

# UC Irvine

## UC Irvine Previously Published Works

### Title

In Defence of Closure

### Permalink

<https://escholarship.org/uc/item/8vt0z6rd>

### ISBN

978-0-367-61231-3

### Author

Pritchard, Duncan

### Publication Date

2022-11-16

### DOI

10.4324/9781003104766-8

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NoDerivatives License, available at <https://creativecommons.org/licenses/by-nd/4.0/>

Peer reviewed

## IN DEFENCE OF CLOSURE

DUNCAN PRITCHARD

*University of California, Irvine*

[dhpritch@uci.edu](mailto:dhpritch@uci.edu)

ABSTRACT. A defence is mounted of the closure principle for knowledge. It is first argued that this principle needs to be specifically understood in a diachronic fashion as a competent deduction principle in order for it to represent an intuitive principle that we would want to endorse. Next, it is claimed that in evaluating putative counterexamples to closure we need to differentiate between cases that employ local error-possibilities and cases that employ global error-possibilities. As regards the former, two strategies for resisting the denial of closure are expounded—in terms of background knowledge and factive reasons, respectively—both of which appeal to a distinction between favoring and discriminating epistemic support. Similarly, as regards the latter, two strategies for resisting the denial of closure are expounded. The first involves maintaining that one can know the denials of radical sceptical scenarios by appeal to factive reasons. The second—which I favor—involves appealing to the Wittgensteinian notion of a hinge commitment. As explained, such commitments are not in the market for knowledge, but neither are they the kinds of commitments that could feature in a closure-based inference. Accordingly, rather than constituting counterexamples to closure, they are instead simply inapplicable to closure-based inferences.

1. The closure principle has been formulated in a range of different ways since it began to be discussed in the 1970s. The initial formulation, due to Fred Dretske (1970), was essentially concerned with the idea that knowledge is closed under entailments, such that one can know what is entailed by what one knows, regardless of whether one's knows or even believes the entailment in question.<sup>1</sup> Dretske famously argued that such a principle did not hold universally.

In particular, he held that closure fails because 'knows that  $\varphi$ ' is not a 'fully penetrating' operator, like 'it is necessary that  $\varphi$ ', which 'penetrates' through to all entailments. Interestingly, however, Dretske argued that it was also not a 'non-penetrating' operator either, like 'it is

surprising that  $\varphi$ , which offers no guarantee of penetration to entailments. Instead, he claimed that ‘knows that  $\varphi$ ’ is a ‘semi-penetrating operator’, such that it ‘penetrates’ through some entailments but not others in a principled fashion. Dretske’s concern was to specify what this principled basis was that prevented certain kinds of operator penetrations to occur.

The proposal that Dretske offers is now very familiar, and can be found in some form in the work of several figures who reject closure in this sense. In broad terms, the idea is that some of our knowledge occurs, entirely legitimately (i.e., not in a way that undermines its epistemic standing), against the background of certain presuppositions holding. It follows that one cannot use this knowledge in order to come to know these presuppositions, on pain of a kind of bootstrapping.<sup>2</sup> Closure-style inferences are thus usually fine—and hence the ‘knows that  $\varphi$ ’ operator penetrates through to entailments—so long as it doesn’t concern inferences to these background presuppositions.

Before we can meaningfully engage with the question of whether the closure principle should be restricted, however, we first need to identify which formulation is at issue. Dretske (2005*a*) himself in more recent work focusses on a reading of the closure principle whereby the entailment is known.<sup>3</sup> This certainly makes closure much more compelling, as I take it there isn’t much of a temptation to regard unknown consequences of our knowledge as being in the market for knowledge purely in virtue of the fact that they are entailed by one’s knowledge. Aside from anything else, there is nothing to ensure that one even believes such consequences, and yet belief is usually thought to be necessary for knowledge.<sup>4</sup>

While the idea that knowledge is closed under known entailment is more plausible, however, it is still prone to obvious counterexamples. In particular, the subject’s belief in the entailed proposition might have nothing whatsoever to do with her knowledge of the antecedent proposition and the relevant entailment. In such a case, why should her knowledge of the entailing proposition and the entailment have any bearing on whether she knows the entailed proposition?

This issue highlights an important point about closure principles, which is that they are trying to capture a specific kind of belief-forming process. In particular, what we are interested in is whether knowledge can be extended via one’s knowledge of entailments to knowledge of the entailed propositions, and that means that we are thinking of the subject’s belief in the entailed proposition as being based on this inference. Since we are now describing an inferential process, this means that we need a diachronic formulation of the closure principle, rather than the kind of synchronic formulation that Dretske was considering. We thus get the following competent deduction articulation of closure that is now standard in the literature:

If  $S$  knows that  $p$ , and  $S$  competently deduces from  $p$  that  $q$ , thereby forming a belief that  $q$  on this basis while retaining her knowledge that  $p$ , then  $S$  knows that  $q$ .<sup>5</sup>

Henceforth, by ‘closure’ we will have in mind this specific formulation.

A few remarks on this formulation are in order. First, the notion of competent deduction in play ensures that the subject isn’t simply making a mistake in undertaking the relevant inference, and hence it is meant to presuppose that the target entailment both holds and is known to hold by the subject. Second, it is here stipulated that the belief in the entailed proposition is based upon the competent deduction from the subject’s knowledge of the entailing proposition, which thereby excludes the issue we just raised where the basis for belief has nothing to do with the subject’s knowledge of the entailing proposition and the entailment. Third, since inferences take place over time, we also need to stipulate that the subject retains her knowledge of the entailing proposition throughout, in order to exclude cases where the subject loses this knowledge during the inference (perhaps because of the inference, in that the inference itself prompts her to doubt that she knows the entailing proposition). Finally, fourth, I want to highlight that the notion of knowledge that is in play here is rationally grounded knowledge. Depending on one’s wider epistemology, rationally grounded knowledge might well be the only kind of knowledge that there is.<sup>6</sup> But we do not need to take sides in that issue. The crux of the matter is that any principle that is concerned with such reflective rational processes as competent deductions is clearly directed at the kind of knowledge that figures in the space of reasons.

