not arguments we ourselves would wish to defend—rather, they shall serve to introduce, and highlight features of, the case we wish to draw attention to.)

2. *Unacknowledged knowledge*

2.1. *The cautious skeptic*

Suppose a cautious skeptic goes around asserting the following sorts of things: she *believes* she is sitting at her desk but doesn't *know* it; she *believes* the room is illuminated, but doesn't *know* it; she *believes* she is writing down her ideas, but doesn't *know* it; etc.. Suppose her perceptual capacities and logical acumen are fine. Her perceptual beliefs are true, and they change accordingly with changes in her environment.

Externalists will surely hold that she *does* know these things despite her denials—because she is appropriately and adequately sensitive to what is going on around her; her beliefs are reliably acquired and safe from error. However, they may allow that she does not *know that she knows* these facts, but only because she does not believe —and, given her misconceptions about knowledge, is not even *in position to believe*— that she knows these facts, and knowledge requires belief.

So, we have KK-failure here, but, arguably, the skeptic herself is at fault.

2.2. *Closure versus KK*

What is less widely appreciated is that the (many) alleged counterexamples to closure —of the sort famously advanced by Dretske (1970, 1971) and Nozick (1981)— may be recast as examples of 'unacknowledged knowledge' instead, thereby providing a case for KK-failure in the same scenarios.

We'll use Dretske's (1971) famous zebra example to illustrate the strategy. Suppose Meena is visiting the zoo, and sees a zebra in an enclosure. Dretske allows, *contra* [EC], that she may know the premises of the following (valid) argument but not be in a position to know its conclusion:

---

[3] The following quotation shows that Dretske himself is not appealing to closure—which he in any case famously rejects (Dretske, 1970, 1971): "Although modest contextualism rejects KK, it does not, taken by itself, deny closure" (Dretske, 2004, p. 176).

[4] To give a bit more detail —a full elaboration is unnecessary here— the argument has Mr Magoo seeing a tree in the distance; what Williamson's [KK]-*reductio* supposedly establishes is that for *some n*, Mr Magoo knows $[\neg T_n]$ *that the tree is not n inches tall,* but he does not know *that he knows* $[\neg T_n]$.

ZEBRA

(1) [Zeb] The animal in the pen is a zebra.

(2) If [Zeb], then [¬CPM] it is not a cleverly painted mule.

(C) Hence, [¬CPM].

According to Dretske, Meena's sighting the zebra suffices for her to know [Zeb], but not for her to know [¬CPM], that what she is looking at is not a cleverly painted mule. This accords with ordinary folk's intuitions: intuitively, to know *that,* she would have to examine the animal up close. And Meena herself would most likely deny that she knows [¬CPM].

Many find the intuitive cost of denying closure too high, however.[5] For example, if one takes ZEBRA to be an [EC]-counterexample, one must surely allow that Meena also knows [¬M] *that the animal in the pen is <u>not a mule</u>* despite not knowing [¬CPM] (that it is not a *cleverly painted <u>mule</u>*). We prefer to pay the costs associated with *endorsing* closure (and make partial payment in § 3.4.)

The pertinent point for present purposes is that the putative closure-counterexamples can be re-modeled as KK-counterexamples. To begin with, closure-endorsers will insist, on the basis of closure, that if Meena *does* know the ZEBRA-premises, then she is simply mistaken in thinking she does not know the conclusion. But, they can then explain her reticence in claiming knowledge of [¬CPM] by putting it down to her *not knowing that she knows* it: she does not know *that she knows* because she does not *believe* that she knows (and knowledge requires belief).[6]

And once it is granted that Meena knows, but does not know that she knows, [¬CPM] in the zoo scenario, our closure-endorsers can then argue that she does not know that she knows [Zeb] either—by consideration of the following argument:

K-ZEBRA

(1) Meena knows [Zeb].

(2) Meena knows (*If* [Zeb], *then* [¬CPM].)

(C) Hence, Meena knows [¬CPM].

Allowing that Meena herself recognizes the validity of [EC], she will recognize that K-ZEBRA is (trivially) valid. And we can surely allow that she knows the second premise—it is not unreasonable to assume that simple *a priori* truths can be known and be known to be known. So, since, as our closure- endorser has it, Meena is not in a position to know that she knows [¬CPM], i.e. (C), her knowledge of [EC] dictates that she does not know the first premise—i.e. we must deny that she knows that she knows [Zeb].[7]

---

5 Contemporary closure-deniers, e.g. Alspector-Kelly (2019) and Kelp (2021), avoid some pitfalls of the original Dretske (1971) and Nozick (1981) accounts—including e.g. Kripke's (2011) notorious Red Barn example. But equally counterintuitive results are inevitable—such as the one we mention here.

