I N T E R V I S T E

Conversation with Duncan Pritchard

by Michel Croce

Duncan Pritchard is Chair in Epistemology at the University of Edinburgh and Director of Edinburgh's Eidyn research centre. He is certainly one of the most influential contemporary epistemologists. His main research interests include topics such as the problem of skepticism, the problem of the nature and value of knowledge, extended knowledge, and social epistemology. He is author and co-author of several books, editor of the International Journal for the Study of Skepticism and Oxford Bibliographies: Philosophy, and principal investigator of several international research projects hosted at Eidyn. His latest book, Epistemic Angst: Radical Skepticism and the Groundlessness of Our Believing, published by Princeton University Press in 2015, provides a novel strategy to address the skeptical challenge. In this

broad interview, Pritchard reveals the theoretical origin of his latest work and illustrates the connections between the main topics he explored in his massive research activity. He also offers interesting anecdotes concerning his personal experience as an academic and insightful advices to those who are willing to pursue a philosophy curriculum.

1. *Dear Duncan, thank you very much for accepting my invitation to tell the readers of APhEx something about your work. We'll have time to discuss in detail various aspects of your philosophy, but let me start with a brief biographical question about your interest in philosophy and in becoming an academic. If I recall correctly, you said you wanted to be an academic, even before deciding to start a degree in philosophy. Am I right? Can you tell us something about your professional dreams at that time?*

DP: Many thanks for the invitation Michel!
Your memory is entirely correct. I must have been a very strange child, because from a very early age (6-7) I somehow formed the idea that I wanted to be a Professor. (I recall the teacher asking the class what we wanted to do when we grew up, and amidst all the usual answers – fireman, policewoman, train driver etc – there was my answer: "Professor"). This is particularly odd given that I don't come from a background that would bring me into contact with Professors, or indeed even people with undergraduate degrees! In fact, I'm the first member of my family to go away to University. I think I liked the idea that a Professor was as high as it went in terms of your knowledge of a subject. I may also have had some vague idea about a "life of the mind", and why this might be attractive (e.g., reading lots of books, as I loved books). Of course, I didn't know what subject I wanted to be a Professor of – that came later.

2. *What brought you to study Philosophy then?*

DP: At some point the general idea of being a Professor grew into the more specific idea of being a Professor of English Literature, and maybe also being a poet on the side (seriously!). I used to write a lot, and even won a few prizes for my writing (including one that was quite prestigious), so I thought I might well have the ability to both teach literature at a high level and also publish too. The way I got into Philosophy was pure luck. My original plan was to try to get into Oxford to study English Literature. But I was an unsophisticated and introverted boy (both are still true of me, alas) from a work-

ing-class background and at a mediocre state school, so my chances were pretty slim. But then a wonderful opportunity presented itself. A part-time job that I had (and which I loved) looking after a local graveyard (mainly cutting the grass, but yes, sometimes digging graves too) looked to be morphing into a more full-time temporary role. The graveyard was very large, and in the middle of a very deprived area of Wolverhampton, near to where I grew up. Working part-time, I could only keep on top of some of the graveyard; the rest was completely overgrown. But the plan was to spend a year getting the whole graveyard under control, and of course once the jungle had been cut back it would be easier to maintain. The thinking was that this would be good for the area, and of course I was just the person to do it. This was going to solve lots of problems for me. I could take a year out and earn a decent wage, enough to finance a trip aboard at the end to do some travelling and make myself a little less unsophisticated. Plus, by taking a year out I could potentially get admitted to Oxford on grades alone, without having to do the interview (which I was dreading, as unlike pupils from private schools I would get absolutely no preparation). So far so good. But then the recession hit Wolverhampton. The factory where my father worked laid off most of its workers, him included, and my mother also lost her job as a school secretary. Worse, the funding for my position collapsed as well, so I was now looking at a year at home with me and my parents all out of work. (I got the grades I wanted by the way, and then some, as I took the so-called 'special level' exams for my subjects too – these were the old Oxbridge entrance papers, and to my knowledge I was the first person to take them at my school, at least as anyone could remember anyway – and I sailed through them as well. But all for naught as it turned out). In any case, I was now facing a crisis. Could I really face living a year in Wolverhampton with my parents and all of us on the dole (much as I love Wolverhampton – it's one of the friendliest places around, for example – it is rather grey and depressing)? Anyway, then I remembered that our School had advised people to apply for University places even if they intended to take a year out, as one could always decline their place. Since I had no intention of going to University that year, I therefore had randomly chosen five northern Universities, and "accepted" the one that replied first, thinking that I would never go there. That turned out to be the University of Hull, and as it happened this was the only place where I had applied to do Philosophy and English Literature as a joint degree rather than just English Literature. So I was now faced with the prospect of a bleak year in Wolverhampton, or actually taking up this place at Hull. I thought about it for a while, and decided to give it a shot. It was an excellent decision. First, because Hull turned out

to be an amazing place. It's a bit like Wolverhampton in fact, in that it is very friendly. But also because I was introduced to philosophy and realised that this was what I was born to do. Yes, I love literature, but what really gets me going is ideas, and I realised that it is the ideas in the literature that draw me in. By the end of the first year I converted to a single honours degree in Philosophy, and never looked back. (Mind you, back then I was a "continental" philosopher, reading lots of Foucault, Deleuze, Nietzsche, Heidegger and so on. I'm still interested in some of these figures as it happens, but these days I'm very much in the analytical mould). I thrived at Hull, and not just academically, as I also got to do other interesting things, such as edit the University newspaper, The Hullfire. In my final year I asked my dissertation supervisor, a very nice chap called T. S. Champlin, who the best analytical philosopher currently working in the UK was (the idea of going abroad never even occurred to me), and he said "Crispin Wright". So I started reading Wright's work, which was an eye-opener for me (especially as it was only in the second half of my degree that I really started to become converted to analytic philosophy). Suitably impressed, I applied to the University of St. Andrews, specifically to do my PhD under Crispin's supervision (which was an *excellent* call).