So construed, closure certainly seems highly intuitive. Given that competent deduction is a paradigmatically rational process, it is hard to see how employing such a process to extract a belief from what one already knows wouldn’t thereby ensure that this belief is also known. We thus have a formulation of closure which is such that it would be highly surprising to discover that it didn’t universally hold. With this in mind, I want to consider two kinds of counterexample that are typically posed against closure. Both concern error-possibilities the falsity of which is entailed by putatively known propositions, but whereas the first kind of error-possibility is local, the second kind is radical (and so of relevance to the problem of radical scepticism). As we will see, it is important in the debate regarding the status of closure to keep local and global error-possibilities apart.

2. Consider Dretske’s (1970) famous counterexample to closure involving zebras and cleverly disguised mules. Our agent, let’s call her ‘Zula’, is looking into a clearly marked zebra enclosure at a regular zoo, and can see the zebra-like creature inside. Does Zula know [Z]: that the creature she is looking at in the enclosure is a zebra? So long as circumstances are as described, and there is no epistemic ‘funny business’ going on, then I think we would treat Zula as knowing this proposition.

Note, however, that [Z] entails [not-CDM]: that the creature Zula is looking at in the enclosure is not a cleverly disguised mule, and presumably Zula will also know that this entailment holds. Accordingly, it seems that Zula can come to know that [not-CDM] via closure by competently deducing it from her knowledge that [Z].

The difficulty, however, is that it doesn't seem that Zula has an adequate epistemic basis for knowing [not-CDM]. In particular, we can stipulate that Zula doesn't have any special expertise that might be relevant in this regard, such as zoological training. Moreover, she hasn't made any special checks of the creature either, such as closely inspecting it at close range. So how then could she possibly be in a position to know [not-CDM]? In particular, it seems that Zula isn't able to discriminate between a zebra and a cleverly disguised mule. And yet it appears to follow from closure that her knowledge of [Z] enables her to come to know, via a competent deduction, [not-CDM]. But how can it be that she can know that it is a zebra and not a cleverly disguised mule if she can't discriminate between the two?

3. Such cases naturally invite the following *relevant alternatives* line. In order to have knowledge, it is not required that one is able to rule out all possible error-possibilities, as that would be to set an unduly austere, infallibilist, requirement on knowledge.<sup>7</sup> All that is required is that one is able to rule out the relevant ones, in some sense of 'relevance' to be determined. The irrelevant alternatives, in contrast, can be reasonably presupposed as false, and thus ignored. Intuitively, the CDM alternative is an irrelevant alternative when it comes to Zula's knowledge (in these conditions) of [Z], which is why she can know the latter even while not being in a position to exclude the former. This is thus why an unrestricted form of closure fails, in that if it were allowed, then one would be able to use closure to come to know an irrelevant alternative the falsity of which is presupposed in, and entailed by, one's knowledge.<sup>8</sup>

There is also a variation on the relevant alternatives line that deserves note, which involves saving closure by going *contextualist*. In particular, given that we are working with a diachronic formulation of closure, it is possible to hold that the very consideration of an irrelevant error-possibility can thereby change the context and thus make that alternative relevant. On this view, Zula can know [Z] even though she doesn't know [not-CDM] just so long as she doesn't entertain the latter. But once she starts to consider the entailment at issue, and hence the CDM alternative that it concerns, then this possibility now becomes relevant and her inability to exclude it ensures that she no longer knows the entailing proposition, [Z]. There is thus no single context in which Zula knows both [Z] and [not-CDM], and hence there is no need (on this score at least) to deny closure.<sup>9</sup>

4. I think we can respond to this kind of case without either denying closure or being obliged to embrace contextualism. The problem arises from the fact that Zula is unable, in her present situation, to perceptually discriminate between a zebra and a cleverly disguised mule. This is why it is puzzling that Zula seems to be in a position, via closure, to know both [Z] and [not-CDM], as this implies that she can tell the difference between the two scenarios. I think what is key here is to recognise that there is a way of knowing the difference which doesn't imply that one can discriminate between the two scenarios at issue. If that's right, then there needn't be anything puzzling about Zula's ability to know both [Z] and [not-CDM], even though she cannot, *ex hypothesi*, perceptually discriminate between them.

Consider the distinction between *favouring* as opposed to *discriminating* epistemic support. What Zula lacks is the discriminative ability to perceptually distinguish between the possibility that the creature before her is a zebra rather than a cleverly disguised mule. Someone with different perceptual abilities, like a zoologist, might be able to undertake such a discrimination, as might someone who has been able to make special checks of the creature. But neither applies to Zula. This needn't mean that Zula is unable to know that the creature is a zebra rather than a cleverly disguised mule, however; it just means that if she does know this it can't be via her perceptual ability to discriminate between the two scenarios.

It is possible to know that one scenario obtains rather than another, even though one cannot discriminate between them, via the possession of favouring epistemic support. In particular, if one has independent reasons for treating the alternative as being highly unlikely, then that will ordinarily suffice to enable one to know that the target scenario has obtained rather than the alternative. One can discount all kinds of implausible alternatives in this way, and thereby come to know that they don't obtain. I know that the person at the mall who looks a bit like Kanye West is not really Kanye West because of a wide range of evidence that I possess which favours this believed scenario over the alternative. Why would Kanye be hanging out at the local mall? Would Kanye be dressed like that? Isn't the real Kanye much shorter than this individual? And so on. This favouring evidence enables me to know that the person is a Kanye doppelgänger and not the real deal, even though I may be unable to perceptually discriminate between a good Kanye look-alike and the real thing.

