THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Skepticism and Epistemic Angst, Redux

Citation for published version:

Pritchard, D 2019, 'Skepticism and Epistemic Angst, Redux', *Synthese*. <https://doi.org/10.1007/s11229-019-02504-2>

Digital Object Identifier (DOI):

[10.1007/s11229-019-02504-2](https://doi.org/10.1007/s11229-019-02504-2)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Synthese

Publisher Rights Statement:

This is a post-peer-review, pre-copyedit version of an article published in Synthese. The final authenticated version is available online at: <https://link.springer.com/article/10.1007/s11229-019-02504-2>

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

For a symposium on *Epistemic Angst*, *Synthese* (forthcoming).

SCEPTICISM AND *EPISTEMIC ANGST*, REDUX

DUNCAN PRITCHARD

University of California, Irvine & University of Edinburgh

ABSTRACT. Part one offers a *précis* of my book, *Epistemic Angst* (Princeton UP, 2015), with particular focus on the themes discussed by the participants in this symposium. Part two then examines a number of topics raised in this symposium in light of this *précis*. These include how best to understand the ‘non-belief’ account of hinge epistemology, whether we should think of our hinge commitments as being a kind of procedural knowledge, whether hinge epistemology can be used to deal with underdetermination-based scepticism, what the status of my acceptance of the closure principle amounts to, whether one’s total evidence in fact supports our hinge commitments, and the nature of the kind of reasoning that Wittgenstein employs when advancing a hinge epistemology. Finally, I offer some remarks on the notion of epistemic risk in the context of the sceptical problematic, and show how this has application to legal epistemology.

1. *PRÉCIS OF EPISTEMIC ANGST*

Epistemic Angst: Radical Skepticism and the Groundlessness of Our Believing (Princeton University Press, 2015) is devoted to resolving the problem of radical scepticism, at least in one of its dominant contemporary guises. Radical scepticism is a puzzle that I’ve struggled with for many years, but I genuinely believe that with this monograph I have made definitive progress on the issue (though of course whether others are convinced is entirely another matter). What is key to my proposed resolution is the idea that what looks like a single problem is in fact two related but distinct problems in disguise. This then enables me to show that two apparently competing anti-sceptical proposals—Wittgensteinian hinge epistemology and epistemological disjunctivism—are in fact, properly understood, complementary parts of a broader solution to the target sceptical problematic.

That the contemporary sceptical problem is two problems in disguise is the central claim of part one of the book. In the contemporary literature, the Cartesian problem of radical scepticism can be found formulated in terms of either an underdetermination principle or a closure principle. On the face of it, the difference of formulation seems not to matter—the conventional wisdom is that these formulations are, if not actually logically equivalent, divergent only in ways that don't really matter to our understanding of the sceptical problem. I argue that this is a mistake. In particular, not only are the two formulations logically distinct, but they are also logically distinct in dialectically important ways.¹ This is because they each turn on a distinct guiding claim—what I refer to as a 'source of scepticism'. For underdetermination-based radical scepticism this is a thesis about the limited nature of the rational support that our beliefs can enjoy, even in the epistemically best case—roughly, that even in epistemically paradigmatic conditions, our reasons for believing what we do about the world around us are compatible with widespread falsity in our beliefs. I call this source of scepticism the *insularity of reasons thesis*. For closure-based radical scepticism, I argue that the source of scepticism is a rather different thesis which concerns not the nature of the rational support that our beliefs enjoy in the epistemically best case, but the scope of rational evaluations. This is the idea that there is no in principle constraint on the scope of these rational evaluations, such that one can legitimately rationally evaluate all of one's commitments at once. I refer to this as the *universality of rational evaluation thesis*.

Once we understand that these two formulations of the radical sceptical problem are not only logically distinct, but also turn on subtly different guiding claims, then this has consequences for how we should think about what a resolution to this problem would look like. In particular, we should not assume that a resolution of the one formulation of the problem would thereby be a resolution to the other. And, indeed, this is just what we discover as the book progresses, for along the way I flag how other traditional responses to radical scepticism—from different types of attributer contextualism, to dogmatism, to alternative versions of hinge epistemology, and so on—struggle to simultaneously handle both problems.

My solution to the sceptical problematic involves bringing together two anti-sceptical proposals that have standardly been thought to be competing. The first is the kind of hinge epistemology set out by Wittgenstein in the impressionistic remarks that make up his final notebooks, published as *On Certainty*.² The second is epistemological disjunctivism, a proposal that I've articulated (e.g., Pritchard 2008; 2012) but which is rooted in the work of John McDowell (e.g.,

1995). As I show in the book, far from being competing, these two anti-sceptical proposals in fact belong together.

Part two of the book is devoted to outlining a plausible hinge epistemology. The general idea that Wittgenstein develops is of there being certain basic visceral commitments—the *hinges*—that one must hold in order for rational evaluations to take place at all, such that the very idea of universal rational evaluation that is at issue in the universality of rational evaluation thesis is simply incoherent. The challenge, however, has always been how to take this basic idea and spell it out in a coherent way that doesn't collapse back into being another version of radical scepticism. Very roughly, the main elements of my own version of hinge epistemology include an account of what makes a commitment a hinge commitment and, relatedly, the thesis that hinge commitments, properly understood, are not beliefs, *at least in a very specific sense of belief*. I highlight this last point, since it is apt to be overlooked (and as we will see is overlooked by some of the papers that make up this volume). I will come back to this issue, but first let me say something briefly about how I understand hinge commitments.

While most accounts of hinge commitments emphasize the apparent relativity of our hinge commitments (to person, place, epoch, circumstances, and so on), I aim to show that all hinge commitments have a common core. For what all of our hinge commitments express is an overarching commitment that I call the *über hinge commitment*—*viz.*, the certainty that one is not radically and fundamentally in error.³ The general idea is that the optimal certainty that is associated with the everyday Moorean certainties (in normal conditions), such as that one has two hands, is an expression of this more general optimal certainty. The apparent variability in our hinge commitments is thus explained in terms of how the general *über hinge commitment*, in concert with different circumstances, and thus a different set of beliefs, gives rise to specific hinge commitments with more concrete propositional contents, such as that one has hands. It is thus the *über hinge commitment* that is presupposed in our believing, with specific hinge commitments a function of this general commitment manifest in concrete ways via one's particular set of beliefs.

As I explain in the book, there are numerous advantages to this account of hinge commitments, even aside from the obvious boon that it demystifies what has hitherto been a rather enigmatic notion. One chief advantage is that it can account for two features of hinge commitments that on the face of it are in direct tension with one another. On the one hand, hinge commitments are completely unresponsive to rational considerations, in the sense that they are commitments that we would retain, and be no less certain of, even if we became aware of the fact that we have no

rational basis for their truth. In particular, our continued certainty in them would be manifest in our actions, so that even if we might claim to doubt them, this ‘doubt’ would be in an important sense fake. On the other hand, however, hinge commitments clearly can change over time, and change in ways that seem to be at least superficially rational. Indeed, the very same proposition can be at one time a hinge commitment and another time an ordinary belief, where this change seems to involve a rational response to changed circumstances. But if hinge commitments are not responsive to reasons in the specific way just described, then how can they change at all, much less change via a rational response to changed circumstances?

My account can explain what is going on here. Take the hinge commitment that one has never been to the moon. What holds this in place is that, given one’s system of beliefs (bearing in mind that we are here distinguishing beliefs and hinge commitments), this commitment is an expression of one’s commitment to the *über* hinge. But one’s beliefs can change over time, and in rational ways. Imagine now a child born today who over a long lifetime becomes aware of the increasing ease of space travel, to such an extent that it becomes the kind of activity that one could partake in without being aware of it. At some point one’s beliefs about space travel will change to such an extent that it is no longer a hinge commitment that one has never been to the moon, but rather becomes instead a belief of the familiar kind (akin to my belief that I’ve never been to certain towns in England). As one’s beliefs change in response to changed circumstances, so which particular propositions manifest the *über* hinge commitment will alter accordingly. We can thus explain how one’s specific hinge commitments can change over time, and in rational ways, even though the commitment is not itself responsive to reasons (in the specific sense set out above).

Or consider a proposition which is at one time a hinge commitment but at other times an ordinary belief. For example, that one has two hands is no longer a hinge commitment in the abnormal conditions where one is coming around in hospital after a major car accident. And this change in one’s commitments seems to be a rational response to a change in circumstances. On my view there is nothing puzzling about this, any more than the person in the previous example who lives long enough to lose his hinge commitment to having never been to the moon. The change in circumstances, to the abnormal conditions specified, will lead to a change in one’s set of beliefs, and hence what specific propositions codify the *über* hinge commitment. (This is why, by the way, I do not follow orthodoxy and talk of hinge *propositions*, since I think the phrasing is apt to mislead—the very same proposition can be at one time a hinge and another time not a hinge; it’s the distinctive *attitude* to the proposition that is important, not the propositional content itself).

As noted above, on my view hinge commitments are not beliefs, at least not in a specific sense of that term. The qualification is important, for while I refer to my proposal as the ‘non-belief’ reading of Wittgenstein, I am also very explicit that I am *not* claiming that it would be incorrect to loosely describe hinge commitments as beliefs in an everyday sense of that term. As I point out, the concept of belief is used in a variety of ways, not just theoretically (e.g., by philosophers, cognitive scientists, and so on on), but also in everyday language. Nonetheless, there is a very specific sense of belief that I claim is particularly relevant to epistemologists, and especially in the context of the radical sceptical paradox. This is belief in the sense of that propositional attitude which is a constituent part of rationally grounded knowledge (or ‘K-apt’ belief as I sometimes call it). Belief in this specific sense has some distinctive properties, not least that it has some base-level conceptual connections to truth and reasons. In particular, to believe a proposition is to have a reasons-responsive commitment to the truth of that proposition at least in this minimal sense—one cannot retain one’s belief in this sense while recognizing that one has no rational basis for the truth of the target proposition. If one continues to be committed to the truth of the target proposition regardless, then one’s propositional attitude is not one of belief in this specific sense, but rather a different propositional attitude entirely (such as wishful thinking, etc.). But since, as noted above, one’s hinge commitments are not affected by the recognition that one has no rational basis for the truth of the target proposition, it follows that they are not beliefs in this sense (even though they count as beliefs in a looser, everyday, sense of the notion).