By the way, there's a postscript to all this that I should mention, not least because it might encourage someone out there to give professional philosophy a shot who might otherwise give it a wide berth. Given my modest background, and in particular the fact that I had absolutely no independent financial means (I worked my way through University), opting for a subject like philosophy was a hell of a risk. Remember that I'm the first person from my family to go away to University – I really should have opted for something with more job security! (Though being a poet is hardly better of course). But I reasoned that I owed it to myself to at least try it, and then further reasoned that if I was going to try it I had to go in 100%. I won't lie: it was bloody difficult, and at one point had it not been for the financial help of a friend of mine, I would have had to quit. I realise in retrospect that I was very lucky to be doing a PhD around the time that British philosophy was starting to become more professional and meritorious (largely because of the RAE, now called the REF, that had been introduced to rate academic subject areas). This meant that no one really cared that much any more that I didn't go to the right school (etc), but focused instead on what one had actually achieved. And I published early, and kept on publishing – that stood me in good stead, and it meant that I managed to get a decent (albeit temporary) academic position straight out of my PhD (cemented further by a grant I received from the Leverhulme Trust, which helped me write *Epistemic Luck*).

Hard work, enthusiasm, and a general bloody-mindedness to refuse to even countenance the worst-case scenario (back to Wolverhampton with my tail between my legs?) did wonders for me back then. And of course the fact that I ended up being paid to do philosophy was incredible to me, and still is, to be frank.

3. *Let's go straight to the core of your research interests. Your philosophical output is impressive and covers many different areas of epistemology, but it's easy to recognize that the problem of radical skepticism is the "fil rouge" of your reflections. You wrote about this topic in your first book, "Epistemic Luck" (2005, OUP), and you went back to the problem in your latest work, "Epistemic Angst" (2015, PUP). How do you conceive the skeptical threat? Why is it so hard to defeat the skeptic?*

DP: There's a line in the new book, *Epistemic Angst*, where I say that the problem of radical scepticism, philosophically speaking, is both my first love and my true love. This is entirely true. It was encountering this problem that made me give up on the idea of being a published author/English Literature academic and focus instead on being a professional philosopher. The problem struck me as both deep and genuine. The former because I don't think one can simply ignore it (see my answer to the question below for more details on this front), and the latter because standard answers to the problem are, in my view anyway, philosophically unsatisfying.

I conceive of the sceptical threat as a *paradox*, since I think this takes the puzzle on its strongest formulation. If you take the sceptical threat as a *position*, such that there needs to be a consistent view known as scepticism that someone can advocate, then I think it becomes much too easy to defeat. Rather, what the sceptic does is demonstrate that there are claims that one is antecedently committed to, and which seem to be mere articles of commonsense, but which are, collectively, inconsistent. She then does nothing else, but merely backs away and watches you squirm at the realisation that you are committed a radical form of revisionism, whereby you need to deny something highly intuitive – though you know not what – in order to remove the contradiction.

The reason why it's so hard to defeat the sceptic (or, as I would prefer, "scepticism", since there need be no sceptic on my view) took me a while to discern. (I genuinely struggled with this one, for many years). I argue that the key to the matter is to realise that what looks like a single problem is in fact two logically distinct problems that are generated by discrete philosophical sources. What is required is thus a response to radical scepticism that is

sensitive to these distinct philosophical sources of scepticism, and which in the process sees the sceptical problematic aright. This is what I call the *biscopic* treatment of the problem, in that we are seeing a dual-faceted problem through two eyes for the first time. ("Biscopic" is an ugly term, of course, and for a long time it was just a place-holder while I came up with a better term. But despite a lot of alternatives suggested by friends, I didn't come across anything more suitable, and so this name stuck). The biscopic treatment of the problem that I offer blends together two apparently competing anti-sceptical proposals: a Wittgensteinian hinge epistemology and epistemological disjunctivism (a view inspired by McDowell's writings). I claim that, properly understood, these proposals are not competing at all, but rather complementary. Indeed, I argue that when combined they offer a complete response to the problem of radical scepticism (of the kind that I am concerned with anyway; there are several sceptical problems), in that they enable us to deal with both aspects of this problem in a cohesive and mutually-supportive manner. Moreover, I claim that the reason why philosophers have hitherto been unconvinced by the anti-sceptical credentials of a Wittgensteinian hinge epistemology or epistemological disjunctivism is that they have assessed these proposals in isolation, which in fact they need to be assessed collectively.

4. *In your latest book you introduced a novel and extremely interesting vocabulary for approaching the skeptical paradox. Can you explain what "epistemic angst" and "epistemic vertigo" are in your view?*

DP: In previous work, such as *Epistemic Luck*, I used the term *epistemic angst* to refer to a general anxiety about one's epistemic position that is generated by engagement with the sceptical problem. I now use it in a more refined way to refer to an anxiety of this sort that is legitimately held. This enables me to draw a distinction between epistemic *angst* and *epistemic vertigo*. The latter is a residual anxiety about one's epistemic position that one cannot shake even once one has discerned a philosophically satisfying response to the sceptical problem (such that epistemic *angst*, understood along more narrow lines, is eliminated). What I wanted to capture was the fact that my particular response to the sceptical problem does not leave everything as it was (even though it is an *undercutting* response – i.e., it demonstrates that what looked like a genuine paradox is in fact nothing of the sort). In particular, the fact remains that although the sceptical problematic is shown to be illegitimate – in that it trades on dubious philosophical claims that should be rejected, rather than on mere commonsense – it is not as if in engaging with