What goes for Kanye West and his look-alike in this scenario can equally apply in the Zula case. If Zula is an ordinary epistemic agent, then we would expect her to have all kinds of reasons at her disposal that would bear on the plausibility of the CDM hypothesis. She knows that this is a reputable zoo, and that reputable zoos don't deceive their clientele. She knows that a deception of this sort would be hard to pull off, and expensive too, which makes it unlikely that anyone would attempt it. She knows that such a deception would be easily spotted by someone, even if it would

trick her, and that there would be penalties for the zoo as a result. And so on. She thus has excellent rational support that favours her belief that the creature is a zebra over the CDM alternative. And if that's right, then she can know both [Z] and [not-CDM] even though she can't perceptually discriminate between these two scenarios.

5. One might object to the foregoing by claiming that it puts the threshold for perceptual knowledge implausibly high. Is it really required for perceptual knowledge that one has the kind of background reasons just noted? But notice that, on closer inspection, this is not a very demanding condition at all.

To begin with, it doesn't follow from the foregoing that in order for Zula to know [Z] she must first assemble a rational basis for excluding such scenarios as CDM. Indeed, there is no reason why she would have even considered such a scenario in normal circumstances. The claim is rather that it is only once Zula attempts to undertake the target competent deduction that such favouring epistemic support becomes relevant. Prior to this error-possibility coming into play, Zula can know [Z] even in the absence of such support, such as by possessing the discriminative capacities that would enable her to tell apart the scenario depicted in [Z] from relevant alternative scenarios (that the creature before her is not a horse, say). In considering the relevant entailment, however, Zula is thereby obliged—as any agent would be when considering an inference of this kind—to consider whether she has grounds to dismiss this alternative, or whether she should instead reconsider her commitment to the entailing belief.<sup>10</sup> This is the point that she should reflect on her background reasons and marshal the supporting favouring epistemic support.

For any ordinary agent, however, such background reasons will be readily available, and hence they ought to be similarly available to Zula. That is, requiring that Zula has these background reasons in these conditions is not an austere epistemic requirement to impose. In contrast, if Zula really is the kind of agent for whom such background reasons are not available—i.e., such that they cannot be summoned by reflection if required, without any further empirical investigation being needed—then it is not so clear that she has knowledge of the entailing proposition in the first place. At the very least, if she knows this proposition, then presumably she does so purely in virtue of exercising some sort of perceptual capacity, and not because her knowledge enjoys rational support.<sup>11</sup> As we noted above, however, our concern in these deliberations over closure is precisely with rationally grounded knowledge.

6. Once the distinction between favouring and discriminating epistemic support has been made explicit, then a further dialectical option becomes apparent. The kind of rational support that one

gains for one's belief via perception is standardly thought to be non-factive, such as considerations about the way things appear (which obviously don't entail how things are). There are epistemological proposals that reject this way of thinking about the reasons that ground perceptual knowledge, however, and maintain instead that a subject's rational support for their perceptual knowledge in epistemically paradigmatic conditions can be factive. In particular, *epistemological disjunctivism* argues that in epistemically paradigmatic conditions one can have perceptual knowledge of a proposition,  $p$ , in virtue of being in possession of the factive reason that one *sees that  $p$*  (where seeing that  $p$  entails  $p$ ).<sup>12</sup>

It would take me too far afield to attempt to defend epistemological disjunctivism here, but I do want to note how treating it as a theoretical possibility bears on Zula-style cases. Epistemological disjunctivism doesn't deny that Zula is unable to perceptually discriminate between a zebra and a cleverly disguised mule in the conditions that she is in. Factive reasons, however, provide a decisive kind of favouring epistemic support for one's belief, in that they entail the falsity of alternative scenarios. In order to be the kind of epistemic agent that would be credited with factive rational support Zula would also be likely to possess the kind of background reasons noted above, but crucially she wouldn't require this background support in order to be able to know both [Z] and [not-CDM], as the factive reason would suffice by itself. There would thus be no impetus in this case to deny closure as opposed to embracing Zula's knowledge of [not-CDM].<sup>13</sup>

7. The general line of reasoning that we have applied to the Zula case can be similarly applied, *mutatis mutandis*, to a range of other putative counterexamples to closure. In particular, it is generally applicable to any such case that essentially appeals, as the Zula case does, to a local error-possibility.<sup>14</sup> More specifically, this response is applicable precisely because the error-possibility at issue, being local, doesn't immediately call into the question the favouring epistemic support that is available to the protagonist.

This point straightforwardly applies when it comes to the background reasons that we noted above. Even if Zula has been given reasons for treating the CDM alternative seriously (i.e., it has been rationally motivated), as opposed to merely having this alternative made explicit to her, it is still possible for Zula to legitimately appeal to background reasons in order to rationally dismiss it, as this background knowledge is not called into question by the CDM error-possibility.

Matters become more complicated when it comes to the factive reasons defended by epistemological disjunctivism. This is because such factive reasons are only available to the subject in epistemologically paradigmatic conditions. Accordingly, it makes a difference whether the CDM error-possibility has been merely raised or rationally motivated, since in the latter case the subject



is no longer in the right conditions in order to possess factive reasons.<sup>15</sup> Nonetheless, so long as this error-possibility is merely raised, then the factive reasons are available to enable Zula to discount this alternative. Moreover, even if it is rationally motivated, it is still possible for Zula to appeal to her background reasons in the way just described in order to dismiss it.

The upshot is that we can accommodate the intuitions in play in Zula-style cases involving local error-possibilities without denying such an intuitive principle as closure (or else going contextualist). So what then, if anything, does provide the impetus for rejecting this principle? I think that for this we need to focus not on cases involving local error-possibilities but on global error-possibilities, since on the face of it at least they call into question both one's background reasons and also that one is in the kind of epistemic conditions that could possibly support factive reasons. In particular, the real case for denying closure concerns the kind of error-possibilities characteristically appealed to in the context of radical scepticism.