6 Skeptics and closure-deniers accommodate people's reluctance to affirm that they know the likes of [¬CPM] by maintaining that they indeed *do not know* it. (Epistemic contextualists take the same line while also attempting to preserve closure). To our mind, the reluctance (or awkwardness) people feel is just as plausibly accommodated by taking it to signify that they *do not believe that they know.* Williamson (2005, p. 227 ff.) likewise argues for KK-violations in such cases by appealing to the knowers' *unpreparedness to use the proposition for practical reasoning* in the pertinent context. We think it simpler to appeal directly to the fact that the subjects do not believe *they know.*

7 We thank a referee for forcing clarification of our reasoning here.

Finally, since putative closure-counterexamples such as ZEBRA can be devised on demand for just about any known proposition, the table-turning anti-KK strategy just outlined yields a case for ubiquitous KK-failure instead.

But this case against [KK] has two, not insignificant, drawbacks. First, it rests on ordinary folk having *misplaced* qualms about what they know: e.g. it is precisely, and only, because Meena mistakenly believes she does not know [¬CPM] that [KK] is violated in the ZEBRA scenario. One might well balk at a case against [KK] which pivots solely on how diffident or cautious subjects are. The strategy would appear to have the consequence that a dithering subject fluctuates between not knowing one minute, and knowing the next, that she knows. The second shortcoming is that one may rightly question whether demonstrating that an individual knows, but does not know that they know, a certain fact is in itself sufficient to refute [KK]: for, [KK] —our version of it at any rate (p. 2)— is violated only if the individual is *not even in a position to* know that fact. Our ditherer, one might think, was always *in a position* to know that she knew.

So, as we proclaimed at the outset, we are not minded to press the foregoing arguments against [KK]. However, the anti-KK strategy we are about to outline is more compelling; and it has the consequence that an *entirely faultless* subject may know she is seeing a zebra, for example, without being in a position to *know that she knows* she is. But, let us give fair warning: it will not suit all closure-endorsers.

3 The case from basic-knowledge bootstrapping

3.1 Strategy

The following case against [KK] rests on a specific type of basic-knowledge situation (call it a *BK-bootstrapping* situation):

— An individual, X, knows some empirical fact, *P* (A)

— *Q* is a proposition which, given the way X acquired her belief in *P*, has to be true in order for X's belief in *P* to count as knowledge (e.g. reliably-acquired, safe from error, indefeasible, non-accidentally true, etc.). Thus:

(X knows *P*) *only if Q* (B)

— X knows (B)—i.e. knows that: (she *knows P*) *only if Q* is true (C)

(For any such *Q*, we'll call [¬*Q*] a *knowledge-defeating possibility* for X's belief *P*).

— But, X is not in a position to know *Q* (D)

(A), (B) and (D) are what our initial characterization of epistemic externalism, i.e. [REL] and [BK], allows: X's knowledge of *P* is basic knowledge; (C) is what makes this a special case of basic knowledge—*special, but commonplace* we contend. For, we think that most people would acknowledge certain possibilities as *knowledge-defeating* possibilities for (just about) any belief they had about the world. (So, our view is that even ordinary folk are really externalists, advocates of basic knowledge, at heart—they just need the right prodding!)

Assuming [EC] (epistemic closure), it is easily demonstrated that [KK] is violated in any BK-bootstrapping situation: X does not know that she knows the pertinent *P*:

The BK-bootstrapping KK-reductio

1. X knows that she knows *P* from (A) and [KK], for reductio
2. X knows that [X knows *P only if Q* is true] from (C)
3. X knows *Q* from (1), (2) and [EC]

4. X doesn't know Q from (D)
5. $\perp$ (3) and (4) contradict
6. So, [KK] is refuted from reductio of (1)

So, [KK] is violated in any situation meeting conditions (A)-(D). We contend that such situations are ubiquitous.