This point is important to how I understand the anti-sceptical credentials of hinge epistemology. The basic idea outlined above offers us a way of rejecting the universality of rational evaluation thesis, but on the face of it this is achieved by rejecting the closure principle. If we cannot undertake universal rational evaluations because of our hinge commitments, then doesn’t this mean that there are consequents of our competent deductions (the hinges) which are unknown, contrary to the closure principle? I argue in part one that the proper formulation of the closure principle to serve the sceptic’s ends is diachronically as a particular kind of competent deduction principle, whereby one gains rationally grounded knowledge in virtue of forming a belief on the basis of a competent deduction from one’s rationally grounded knowledge (while retaining one’s rationally grounded knowledge of the entailing proposition).⁴ So construed, however, it is important to the principle that it is concerned with the acquisition of a belief (in the K-apt sense) via a paradigmatically rational process (competent deduction). Both conditions fail to be met on my view, but let’s focus on the first. If our hinge commitments are not beliefs in the specific sense set out,

and yet closure, properly formulated to play its intended role in the sceptical problematic, concerns the acquisition of belief in just this sense, then it follows that by appealing to hinge commitments to reject the universality of rational evaluations we need not thereby deny the closure principle.

This is a dialectically important claim. The idea that we can put clear blue water between these two theses, such that the former can be denied without thereby denying the latter, helps to establish the *undercutting* credentials of hinge epistemology as an anti-sceptical proposal (as opposed to *overriding*, or revisionist, forms of anti-scepticism). By this I mean that we have a way of dealing with the sceptical problem such that it shows that what looks like a paradox, drawing only on commonsense claims that we would all otherwise accept, is in fact nothing of the sort, but rather arises out of theoretical claims masquerading as commonsense. Moreover, once these commonsense claims are denied, then the paradox disappears. In terms of the closure-based formulation of the sceptical puzzle, it turns out that it is not closure that is the problem. Properly understood, the closure principle doesn't generate the advertised sceptical conclusion, since one does not believe one's hinge commitments, much less acquire them via competent deductions. The key to the closure-based formulation of the sceptical paradox is rather the underlying claim regarding the universality of rational evaluations, and a hinge epistemology shows us why this ought to be rejected, and also (relatedly) how it is not rooted in our everyday practices.

Part three of the book is devoted to epistemological disjunctivism. Hinge epistemology, properly unpacked, can be employed to deal with closure-based radical scepticism, but it has no obvious bearing on underdetermination-based radical scepticism. That rational evaluations are essentially local doesn't tell us anything about the nature of the rational support that one's beliefs enjoy in epistemically paradigm conditions. In particular, perhaps rational evaluation is essentially local, *contra* the universality of rational evaluation thesis, *and* also essentially insular, in line with the insularity of reasons of thesis? A path to radical scepticism would thus be left open. Given our account in part one of how these two formulations of radical scepticism differ, this result should hardly be a surprise.

This is where epistemological disjunctivism enters the picture. In summary, this proposal holds that the rational support that our perceptual beliefs can enjoy in epistemically paradigm conditions is *factive*, in the sense that it entails the truth of the target proposition believed. This is a radical thesis, and hence requires defense, which I offer in this part of the book.⁵ If true, however, then epistemological disjunctivism directly neutralizes underdetermination-based radical scepticism by undermining the insularity of reasons thesis. Moreover, like hinge epistemology, it does so by

undercutting this formulation of radical scepticism. For according to epistemological disjunctivism, the insularity of reasons thesis, far from being a commitment that naturally arises out of our everyday concepts, in fact trades on contentious theoretical ideas that we should dispense with. Indeed, a core claim of epistemological disjunctivism is that these factive reasons are central to our ordinary epistemic practices, and that we are only led to reject this quotidian picture because we are led astray by dubious theoretical reasoning. There is thus no sceptical paradox that trades on this thesis.

While epistemological disjunctivism is effective at dealing with underdetermination-based radical scepticism, it offers us an awkward, at best, purchase on closure-based radical scepticism. In particular, if one tries to respond to the latter by appealing only to epistemological disjunctivism, then one seems to be led to endorsing some quite uncomfortable epistemic claims. The key to resisting these consequences is to recognize that epistemological disjunctivism is only meant to be the antidote to underdetermination-based radical scepticism, just as hinge epistemology is only meant to be the solution to closure-based radical scepticism. Just as these two formulations of the sceptical problem are logically and dialectically distinct, so it now becomes clear that these two anti-sceptical proposals, far from competing, are in fact not only compatible but mutually supporting.

Explaining why this is so is the task of part four of the book. Their compatibility becomes clear once we realize the very different kinds of epistemic claims that they are offering. While hinge epistemology is concerned with demonstrating the logical limitations of the scope of rational evaluation, epistemological disjunctivism is devoted to explaining how the rational support that one's beliefs enjoy in epistemically paradigm cases can be of a particularly robust kind. These are distinct theoretical concerns, with no obvious tension between them once they are unpacked.

But the more important issue is the way that these two proposals are mutually supporting. The basic idea is that each proposal is more plausible when combined with its sister view. So, it is easier to live with the essential locality of rational evaluation if one is also able to demonstrate that paradigm cases of perceptual rational support are factive. And it is easier to live with the idea that paradigm cases of perceptual rational support are factive if one embraces the essential locality of rational evaluation (i.e., because one is not thereby committed to the epistemic immodesty of supposing that one can have a factive rational basis for dismissing radical skeptical hypotheses). Note too the extent to which these proposals are in the same spirit, in that they both offer us undercutting treatments of the skeptical "paradox."

I call the unified defense against radical skepticism the *biscopic* proposal. It is, admittedly, an ugly name—I tried to find a better moniker but without success. But, despite its ugliness, it does convey the bare essentials of the proposal. This is that we have been looking at this problem through, as it were, only one eye—only one eye at a time anyway—and that we need to use both of our philosophical “eyes” in order to see the problem aright. Only then can we gain the right perspective on the problem and thereby recognize what the correct solution to the sceptical puzzle must be.

The book closes with a discussion of a related, though importantly distinct, kind of sceptical anxiety, one that remains even once the sceptical problem has been dissolved (and hence the source of the epistemic *angst* has been removed). I call this psychological state *epistemic vertigo* (or *epistemic acrophobia*, if one wants to be pedantic), in order to capture the idea that it is essentially a kind of phobic reaction to one’s epistemic predicament. Just as one can suffer from vertigo when high up, even while fully recognizing that one is not in any danger, so I think that even after the problem of radical skepticism has been resolved, and hence the epistemic *angst* generated by this problem is defused, it can nonetheless be the case that one feels a residual unease about one’s epistemic situation. The reason for this disquiet is embedded in the Wittgensteinian account of the structure of rational evaluation itself. For what Wittgenstein alerts us to is how one’s apparently very ordinary commitments—such as that one has two hands—can be playing a quite striking role in the system of rational evaluation. Wittgenstein wrote that our hinge commitments “lie apart from the route travelled by inquiry.” (Wittgenstein 1969, §88) He means that the question of their rational standing simply never arises in normal conditions, and so we are unaware that these ordinary commitments play an extraordinary epistemic role. Once one has inquired into their rational standing, however—and the stimulus for this inquiry will almost certainly be philosophical in nature—then it is hard not to continue to be struck thereafter by their peculiarity. To employ a phrase memorably used by Stanley Cavell (1988) in this context, that which is most ordinary becomes ‘uncanny’ once it is made explicit to one.

Another way of putting this point is that while in everyday life we do not take it as given that universal rational evaluations are possible—indeed, we don’t consider the issue at all—neither do we recognize that they are impossible. That’s not to say that we don’t recognize that our everyday practices of giving reasons for and against particular claims is local, as we surely do recognize this, at least implicitly. The crux of the matter is rather that our practices of rational evaluation, while local, also seem to be entirely open to indefinite broadenings of scope. That is, there seems no inherent

limits to the scope of rational evaluation, even if in practice it is always local in nature. That there is such an inherent limit—that a fully general rational evaluation, one that encompassed even our hinge commitments, is impossible—is a philosophical discovery. Moreover, in discovering it, we also realize that our everyday epistemic practices disguise this fact. It is thus unsurprising, then, that even once epistemic *angst* has been removed, epistemic vertigo might well remain, for we now have a perspective on our practices of rational evaluation that is in a certain sense *unnatural*. We have, as it were, epistemically “ascended” and adopted a vantage point that we would not normally adopt. From this unnatural vantage point, epistemic vertigo is a natural response. My point is that one can accept that there is a genuine phenomenon of epistemic vertigo without thereby conceding anything of substance to the radical skeptic. Epistemic *angst* is averted—this is no skeptical solution of radical skepticism. But as with any engagement with a deep philosophical problem, things are not left entirely as they were before.⁶

2. SCEPTICISM AND *EPISTEMIC ANGST*, REDUX

With the foregoing in mind, let me comment on the papers in this special issue. I will start with the papers that are explicitly devoted to attacking the positions that I set out in the book. In many cases, it seems that the concerns raised primarily relate to a misunderstanding of what I am claiming. For example, the paper by Mona Simion, Johanna Schnurr, and Emma Gordon (‘Epistemic Norms, Closure, and No-Belief Hinge Epistemology’) trades on a fairly straightforward misreading of my claim that our hinge commitments are not beliefs. As I make clear in the book (and further explain above), I am not arguing that our hinge commitments are not beliefs in an everyday sense of the term, which is what Simion, Schnurr and Gordon clearly suppose I am claiming. Quite the contrary, I hold that they *are* beliefs in this loose, inclusive sense, such that belief just signals a general endorsement of the target proposition. *Obviously* these commitments are beliefs in this sense, as one is optimally certain of them. My claim is much narrower, in that I hold only that there is a particular sense of belief that is relevant to epistemology (and the problem of radical scepticism in particular)—i.e., the propositional attitude that is a constituent part of rationally grounded knowledge—and that our hinge commitments are not beliefs in this specific sense.

Since the bulk of Simion, Schnurr and Gordon’s critique turns on the idea that I am maintaining the more radical thesis—*viz.*, that hinge commitments are not beliefs in any sense,

including in the everyday sense—it completely misses its target. For example, consider this passage at the very start of the paper:

“This theory [*i.e.*, *non-belief hinge epistemology*] entails substantial fallibility in our self-ascription of attitudes: the propositions classified as hinges are ones many of us take ourselves to believe. Why should we buy into this error theory?” (Simion, Schnurr & Gordon, 1-2)

But since I explicitly argue that our hinge commitments are beliefs in an everyday sense of the term, then there is no need at all for an error theory to explain away widespread misascription of attitudes, as there is no widespread misascription in the first place. Relatedly, Simion, Schnurr and Gordon claim that my argument for the non-belief version of hinge epistemology trades on the following claim:

“Doxastic Responsiveness (DR): Necessarily, a mental state is a belief only if it is responsive to rational considerations.” (Simion, Schnurr & Gordon, 2)

But I obviously do not endorse (DR), since my claim is only about beliefs in a very specific sense—it’s perfectly compatible with my proposal that beliefs in general might be unresponsive to rational considerations and yet be *bona fide* beliefs nonetheless.