8. Let's follow convention and focus on the radical sceptical hypothesis that one is a disembodied brain-in-a-vat being 'fed' deceptive experiences [BIV], such that most of one's beliefs are false. Let's now imagine a subject who is in entirely normal epistemic conditions, and so is not a victim of the BIV hypothesis, nor is such a scenario even modally close. Our subject believes something mundane that entails [not-BIV], such as that she is presently sitting down [E]. We would naturally attribute knowledge to our subject of [E]. And yet, if she does know [E], then since this obviously entails [not-BIV], it follows that our subject can reason, via closure, to knowledge of [not-BIV].

Crucially, however, it is widely accepted that it is at least problematic to treat subjects as having knowledge of the denials of radical sceptical hypotheses. How could our subject come to know [not-BIV]? Given that BIV depicts a scenario that is *ex hypothesi* indistinguishable from the subject's ordinary experiences, what would she possibly appeal to in order to exclude it? In particular, the subject lacks a perceptual capacity to discriminate between everyday circumstances and the BIV alternative. Moreover, given the radical nature of the BIV hypothesis, it also calls the subject's background reasons into question too, and hence it does not seem appropriate for the subject to appeal to favouring epistemic support via this route. So how then can it be possible for our subject to know both [E] and [not-BIV], given that she has no way of distinguishing between the two scenarios at issue, either directly via a perceptual capacity or indirectly via appeal to favouring epistemic support? Radical sceptical arguments of this kind thus pose a challenge for closure, of a kind that is distinct from the difficulty that is presented by Zula-style cases involving local error-possibilities.

9. So how might we deal with this challenge to closure? I want to explore two options. The first is

to revisit the idea of favouring epistemic support that we saw could be effective in dealing with Zula-style cases. This avenue might seem unpromising given the point we just made about how radical sceptical error-possibilities call into question one's background reasons. Recall, however, that there was a second way of appealing to favouring epistemic support available, which concerned the factive reasons defended by epistemological disjunctivism. According to this proposal, if one is in epistemically paradigm conditions, then one can have available to one factive reasons in support of one's perceptual knowledge. As we noted above, such factive reasons would provide one with a decisive form of favouring epistemic support to enable one to know the entailed proposition even when one lacks the relevant discriminative powers to exclude the target error-possibility. Could this strategy be applied to the radical sceptical case?

One might initially find such an approach suspect. If the radical sceptical error-possibilities call into question one's background reasons, then why don't they also undermine one's factive reasons? In particular, why don't they entail that one is not in the kind of epistemically paradigm conditions that are the prerequisite of being in possession of factive reasons?

There is, however, an important factor that is relevant here, which is whether the target error-possibility has been rationally motivated, or whether it is merely being considered by the subject. We noted above that different epistemic burdens are imposed on the subject in the different cases, especially with regard to factive reasons. In particular, if Zula is provided with good reasons for thinking that there is the kind of deception taking place that would motivate the CDM alternative, then it's unlikely that she is in the kind of epistemically paradigm conditions that would be needed for her perceptual belief to enjoy factive rational support. If that's right, then wouldn't that point also neutralise the appeal to factive reasons in response to radical scepticism?

Notice, however, that radical sceptical error-possibilities are precisely not rationally motivated. For one thing, it is hard to see what possible motivation could be given. What could possibly function as a reason for thinking that one's beliefs are radically in error, such that it wouldn't itself be called into question by the very error-possibility at issue? In any case, it is certainly not part of the employment of radical sceptical hypotheses as part of the radical sceptical puzzle to rationally motivate them. This would bring with it empirical commitments that the radical sceptic would not want to be saddled with. Moreover, radical scepticism of this kind, at least when presented in its strongest form, is a putative *paradox*, in that it is meant to be exposing deep tensions within our epistemological concepts.<sup>16</sup> In presenting a paradox, however, one is meant to draw only on commitments that we all (supposedly) share, and not advance controversial claims, much less controversial empirical claims, as that would undermine the thesis that what is under consideration is a genuine paradox. That's why there is no contention as part of the radical sceptical paradox that we have any reason for thinking that we are the victim of a radical sceptical

hypothesis. The thesis is just that such a scenario is possible and that we cannot exclude it, but that our failure to exclude it conflicts with our everyday knowledge (where our commitment to closure is playing a large part in setting up that conflict).

If radical sceptical hypotheses are not rationally motivated, however, but merely raised, then that alters the dialectical situation somewhat. In particular, there is no basis for treating the subject as not being in epistemically paradigm conditions, such that factive reasons are unavailable to her. Moreover, notice that the factive reasons advocated by epistemological disjunctivism are different to the sort of background reasons that might be appealed to in the context of radical scepticism. In the BIV case, for example, the latter might consist of considerations concerning the state of current technology that run counter to the BIV hypothesis. But what real purchase do these considerations really have on the plausibility of the BIV scenario? After all, these are considerations that are entirely compatible with the truth of the BIV hypothesis.