the sceptical problem we can forget the inherent limitations of our epistemic position that it exposes; there is no return to a state of epistemic innocence. These limitations are not themselves a source for scepticism, as I explain in the book, but they are hidden from view in our everyday practices, which never fully engage with radical sceptical concerns. The upshot is that while those who have never been occupied with the sceptical problematic can rightly ignore it, those who have engaged with it will be left with an unnatural perspective on one's epistemic position that is liable to generate this vestigial anxiety, even though one is fully aware that the sceptical problem has been exposed as dubious. This is why I call it epistemic vertigo (or epistemic *acrophobia*, if one wishes to be pedantic). The point is that when we attempt a certain kind of philosophical perspective on our epistemic position – when we epistemically *ascend*, as it were – we can feel an anxiety about the security of our epistemic position even while being fully aware that there is no epistemic danger (not from radical scepticism anyway). Our reaction is thus in a sense *phobic*, just as one who is high up might fear falling even though she is fully aware that there is no risk of falling. (Epistemic vertigo is thus an *alief*, rather than a belief, to use Gendler's useful terminology). I think it is important to capture this aspect of the phenomenology of our engagement with scepticism, in order to show that not all epistemic anxiety is legitimate, even if it might be in a certain sense inevitable. (My approach to radical scepticism is broadly in line with Cavell's on this score, in that I think the deep truth in scepticism is that our relationship to the world as a whole is not one of knowing. Our ordinary practices hide this point, a point which an engagement with scepticism reveals. But it is not itself a source of scepticism, properly understood, but merely an inevitable feature of what it is to be a knowing subject. Still, the "unnatural" philosophical perspective from which one recognises this point is dbound to bring with it the residual anxiety of epistemic vertigo).

5. *I really like your approach, because it makes it clear that the skeptical threat challenges our epistemic and psychological states at the same time. Let's dive into your analysis of radical skepticism. You distinguish between two distinct versions of skepticism, i.e. closure-based skepticism and underdetermination-based skepticism. What's the difference between them? And why is so important to keep them distinct?*

DP: For a long time these formulations of the sceptical problem were taken to either be logically equivalent or else dialectically equivalent (in the sense that what logical differences there were, they didn't matter to our dealings

with the sceptical problem). In an early paper of mine – «The Structure of Sceptical Arguments», *Philosophical Quarterly* (2005) – and in *Epistemic Luck*, I noted that in fact they are logically distinct sceptical arguments, to the extent that one could not expect that a response to the one problem would thereby be a response to the other (so in that sense they are not dialectically equivalent either). I didn't appreciate how important this result was at the time, but only came to see its significance later as I tried to work out how a Wittgensteinian hinge epistemology might relate to epistemological disjunctivism. As I came to recognise, these proposals were in fact engaging with distinct sources of scepticism, each of which was associated with one of these formulations of the sceptical argument.

In particular, it came to me that the sceptical appeal of the closure-based sceptical argument trades on what I call the *universality of rational evaluation*. This is the idea that there are no in principle constraints on the extent of one's scope of rational evaluations. This looks like a harmless claim, since why would there be any *in principle* constraints of this kind? (Of course there are practical constraints, like time, imagination, limited resources, and so on). This idea crops up in the closure-based argument in terms of how it seems harmless to transition from everyday local rational evaluations to, via closure-style inferences, the kinds of wholesale rational evaluations that involve radical sceptical hypotheses.

In contrast, the underdetermination-based sceptical argument trades on a distinct source of scepticism, what I refer to as the *insularity of reasons* thesis. This is the idea that the rational support we have for our knowledge, even in the very best case, is always compatible with our beliefs being massively in error. It is only with this claim in play that one can use the underdetermination principle to generate a sceptical conclusion.

I argue that these two sources of scepticism are both completely distinct – in that one could reject/endorse the one without thereby being obliged to reject/endorse the other – and that despite their surface plausibility they are also both highly dubious. In particular, I maintain that they are not derived from our everyday epistemic practices, but are rather the result of questionable philosophical theorising, and hence should be rejected. This is important since it enables us to construct an undercutting response to the problem of radical scepticism, one that contends that what looks like a genuine paradox is in fact nothing of the sort. In particular, we can respond to both the closure and underdetermination-based formulations of the sceptical problem without having to deny either of the epistemic principles in play. We are thus able to deal with the sceptical problem without having to resort to philosophical revisionism. This is only possible, however, because we

have differentiated these two approaches to radical scepticism and recog-nised how they trade on distinct sources of scepticism. The key to the scep-tical problem is thus to realise that it is in fact two problems in disguise, and as such requires the kind of *biscopic* treatment outlined above.

6. *Now you have to tell us about your epistemological cure for the epistem-ic angst. Which strategy do you adopt against the two forms of skepticism?*

DP: As noted above, I opt for a biscopic treatment of the problem, which is sensitive to the fact that this is in fact two logically distinct sceptical prob-lems in disguise, which in turn trade on two distinct sources of scepticism.