Simion, Schnurr and Gordon spend a lot of the paper trying to make sense of why I hold that hinge commitments are not beliefs. This is what they end up claiming:

“We take it, though, on a maximally charitable reading, that what Pritchard has in mind is something like the following: hinges are not the proper subject of rational evaluation, which is why they are also not responsive to reasons.” (Simion, Schnurr & Gordon, 5)

As ‘maximally charitable’ readings go, this is a rather strange gloss on my view, not least because it doesn’t correspond to what I explicitly say is my reason for why hinge commitments are not beliefs. Here is the relevant passage, which Simion, Schnurr and Gordon quote, but then proceed to set to one side:

“[T]here are certain minimal, but constitutive nonetheless, connections between belief and truth such that a propositional attitude that didn’t satisfy them simply would not count as a belief, but would be a different propositional attitude entirely.

In particular, it makes no sense, for example, for there to be an agent who believes that *p* while taking herself to have no reason whatsoever for thinking *p* to be true. *A fortiori*, it does not make sense for an agent to believe that *p* while taking herself to have overwhelming reasons for thinking that *p* is false.” (Pritchard 2015, 90)⁷

So my claim is that a propositional attitude that didn’t pass this test wouldn’t be a belief, in a very specific (K-apt) sense of belief, and I further claim that hinge commitments don’t pass the test. So

yes, hinge commitments are not the proper subject of rational evaluations, but that's not the reason why they are unresponsive to reasons, much less is it why they aren't beliefs. They aren't beliefs because they don't pass the test for (K-apt) belief that I describe. And that's not some general claim about reasons-responsiveness that needs further unpacking, but rather a very concrete conceptual thesis about what I claim is a necessary condition for belief in the relevant (K-apt) sense.

Do Simion, Schnurr and Gordon think that this test is faulty? Alternatively, do they think that our hinge commitments satisfy the test? It's hard to answer the former question, since their focus is not on belief in the (K-apt) sense that I have set out, but rather belief in general (and of course I don't hold that this test applies to belief in general). So, for example, they note that beliefs can be the result of such factors as "a bump on the head" (5), and hence might have nothing whatsoever to do with reasons. In terms of the everyday notion of belief, however, which just concerns a general endorsement of a proposition, then that's fine by me, but this clearly isn't the kind of belief that I'm focused upon (which I would claim clearly cannot be arrived at merely through a 'bump on the head').

Alternatively, in terms of the question of whether hinge commitments satisfy the belief test I lay down, the critical thought seems to be that one can doubt one's hinges, and hence there must be a sense in which they are responsive to reasons in the relevant sense:

"The impressionable teenage moviegoer who mistook *The Matrix* for a documentary can ask herself "Is it rational to believe the world as I think of it exists?" Both of these are instances of evaluating a hinge commitment for rationality." (Simion, Schnurr & Gordon, 6)⁸

The problem with appealing to cases like this is that I'm quite clear in the book that what's at issue is *actual doubt*, and not merely going through the motions of doubt. Descartes may say that he doubts all his beliefs at once, but if one endorses the general Wittgensteinian line, then one won't be inclined to take such assertions at face value. In particular, the Wittgensteinian claim is that one's certainty in the hinges is manifest in one's actions, such that even if one declares that one is doubting the hinge, one's behavior indicates otherwise as it reveals one's visceral certainty. Accordingly, citing cases of people who claim to be doubting a hinge commitment is not sufficient to settle the matter; one needs to further claim that this is a genuine doubt in the relevant sense, and frankly I'm not convinced.

In any case, if there is genuine doubt in this case, then all that shows is that it isn't a hinge commitment. Remember that most hinge commitments—what I call the 'personal' hinge commitments—are only hinges in normal conditions. As I noted above, when I wake up in hospital

after a serious car accident, it might no longer be a hinge commitment of mine that I have two hands; indeed, it might make perfect sense to confirm that my hands are there by looking for them. The very same proposition can function at one time as a hinge and another time as an ordinary belief (which is precisely why I talk of hinge commitments rather than hinge propositions). So it is not even enough to find a propositional content that ordinarily functions as a hinge and have some person genuinely doubt it. Rather one needs to find a proposition functioning as a hinge commitment but which is also subject to genuine doubt. But I think this is going to be hard to find. Does anyone with hands genuinely doubt—i.e., as evident in their actions, rather than merely going through the verbal motions of doubt—that they have hands *in normal conditions*?

It's hard to say, but I think part of the reason why Simion, Schnurr and Gordon fail to understand this point is because they don't have my proposal as a whole in view. Their focus is on the non-belief aspect of the thesis (though as noted they misunderstand this aspect too), but one also needs to understand this part of the proposal in light of the other moving parts. In particular, my account of how the über hinge commitment and the personal hinge commitments are related to one another is very important to my account. For one thing, it can explain how one's personal hinge commitments can change over time, and change in ways that are rational. Simion, Schnurr and Gordon (e.g., 7) are very puzzled by the idea that one's hinge commitments can change, but that's because they don't even introduce the distinction between über and personal hinge commitments that is crucial to understanding how this works. As explained above, it is this distinction that enables me to explain how there can be rational ways in which one's personal hinge commitments change over time, as one's beliefs change in response to new information and so one's overarching über hinge commitment (which never changes) is codified by different personal hinge commitments. But since Simion, Schnurr and Gordon never even mention the distinction between über and personal hinge commitments, even though it's central to my account of hinge commitments, it's no wonder that they are puzzled by the idea that on my view one's (personal) hinge commitments can change over time.

In his contribution to the volume, Xiaoxing Zhang ('Closure, Deduction and Hinge Commitments') also makes the key interpretative error of supposing that the non-belief version of hinge epistemology claims that hinge commitments are not beliefs in an everyday sense of the term (i.e., as opposed to them rather not being beliefs in the very specific sense that I outline). He also overlooks the concrete reason that I explicitly offer for why hinge commitments are not beliefs in the manner that I have in mind (i.e., the simple test set out above, the conceptual point that beliefs

in this sense are not an endorsement of the truth of the target proposition that would survive the recognition that one has no rational basis for the truth of this proposition). Most of Zhang's paper is devoted to analyzing different ways that hinge commitments might be unresponsive to reasons without ever actually engaging with this specific claim at all, even though it is the very reason I offer for taking the line that I do.

One issue that Zhang (10) raises is that the fact that the hinges are already endorsed doesn't prevent them from being reevaluated via competent deductions, in line with the closure principle. As a general point about propositional attitudes, this is of course true, but it simply doesn't apply to hinge commitments as I describe them. Remember that I claim that all our hinge commitments are manifestations of the über hinge commitment, which is a general overarching visceral certainty that one is not radically mistaken. It is not an option to believe one's hinge commitments on a different basis, and in this sense they are very unlike other propositional attitudes. I might decide to no longer endorse the truth of a certain proposition, for example, and then endorse that same content on a different basis, and there is nothing incoherent about that. So, Zhang reasons, why can't one's hinge commitments become reevaluated via a competent deduction, and thereby become a (K-apt) belief in the usual way (and hence be in the market for knowledge)?

Crucially, however, I argue that such optionality of commitments is not available in the case of hinge commitments. If Zhang is unconvinced by this claim then he needs to argue against it directly; merely noting that other propositional attitudes can be changed at will doesn't by itself have any implications for how we should understand our hinge commitments. The more important point, however, is that basis-relativity is explicitly built into my account of closure, and so far as I can tell Zhang doesn't object to this aspect of the formulation. But of course it then follows that if one's hinge commitment to the relevant proposition is not based on the competent deduction—which is precisely what I am claiming—then the closure principle simply cannot have application in this case.

Zhang's failure to note that I do regard hinge commitments as beliefs in a broad sense of the term might explain why he fails to realize why the critical line he offers won't stick. For it is already the case on my view both that one is optimally certain of one's hinge commitments (and hence committed to their truth) and that one can come to evaluate them on the basis of a competent deduction and thereby come to explicitly endorse them on this basis. On the former point, Zhang talks about hinge commitments as if, not being beliefs in a very specific sense, one does not endorse the truth of the target proposition (when of course one not only endorses the truth of one's hinge commitments; one is also optimally certain of these propositions).⁹ On the latter point, there is

nothing in the notion of a hinge commitment that excludes, for example, one adopting the propositional attitude of acceptance towards it. But that won't be relevant to the basis-relativity of the closure principle, which demands that one is able to adopt a very specific propositional attitude—belief, in the K-apt sense—and that propositional attitude is, I argue, not available for our hinge commitments.

Now, as noted above, one's hinge commitments can change over time and in response to different conditions (especially a shift to abnormal conditions). So there is a further sense in which one can reevaluate one's hinge commitments and thereby come to adopt a different propositional attitude toward them. Take the case noted above of the person who awakes after a serious car accident who now genuinely doubts that he has hands. Is this not an example where the subject is reevaluating their hinge commitments as a result of some kind of competent deduction? I doubt that there is any competent deduction on display here, but it doesn't matter for our purposes as the resulting propositional attitude is not a case of someone now believing a hinge commitment anyway, but rather someone who has lost the target hinge commitment and now acquires a belief in the proposition in question. Clearly this is not going to be a problem for my view.

But let us turn to Zhang's other main critical claim, which concerns the following scenario:

“(CHAOS) Sam has grown up in a world that is internally identical to ours. One day, as he wakes up, he sees his hands miraculously disappear. His bed turns into a river, on which people walk toward him, calling him by a different name. After a while, the world returns to normal. The same kinds of events recur several times, differing in their details. Eventually, not a single physical object that Sam could discern remains.” (Zhang, 25)

This is what Zhang says about this case:

“[...] it can be rationally acceptable for Sam, the victim of this chaotic situation, to believe that his opinions about the world are systematically mistaken. Suppose that he does: Sam reacts to his evidential situation by conceding that he has been systematically mistaken about the world. The über hinge commitment, therefore, is revised.” (Zhang, 25)

This is a very puzzling claim to make, particularly in the context of attacking a Wittgensteinian hinge epistemology. After all, this is precisely the kind of scenario that Wittgenstein is claiming is impossible. It is held to be a philosophical fiction, the kind of thing that seems plausible in the abstract but which doesn't stand up to closer scrutiny.