The factive reasons that epistemological disjunctivism advances are very different in this regard. For what epistemological disjunctivism is precisely challenging is a key presupposition of the radical sceptical paradox to the effect that the rational support one's beliefs enjoy, even in epistemically paradigm conditions, must always be non-factive. That is, radical scepticism is trading on the fact that even in ideal epistemic conditions the rational support one's beliefs enjoy is compatible with those beliefs being massively false. If that's right, then even if the radical sceptical hypotheses are not only false but also modally far-fetched, and even if no rational basis has been offered (or could be offered) for taking them seriously, it would still remain that one lacks a rational basis for excluding them. In contrast, however, if one grants that one's perceptual knowledge can be supported by factive reasons when one is in ideal epistemic conditions, and radical scepticism has offered no rational basis for thinking that one is not in such conditions, then it is open to one to appeal to one's factive reasons to exclude radical sceptical scenarios. In particular, notice that one is now in a much better dialectical position than one would be if one could only appeal to background reasons. For unlike one's background reasons, the factive reasons that support one's perceptual knowledge actually entail the falsity of radical sceptical hypotheses.<sup>17</sup>

The upshot is that epistemological disjunctivism offers one way of dealing with the challenge to closure presented by the radical sceptical paradox. Rather than rejecting closure, one can instead embrace the idea that agents can have knowledge of the entailed proposition. This is not in virtue of having some mysterious perceptual capacity that enables them to discriminate between everyday life and the sceptical alternative, but rather due to the possession of favouring epistemic support. Moreover, unlike one's background reasons, this favouring epistemic support does speak to the radical sceptical hypothesis at issue, by entailing its falsity (as opposed to being entirely compatible with its truth). Whether one finds such a strategy plausible will, of course,

depend on whether one finds epistemological disjunctivism credible (and we haven't argued for this thesis here). Nonetheless, the point is that on this epistemological proposal there is no special challenge raised for closure in the context of radical scepticism.<sup>18</sup>

10. While I hold that epistemological disjunctivism has an important part to play in responding to the problem of radical scepticism, ultimately I don't think that it offers the right response to closure-based radical scepticism.<sup>19</sup> The reason for this is that I'm independently convinced by Wittgenstein's account of hinge commitments. In essence, this is the idea that it is in the nature of our system of rational evaluation that it presupposes that one has certain basic 'hinge' commitments which are themselves immune to rational evaluation. It follows that one does not have specific reasons for regarding one's hinge commitments as true. Moreover, these hinge commitments are not like hypotheses or assumptions, which one might endorse without having any commitment to their truth, but are rather claims that one is optimally certain of. Indeed, far from being concerned with theoretical theses, they instead primarily consist of entirely mundane claims, such as that one has two hands or that one's name is such-and-such. One's visceral certainty (in normal conditions) in these claims is manifest in one's actions, such that genuine (as opposed to merely performative) doubt of them is simply impossible.

There are various accounts of the Wittgensteinian notion of a hinge commitment available in the literature, though I take the former description to be broadly common ground among most of these views. Where they tend to depart from each other is in terms of what they add to this basic account.<sup>20</sup> On one prominent proposal, for example, our hinge commitments can be in the market for knowledge even despite the hinge role that they play in our rational evaluations. For although we cannot have any rational basis for thinking that they are true, there is nonetheless a kind of *strategic* rational basis for their truth—an epistemic entitlement, as some commentators put it—that accrues to them in virtue of how they are constitutive of our system of epistemic rationality. As we might put the point, if in order to be an epistemically rational subject at all one is required to have hinge commitments, then it must be epistemically rational to believe one's hinge commitments even though, by definition, one cannot have a reason for thinking that these commitments are true. Accordingly, this opens up a way for them to count as knowledge, even though one has no reason for regarding them as true.<sup>21</sup>

11. I don't find such an approach plausible, however. There are various reasons for this, but I want to focus on one core concern, which is that this way of thinking about hinge commitments doesn't take on board the nature of the propositional attitude in play. Recall that our hinge commitments exhibit an optimal certainty that is manifest in our actions, even while being

divorced from any rational basis for their truth. Wittgenstein characterises this arational certainty as being visceral, ‘animal’; it is simply there ‘like our life’.<sup>22</sup> Wittgenstein thus regards our hinge commitments as involving a distinctive kind of propositional attitude.<sup>23</sup> Given how broad our folk notion of belief is—which takes in a wide range of propositional attitudes, including acceptances and faith-like commitments, as well as the kind of propositional attitude that is involved in having knowledge—one could truly describe our hinge commitments as beliefs.<sup>24</sup> But to do so glosses over important distinctions that are relevant here.

In particular, our hinge commitments are not assumptions or hypotheses, such that one might be in some sense committed to them (e.g., committed to act as if they are true), without thereby being committed to their truth. For the visceral certainty we have for our hinge commitments ensures that they are incompatible with agnosticism about the truth of the target proposition. More importantly for our purposes, our hinge commitments are also not beliefs in the specific sense of belief that is most relevant to epistemology—*viz.*, that propositional attitude which is a constituent part of rationally grounded knowledge. Belief in this sense—*K-apt belief*, if you will—bears certain basic conceptual connections to reasons and truth that set it apart from other propositional attitudes that might fall under the broad folk conception of belief, including our hinge commitments. Here is one key consideration in this regard, which is that *K-apt* beliefs are such that one cannot retain one’s *K-apt* belief that  $p$  while recognising that one has no rational basis for the truth of  $p$ . If one retains one’s commitment to  $p$  even while recognising this lack of a rational basis, then whatever one’s propositional attitude in this regard, it cannot be a *K-apt* believing, but must instead be some other propositional attitude (such as a wishful thinking that  $p$ , for example).

Crucially, however, one’s hinge commitments precisely fail this test for *K-apt* belief, in that they are all-out commitments to the truth of  $p$  that one retains even when one comes to appreciate that, as hinge commitments, one lacks a rational basis for them.<sup>25</sup> That is, while one might profess doubt of one’s hinge commitments—as part of a philosophical investigation, for example—this will not be genuine doubt of them, as will be apparent from how one’s actions continue to manifest one’s hinge certainty. That our hinge commitments are not *K-apt* beliefs is important, since it highlights that they are simply not in the market for knowledge, as our commitment to them is not the kind of propositional attitude that could function as a constituent part of rationally grounded knowledge.<sup>26</sup> It follows that attempts to show that they can be known via appeal to broadly strategic epistemic considerations are not going to work. For even if it is true that there is a positive epistemic story to tell about our hinge commitments on this front, the fact remains that our hinge commitments are simply not the kind of propositional attitude that could amount to

knowledge.