I argue that a Wittgensteinian hinge epistemology, properly understood, is the antidote to closure-based radical scepticism, in that it undercuts the universality of rational evaluation thesis. In a nutshell, what Wittgenstein demonstrates in *On Certainty*, I argue, is how it is a requirement of being a rational subject at all – i.e., someone who has beliefs, who can offer reasons in support of, or against, particular propositions, and so on – that one has a body of arational hinge commitments. As Wittgenstein emphasises time and again, it is a matter of logic that rational evaluation is structured in this way, such that it has these arational commitments at its heart. One consequence of this claim is that the very idea of a fully general rational evaluation – whether negative, as in the case of the sceptic, or positive, as in the case of the traditional anti-sceptic –is simply incoherent, since there could be no such thing. All rational evaluation is instead in its nature essentially local. (Note that this is not a limitation on our rational evaluations, any more than the fact that one cannot invent a circle-square is a limitation on one's crea-tivity). This means that it is the universality of rational evaluation thesis that is the joker in the pack in this regard, and must be rejected. Wittgenstein ex-poses how this is in fact a philosophical fiction, one that is not rooted in our everyday epistemic practices at all. Crucially, however, I also claim that if we take Wittgenstein's proposal seriously – and in particular embrace his claim that our hinge commitments are essentially "animal", "visceral" and so on – then one can reject the universality of rational evaluation thesis without thereby having to deny the closure principle. We thus get an under-cutting – and hence completely non-revisionist – response to closure-based radical scepticism.

In contrast, I argue that epistemological disjunctivism, again properly understood, is the antidote to underdetermination-based radical scepticism. According to this proposal, in paradigm cases of perceptual knowledge the rational support one's belief enjoys can be *factive*, where this means that it

entails the proposition that it is rational support for. In particular, I claim that in paradigm cases of perceptual knowledge, one's rational support for believing that *p* can be that one *sees that p*, where seeing that *p* entails *p*. Such a proposal has, until recently, simply found no home in contemporary epistemology, since it was widely accepted that if rational support was reflectively accessible then it couldn't be factive (and, conversely, that if epistemic support was factive, then it couldn't be reflectively accessible rational support). I argue that the philosophical reasons offered for rejecting epistemological disjunctivism are all highly dubious on closer inspection, and moreover that factive reflectively accessible reasons are rooted in our everyday epistemic practices. (Key to this argument is a distinction I draw between favouring versus discriminating epistemic support. I maintain that all epistemologists should endorse this distinction, but that it is particularly helpful for epistemological disjunctivism). It follows that our rejection of epistemological disjunctivism was based on faulty philosophical reasoning and hence that we should embrace our everyday practices that have factive reflectively accessible reasons at their heart. With such a picture in place, however, we are in a position to reject the insularity of reasons thesis – which is logically incompatible with epistemological disjunctivism – and hence undercut underdetermination-based radical scepticism.

If one treats either of these proposals as being a single silver bullet to defeat radical scepticism, however, then they start to run into difficulties. A Wittgensteinian hinge epistemology has absolutely nothing to say about whether reasons can ever be factive, and as such it is silent about the insularity of reasons thesis. But with underdetermination-based scepticism still in play, it can be hard to see how a response to radical scepticism that insists on the essential locality of rational evaluations is any comfort. Epistemological disjunctivism, in contrast, *can* be applied to closure-based radical scepticism – indeed, I had a go at this strategy in an earlier monograph, *Epistemological Disjunctivism* (Oxford UP, 2012) – but it gains an awkward purchase, at best, on this formulation of the problem. For epistemological disjunctivism is now required to maintain not just that one can have rational support that decisively favours one's everyday beliefs over incompatible alternatives, but that one has, in addition, a factive rational basis for knowing the denials of radical sceptical hypotheses. That claim is hard to stomach.

But once one combines these proposals the difficulties that afflict each alone disappear. For one thing, notice that these theories are not competing anti-sceptical stances, as many have hitherto supposed, as they have very distinct domains of application. A hinge epistemology, properly understood, only tells us something interesting about the structure of rational evalua-

tions, but it is entirely silent on whether, in particular, rational support can ever be factive. Epistemological disjunctivism, properly understood, is only a claim about a paradigm kind of knowledge (and possibly about paradigm cases of knowledge more generally, though that is an argument for another day), and not a thesis about knowledge in general. It is thus entirely compatible with the idea that rational evaluations might be essentially local.

These proposals are also mutually supporting when combined. It is easier to live with the idea of local rational evaluations if one can be assured, via epistemological disjunctivism, that there can be factive rational support (and hence that underdetermination-based radical scepticism is blocked). And it is easier to live with epistemological disjunctivism if it is not committed to there being factive rational support for the denials of radical sceptical hypotheses. But this is precisely one consequence of combining the views, since such a result is ruled-out by the fact that our commitments to the denials of radical sceptical hypotheses are hinge commitments.

Finally, notice how both responses are entered in the same anti-sceptical spirit, in that they are both undercutting responses to the problem. In particular, they both argue that what looks like a genuine paradox is in fact nothing of the sort, in that it in fact illicitly trades on dubious philosophical claims that should be rejected. We thus get a complete response to radical scepticism, one that is sensitive to its dual nature, and which offers an undercutting dual response as a result. Such are the advantages of taking a biscopic line on radical scepticism.

7. *I want to go back to your first book for a moment, to understand better how the problem of radical skepticism fits into the whole philosophical picture you built in these years. First of all, am I right if I say that "Epistemic Luck" contains the main ingredients of your epistemology? I want to mention, among others: skepticism, Ludwig Wittgenstein, epistemic luck, and virtue epistemology.*

DP: Yes, I think that's right, though I hadn't realised at the time what a rich seam I was mining. (Had I realised this, the book probably wouldn't have been quite so ambitious, and hence would probably not have had the influence it did. As Cate Elgin once memorably argued, there can be sometimes be an epistemic efficacy in being ignorant!).