3.2. Examples

Example 1: LIAR

Suppose Meena meets someone at a party who introduces himself as Dilip and, in the course of the ensuing conversation, reveals he is a physicist. She believes him. There is no deception here. We can allow, then, that Meena knows [P] *that Dilip is a physicist.* Further, given her grounds for [P], we can also assume Meena recognizes that she does not *know* [P] if [L] *Dilip is in fact a convincing compulsive liar.*[8] In our jargon, [L] is a knowledge-defeating possibility for Meena's belief [P]—and she knows it. But, she surely does not know [$\neg L$] on the basis of their brief conversation.

So, consider the following valid argument:

LIAR

(1) Meena knows [P] that Dilip is a physicist.
(2) (Meena knows [P]) only if [$\neg L$] Dilip is not a convincing compulsive liar.
(C) Hence, [$\neg L$] Dilip is not a convincing compulsive liar.

The case against [KK] is straightforward: we have allowed that Meena knows LIAR(2); if she also knew LIAR(1), closure ([EC]) would dictate that she is in a position to know what is entailed by them, namely, [¬L]—which she doesn't. And Meena cannot come to know [$\neg L$] by way of deducing it from the premises; for, in the given circumstances, the reasoning exemplifies a form of *bootstrapping:* in affirming LIAR(1), Meena has *already* taken the falsity of [L] *as granted,* as it were, albeit tacitly. So, Meena does not know LIAR(1), period. Hence, Meena knows, but does not know that she knows, [P]. As it happens, Dilip is a wonderful, honest person; but, she simply is not yet in a position to know that.

Before proceeding to the second example, let us spell out the *formal* differences between ZEBRA, K-ZEBRA and LIAR just to highlight what is distinctive about our bootstrapping case against KK, even though all three assume [EC], epistemic closure.[9]

— ZEBRA is simply an instance of *modus ponens:* [$A, A \rightarrow B \models B$]; the externalist case against KK we floated in this case granted that Meena knows the premises, and, hence, by [EC], that she knows the conclusion, B; it accommodates her own awkwardness about affirming that she knows B by proposing that she does not *believe* that she knows B, and, hence, since knowledge is factive, that she does not *know that she knows B*.

— *K*-ZEBRA basically *piggybacks* on the above stance; it is of the form [$K(A), K(A \rightarrow B) \models K(B)$]. This is an *instance* of [EC]; so, *assuming Meena herself knows* [EC], she knows that *K*-ZEBRA is trivially valid. Meena's knowledge of [EC] therefore dictates that if she knows *K*-ZEBRA's premises, she is at least in a position to know that she knows B, which

[8] We emphasize 'know ' to highlight that [P] itself is not at odds with [L]; it is Meena's *knowing* [P] which is. The same point applies to the arguments we'll be considering later.

[9] Again, we thank a referee for pressing the need for such clarification.

our externalist maintains is *not* the case. So, our externalist argues, one must deny that Meena knows both premises—the premises in this case being that she knows A and that she knows $[A \rightarrow B]$. Next step: since $[A \rightarrow B]$ in this specific case is an *a priori* truth, it is granted that she knows this, and also that she knows that she knows this. We are thereby compelled to deny that she knows the first premise—i.e. to deny that she *knows that she knows A*.

— LIAR is of the same logical form as ZEBRA (i.e. *modus ponens*), just with 'A' replaced by '$K(A)$': $[K(A), K(A) \rightarrow B \models B]$. But, unlike ZEBRA, this is a bootstrapping argument for Meena. So, our externlaists takes there to be no grounds for maintaining that Meena knows the conclusion, B, to be true. Hence, as with K-ZEBRA, we are led to deny that she knows both premises; and, again, since the second premise appears *a priori* true, we can grant that she knows it to be true. In which case, she does not know the first premise, i,e, she does not *know that she knows A*.

Let us return to our other examples of BK-bootstrapping.

Example 2: HOLIDAY

Here is different sort of example illustrating that the reliability (or safety) of a belief depends on far more than its *provenance.* Suppose Meena has just booked her flight for a holiday in the Bahamas a month hence. Presumably, she *may* then know [B] *that she will be in the Bahamas the following month*. But, as she herself will grant, she knows [B] only if [$\neg A$] *she does not come close to having a serious accident days before her departure*. However, she surely is not in a position to know the latter. So, we have another counterexample to [KK]: given that Meena knows that *her knowing* [B] is incompatible with [A], if she *knew that she knew* [B] as well, [EC] would dictate that she *is* in a position to knew [$\neg A$]—which she isn't; so Meena does not know that she knows [B].[10]