12. This issue is relevant to our present concerns because of how our commitment to the denials of radical sceptical hypotheses seem to be hinge commitments. That is, the general form of a hinge commitment is that it is the kind of consideration that would reveal wholesale and fundamental error on one's part, such as being wrong about whether one has hands in normal conditions (i.e., as opposed to the abnormal conditions of waking up after a serious car accident).<sup>27</sup> If that's right, then it follows that one is hinge committed to the falsity of radical sceptical hypotheses, as they directly concern the possibility of wholesale and fundamental error. Accordingly, if such propositions are not in the market for knowledge, then that seems to pose a challenge for the closure principle.

One response to this puzzle might be to simply accept that a Wittgensteinian hinge epistemology entails the denial of closure, and some commentators have embraced that line.<sup>28</sup> Another alternative is to contend that one can have knowledge of one's hinge commitments and thus that one doesn't need to reject closure. Given the foregoing, one might think that I am embracing the first option over the second, but in fact I think that this is a false dilemma, in that our inability to know our hinge commitments is consistent with the closure principle.

In order to see why this is the case, we need to remember how we formulated that principle above. Recall that it was argued that if we are to regard closure as a compelling principle that isn't susceptible to straightforward counterexamples, then it is crucial that it is understood diachronically as a competent deduction principle. In particular, via closure-based inferences one thereby comes to acquire a belief in the entailed proposition that is based on the competent deduction. Since the idea is that closure delivers rationally grounded knowledge, we are clearly concerned with the acquisition of K-apt belief in the entailed proposition. We have seen, however, that one simply cannot have a K-apt belief in one's hinge commitments, as the visceral certainty involved in the latter excludes this possibility.<sup>29</sup> Accordingly, there is no need to deny the closure principle in order to accommodate the unknowability of our hinge commitments, since there cannot be a closure-based inference involving our hinge commitments that generates the target counterexample. Rather than there being exceptions to the closure principle, it is instead simply inapplicable to our hinge commitments.

13. The challenge posed by local and global error-possibilities to the closure principle is thus found to be illusory. To begin with, we have seen the importance of properly characterising what this principle amounts to, and thus what we are meaning to defend, as not every formulation of this principle is even *prima facie* plausible. Next, we need to distinguish between local and global

error-possibilities, as we have seen that they pose distinct issues for the closure principle, and hence the response we should mount is similarly different. In the case of local error-possibilities, we saw how favouring epistemic support can enable a subject to have knowledge of the entailed proposition in the relevant cases, either by appealing to background knowledge or by appealing to the factive reasons advocated by epistemological disjunctivism. In the case of the global error-possibilities raised as part of the problem of radical scepticism, the background knowledge is no longer applicable because it is called into question by the error-possibility in play (in contrast to scenarios involving local error-possibilities). Nonetheless, there are still ways of responding to this version of the challenge. On the one hand, one can directly appeal to one's factive reasons in order to maintain that one can have knowledge of the entailed proposition. On the other hand, one can opt for my favoured route of maintaining that one is hinge committed to the denials of radical sceptical hypotheses, and that such commitments are simply not in the market for knowledge. More specifically, one is unable to have a K-apt belief in one's hinge commitments of the kind that is relevant to a closure-based inference. Accordingly, far from closure failing in this case, closure is simply inapplicable. The upshot is that regardless of whether the error-possibility is local or global, there is no challenge to the closure principle, properly understood.<sup>30</sup>

## REFERENCES

- Cohen, S. (1984). 'Justification and Truth', *Philosophical Studies* 46, 279-96.
- (1988). 'How to be a Fallibilist', *Philosophical Perspectives* 2, 91-123.
- (1999). 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives* 13, 57-89.
- (2000). 'Contextualism and Skepticism', *Philosophical Issues* 10, 94-107.
- (2002). 'Basic Knowledge and the Problem of Easy Knowledge', *Philosophy and Phenomenological Research* 65, 309-29.
- Coliva, A. (2010). *Moore and Wittgenstein: Scepticism, Certainty, and Common Sense*, London: Palgrave Macmillan.
- (2015). *Extended Rationality: A Hinge Epistemology*, London: Palgrave Macmillan.
- DeRose, K. (1995). 'Solving the Skeptical Problem', *Philosophical Review* 104, 1-52.
- Dretske, F. (1970). 'Epistemic Operators', *Journal of Philosophy* 67, 1007-23.
- (1981). 'The Pragmatic Dimension of Knowledge', *Philosophical Studies* 40, 363-78.
- (2005a). 'The Case Against Closure', *Contemporary Debates in Epistemology*, (eds.) E. Sosa & M. Steup, 13-26, Oxford: Blackwell.
- (2005b). 'Reply to Hawthorne', *Contemporary Debates in Epistemology*, (eds.) E. Sosa & M. Steup, 43-6, Oxford: Blackwell.
- Hawthorne, J. (2005). 'The Case for Closure', *Contemporary Debates in Epistemology*, (eds.) E. Sosa & M. Steup, 26-43, Oxford: Blackwell.
- Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy* 74, 549-67.
- McDowell, J. (1995). 'Knowledge and the Internal', *Philosophy and Phenomenological Research* 55, 877-93.
- McGinn, M. (1989). *Sense and Certainty: A Dissolution of Scepticism*, Oxford: Blackwell.
- Moyal-Sharrock, D. (2004). *Understanding Wittgenstein's On Certainty*, London: Palgrave Macmillan.
- Neta, R., & Pritchard, D. H. (2007). 'McDowell and the New Evil Genius', *Philosophy and Phenomenological Research* 74, 381-96.
- Nozick, R. (1981). *Philosophical Explanations*, Oxford: Oxford University Press.
- Pritchard, D. H. (2007). 'How to be a Neo-Moorean', *Internalism and Externalism in Semantics and Epistemology*, (ed.) S. Goldberg, 68-99, Oxford: Oxford University Press.
- (2008). 'McDowellian Neo-Mooreanism', *Disjunctivism: Perception, Action, Knowledge*, (eds.) A. Haddock & F. Macpherson, 283-310, Oxford: Oxford University Press.
- (2010). 'Relevant Alternatives, Perceptual Knowledge and Discrimination?', *Noûs* 44, 245-68.
- (2011a). 'Epistemological Disjunctivism and the Basis Problem', *Philosophical Issues* 21, 434-55.
- (2011b). 'Wittgenstein on Scepticism', *The Oxford Handbook on Wittgenstein*, (eds.) O. Kuusela & M. McGinn, 521-47, Oxford: Oxford University Press.
- (2012). *Epistemological Disjunctivism*, Oxford: Oxford University Press.
- (2014). 'Sceptical Intuitions', *Intuitions*, (eds.) D. Rowbottom & T. Booth, 213-31, Oxford: Oxford University Press.
- (2015). *Epistemic Angst: Radical Skepticism and the Groundlessness of Our Believing*, Princeton, NJ, Princeton University Press.
- (2017). 'Wittgenstein on Hinge Commitments and Radical Scepticism in *On Certainty*', *Blackwell Companion to Wittgenstein*, (eds.) H.-J. Glock & J. Hyman, 563-75, Oxford: Blackwell.
- (2020). 'Epistemic Vertigo', *The Philosophy and Psychology of Ambivalence: Being of Two Minds*, (eds.) B. Brogaard & D. Gatzia, ch. 7, London: Routledge.
- (2021). 'Ignorance and Inquiry', *American Philosophical Quarterly* 58, 111-23.