8. *Can you tell us how you got to Epistemic Angst from Epistemic Luck? What is the main progress of the original view?*

DP: My stance on radical scepticism has changed quite significantly, as I was never quite persuaded of the line I take in *Epistemic Luck* (though I still think I did a decent job of defending it). As noted above, although *Epistemic Luck* contains a discussion of the logical differences between closure-based and underdetermination-based radical scepticism, I didn't appreciate at the time how important this point was to our understanding of the sceptical paradox. There are several other key ingredients that mark the transition from the first book to the most recent. One of them is my articulation of epistemological disjunctivism (e.g., in my 2012 Oxford UP monograph, *Epistemological Disjunctivism*). The other is the very specific non-belief rendering I now give of a Wittgensteinian hinge epistemology, which enables me to retain the closure principle. The third is an important distinction that I draw between favouring and discriminating epistemic support (see, e.g., «Relevant Alternatives, Perceptual Knowledge and Discrimination», *Noûs* (2010)), and which underpins my argument for epistemological disjunctivism. Finally, I also came to recognise that the project of analyzing knowledge and the project of dealing with radical scepticism are best kept apart (rather than run together, as I did in *Epistemic Luck*). Putting all these points together, one ends up with the biscopic treatment of radical scepticism in *Epistemic Angst*. It took a long time to get there, but I got there in the end!

9. *Since we moved from specific questions to the broad picture, let me ask you also about your 2012 "Epistemological Disjunctivism". You told us about this view as a component of your biscopic solution to radical skepticism, but I want to understand better the scope of this project. Should we endorse epistemological disjunctivism only to account for perceptual knowledge? Or, do you think it's plausible to be disjunctivist about other sources of knowledge?*

DP: Yes, my view – although I haven't argued for this in print as yet – is that if epistemological disjunctivism is plausible then it ought to be generally applicable to paradigm cases of knowledge, at least when that knowledge comes from core sources, such as perception, memory and so on. The focus on perceptual knowledge is merely because I think this is the cleanest case, and hence it's the best case on which to mount a defence of epistemological disjunctivism as a viable position (bearing in mind that, until quite recently, the position was thought simply unavailable as a theoretical option by contemporary epistemologists). That said, it may well be enough as far as the sceptical problem that I am concerned with goes if only paradigmatic per-

ceptual knowledge enjoys factive rational support. If that's right, then the extension of epistemological disjunctivism to other kinds of knowledge need not carry any anti-sceptical load, even though it would be an epistemologically important result.

10. *Let's talk about another important part of your work: the theory of knowledge that you named "anti-luck virtue epistemology". As you explain in the first investigation of "The Nature and Value of Knowledge" (2010, OUP), your account challenges robust virtue epistemology, i.e. those views that attempt to "exclusively analyse knowledge in terms of a true belief that is the product of epistemically virtuous belief-forming process" (38). Can you briefly tell us what are the limits of robust virtue epistemology and how your anti-luck virtue epistemology goes beyond them?*

DP: It's interesting that whereas in *Epistemic Luck* my account of knowledge and my response to radical scepticism went hand-in hand, in subsequent work I saw the virtue of keeping them apart. In particular, I don't think that we need to offer a complete theory of knowledge in order to resolve the sceptical problem – indeed, my response to the sceptical problem in *Epistemic Angst* doesn't appeal to any specific theory of knowledge. Moreover, one doesn't need to resolve the sceptical problem in order to offer a complete theory of knowledge (as *The Nature and Value of Knowledge* attests, as it hardly mentions the sceptical problem).

The backdrop of my defence of anti-luck virtue epistemology was my desire to somehow make the kind of *robust virtue epistemology* – as defended by Ernest Sosa, John Greco, Linda Zagzebski and others – into a defensible position. I really liked the elegance of this proposal, and in particular how it seemed to offer a very straightforward account of the value of knowledge in terms of knowledge being a particularly cognitive kind of achievement, but the more I examined the view the more I came to realise that it simply wasn't tenable. In a nutshell, robust virtue epistemology understands knowledge as being a cognitive success that is primarily attributable to cognitive ability (which in turn generates the analogy with achievements, since they are plausibly to be understood as success that is primarily attributable to ability). The problem such a proposal faces is that it cannot account for a phenomenon that I refer to as *epistemic dependence*. This concerns the way in which whether one's cognitive success amounts to knowledge can significantly depend upon factors outwith one's cognitive agency.

Epistemic dependence has both a positive and a negative aspect. *Positive epistemic dependence* is when one exhibits a level of cognitive agency that would not normally be sufficient for knowledge, but one counts as knowing regardless because of factors outwith one's cognitive agency. *Negative epistemic dependence*, in contrast, is when one exhibits a level of cognitive agency that would normally easily suffice for knowledge, but one does not count as knowing due to factors outwith one's cognitive agency.

A good example of the former are testimonial cases where one for the most part trusts the word of an informant, in an environment that is epistemically friendly. Mere gullibility is never a route to knowledge (*contra* Lackey, for example, I don't think that one can gain testimonial knowledge merely by believing an informant, no matter how friendly the environment), but crucially in the right epistemic conditions one can gain testimonial knowledge by exhibiting very little by way of cognitive agency (e.g., just by being reasonably discriminating about who one would ask, and what one would believe). Such knowledge is not primarily attributable to one's cognitive agency, however, but if anything to one's informant's cognitive agency (or, if you like, the 'cognitive whole' of subject-and-informant).

A good example of the latter are cases of purely *environmental epistemic luck*. I introduced this distinction to enable me to contrast standard cases of veritic epistemic luck that involve something intervening between belief and fact (*intervening epistemic luck*, as I call it), and those cases of veritic epistemic luck where the luck in play is purely environmental, and nothing intervenes. Setting aside specifically epistemic cases for a moment, we can see this distinction in play across a range of scenarios. Imagine first a skilful archer who fires her arrow at the target and who hits the target, but where the reason she hits the target is that a dog has jumped up mid-flight, caught the arrow and placed it in the target. This would be a case of intervening luck, where the agent's success is such that it could have easily been a failure due to something intervening to ensure the success. Contrast this case with one where nothing gets in the way, but the archer is skilfully firing her arrow in an environment where conditions are such that it is nonetheless still the case that her success could very easily have been a failure (e.g., a dog very narrowly failed to catch the arrow mid-flight, and had he caught it he would have run off with the bolt). This would be a case of purely environmental luck. A case of environmental *epistemic* luck would be a barn façade scenario, where the agent really is seeing a genuine barn before her, but happens to be in barn façade county where most of the things that look like barns are nothing of the kind (such that her cognitive success could easily have been a cognitive failure).