Example 3: MEETING

Our third candidate KK-counterexample is a variation of ZEBRA (§ 2.2). It suggests that beliefs have to be *indefeasibly justified,* not just reliably-acquired or safe from error, in order to qualify as knowledge. Consider the earlier zebra-scenario: Meena sees a zebra at the zoo and thereby comes to know [Zeb] *that the animal in the pen is zebra.* But, presumably, her belief would not have qualified as knowledge if e.g. [Meet] *there had been a recent meeting of the zoo's governors where the possibility of deploying mules cleverly painted to look like zebras was seriously considered*—(there is no need to delve into possible explanations for such a meeting). We allow that Meena knows that *her knowing* [Zeb] is incompatible with [Meet], but that she is in no position to know that the latter is false. As before, given that Meena knows [Zeb], and knows that her knowing this is incompatible with [Meet], she cannot also know that she knows [Zeb], for by closure this would dictate that she is in a position to know [¬Meet], which she isn't.

The diverse variety of illustrative examples we have just considered would seem to suggest that BK-bootstrapping arguments could be devised for just about any known facts about the world. If our case against [KK] is sound, KK-failure is *rife;* knowledge *in general is opaque:* there are hardly any known facts about the world which are known to be known.

3.3. A disparity in need of explanation

However, an obvious anomaly needs to be addressed: why should someone who maintains that Meena *knows* [¬CPM] (that the animal in the pen is not a cleverly painted mule) in the zoo scenario —*despite* her own protestations to the contrary one might add— *deny* that she knows [¬Meet] (that the zoo's governors have not met recently to consider the deployment of

[10] As a referee points out, some may question whether one can have knowledge of future events at all; so here is a variation on HOLIDAY involving knowledge of the *past.* Decades after Meena's trip to the Bahamas, her son, Raju, comes across her booking among her belongings; we can surely allow that Raju knows [B] *that Meena visited the Bahamas;* but Raju realizes that he *knows* [B] only if [$\neg C$] *Meena did not come close to canceling the trip in order to attend a wedding,* which he presently is not in a position to know. In this BK-bootstrapping situation, Raju knows, but does not know that he knows, [B].

cleverly painted mules) in the same setting? Our stance is apt to seem even more puzzling given that closure-endorsers must maintain that Meena *knows* —because it is entailed by [Zeb]— [¬BIV]: *that she is not a deceived brain in a vat in a zebra-less world.* So, why the disparity? Why not take Meena to know [¬Meet], even if she does not believe that she does?

Closure-endorsers will indeed have diverging responses here. We shall briefly demarcate what we take to be a feasible stance that underwrites our treating the likes of [¬CPM] and [¬BIV] differently from the conclusions of our BK-bootstrapping arguments in § 3.2.

4. Epistemic priority, transfer of warrant and collateral knowledge

We shall distinguish three kinds of closure-endorsers by their take on putative counterexamples such as ZEBRA.

Knowledge by deduction

(A) First, there are those who allow that Meena can *come to know (acquire knowledge that)* [¬CPM] holds by way of *competently deducing it* from [Zeb], which (by hypothesis) she knows on the basis of her sighting.

We would put Williamson (2000, p. 117) and Hawthorne (2013, p. 29) in this camp, since they appear to maintain that one can, for any valid inference, *acquire* knowledge of the conclusion by way of *competent deduction.* Other possible (A)-listers are Beebee (2001) and Markie (2005), who take the arguments from [Zeb] to [¬CPM] and [¬BIV] to *beg the question* against those who doubt or question the conclusions, but, apparently, nevertheless allow the 'inferential-route', as it were, to knowledge of those conclusions.

(A)-type closure-endorsers may resist our case against [KK], insisting that Meena can indeed come to know the conclusions of the arguments considered in § 3.2 by way of deduction—so long as she does know the pertinent premises. But to insist this just begs the question against our contention that these are *bootstrapping* arguments, which, intuitively, can give no succor to the conclusions in the specified circumstances.

Epistemic entitlement

(B) The second type of closure-endorser *denies* that Meena can acquire knowledge of [¬CPM] by way of deduction from the premises of ZEBRA, on the grounds that the *warrant* for the first premise *'fails to transfer'* to the conclusion; but, rather than *rejecting* closure on these grounds as Dretske (1970, 1971) does, they maintain that Meena *'already' knows* [¬CMP] to be true.