I'm not altogether clear why Zhang believes that these scenarios where one is supposed to genuinely doubt the über hinge commitment are coherent, but the following passage is, I think, very revealing:

“The fact that the entire building would tragically collapse if the foundation were undermined does not logically entail that the foundation itself is protected from attack. We live in an internally coherent world, but the basis of our epistemic system is not *ipso facto* proof against extreme counterevidence.” (Zhang, 25-26)

The thought in play here seems to be that the view that one cannot coherently doubt the über hinge commitment is equivalent to maintaining that this claim is guaranteed to be true (that it is ‘protected from attack’). But this is a bizarre reading of hinge epistemology. *Of course* the über hinge commitment can be false! And if it can be false, then *of course* there can all sorts of things that might happen that would indicate that it is false. Indeed, one can in a (somewhat tenuous) sense even doubt the über hinge commitment, but the point that Wittgenstein makes is that if this is genuine, then one has exited the realm of reasons altogether and begun a path towards insanity. It is the idea that one genuinely doubts such a thing and yet continues to be a citizen of this realm that is held to be impossible, but that is no guarantee at all of ‘protection from attack’.

Moreover, if Zhang’s claim in this quotation were sound, then it would mean that we *would* have a rational basis for the über hinge commitment—what Zhang refers to as the logical basis for supposing that ‘the foundation itself is protected from attack’. But I’ve explicitly argued that such a thing is impossible, which should have given Zhang pause to doubt whether he has understood the proposal correctly. Indeed, on my view it is key to hinge epistemology that one has to recognize that there are no fundamental protections of this kind. This is precisely why the “difficulty is to realize the groundlessness of our believing.” (Wittgenstein 1969, §166) It is also part of the reason why on my view hinge epistemology generates a kind of epistemic anxiety, albeit one distinct from epistemic *angst*—*viz.*, epistemic vertigo.

There are other oddities to Zhang’s stance in this regard. At one point he talks about how one could rationally take an equivocal stance in response to (CHAOS), such that it would be rational both to retain the über hinge commitment and to reject it:

“Situations in which we can be rational in both ways often occur when our evidence is indeterminate regarding the truth of a proposition or when we lack enough practical reasons to make a decision. With inconclusive evidence, for instance, a detective can suspect someone is a criminal, but he can also consider him innocent. In view of the detective’s evidential situation, both attitudes can be rational.” (Zhang, 28)

But notice that this scenario is only coherent because we are talking about propositional attitudes, like suspicion, which can be adopted towards contradictory propositions. The propositional attitudes that concerns us, however, are belief in the K-ap_t sense and hinge commitments, and clearly neither

cannot be adopted towards contradictory propositions. What would it even mean to endorse the truth of both p and not- p in either of these senses (e.g., to be optimally certain of the truth of both of them)?

Matters get stranger when Zhang explicitly discusses the nature of one's endorsement of the über hinge commitment in terms of the strategic reasoning of epistemic rationality. He seems to think that hinge epistemology is committed to the idea that there is a strategic epistemically rational basis for avoiding the 'epistemic suicide' (Zhang's terminology) of doubting the über hinge commitment. But of course this way of thinking about our über hinge commitments is entirely alien to the Wittgensteinian picture, at least as I describe it anyway, in that it puts über hinge commitments into the system of rational evaluation rather than being presuppositions of it.¹⁰

In 'Strange Bedfellows: On Pritchard's Disjunctivist Hinge Epistemology', Annalisa Coliva offers a characteristically nuanced treatment of the book, one that reflects her own distinctive contribution to the hinge epistemology literature.¹¹ I'll take her main critical points in turn.

Coliva begins by taking issue with the way that I argue that a Wittgensteinian hinge epistemology is able to retain the closure principle. As explained above, there are three main elements to my argument here. The first is that the closure principle must be formulated in a certain kind of way—essentially, as a competent deduction principle where the (K-apt) belief in the consequent is based on this deduction—in order for it to support the closure-based formulation of the sceptical argument. In particular, it is only so formulated that the principle looks like the kind of claim that one can only reject on pain of absurdity, which would necessitate that any anti-sceptical proposal that involved its rejection would be clearly revisionary, and hence not undercutting. The second is that closure is specifically concerned with the acquisition of a belief in the specific K-apt sense outlined above, which is formed on the basis of this paradigmatically rational process. The third is that one's hinge commitments are not (K-apt) beliefs, much less are they (K-apt) beliefs that are formed on the basis of such rational processes. So the closure principle doesn't need to be rejected; in particular, rejecting the universality of rational evaluation, as Wittgenstein urges us to do, doesn't lead to denying the closure principle.

One point that Coliva makes is that my formulation of closure in terms of competent deduction is simply what other commentators—most notably Crispin Wright (e.g., 2004)—refer to as 'transmission'. The idea is that we let 'closure' refer to the early formulations offered by, for example, Robert Nozick (1981) whereby it only concerns the idea that knowledge is closed under known entailments. Transmission, in contrast, is the more specific—and more demanding—idea

that the rational support transmits across the entailment to be rational support for one's belief in the consequent.

That was Wright's (2004) proposal at any rate. Oddly, it seems that Coliva, while insisting on this terminological distinction between closure and transmission, in fact understands it slightly differently. Here is her description:

“Closure for justification would be respected and yet it would not thereby give one a first justification to believe Q; nor would it enhance any antecedent justification one would have for Q already. Transmission, in contrast, would be the generative principle that would either give one a first justification to believe Q, or enhance one's antecedent justification for it.” (Coliva, 5-6)

But there is no obvious notion of ‘transmission’ in play in Coliva's formulation of it, as there's no demand that the rational support (or ‘justification’, if that's a different notion) is actually transferring across the entailment. In any case, I'm not convinced that what I'm calling the closure principle is just the transmission principle in disguise, whether on Wright's or Coliva's reading. It's certainly different from Wright's, as there is nothing in my formulation of the closure principle that demands that the original rational support for believing the entailing proposition is thereby one's rational support for believing the entailed proposition. The critical issue is rather *basing—viz.*, that the belief in the entailed proposition is based on this competent deduction. Moreover, it doesn't seem to quite correlate with Coliva's notion of transmission either. There's nothing in my formulation of closure that makes it an essentially generative principle. One might already know the consequent, but then come to know it on a different basis, in which case no new knowledge is generated. (One's epistemic basis needn't be thereby improved either. Perhaps the previous basis for belief becomes discounted in the process, along with whatever epistemic support it was thought to provide). Again, what is important about the reformulation of closure that I endorse—which follows earlier work by Timothy Williamson (2000, 117) and John Hawthorne (2005, 29)—is that it involves a K-apt belief in the consequent that is based on the competent deduction in question.

Furthermore, I'm not persuaded that this really is merely a terminological dispute, as Coliva maintains. Remember that I argue for the claim that the closure principle *needs* to be formulated this way if it is to play the role in the sceptical problematic that it is meant to. Relatedly, I argue that other formulations of this principle—most notably the basic formulation in terms of knowledge being closed under known entailment that Coliva is contrasting with transmission—are simply not credible. So why bother keeping them in the debate under the heading of the ‘closure’ principle? This is especially so given that the transmission principle, as just explained, isn't obviously the right

principle to focus upon in this regard anyway, and certainly not equivalent to the formulation of the closure principle that I offer. I think we would be better dispensing with the closure/transmission distinction altogether, and simply focusing on that formulation of the closure principle that is the one operative in the sceptical problematic.

Coliva is also unsure what I have in mind when I talk about the propositional attitude that is at issue in our hinge commitments (even though I devote a whole chapter of the book to this very endeavor). In particular, she notes that it isn't clear how my notion of a hinge commitments differs from that of an acceptance of a proposition, which Coliva (7) understand as "a propositional attitude in which the subject takes or holds the propositional content to be true." In fact, I do explicitly discuss the relationship between hinge commitments and acceptances in the book. In particular, I argue that acceptance is not the relevant propositional attitude because one can accept a proposition as true while nonetheless being agnostic about the truth of that proposition. (Pritchard 2015, 82) But of course hinge commitments cannot be compatible with agnosticism about the truth of the target proposition, as one is optimally certain of it. So it ought to be clear why I think that our hinge commitments cannot be mere acceptances (just as it should be clear why it doesn't correspond to any of the various other propositional attitudes that I distinguish from our hinge commitments, such as aliefs, trustings, including rational trustings, hypotheses, and so on; remember that I claim that hinge commitments are *sui generis*).

But let us turn to more significant issues. Coliva objects to my way of dealing with closure-based scepticism as follows:

"For, if "There is an external world" is a proposition and, as Pritchard holds, it is knowledgeably entailed by "Here is a hand", how come that we can know that we have a hand and yet not know that there is an external world? In fact, how come that we can know we have a hand and yet not even be allowed to form a belief with respect to "There is an external world", which is a proposition we know to be entailed by "Here is a hand"? Merely insisting that it is a commitment would sound *ad hoc*." (Coliva, 7-8)

To begin with, I should note that 'there is an external world' is not a hinge commitment on my view (and neither is the closely related 'there are physical objects'). This is because I follow Wittgenstein (in the first notebook that makes up *On Certainty*) in treating this claim as simply nonsense—it fails to coherently pick out any proposition at all, unlike a genuine hinge commitment.¹² But let that issue slide, as all that matters is that some hinge commitments have the supposedly problematic consequence that Coliva alleges, and I agree that they do.

What, exactly, is *ad hoc* about the proposal that I offer? Remember that I don't simply insist that we lack a belief, in the K-apt sense, when it comes to hinge commitments, but *argue* for this claim. Moreover, I argue for my account of the nature of hinge commitments. Coliva doesn't explicitly reject any of these arguments, much less argue against them. With this in mind, the charge that the account is *ad hoc* looks like the mere absence of a contrary argument (which itself puts it in the vicinity of a more plausible *ad hoc* charge).

Here is how Coliva expands on her concern:

"I would grant Pritchard that we do not typically come to hold that there is an external world by going through an inference such as the previous one. That inference, that is, is not generative of knowledge. But why cannot it at least generate belief, *ex post*, if it is a valid inference, by Pritchard's lights? The problem is similar to the one usually raised against other hinge epistemologies, and known by the name of "alchemy". Suppose you have a hinge commitment in "There is an external world", acquired one way or another (Pritchard does not say how), which allows you to have a justified belief (or even knowledge) that there is your hand where you see it. Now, consider realizing that "Here is a hand" entails "There is an external world". Why shouldn't this inference allow you to bolster your initial commitment in "There is an external world" and turn it into a belief (or even into a justified, or knowledgeable belief)? The answer is not clear and more would need to be said." (Coliva, 8)

As previously noted, I explicitly reject the idea that 'there is an external world' is a hinge commitment. Coliva also says that I don't explain how our hinge commitments are acquired—even though I say a great deal, as previously noted—but let's set this to one side as well. Worse, Coliva seems to believe that I hold that one can have knowledge, in normal conditions, that one has hands (the entailing proposition in her envisaged closure-style inference), even though I am quite clear that this is a hinge commitment that is not in the market for knowledge. Let's put all these misunderstandings of my position to one side.