- (Forthcoming). ‘Moderate Knowledge Externalism’, *Externalism About Knowledge*, (ed.) L R. G. Oliveira, Oxford: Oxford University Press.
- Pryor, J. (2000). ‘The Skeptic and the Dogmatist’, *Noûs* 34, 517-49.
- Schönbaumsfeld, G. (2016). *The Illusion of Doubt*, Oxford: Oxford University Press.
- Stevenson, L. (2002). ‘Six Levels of Mentality’, *Philosophical Explorations* 5, 105-24.
- Stine, G. C. (1976). ‘Skepticism, Relevant Alternatives, and Deductive Closure’, *Philosophical Studies* 29, 249-61.
- Stroud, B. (1984). *The Significance of Philosophical Scepticism*, Oxford: Clarendon Press.
- Unger, P. (1971). ‘A Defence of Skepticism’, *Philosophical Review* 80, 198-219.
- (1975). *Ignorance: A Case for Scepticism*, Oxford: Clarendon.
- Vahid, H. (2011). ‘Externalism/Internalism’, *Routledge Companion to Epistemology*, (eds.) S. Bernecker & D. H. Pritchard, 144-55, London: Routledge.
- Williams, M. (1991). *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*, Oxford: Blackwell.
- Williamson, T. (2000). *Knowledge and its Limits*, Oxford: Oxford University Press.
- Wittgenstein, L. (1969). *On Certainty*, (eds.) G. E. M. Anscombe & G. H. von Wright, (tr.) D. Paul & G. E. M. Anscombe, Oxford: Blackwell.
- Wright, C. (2004). ‘Warrant for Nothing (and Foundations for Free)?’, *Proceedings of the Aristotelian Society* (supp. vol.) 78, 167-212.