Crucially, however, cases of environmental luck are such that the subject's success *is* primarily attributable to her agency – think, for example, of the archer just noted who successfully hits the target without any interference, but whose bolt could have easily been taken by the dog. This means that cases of environmental *epistemic* luck ought to count as knowledge too, and yet since they are also cases of veritic epistemic luck that's not very plausible (and if you think that it is plausible, see what I say about epistemic risk below). Moreover, although various people have tried to get out of this problem by tweaking the robust virtue-theoretic account of knowledge, I've since argued, in a joint paper with Jesper Kallestrup («Virtue Epistemology and Epistemic Twin Earth», *European Journal of Philosophy* (2014)), that such tweaks are ultimately hopeless. In particular, even if we keep fixed everything that might remotely be relevant to an attribution of cognitive agency, one can still generate cases of positive epistemic dependence via appeal to environmental epistemic luck, and so robust virtue epistemology is untenable.

Moreover, notice that these two objections to robust virtue epistemology – via cases of positive and negative epistemic dependence – are particularly problematic for the view in that they pull the position in two opposing directions. In order to deal with positive epistemic dependence, the temptation is to raise the bar for knowledge in order to exclude cases of environmental epistemic luck. But in order to deal with negative epistemic dependence, the temptation is to lower the bar, in order to allow for knowledge to be had in epistemically friendly environments even though very little cognitive agency is on display. Put the two objections together, and the position is pulled asunder.

That's the negative story, but there's also a positive story to tell as well. Although I was very disappointed to find that I couldn't make robust virtue epistemology work, it dawned on me that the source of the problem here was a failure to recognise that knowledge is in fact answering to two distinct, albeit overlapping, considerations. The first is the need to eliminate high levels of epistemic luck (or, as I put it these days, high levels of epistemic risk). The second is the need to accommodate the sense in which when one knows it is down to one's cognitive agency, in some significant way, that one is cognitively successful. These intuitions can look like two sides of the same coin, in that what it usually takes for one's cognitive success to be a manifestation of one's cognitive agency is that it be non-lucky, and what it usually takes for one's cognitive success to be non-lucky is that it is a manifestation of one's cognitive agency. But while they overlap, they are not the identical. Indeed, I argue that they come apart in both directions,

in that there can be cases where the cognitive success is non-lucky in ways that have nothing to do with one's manifestation of cognitive agency, and there can be cases where the cognitive success is primarily attributable to one's cognitive agency and yet is nonetheless a matter of luck (environmental epistemic luck is a case in point on the latter front). Once we recognise this point, we realise that what we are really after is not a theory of knowledge that exclusively understands knowledge in terms of the exclusion of epistemic luck or in terms of the manifestation of cognitive agency, but rather a proposal that can accommodate both constraints and how they intersect. I call the resulting proposal *anti-luck (/risk) virtue epistemology*. Roughly, it holds that knowledge is a safe cognitive success where the safety of one's cognitive success is at least significantly (but need not be primarily) attributable to one's manifestation of cognitive agency. I argue that this is a complete theory of knowledge, in that it is able to deal with all the problem cases. It also casts light on a number of interesting issues in epistemology, such as the relationship between knowledge and cognitive achievements, the relationship between knowledge and understanding, the relationship between propositional knowledge and ability knowledge, the value of knowledge, and so on.

11. *Two of the several externally funded Eidyn research projects you lead as Director of this research centre based in Edinburgh are about the virtue of intellectual humility. What did you discover about intellectual humility through these projects? And why is this non-classical virtue becoming so important nowadays?*

DP: One thing I discovered in engaging with all this wonderful new work on intellectual humility is that my conception of it is largely at odds with everyone else's. In particular, I think it is standard to regard intellectual humility as at least involving some kind of ownership of one's intellectual limitations, and possibly even having a downgraded conception of one's intellectual abilities. I think the latter is completely at odds with intellectual humility being an intellectual virtue (which I think it is), since I don't think intellectual virtues can essentially involve inaccuracy. But the former is also problematic. In particular, I hold that intellectual humility is best characterised in terms of having certain other-regarding – better, other-respecting – dispositions, and not in terms of some sort of distinctive attitude that one takes to oneself and one's intellectual abilities. For example, it consists in a willingness to listen carefully to others, to think about what they have to say, to present one's reasons and listen to their counterarguments, and so on.

This conception of intellectual humility has important ramifications for a number of debates, not least the debates about the epistemology of disagreement. It is often taken as given in these debates that if one, on reflection, responds to a disagreement with one's epistemic peers by sticking to one's opinions, then one is thereby being dogmatic, and hence lacking in intellectual humility. But I don't think that this follows at all. If one reflects on the matter and one is genuinely convinced, then I think one owes it to oneself to maintain one's opinions. But that's entirely compatible with being intellectually humble, since on my view that is about manifesting certain other-respecting dispositions towards one's disputant, and not in terms of any distinctive attitude one takes to one's own beliefs. (And note that these dispositions must really be rooted in a respect for others – merely going through the motions of respecting another's point of view is not sufficient). More generally, I think this conception of intellectual humility is very important to public debate. We need to learn how to respect others – where I think this is ultimately rooted in a genuine care for other people – while also recognising that this is entirely compatible with one having genuine convictions (e.g., political, ethical, etc).