The more fundamental question that Coliva raises is why one can't undertake a closure-style competent deduction and thereby come to believe, in the K-apt sense, a proposition that is functioning as a hinge commitment. Coliva talks here as if I side-step this question, when in fact it is central to my proposal to explain why this is case. As with previous commentators, I suspect that the reason Coliva doesn't understand the line that I take is that she is ignoring the important qualification I make about our hinge commitments not being beliefs in a very specific (K-apt) sense of the term (after all, I don't dispute that our hinge commitments can be beliefs in the loose everyday sense of the term). Tellingly, when she describes my view, she doesn't mention this distinction, nor does it figure in her account of her critique. And yet it is key to understanding why I hold that one cannot acquire (K-apt) belief in a proposition that one is hinge committed to as a result of a closure-style competent deduction. By failing to even mention this part of my view, even

though it is central, Coliva is able to make her concerns sound credible, when in fact they don't engage with the details of what I actually propose at all.

Coliva is also puzzled about how I deal with abominable conjunctions, whereby one claims (rationally grounded) knowledge of the first conjunction but claims that one lacks such knowledge of the (clearly entailed) hinge commitment. But I am not sure how abominable such conjunctions are on my view. In particular, doesn't the awkwardness of the assertion rely entirely on one not clarifying what one means? In particular, once it is made clear why the second conjunction is unknown—i.e., because it is a hinge commitment that isn't in the market for knowledge—then the assertion ceases to be puzzling, especially since this clarification cancels a weak implication that one doesn't endorse the truth of the second conjunct (or is at least unwilling to commit to it).

A further objection that Coliva raises concerns my claim that a Wittgensteinian hinge epistemology is only applicable to the closure-based formulation of the sceptical problem, and that one thus needs to supplement it with epistemological disjunctivism (as part of a bispic anti-sceptical proposal) in order to deal with the underdetermination-based formulation of the sceptical problem. Coliva thinks that this move can be resisted, on the grounds that there are resources in a Wittgensteinian hinge epistemology to block underdetermination-based radical scepticism. Here is Coliva:

“[...] by means of hinges such as [...] ‘Our senses are broadly reliable’, by having a hand-like experience, absent defeaters, one would thereby possess a justification for ‘Here is a hand’. That is, thanks to the hinge, one would be entitled to take one's perceptual experience at face value as favoring ‘Here is a hand’ (as opposed to its skeptical counterpart, e.g., ‘I am a handless BIV hallucinating having a hand’). Thus, there is a sense in which, thanks to hinges, our perceptual experiences provide us with a better rational support than their subjectively indistinguishable skeptical counterparts, such as BIV-experiences. That is, thanks to hinges, they provide us with ‘decisive’ reasons in favor of the corresponding ordinary empirical propositions.” (Coliva, 10)

This proposal is only credible if the Wittgensteinian hinge epistemology in question is the one that Coliva proposes, however. On my view, after all, our hinge commitments lack any rational standing whatsoever, and hence cannot possibly play a role in ensuring that our everyday beliefs are rationally favored over manifestly indiscriminable sceptical alternatives. In contrast, the line that Coliva sets out here does make sense on her ‘extended rationality’ line, whereby the hinge commitments do enjoy a positive epistemic status, in virtue of their framework role in epistemic rationality, but just not one that is evidential. If that's right, then it's far more coherent to suppose that this epistemic standing can play the kind of role that Coliva describes, such that we have the desired favoring rational support. (Although even then I do not see how such support can be ‘decisive’ as Coliva

maintains in this quotation; only the epistemological disjunctivist can offer that, via their appeal to factive reasons).

This issue thus comes down to the clash between our respective accounts of hinge epistemology. It is very unfortunate that Coliva's (2015) important monograph setting out her view appeared too late to be covered in *Epistemic Angst*.¹³ It would take me too far afield to contrast our positions here, but the reader can probably work out the main points where we diverge. In particular, I do not think of our hinge commitments as playing some strategic role within epistemic rationality as Coliva does, as if they are reasonable assumptions. Instead I take seriously what Wittgenstein says about the animal, visceral nature of these commitments, an all-out conviction in the truth of the target proposition but one of an entirely arational, and certainly non-strategic, nature. Nonetheless, Coliva's proposal is important, and marks a significant alternative to the one that I set out. I hope to have the opportunity before long to offer a more detailed critical response to her view.¹⁴

In 'Hinge Epistemology and the Prospects for a Unified Theory of Knowledge', John Greco offers the fascinating proposal that our hinge commitments should be construed as a sub-set of procedural knowledge. His intriguing idea is that by thinking of our hinge commitments in this way we are able to integrate them within a wider virtue-theoretic epistemology, something which he claims would be an important philosophical result.

The first point that I want to make is that in principle I could accept a lot of what Greco proposes, in that his idea is not strictly inconsistent with my account of hinge epistemology. My non-belief account of hinge epistemology involves making the claim that our hinge commitments do not amount to rationally grounded propositional knowledge, but that's compatible with them being some other kind of knowledge. In particular, since procedural knowledge is usually modelled along non-propositional lines, in the manner of knowledge-how, there could be a case for arguing that hinge commitments are a kind of knowledge, but just not the kind of knowledge at issue in the sceptical problematic (which I maintain is focused on rationally grounded knowledge).

Nonetheless, while I could in principle accept Greco's proposal, I'm inclined to think that our relationship to our hinges is not knowledge in any sense. In this I follow Cavell, who I think accurately captured Wittgenstein's central insight in this regard (which he also attributed to Heidegger) as being that our fundamental relationship to the world is not one of knowing. Admittedly, Cavell (1969, 241) added a famous qualification to this claim: "anyway, not what we think of as knowing", which leaves open that our hinge commitments might be knowledge in *some*

sense. (Arguably, Cavell's notion of acknowledgement, which he thought was applicable in the case of our hinges, is a kind of knowledge, albeit very different from our everyday propositional notion). I'm afraid that I think Cavell's qualification merely muddies the issue, which I take to be very clear. Let me explain.

At one point in his discussion, Greco notes that procedural knowledge can be made explicit, at which point it becomes propositional knowledge. The subject thus has both propositional and procedural knowledge at the same time. But this is precisely what I claim is not possible in the case of our hinge commitments. This isn't because we can't make them explicit, as obviously we can (albeit we tend to do so within an artificial context of inquiry, one that tends to obscure their real nature; it is also the making explicit of the hinge commitments that leads to the intellectual anxiety that I call epistemic vertigo). Rather the issue is that in making them explicit they do not thereby become beliefs, at least in the sense that is relevant to rationally grounded knowledge. I don't have a problem with hinge commitments being propositional commitments—here is where I diverge with the non-propositional account of hinge epistemology—but I do think it is mistaken to think of them as being beliefs in the sense that is relevant to rationally grounded knowledge.¹⁵ As explained above, they can never be that, and so they are not in the market for such knowledge.

That leaves open that our hinge commitments might nonetheless be non-propositional procedural knowledge of some kind. One concern one might have with this idea is that non-propositional procedural knowledge can be made explicit to the subject, at which point it becomes in the market for being rationally grounded knowledge. Since I hold that the latter is not possible in the case of our hinge commitments, even though I grant that one can become (propositionally) aware of these commitments, then clearly I ought to be suspicious of the idea that they could amount to procedural knowledge in the first place. In particular, it would be important to any development of this idea that the procedural knowledge is only made explicit to the subject in a way that doesn't put it in the market for knowledge (as presumably it wouldn't be in the case of a hinge commitment, for the reasons we have previously explored).

This brings me to a more fundamental concern. It seems to me that it is just that not plausible that our hinge commitments can be knowledge in any sense, including non-propositional knowledge (although I never discuss this possibility in the book itself, as I think our hinge commitments are propositional in form), at least once we are clear about how they are acquired and maintained. What is in the background here is a claim that I share with Greco, albeit in a slightly different form, which is that I hold that knowledge is attributable to cognitive agency in some

significant way.¹⁶ But it seems that our hinge commitments are not obviously attributable to our cognitive agency at all. Indeed, I think this becomes clear once we reflect on the basis for these commitments. As Wittgenstein makes explicit in his discussion of them, they simply don't arise in the way ordinary beliefs are acquired. It comes as a philosophical discovery that one's commitment (in normal conditions) to having hands is not based on one's experiences of one's hands, for example. These commitments are, instead, just there, "like our life" as Wittgenstein (1969, §559) puts it at one point. If that's right, then there's no substantive sense in which we can think of these commitments, where true, as being such that that this cognitive success is attributable to one's cognitive agency. Accordingly, that excludes them from being knowledge.

Now one might object that all this shows is that they do not amount to *propositional* knowledge, either of a rationally grounded variety (of the kind that is the focus of radical scepticism, and hence my book), or of a more general kind that doesn't presuppose rational grounding. One's cognitive success can be attributable to one's cognitive agency in ways that don't presuppose rational grounding, after all, as virtue epistemologists like Greco have pointed out. But why could this knowledge not be procedural knowledge as Greco suggests, where this is contrasted with propositional knowledge (and assuming now, for the sake of argument, that it doesn't become propositional knowledge merely by the subject becoming appropriately aware of it, as Greco proposes)?

The crux of the matter is that even procedural knowledge bears some appropriate connection to manifestations of cognitive agency, as even Greco would admit. And yet our hinge commitments are not in any way manifestations of cognitive agency. Indeed, to think of them in these terms is to misunderstand their role in our cognitive economy. These are commitments that enable certain epistemic practices to occur, rather than being themselves epistemic standings. Accordingly, I think we would be wise to stick to the picture on which our hinge commitments do not amount to knowledge in any sense, including in a non-propositional sense.

A final comment that I would like to make concerns Greco's remarks about the putative merits of an integrated epistemology. It's not clear to me why the more general epistemology that I set out is any less integrated in virtue of it containing unknowable hinge commitments than Greco's proposed theory. Like Greco, I offer a completely general account of knowledge in broadly virtue-theoretic terms. (The details differ, but we can ignore them for our purposes).¹⁷ It's just that once we specify what these conditions are then it becomes clear that our hinge commitments do not qualify. What is unintegrated about that? It's almost as if Greco thinks an integrated account must be one

that treats every putative case of knowledge as *bona fide*, but that's surely not the kind of integration that a theoretician should be exclusively interested in.