Davies (2004) and Wright (2004) are prominent (B)-listers; they attempt to preserve closure by maintaining that Meena is *'epistemically entitled'* to accept [¬CPM], that the animal she sighted is not a cleverly painted mule. In their view (albeit for different reasons), Meena's knowledge of [Zeb] *presupposes* (depends on, and is partly explained by, the fact) that she knows [¬CPM]—the latter item of knowledge is *epistemically prior,* so to speak.

But, to be blunt, we consider the appeal to 'epistemic entitlement' —and we include the invocation (of knowledge) of so-called *hinge* propositions here— to be wishful thinking. The fact that we would have to rule out knowledge *tout court* in a certain area of discourse, a 'region of inquiry' as it were, unless a proposition *p* is assumed to be true, does not, in our view, simply *entitle* us —as Wright (2004) appears to maintain— to take *p* to be true, let alone to take ourselves to *know* that is.

Collateral knowledge

Our misgivings about these approaches lead us to the third closure-preserving camp—though we fear we may be the only occupants as yet:

(C) This camp sides with (A)-listers in taking Meena's knowledge of [¬CPM] to be explained by —to stem ('flow') from— her knowledge of [Zeb] rather than, as (B)-listers have it, the other way around. But it sides with (B)-listers in denying that Meena can *acquire* knowledge of [¬CPM] by way of *deduction* from ZEBRA's premises. Rather, it takes her knowledge of [¬CPM] to emerge *in sync* with, at the same time as, her coming to know [Zeb]; but, this is not a case of *inferential*-knowledge.[11]

Crudely, we think of Meena's knowledge of [¬CPM] —and let us not forget [¬BIV] (that she is not a deceived brain in vat in a zebra-less world)— as *collateral knowledge:* she attains this knowledge *by virtue* of coming to know [Zeb]. To give a more mundane example: Meena's knowing [Zeb], given the concepts she possesses, dictates her knowing a host of other facts—e.g. that the animal in the pen is not a mule; that it is not a kangaroo; not a giraffe; etc. These latter items of knowledge are not, in our view, known by way of inference from [Zeb]; for, Meena's belief that the animal is zebra *comes with* the (albeit tacit) belief that it is not a mule, not a kangaroo, not a giraffe, and so on. She likewise acquires more esoteric knowledge of the likes of [¬CPM] and [¬BIV] in our view, since her understanding of what it is to be a zebra, given other concepts she possesses, makes it rationally impossible for her not (to be disposed) to believe these other propositions if she believes [Zeb].[12]

Of course, some story will have to be told about what makes her knowledge of [Zeb] different, epistemically prior or more basic; we hazard that it is distinguished by the fact that her belief in [Zeb], unlike her belief in [¬CPM] or [¬BIV], meets some sort of 'sensitivity' condition—but it would take us too far afield to detail or defend a specific account here.[13]

By contrast, coming back to the disparity 'in need of explanation', [¬Meet] is *not* a consequence of [Zeb]—as we have stressed, it is Meena's *knowing* [Zeb] which entails [¬Meet]. And, as we have argued, since her knowing [Zeb] *(a priori) depends* on the truth of [¬Meet], she cannot come to know that conclusion by way of deducing it from the premise that she knows [Zeb]—for, this would be a case of epistemic bootstrapping. So, we deny that Meena knows [¬Meet] because, in our eyes, there are *no satisfactory explanation* of how she attains this putative knowledge. More generally, we distinguish between the (obvious) consequences of 'worldly' propositions one knows, such as [¬CPM] and [¬BIV] —these comprise collateral knowledge one has— and the consequences of one's *knowing* such worldly propositions, which *ipso facto* are not known if they conform to our bootstrapping template in § 3.1 (p. 7).

It is incumbent on any externalist who still maintains that Meena really is in a position to know the conclusions of the bootstrapping arguments in § 3.2 —e.g. some steadfast A- or B- lister— to explicate the sorts of 'reliability' (safety, defeasibility) conditions they take to underpin Meena's knowledge of [Zeb], say, but which she is not in a position to know. Unless some such condition is advanced, we invite clarification on what makes them externalists at all—for, presumably, the existence of basic knowledge will not be key.

Conclusion

We can summarise our final position by comparison with the remarks from Williams (1991) we quoted earlier (p. 2). We wholly agree with him that "*no* item of knowledge guarantees *full* knowledge of its own reliability conditions" (Williams, 1991, p. 348, our emphases). But we would go further: not only are there reliability conditions which the

---

[11] We note a divergence from Steup (2019), who defends an analogous view: (roughly) that one attains knowledge of the *reliability* of the source of one's knowledge of *P as one comes to know P;* but we shy away from this view, since it is one's *knowing P,* not *P* itself, which entails the reliability the source. Nevertheless, he may count as a (C)-lister by our characterisation.