12. *We haven't talked about your work in applied epistemology so far, but I think it's important that the readers of APhEx get to know about what you and other members of Eidyn are doing in this field. As Director of this research centre, could you tell us a bit about the projects and the activities hosted by Eidyn? You are also doing research in sub-fields such as social epistemology, epistemology of education, legal epistemology. What are the prospects for applied epistemology?*

DP: One of the distinctive features of *Eidyn* is its commitment to outreach and, relatedly, to developing impact from the research projects that we host. We've had a lot of success on this front, and it's something that I'm very proud of. Since you ask about applied epistemology in particular, let me tell you about one specific project that we've run which is devoted to transforming prison education (see http://eidyn.ppls.ed.ac.uk/project/philosophy-prisons). This involved collaboration with the Scottish Prison Service and colleagues at New College Lanarkshire and the University of Edinburgh's Moray House School of Education. The goal was to run a study that explored how the teaching of philosophy in prisons could help develop prisoners' intellectual character (drawing on, in the process, some work I've done on the role of intellectual virtues in educational contexts). To this end we used an off-line version of *Eidyn's* very popular *Introduction to Philosophy*

MOOC (= massive open online course) and a particular pedagogical approach to teaching philosophy known as Community of Philosophical Inquiry, or CoPI for short. Throughout the study we were looking for signs of the development of intellectual character, such as the manifestation of particular cognitive traits – for example, intellectual tenacity, creativity in problem-solving, the ability to articulate one's reasoning, responsiveness to the reasons offered by others, and so on. The results were far more impressive than we could have hoped, as there was clear evidence that the prisoners not only developed their intellectual characters, but that they also developed important social and interpersonal skills (e.g., self-esteem and respect for others). This is a good example not just of applied epistemology, but also impact, in that this project has transformed the provision of prison education in at least two prisons in Scotland, with scope to eventually change prison education across the sector (there are plans for follow-up projects aimed at achieving just that).

The foregoing in itself offers a flavour of why one might profit from engaging in applied epistemology. But even where there is no obvious impact of this kind, I still think that applied epistemology is worth exploring. In particular, it is very useful to take general ideas in epistemology and see how they play out within an applied domain. Recently, for example, I've been taking my new work on epistemic risk (described below) and applying it to two issues in the philosophy of law, concerning the notion of legal evidence and the question of how much legal error is compatible with a just legal system. I've also been doing a lot of work on the epistemology of education, such as concerning what the epistemic goals of education are, and how the intellectual virtues figure in those goals (this work obviously connects with the philosophy in prisons project just described). And I have an ongoing interest in trying to work out how general themes in epistemology play out within the specific domain of religious belief. In all these cases I think one's understanding of both the specific domain in question and of epistemology in general is enhanced, in that the flow of understanding goes in both directions. This why when I teach epistemology I increasingly try to bring in at least some applied element into the discussion, in order to get students to see that epistemology isn't just about dealing with radical sceptics or defining knowledge, but can also be about something as practical as how to understand legal evidence.

13. *And what about the completed project on Extended Knowledge? It sounds as it is a great interdisciplinary topic, involving general epistemology, social epistemology and philosophy of mind.*

DP: The basic idea behind *Eidyn's* Extended Knowledge project was to systematically explore – incredibly, for the first time – the epistemological ramifications of the extended cognition programme in the cognitive sciences (see http://eidyn.ppls.ed.ac.uk/project/extended-knowledge-2013-15-0). Edinburgh was the ideal place to host such a project, given that our own Andy Clark has been the leading figure when it comes to work on extended cognition. I wrote what I think is the very first paper exploring the epistemological implications of extended cognition («Cognitive Ability and the Extended Cognition Thesis», *Synthese* (2010)), but since then the floodgates have opened on this score, largely due to our very successful project. In particular, we didn't just examine the epistemological ramifications of a technological or environmental cognitive extension, but also the possibility of group, or social, cognition. The result has been a wealth of publications, not least a very long survey paper written by myself and my collaborators on the project ('Varieties of Externalism' (with J. A. Carter, J. Kallestrup & O. Palermos), *Philosophical Issues* (2014)), and two volumes of papers, containing a range of interdisciplinary perspectives on the issues, that are forthcoming with Oxford UP (*Extended Epistemology* and *Socially Extended Epistemology*).

14. *Before concluding this interview with a few curiosities, let me ask you what are you working on at the moment. I saw that you recently published on the topic of epistemic risk («*Epistemic Risk*», Journal of Philosophy (forthcoming)). What is epistemic risk? What's the relationship between luck and risk in your view?*

DP: When I first started developing my anti-luck epistemology methodology as part of *Epistemic Luck*, I was proceeding on the basis that there must be a philosophical account of the nature of luck out there. To my surprise, there wasn't, in that it was mostly just treated as an undefined primitive. This is particularly remarkable given that the notion appears across a number of philosophical debates, such metaphysics (e.g., causation), political philosophy (e.g., just deserts), ethics (e.g., moral luck), and epistemology (e.g., epistemic luck). Usually philosophers cannot help but define everything, so it's surprising to see that they let this one slip through the net. On the plus side, however, it turned out that there was a wealth of interesting empirical work on luck and risk ascriptions that was helpful to developing my own account of luck, which ended up going down modal lines. In *Epistemic Luck*, I put the material on risk to one side, as I tended to assume, in

common with the empirical literature, that the differences weren't that important (the folk often tend to run these notions together, which is why they tend to appear side-by side in the empirical literature).