Now Greco can respond here that the issue for him is about the integration of propositional and non-propositional knowledge, rather than an integration of different kinds of (putative) propositional knowledge. That's a fair point. But notice that making this explicit rather blunts the methodological force of the integration claim. After all, I can also account for the arational nature of our (propositional) hinge commitments via my account of propositional knowledge, so for my view there is no need to consider an integration with non-propositional knowledge. As it happens, I'm inclined to think (though the matter isn't settled in my mind) that the kind of framework that I apply to propositional knowledge—in terms of an intertwined virtue-theoretic and anti-luck/risk condition—is applicable, at least with suitable modifications, to all kinds of knowledge.¹⁸ If that's right, then I might be in a position to offer the kind of unified account of propositional and non-propositional knowledge that Greco is urging. But the crucial point remains that one does not need theoretical integration of this kind in order to make sense of hinge commitments, at least if my account of them is correct.

In 'An Evidentialist Account of Hinges', Ram Neta offers a fascinating twist on the hinge epistemology line. While endorsing some of the main features of the hinge epistemology line, not least that we have such commitments at the core of our rational practices, he claims that one's total evidence necessarily supports one's über hinge commitment, and hence that it amounts to knowledge. So far from being essentially unknown, as I claim, one's über hinge commitment in fact enjoys a splendid evidential status, more than sufficient for knowledge.

Neta makes some curious readjustments when he states my hinge epistemology line. One is to reinterpret claims about reasons into claims about evidence. Neta clearly thinks that this is harmless, but personally I worry that there is in fact quite a lot of divergence between these notions, at least when it comes to the kind of fundamental epistemic questions that a hinge epistemology raises. In particular, when hinge epistemology trades in reasons it is clearly thinking of them as considerations which, to put the point as I've done elsewhere, are *normative*, *motivational*, and *explanatory* (though the contents that fall under these categories do not always align as one might expect).¹⁹

Normative epistemic reasons are objectively good reasons for regarding the target proposition as true, and they can be either *decisive*, in the case of the factive reasons at issue in epistemological disjunctivism, or they can be *suggestive*, as applies to non-factive reasons. It is the normative element

of reasons that makes them genuine, as opposed to merely apparent, epistemic reasons. As the name suggests, *motivational epistemic reasons* are in fact one's reason for believing that p , and they are primarily what is at issue in the kind of diaphanous response to radical scepticism that I offer (though the other two kinds of reasons are important too, as I will explain). Clearly, a consideration can be one's reason for believing that p without being an objectively good reason for regarding p as true, as when a subject believes that p for lousy reasons (e.g., because she desires that p is true, assuming this is ever a route to belief, at least in the specific sense that concerns us). Going in the other direction, something can clearly be a normative epistemic reason in support of p without it thereby being your motivational epistemic reason to believe that p (as when experts testify that p , but one believes that p because one desires that p is true).

Finally, *explanatory epistemic reasons* are considerations that explain why one believed that p . Such reasons need not line up with normative or motivational epistemic reasons. That they don't line up with normative epistemic reasons should be straightforward. If there are objectively good reasons for treating p as true, but these reasons are not the subject's motivational epistemic reasons, then one would hardly try to explain why the subject believed that p by appeal to the normative epistemic reasons. Putting the point this way, however, might lead one to suspect that one's motivational epistemic reasons and one's explanatory epistemic reasons are just the same thing, but this would be a mistake. In cases where one is radically in error, one's motivational for believing as one does might be very different from what in fact explains what one believes what one does. In short, in the epistemic good case one's normative, motivational and explanatory reasons line up, but in the epistemic bad case, where one is radically in error, these come apart, especially in terms of one's normative and explanatory reasons.

In any case, I digress—the details, while philosophically interesting, are not relevant to our current concerns.²⁰ My point in the foregoing is that the very talk of reasons trades on a rich account of how these considerations feature in one's rational economy. Evidence is a very different notion, however. Something can be evidence without one even being aware of it, and in this sense it manifestly comes radically apart from the notion of a reason that hinge epistemology is employing (which I take to be primarily motivational). It's not at all clear how Neta is thinking of evidence, since he never tells us, but it is a rather crucial detail to understanding his proposal. In particular, what exactly is one's 'total evidence', this crucial notion that Neta's whole enterprise turns upon?

Neta seems to presuppose that since we have a general grasp on the notion of evidence, then we thereby have a general grasp on the notion of one's total evidence. Isn't this just the *aggregate* of

one's evidence, after all? But a moment's thought indicates that this line of thought, while deceptively simple, in fact can't be taken at face-value. Here is one issue: one's evidence isn't always available to one, and Neta's argument clearly presupposes that he thinks of total evidence in such a way that one might not be aware of what one's evidence is. Indeed, he makes this point explicit when he endorses the view, which he attributes to G. E. Moore, that "my total evidence is more evidence than I know to state or attend to." (15) But then clearly one's total evidence doesn't play any explanatory or motivational role in terms of my beliefs, and if so that makes it a very different entity to that involved when hinge epistemologists talk of reasons. For even if one's über hinge commitment is supported by one's total evidence in the way that Neta alleges (though I have doubts about that), this is not an evidential standing that has any bearing on whether one has knowledge of this proposition for the simple reason that one doesn't base one's beliefs on this evidential basis—it is never one's motivational epistemic reason, even by Neta's own lights. It follows that even if one's total evidence in some way supports one's conviction in the über hinge commitment, this commitment will remain unknown regardless, contrary to what Neta is claiming.

Another issue with the idea of our total evidence is that our everyday notion of evidence is highly context-sensitive. I know what it means for something to be good evidence in a court of law, I know what it means for something to be good evidence in an informal situation (like whether to trust a colleague's testimony), I am aware of specialised kinds of evidence that are operative in particular scientific contexts, and so on. But there is no obvious way in which these everyday notions of evidence aggregate; in fact, they seem to largely be in conflict. Rumour and hearsay might well count as evidence in an informal situation—we certainly often treat them as such anyway—but they don't count as evidence in a court of law. And even if one pushes back against this claim by arguing that we ought not to treat rumour and hearsay as evidence in everyday contexts as they are not appropriately connected to the truth, it will not be hard to come up with types of evidence that clearly are appropriately connected to the truth but which are inadmissible in legal contexts (e.g., evidence concerning past convictions). So what sense can there be in the idea of one's total evidence?²¹

More generally, notice that this technical notion of total evidence doesn't obviously have a counterpart in terms of reasons. Even if there is a harmless way of thinking of reasons as evidence, and *vice versa* (though in fact I very much doubt that there is), what the devil would one's 'total reasons' be, exactly? Again, then, we find that Neta's conversion of reasons into evidence is not

nearly as harmless as he supposes. And yet it seems that his critique of hinge epistemology turns precisely on this conversion.

In ‘Closure, Credence and Rationality: A Problem for Non-Belief Hinge Epistemology’, Matt Jope’s critique of my view turns on a quotation that he extracts from *Epistemic Angst*. Here is what Jope quotes:

“Propositions about which one is optimally certain cannot, according to Pritchard, have a rational basis because “to conceive of [*a*] proposition as rationally grounded is to suppose that the rational grounds are *more certain* than the proposition itself” (2016a, 65; *emphasis added by Jope*).” (Jope, 5)

Since the optimally certain propositions in question here are the hinge commitments, Jope extracts the following principle which he claims my view is committed to:

“(PG1) One proposition *p* can count as a reason for believing another proposition *q* only if *p* is more certain than *q*.” (Jope, 7)

Jope argues that this claim is crucial to my Wittgensteinian defense of hinge commitments, and then proceeds to extract various problematic consequences from this principle. He therefore concludes that my proposal is in jeopardy.

Crucially, however, Jope is simply wrong that I’m committed to (PG1). Indeed, I have no such commitment, and hence whatever problematic consequences there are from this claim are irrelevant to my view. We can understand why by going back to the original passage that Jope quotes from *Epistemic Angst*. This is in a part of the book that is introducing Wittgenstein’s proposal, before I begin to set out my own positive account of hinge commitments. In particular, it is a commentary on a remark that Wittgenstein makes in *On Certainty* (§250). The full quotation will make this clearer:

“Wittgenstein is suggesting [*in this example*] that to conceive of this proposition as rationally grounded is to suppose that the rational grounds are more certain than the proposition itself, which of course is *ex hypothesi* impossible since the proposition at issue is held to be optimally certain.” (Pritchard 2015, 65)

As the full quotation clarifies, my comment is aiming to unpack a remark that Wittgenstein makes regarding the rational status of a particular proposition. It is thus odd for Jope to present this claim as being offered in support of my positive proposal. This wouldn’t matter if it turned out that my view was in fact committed to (PG1). Crucially, however, it isn’t.

So what is going on here? Let me begin with Jope’s claim that a commitment to (PG1) is what is motivating my defense of hinge commitments, as I think that if we can understand why this is false, then this will go a long way towards understanding what is at issue here. I think Jope has

made a mistake that a lot of epistemologists make who engage with Wittgenstein indirectly, via other commentators, which is to fundamentally misunderstand the nature of how he is arguing for his proposal. In particular, Wittgenstein is not advancing general claims about the nature of rational support and then arguing that when these are applied to our hinge commitments we find that they are essentially arational. Rather, his commentary on cases is meant to get us to recognize the peculiar nature of these commitments and how they stand aside from our normal rational practices. That is, we might suppose, when conducting a philosophical examination of our commitments, that these ‘Moorean’ commonsense certainties function just the same as our other empirical commitments, and hence should be subject to rational evaluation in the usual way. But this is the very picture—deceptively harmless though it might seem at first blush—that Wittgenstein wants to disabuse us of. He does this not by offering general philosophical claims, but rather by focusing on particular cases and getting us to see that our relationship to the commitments in questions is precisely not of a rational kind as it is with other commitments. That’s what’s going on in the passage that I was commenting on (and which Joep quoted). Wittgenstein is appealing to the special certainty that attaches to these commitments to get us to see that they are not part of the usual rational fabric of our lives as our other everyday commitments are. For example, when we realize that our optimal certainty that we have hands (in normal circumstances) is such that we would be more inclined to doubt our eyes than the existence of our hands if we didn’t see them, then we realize that this claim is playing a very different role to our other commitments.