[12] It should be clear from these remarks that we definitely are not denying that one may *extend* one's knowledge by way of competent deduction in other sorts of cases.

[13] But, to give a heads up, we have in mind something along the lines of accounts proposed by Roush (2005, 2012); Ramachandran (2015).

knower is not in a position to know, we reckon the knower herself will be aware of some such condition. For example, we took it for granted that Meena herself would realize that her knowing [*P*], that Dilip is a physicist, in the LIAR scenario *depends on* [¬*L*], Dilip's not being a convincing complulsive liar.

Secondly, we agree with Williams's view that the ramification of there being reliability conditions which the knower is not a position to know, given externalism, is that "knowing *does not guarantee* knowing that one knows" (Williams, 1991, p. 348, our emphasis); but, again, we would go further: *if* our BK-argument is sound, the existence of such conditions guarantees that one never (at any rate, *hardly ever*) knows that one knows—as far as 'wordly' knowledge is concerned, at least.

We have presented a (neglected) case for widespread KK-failure from the perspective of a certain (neglected) externalist position. We submit that there is, at any rate, a case to answer. We close with the following question: *does it actually matter if we generally don't know that we know?* We are not sure what of significance is lost if the present case is accepted.

REFERENCES

Alspector-Kelly, M. (2019). *Against Knowledge Closure.* Cambridge University Press.

Beebee, H. (2001). Transfer of warrant, begging the question, and semantic externalism. *Philosophical Quarterly,* 51(204), 356–74.

Cohen, S. (2002). Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research,* 65(2), 309–329.

Davies, M. (2004). Epistemic entitlement, warrant transmission and easy knowledge. *Aristotelian Society Supplementary Volume,* 78(1), 213–245.

Dretske, F. (1970). Epistemic operators. *Journal of Philosophy,* 67(24), 1007–1023.

Dretske, F. (1971). Conclusive reasons. *Australasian Journal of Philosophy,* 49(1), 1–22.

Dretske, F. (2004). Externalism and modest contextualism. *Erkenntnis,* 61(2-3), 173–186.

Hawthorne, J. (2013). The case for closure. In Steup, M. and Turri, J. (Eds.), *Contemporary Debates in Epistemology.* Wiley-Blackwell, 26-43.

Kelp, C. (2021). *Inquiry, Knowledge, and Understanding.* Oxford University Press.

Kripke, S. A. (2011). Nozick on knowledge. In Kripke, S. A. (Ed.), *Philosophical Troubles: Collected Papers,* Volume 1. Oxford University Press, 162-224.

Markie, P. J. (2005). Easy knowledge. *Philosophy and Phenomenological Research,* 70(2), 406–416.

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

Ramachandran, M. (2015). Knowing by way of tracking and epistemic closure. *Analysis,* 75(2), 217–223.

Roush, S. (2005). *Tracking Truth: Knowledge, Evidence, and Science.* Oxford University Press.

Roush, S. (2012). Sensitivity and closure. In Becker, K. and Black, T. (Eds.), *The Sensitivity Principle in Epistemology.* Cambridge University Press, 242-268.

Steup, M. (2019). Easy knowledge, circularity, and the puzzle of reliability knowledge. *Episteme,* 16(4), 453–473.

Williams, M. (1991). *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism.* Blackwell.

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

Williamson, T. (2005). Contextualism, subject-sensitive invariantism, and knowledge of knowledge. *Philosophical Quarterly,* 55(219), 213–235.

Wright, C. (2004). Warrant for nothing (and foundations for free)? *Aristotelian Society Supplementary Volume,* 78(1), 167–212.

Murali Ramachandran is Associate Professor of Philosophy at TSLAS, Thapar University. He specializes in epistemology, philosophical logic and metaphysics, and has published on counterpart theory, the theory of descriptions, counterfactual accounts of causation, paradoxes, and Williamson's anti-luminosity arguments.

Address: Thapar School of Liberal Arts and Science, Thapar University, Bhadson Road, Patiala, Punjab 147004, India. E-mail: murali.ramachandran@thapar.edu – ORCID ID: https://orcid.org/0000-0001-5611-518X