## NOTES

- <sup>1</sup> Note that in what follows I will be restricting my attention to single premise formulations of closure.
- <sup>2</sup> As Dretske (2005*a*) puts the point in more recent work, our knowledge doesn't penetrate to the 'heavyweight implications' of what is known, where these concern the legitimate presuppositions of our knowledge.
- <sup>3</sup> See also Nozick's (1981) influential denial of the closure principle, which also treats knowledge as being closed under known entailments.
- <sup>4</sup> At the very least, it is usually thought that there is a belief-like propositional attitude that is a constituent of knowledge. (The reason for this *caveat* will become apparent below).
- <sup>5</sup> This is essentially the formulation of the closure principle put forward by Williamson (2000, 117) and Hawthorne (2005, 29).
- <sup>6</sup> That's at least one way of thinking of the epistemic externalism/internalism distinction as it is applied to knowledge, such that epistemic externalism about knowledge maintains that there can be knowledge that is not rationally grounded, whereas epistemic internalism about knowledge denies this. See Pritchard (*forthcoming*) for more on this point. For further discussion of the epistemic externalism/internalism distinction more generally, see Vahid (2011).
- <sup>7</sup> Infallibilism of this kind is usually treated as an implausible position, though for an exception see Unger (1971; 1975).
- <sup>8</sup> Dretske (1970; 1981) himself advanced this relevant alternatives line. For some useful discussions of Dretske in this regard, see Stine (1976) and Cohen (1988).
- <sup>9</sup> Strictly speaking, contextualism of this kind should be expressed metalinguistically, as being concerned with 'knows', but we can set aside this complication for our purposes. For some key defences of contextualism of this variety, see Cohen (1988; 2000; 2003), DeRose (1995), and Lewis (1996).
- <sup>10</sup> Note that I'm not suggesting here that an error-possibility figuring in a competent deduction is the only way to make it relevant in the applicable sense (i.e., such that one needs to exclude it in order to have knowledge of the target proposition), only that this is sufficient to make it relevant. For example, just as there can be normative epistemic defeaters, which one needs to exclude even if one isn't aware of them, so there can be error-possibilities that are relevant in the applicable sense, and so must be excluded, even if one isn't aware of them (on account of the fact that one ought to be aware of them, just as with normative epistemic defeaters). (Note too that I'm also not suggesting that merely being aware of an error-possibility makes it relevant in the applicable sense, though for simplicity in what follows I will proceed as if this does suffice. Given that such a concession can only benefit my opponent, it shouldn't be controversial). For further discussion of these issues, see Pritchard (2010). See also endnote 11.
- <sup>11</sup> This is the line I advance in Pritchard (2010) as part of what I call a 'two-tiered' account of perceptual knowledge. In particular, while one can have knowledge purely in virtue of exercising certain perceptual capacities, once one entertains (indiscriminable) error-possibilities like CDM, and so infers their denial, then favoring reasons are required. Such an account is thus distinct from both *conservatism* and *dogmatism* in this regard, at least as those two proposals are applied to local error-possibilities like CDM (i.e., as opposed to the kind of radical sceptical scenarios we will consider below). It is distinct from the former since although it does demand the possession of epistemic support that bears on the entailed proposition insofar as one undertakes the relevant closure inference, it doesn't deny that one can know the entailing proposition without it. And it is distinct from dogmatism in that it denies that one can know the entailed proposition via competent deduction closure absent the relevant favoring epistemic support. For a key defence of conservatism, see Wright (2004). For a key defence of dogmatism, see Pryor (2000).
- <sup>12</sup> Epistemological disjunctivism is rooted in the work of McDowell (e.g., 1995), but see, especially, Pritchard (2012) for a full development and defence of this proposal.
- <sup>13</sup> Favoring reasons of the more familiar kind would still be required in non-paradigmatic epistemic conditions though, such as where the subject is in a context in which the CDM error-possibility has been rationally motivated (and hence counts as a misleading defeater).
- <sup>14</sup> Such as, for example, the hypothesis that the table one is looking at is not red but rather white and illuminated by red light—see Cohen (2002, 313).
- <sup>15</sup> See Pritchard (2011*a*; 2012, part 1) for further discussion of what such epistemologically paradigm conditions involve, and in particular why they include both objective factors (such as whether the error-possibility at issue is objectively likely to occur) and subjective factors (such as whether the error-possibility at issue has been rationally motivated, even if it is in fact objectively unlikely to occur). Note that it follows from how we have set-up the Zula case that the CDM alternative is objectively unlikely, as that's required to ensure the security of the intuition that she knows [Z].
- <sup>16</sup> I defend the idea that radical scepticism is best understood as a (putative) paradox in more detail in Pritchard (2014; 2015, part 1). See also Stroud (1984) for an influential defence of this idea.
- <sup>17</sup> As I discuss in Pritchard (2012, part 3), there might be dialectical limitations on the actual citing of factive reasons in a debate with the radical sceptic, but the point is that from a purely epistemological point of view there is nothing amiss in treating factive reasons as providing the relevant favoring epistemic support.
- <sup>18</sup> See Pritchard (2012, part 3) for a fuller development of such a 'neo-Moorean' response to radical scepticism. See also Pritchard (2007; 2008).

<sup>19</sup> I think the real anti-sceptical import of epistemological disjunctivism lies in responding to what I call underdetermination-based radical scepticism, which turns on a difference epistemic principle to closure, and generates what I claim is a logically distinct radical sceptical argument. See Pritchard (2015, parts 1 & 3) for the details.

<sup>20</sup> For some of the key discussions of a Wittgensteinian hinge epistemology, see McGinn (1989), Williams (1991), Moyal-Sharrock (2004), Wright (2004), Coliva (2010; 2015), and Schönbaumsfeld (2016). I summarize my own views in this regard in Pritchard (2015, part 2). For an overview of this debate, see Pritchard (2011*b*; 2017).

<sup>21</sup> The primary defence of this line can be found in the work of Wright (e.g., 2004). See also Coliva (2015) for an important development of this general type of proposal.

<sup>22</sup> See Wittgenstein (1969, §359 & §559).

<sup>23</sup> This is why I refer to our hinge *commitments* rather than, as is more common in the literature, describing them as hinge *propositions*. The latter locution prompts us to focus on the particular propositional content that is in play, but that's not what is distinctive about our hinges, which is rather the special kind of propositional attitude that we bear to these contents in appropriate conditions.

<sup>24</sup> For a useful article exploring the different kinds of propositional attitudes that are typically treated as beliefs, see Stevenson (2002).

<sup>25</sup> See Pritchard (2015, part 2) for further discussion of the propositional attitude involved in our hinge commitments. In particular, this attitude is *sui generis*, in that it is distinct from other kinds of propositional attitude in the vicinity of which we are familiar, such as K-apt beliefs, acceptances, hypotheses, assumptions, aliefs, and so on. As I also explain, on my account of hinge commitments there is a sense in which they can be indirectly responsive to rational considerations even if they are not directly responsive, in virtue of how the content of one's hinge commitments is determined by one's set of K-apt beliefs, and these are in principle responsive to rational considerations in the usual way.

<sup>26</sup> It follows that although our hinge commitments are unknown, this does not constitute an epistemic failing on our part, as they are not in the market for knowledge. Our inability to know our hinge commitments is thus akin to our inability to draw a circle-square. Relatedly, one is not ignorant of one's hinge commitments either, even though they are unknown (and not truly believed either, at least in the specific K-apt sense)—see Pritchard (2021).

<sup>27</sup> Indeed, on my account of hinge commitments, there is really only one overarching hinge commitment that we all share—what I refer to as the *über hinge commitment*—which is the primitive certainty that one is not radically and fundamentally in error. All of one's specific hinge commitments are held to be manifestations of the über hinge commitment. See Pritchard (2015, part 2).

<sup>28</sup> See, for example, Coliva (2015).

<sup>29</sup> Relatedly—although I haven't had space to explore this contention here—one cannot acquire one's hinge commitments via rational processes such as competent deduction anyway. See Pritchard (2015, part 2), where I discuss the distinctive way in which one's hinge commitments are acquired.

<sup>30</sup> I am grateful to Yuval Avnur and Crispin Wright for helpful discussion of the issues covered in this chapter.