Now I don’t doubt that this methodology may be a great puzzle to many analytical philosophers. They are looking for some kind of logical argument that demonstrates a particular philosophical conclusion. Accordingly, when they see a discussion of one of Wittgenstein’s cases they suppose that some general thesis is being argued for on the basis of the case, with the appropriate entailment drawn. But this isn’t what Wittgenstein is doing.²² In particular, he is not offering us a deductive argument for the existence of hinges at all, at least in any traditional sense of that terminology anyway. The idea is rather that our philosophical puzzlement is arising out of a certain picture that Wittgenstein is trying to show us is problematic. That picture, while presented philosophically, can look as if it is simply describing our ordinary commitments, but Wittgenstein’s point is that once we attend to the details we realize that there are theoretical claims lurking in the background here, masquerading as commonsense. Once they are removed, then the philosophical puzzlement disappears. And what needs to be eradicated is a conception of our system of rational evaluation that treats our hinge commitments as being just like any other commitment we have, such

that what differences there are between the two are merely one of degree rather than kind. The crux of the matter is that Wittgenstein is not arguing for the existence of our arational hinge commitments and the special role they play in our rational practices, but merely showing that they are there. So the idea that I motivate hinge commitments, much less that Wittgenstein motivates hinge commitments, via an appeal to a general epistemic principle like (PG1) is simply a misunderstanding—albeit, I grant, a reasonable misunderstanding—of what the Wittgensteinian hinge epistemology is claiming.

The final paper in this special issue, ‘In the Space of Reasonable Doubt’, by Ulrike Hahn and Marion Vorms, covers very different ground to the rest. They are still concerned with scepticism, but more in the applied setting of the notion of reasonable doubt as it operates in the context of the criminal trial, where this is meant to give us a model to understand reasonable doubt more generally. They offer a decision-theoretic account of this notion, and claim that that there is no absolute (i.e., decision-independent) notion of ‘reasonable doubt’, such that reasonable doubt cannot be accounted for only in terms of degrees of belief and probabilistic threshold.

The authors’ last claim that reasonable doubt cannot be accounted for in terms of a probabilistic threshold is of most interest to me, since it intersects most closely with my own work, including themes relevant to *Epistemic Angst*. While I agree with this claim, I think that there is a more pressing way of making the point. The nub of the matter is that epistemic risk, like risk more generally, is a modal notion, and hence is not amenable to a straightforwardly probabilistic account.

In terms of the problem of radical scepticism, this issue is most often put in terms of the related notion of epistemic luck rather than epistemic risk.²³ Luck and risk are closely affiliated notions, in that risky events are characteristically lucky events, and *vice versa*. It’s lucky that one survived that car accident, and one was also at a high level of risk of being killed, and so on. Part of the challenge of responding to radical scepticism is to show that, for all the radical sceptic has argued, it is not a matter of luck that our beliefs are true, where that means that we are also not at a high epistemic risk of being in error. For our purposes we can treat the notions of luck and risk as interchangeable—even though there are some important differences—and focus on epistemic risk.²⁴

Risk, like luck, is a modal notion since it concerns the modal closeness of target risk events. What makes juggling with dynamite risky is that this is an activity where the relevant risk event—being blown up—is modally close. One consequence of this point is that probabilities will usually be at best a good guide to risk, rather than the full story. For while low-risk events will tend to be modally far-off and also low probability, and high-risk events will tend to be modally close and also

high probability, this is not universally the case, and where modality and probability diverge, our judgements about risk (like our judgments about luck) track the modal closeness of the risk event over its probabilistic likelihood. That probability and modality come apart in this way is manifest in lottery-style events. Given that one has purchased a ticket for a fair lottery, that one's ticket is the winner is both modally close—since one could easily be the winner; just a few colored balls need to fall in the right configuration—and yet astronomically probabilistically unlikely. We can make this point vivid in terms of risk by considering a lottery where the 'prize' is not a positive event like a normal lottery event, but rather a risk event, such as being killed. We now have a scenario where one is at great risk of being killed—since one could very easily be killed, given the circumstances—but where nonetheless in terms of the probabilities it is astronomically unlikely that one is killed. The modal account of risk generates the right result—that one is at high risk of being killed—unlike the probabilistic account (which predicts that one is at low risk of being killed).

A good response to the sceptical problematic should demonstrate that the radical sceptic has not provided us with any grounds for supposing that the risk event in question—that we are radically and fundamentally in error—is modally close, such that we might reasonably conclude that we are not subject to high levels of epistemic risk. This is exactly what my biscopical proposal offers. On the one hand, the hinge epistemology component of my proposal entails that a certain kind of anti-scepticism—whereby one undertakes a universal rational evaluation of one's commitments and finds them in order—is impossible. That might seem like a negative result, except that it comes hand-in-hand with a demonstration that the sceptical attempt to rationally evaluate all our commitments at once and find them wanting is also incoherent. So there are no sceptical grounds for epistemic risk (though, as I explained above, there is a basis for epistemic vertigo, albeit where this is a distinct phenomenon to the sceptical problematic, not least because there is no epistemic risk at issue). On the other hand, there is also the epistemological disjunctivist component of the biscopical proposal which demonstrates that one can have factive rational support for one's beliefs in epistemically paradigmatic conditions. And since the radical sceptic doesn't offer any rational basis for thinking that we are not in such conditions, it follows that for all the radical sceptic has shown we are not subject to high levels of epistemic risk; on the contrary, we can reasonably suppose that our beliefs enjoy factive rational support.

How does all this relate to Hahn and Vorms's paper? Well, the framework regarding epistemic risk that I have just set out, while applied here to understanding the problem of radical scepticism, can also be applied to the legal case (and many other cases besides).²⁵ In particular, just as

epistemic risk cannot be straightforwardly understood in terms of probability, the same goes for other relevant legal judgements. Indeed, this is a positive result, since it helps us to cast light on certain puzzles in legal theory. For example, it is often noted that it seems wrong to deliver a legal judgment against someone based on merely statistical evidence alone, even when the statistical evidence is very strong. This is most clear in civil cases, since here the requirement is merely that the balance of evidence be in one's favor.

Imagine, for example, that it is agreed by both sides that the injured party in a civil suit was knocked down by a bus at a particular time and on a particular day. Suppose further, that it is a matter of fact that most buses in the city are operated by a particular bus company (more than 80%, say), such that it is overwhelmingly likely that it was this bus company that was responsible for the injury. Even so, if there is no evidence to show that the person was knocked down by one of the buses from this company, doesn't it seem unfair to find against the bus company in this case?

Moreover, notice that it does seem appropriate to find against the bus company where the overall balance of probabilities is much weaker, but where the evidence in question is not merely statistical. Let's say that very few of the buses operating in that town are run by this company, but there are several independent witnesses who have testified that it was one of their buses that caused the injury. That evidence does seem sufficient to find the company liable, even though the likelihood that it was one of their buses that was involved may not that be much higher than 50% (and may be much less than 80+%).²⁶

We can explain what is going on here by noting that the merely statistical evidence, while being strong probabilistically speaking, is unsafe. Given this evidence, one's judgment that the bus company is genuinely responsible for the injury could easily be false. That is, there are close possible worlds where one forms this judgment on the same basis and yet one's belief is false (i.e., where one of the buses operated by another company is in fact responsible). There is thus a high level of legal risk that one makes the wrong judgement. In contrast, the probabilistically weaker evidence in the second case is safe. Given that several independent witnesses have testified that the bus company was responsible for the injury, this judgement could not easily be false. That is, one would need to consider possible worlds fairly remote from the actual world to find one where one forms one's judgment on the same basis (the testimony of the witnesses) and one's judgement is false (e.g., one would have to imagine some kind of conspiracy is occurring).

So there is in fact a very straightforward way of understanding why focusing on mere probabilities is problematic in the legal case.²⁷ Note that I'm not suggesting that this is an alternative

to the nuanced proposal that Hahn and Vorms set out; indeed, it might plausibly be viewed as supplemental to it. Either way, I don't think we have an adequate handle on this issue unless we bring in the notion of risk, and in particular its modal character.²⁸

REFERENCES

- Burge, T. (2010). *Origins of Objectivity*, Oxford: Oxford University Press.
- Cavell, S. (1979). *The Claim of Reason: Wittgenstein, Skepticism, Morality, and Tragedy*, Cambridge, MA: Harvard University Press.
- (1988). *In Quest of the Ordinary: Lines of Skepticism and Romanticism*, Chicago, IL: University of Chicago Press.
- Cohen, L. J. (1977). *The Probable and the Provable*, Oxford, Oxford University Press.
- Coliva, A. (2010). *Moore and Wittgenstein: Scepticism, Certainty, and Common Sense*, Basingstoke, UK: Palgrave Macmillan.
- (2015). *Extended Rationality. A Hinge Epistemology*, Basingstoke, UK: Palgrave Macmillan.
- Enoch, D., Spectre, L., & Fisher, T. (2012). ‘Statistical Evidence, Sensitivity, and the Legal Value of Knowledge’, *Philosophy & Public Affairs* 40, 197-224.
- Greco, J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*, Cambridge: Cambridge University Press.
- (This Volume). ‘Hinge Epistemology and the Prospects for a Unified Theory of Knowledge’, *Synthese*.
- Jope, M. (This Volume). ‘Closure, Credence and Rationality: A Problem for Non-Belief Hinge Epistemology’, *Synthese*.
- Hahn, U., & Vorms, M. (This Volume). ‘In the Space of Reasonable Doubt’, *Synthese*.
- Hawthorne, J. (2005). ‘The Case for Closure’, *Contemporary Debates in Epistemology*, (eds.) E. Sosa & M. Steup, 26-43, Oxford: Blackwell.
- Helmreich, J., & Pritchard, D. H. (Forthcoming). ‘Due Care and Modal Risk: A Solution to the Problem of Aggregative Permissibility’, *Truth and Trials: Dilemmas at the Intersection of Epistemology and Philosophy of Law*, (eds.) Z. Hoskins & J. Robson, London: Routledge.
- McDowell, J. (1994). ‘Knowledge by Hearsay’, *Knowing from Words: Western and Indian Philosophical Analysis of Understanding and Testimony*, (eds.) B. K. Matilal & A. Chakrabarti, 195-224, Dordrecht, Holland: Kluwer.
- (1995). ‘Knowledge and the Internal’, *Philosophy and Phenomenological Research* 55, 877-93.
- Moyal-Sharrock, D. (2004). *Understanding Wittgenstein’s On Certainty*, London: Palgrave Macmillan.
- Neta, R. . (2002). ‘S Knows that P’, *Noûs* 36, 663-89.
- (2003). ‘Contextualism and the Problem of the External World’, *Philosophy and Phenomenological Research* 66, 1-31.
- (This Volume). ‘An Evidentialist Account of Hinges’, *Synthese*.
- Nozick, R. (1981). *Philosophical Explanations*, Oxford: Oxford University Press.
- Pritchard, D. H. (2005a). *Epistemic Luck*, Oxford: Oxford University Press.
- (2005b). ‘The Structure of Sceptical Arguments’, *Philosophical Quarterly* 55, 37-52.
- (2007). ‘Anti-Luck Epistemology’, *Synthese* 158, 277-97.
- (2008). ‘McDowellian Neo-Mooreanism’, *Disjunctivism: Perception, Action, Knowledge*, (eds.) A. Haddock & F. Macpherson, 283-310, Oxford: Oxford University Press.

- (2010). ‘Epistemic Relativism, Epistemic Incommensurability and Wittgensteinian Epistemology’, *The Blackwell Companion to Relativism*, (ed.) S. Hales, 266-85, Oxford: Blackwell.
- (2012a). ‘Anti-Luck Virtue Epistemology’, *Journal of Philosophy* 109, 247-79.
- (2012b). *Epistemological Disjunctivism*, Oxford: Oxford University Press.
- (2014). ‘The Modal Account of Luck’, *Metaphilosophy* 45, 594-619
- (2015a). ‘Anti-Luck Epistemology and the Gettier Problem’, *Philosophical Studies* 172, 93-111.
- (2015b). *Epistemic Angst: Radical Scepticism and the Groundlessness of Our Believing*, Princeton, NJ: Princeton University Press.
- (2015c). ‘Risk’, *Metaphilosophy* 46, 436-61.
- (2016). ‘Epistemic Risk’, *Journal of Philosophy* 113, 550-71.
- (2017). ‘Anti-Risk Epistemology and Negative Epistemic Dependence’, *Synthese* [Online First, DOI: <https://doi.org/10.1007/s11229-017-1586-6>].
- (2018a). ‘Aesthetic Risk’, *Think* 17, 1-14.
- (2018b). ‘Legal Risk, Legal Evidence, and the Arithmetic of Criminal Justice’, *Jurisprudence* 9, 108-19.
- (2019). ‘Wittgensteinian Epistemology, Epistemic Vertigo, and Pyrrhonian Scepticism’, *Epistemology After Sextus Empiricus*, (eds.) J. Vlasits & K. M. Vogt, ch. 8, Oxford: Oxford University Press.
- (Forthcominga). ‘Anti-Risk Virtue Epistemology’, *Virtue-Theoretic Epistemology: New Methods and Approaches*, (eds.) J. Greco & C. Kelp, Cambridge: Cambridge University Press.
- (Forthcomingb). ‘Epistemic Vertigo’, *Ambivalence*, (eds.) B. Brogaard & D. Gatzia, London: Routledge.
- (Forthcomingc). ‘Shadowlands’, *The New Evil Demon: New Essays on Knowledge, Justification and Rationality*, (ed.) J. Dutant, Oxford: Oxford University Press.
- (Forthcomingd). ‘Wittgenstein’s *On Certainty* as Pyrrhonism in Action’, *Wittgensteinian (adj.): Looking at Things From the Viewpoint of Wittgenstein’s Philosophy*, (eds.) N. de Costa & S. Wuppuluri, Dordrecht, Holland: Springer.
- Pritchard, D. H., Millar, A., & Haddock, A. (2010). *The Nature and Value of Knowledge: Three Investigations*, Oxford: Oxford University Press.
- Schönbaumsfeld, G. (2016). *The Illusion of Doubt*, Oxford: Oxford University Press.
- Simion, M., Schnurr, J., & Gordon, E. (This Volume). ‘Epistemic Norms, Closure, and No-Belief Hinge Epistemology’, *Synthese*.
- Thomson, J. J. (1986). ‘Liability and Individualized Evidence’, *Law and Contemporary Problems* 49, 199-219.
- Williams, M. (2004). ‘Wittgenstein’s Refutation of Idealism’, *Wittgenstein and Scepticism*, (ed.) D. McManus, 76-96, London: Routledge.
- 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. J. G. (2004). ‘Warrant for Nothing (and Foundations for Free)?’, *Proceedings of the Aristotelian Society* (supp. vol.) 78, 167-212.
- Zhang, X. (This Volume). ‘Closure, Deduction and Hinge Commitments’, *Synthese*.

NOTES

¹ This claim is rooted in earlier work, especially Pritchard (2005*b*).

² Wittgenstein (1969).

³ Just to clear, in calling the über hinge a *commitment* I am not suggesting that it is something that one needs to have ever explicitly considered. That one has this commitment is rather manifest in one's actions.

⁴ This particular formulation of the closure principle is advocated by Williamson (2000, 117) and Hawthorne (2005, 29).

⁵ I offer a more extensive articulation and defence of epistemological disjunctivism in Pritchard (2012*b*).

⁶ I further discuss the phenomenon of epistemic vertigo, and its philosophical implications, in Pritchard (2019; *forthcomingb*, *forthcomingd*).

⁷ Note, by the way, that this passage comes immediately after a passage where I make clear that I'm only talking about belief in the specific sense of that propositional attitude that is a constituent part of rationally grounded knowledge, rather than belief in general, so it really is very odd that Simion, Schnurr and Gordon overlooked this qualification.

⁸ Incidentally, Simion, Schnurr and Gordon regularly cite the existence of the external world as an example of a hinge commitment on my view. In fact, I agree with Wittgenstein that it isn't obviously a contentful claim in the first place, which would mean there's no straightforward way in which we can be committed to it, much less hinge committed to it. To be fair to the authors, however, this point appears in an endnote, and so is easily missed—see Pritchard (2015, 201-02). (Indeed, Coliva also misses this point in her commentary—see endnote 12). As I explain there, on this exegetical point, I follow Williams (2004).

⁹ At one point Zhang (21) is puzzled by my claim that we can still endorse the propositions that function as our hinge commitments, given that I claim that they are not beliefs. But what is puzzling about this? I don't deny that we think that our hinge commitments are true, after all; indeed, I maintain that we are optimally certain of them, so we clearly regard them as true. Moreover, I don't even deny that they are beliefs in a loose everyday sense of the notion. So where is the puzzle? One would only find this puzzling if one has set aside the distinction between belief in general and the more specific notion of belief that I claim is lacking in the case of our hinge commitments.

¹⁰ I think Zhang may be here confusing my account of hinge epistemology with the very different interpretations offered by such figures as Wright (e.g., 2004) and Coliva (2015).

¹¹ See, in particular, Coliva (2010; 2015).

¹² As noted above—see endnote 8—to be fair to Coliva this point appears in a footnote, and so it is easily missed; see Pritchard (2015, 201-02).

¹³ The same goes for another important monograph relevant to hinge epistemology—Schönbaumsfeld (2016)—which appeared just after *Epistemic Angst*.

¹⁴ I should note that Coliva closes by raising an objection to metaphysical disjunctivism—which, plausibly, epistemological disjunctivism is at least dialectically committed to, as I've noted elsewhere (e.g., Pritchard 2012)—that has been forcefully put by Burge (2010). This objection trades on claims from the cognitive science of perception. While I am aware of this line of argument, I must confess that I lack sufficient knowledge of the relevant empirical work to adequately assess it.

¹⁵ For an important representative work defending the non-propositional reading of hinge epistemology, see Moyal-Sharrock (2004).

¹⁶ In terms of propositional knowledge, I express this claim in terms of a view that I call *anti-luck virtue epistemology*—or, more recently, *anti-risk virtue epistemology*, though the differences between the two formulations don't matter for our purposes—for the former, see, e.g., Pritchard, Haddock & Millar (2010, ch. 1-4) and Pritchard (2012*a*); for the latter, see, e.g., Pritchard (2016; *forthcominga*). Although this departs from the kind of robust virtue epistemology that Greco has defended—see, e.g., Greco (2010)—we are both defending a core claim about the necessity of a virtue condition on knowledge.

¹⁷ See endnote 16 for more on the differences.

¹⁸ For details about the anti-luck/risk and virtue-theoretic elements of my account of knowledge, see endnote 16.

¹⁹ See especially Pritchard (*forthcomingc*).

²⁰ For further discussion of how these three types of reason come apart in different conditions—and why this is important, especially for our understanding of epistemological disjunctivism—see Pritchard (*forthcomingc*).

²¹ Moreover, the notion of total evidence is a very strange one for someone like Neta to endorse, given that one of the proposals that he is famous for is a distinctive kind of contextualism which is concerned with evidential standing. On this proposal, the scope of one's evidence can shift from one (attributer) context to another, such that one might be truly ascribed quite robust evidence relative to one context and yet at the same time be falsely ascribed that level of evidence relative to another. It is difficult to see how one squares this position with a view (presumably acontextual) about total evidence. See Neta (2002; 2003).

²² Interestingly, my former self made the very same exegetical mistake, as at one point I attempted to formulate Wittgenstein's reasoning as an argument appealing to general epistemological claims—see Pritchard (2010). I subsequently came to recognize that this reflected a misunderstanding of what Wittgenstein is trying to do.

²³ See especially my earlier monograph, Pritchard (2005*a*).

²⁴ For more on the modal account of luck, see Pritchard (2014). For a comparative discussion of the notions of luck and risk, see Pritchard (2015*c*). For my account of epistemic risk, specifically—including how this notion comes apart from the closely related notion of epistemic luck—see Pritchard (2016; 2017; *forthcominga*).

²⁵ I apply it to the legal case in Pritchard (2015*c*; 2018*b*) and Helmreich & Pritchard (*forthcoming*). For a further application to aesthetics, see Pritchard (2018*a*).

²⁶ There are various discussions of this kind of problem in the legal literature (also known as the 'gatecrasher' problem). See, for example, Cohen (1977), Thomson (1986), and Enoch, Spectre & Fisher (2012).

²⁷ In fact, I argue that understanding the nature of legal risk can help us to resolve some other problems in legal theory, such as what level of fallibility is permissible in the criminal trial (Pritchard 2018*b*) and a puzzle regarding the proper requirements of due care (Helmreich & Pritchard *forthcoming*).

²⁸ This symposium arose out of a conference devoted to *Epistemic Angst* held at the Sorbonne in 2017. I am very grateful to the organizer of this event, Jean-Baptiste Rauzy, and also the participants, especially the speakers. Thanks also to the editors of this special issue—J. Adam Carter, Gregoire Leffetz, and Guillaume Dechauffour—and to all the contributors. Finally, thanks to Ram Neta for helpful discussion.