A few years ago, however, I came to realise that the differences between luck and risk were more significant than I previously thought, and in fact could help me resolve some problems with my theory of knowledge. There was one lacuna in my theory of knowledge in particular that always bugged me. According to my modal theory of luck, lucky events are, roughly, those events that obtain in the actual world but which could easily have not obtained (i.e., which obtain in close possible worlds, where the initial conditions for the target event are kept fixed). When I plugged this into an anti-luck epistemology, I drew the consequence that a true belief that was not subject to veritic epistemic luck would not just be true in the actual world, but would continue to be true across close possible worlds where the belief continues to be formed on the same basis as in the actual world (the belief is thus *safe*, to use the common parlance, though note that anti-luck epistemology generates a particularly rendering of safety). But that's not quite licensed by the theory of luck in play. Strictly speaking, a true belief would be veritically lucky if *either* one formed a false belief on the same basis in close possible worlds, *or* one failed to form a true belief on the same basis in close possible worlds (both eventualities, after all, are "non-obtainings of the target event"). Crucially, the second kind of eventuality doesn't seem to be at all incompatible with knowledge, in the way that veritic epistemic luck is. In *Epistemic Luck*, I convinced myself that this wasn't an issue because of the basis-relativity involved, in that provided we kept the basis fixed then there ought not to be case where the second eventuality can in practice come apart from the first eventuality. But still it bugged me, and bugged me even more once I started to think of cases that would exploit this loophole.

Oddly, no-one else spotted this lacuna, though I expected someone to. This meant that I was grappling with a problem for my view that only I recognised as a problem, which is rather unusual. More recently, I've come to realise that there was a second unmotivated aspect of my anti-luck epistemology, which was the way that I dealt with safe beliefs in necessary propositions. Necessary propositions are thought to be a problem for safety-style views because if one happens on a belief in a necessary proposition, no matter how epistemically suspect the basis for belief is, then of course there are no possible worlds, close or otherwise, where one forms the belief on the same basis and ends up with a false belief (so the belief is safe by default). My solution to this, which I thought fell-out of my anti-luck epistemology

methodology, was to insist that what we are interested in is the doxastic output of a certain actual basis for belief. So while this might lead to the belief that *p* in the actual world, we shouldn't just focus on close possible worlds where the subject continues to believe that *p* on the same basis, but rather close possible worlds where that same basis leads to belief (whether *p* or otherwise). In this way, we can explain why, for example, tossing a coin is not a route to safe belief that 2+2=4, for although there is no close possible world where one falsely believes that proposition on the same basis, there are close possible worlds where the same basis leads to false belief (e.g., that 2+2=5). Such a manoeuvre seems very sensible, but why is it motivated by the modal account of luck? That is far from obvious. After all, the modal account of luck just talks about the non-obtaining of the target event, and doesn't place any further restrictions on what counts as a lucky event.

Focussing on risk, specifically, helps us to resolve these problems, and in a principled way. It also helps us to better deal with some other challenges to the view. There are a number of key differences between luck and risk. Here are three. First, luck has a bivalence that risk tends to lack, in that luck can be both good and bad, but risk is usually (though not always) bad. Second, when we make evaluations of risk we pick out a specific risk event. In boarding a plane, for example, the risk event might be crashing, or hitting turbulence, or having a panic attack mid-flight, and so on. What risk event we focus on determines the level of risk in play (perhaps the panic attack is high risk, since this is an event which is modally close, but the plane crashing is low risk, since this is an event that isn't modally close). Luck lacks this feature, in that ascriptions of luck usually just get us to focus on what is happening with regard to the target actual event in the modal environment. Finally, third, luck is essentially backwards-looking, while risk is essentially forwards-looking. Suppose you narrowly avoid being hit by a sniper's bullet. Looking backwards, you will say that you were lucky to be alive. Looking forwards, you will say that you were at high risk of being hit by the bullet.

The most significant of these differences when it comes to resolving the lacunas in anti-luck epistemology is the second. The point is that we need a modal account of risk, one that is applied, in the epistemic case, such that the target risk event is the formation of false belief on the same basis. With that in place, we can now explain why mere non-belief in the target proposition on the same basis wouldn't count as knowledge-undermining epistemic risk, since it doesn't concern the relevant risk event. Moreover, we can motivate the manoeuvre introduced to deal with safe belief in necessary propositions. For notice that our natural formulation of the risk event does not

trade specifically on the proposition actually believed, but rather concerns any belief that results from that same basis.

The move to an anti-risk epistemology also has other advantages. Recently it has become fashionable to argue that lucky knowledge might not be so bad – perhaps we do have some *bona fide* knowledge that is subject to veritic epistemic luck? I don't share this intuition, but I can at least see that there seems nothing obviously contradictory about lucky knowledge. Things change, however, once we put the emphasis on risk. Can we really make sense of the idea of genuine knowledge that is subject to a high degree of epistemic risk? That certainly does sound contradictory. Moreover, our account of luck and risk and how they relate to one another presents us with a diagnosis of why we have tended, post-Gettier, to focus on knowledge-undermining epistemic luck rather than epistemic risk. A key feature of Gettier-style cases, after all, is that the belief in question is granted as being true. This naturally prompts us to take a backwards-looking approach to the issue, and hence employ the language of epistemic luck. But we could just as easily have taken a forwards-looking approach and dealt instead in terms of epistemic risk. Moreover, I contend that it is the notion of epistemic risk that has the whip-hand here, in that we only care about the elimination of high levels of epistemic luck because we care about the elimination of high levels of epistemic risk, rather than *vice versa*. We were thus wrong to have fixated so much on epistemic luck, and in particular to have fixated on this notion at the expense of the more fundamental epistemological notion of epistemic risk. We thus move from an anti-luck epistemology to an anti-risk epistemology, and more generally from an account of knowledge in terms of anti-luck virtue epistemology to one that is an anti-risk virtue epistemology.

Incidentally, I think the modal account of risk that I defend has lots to offer to a number of philosophically important debates. For one thing, I think it is far preferable to the probabilistic account of risk that is dominant in the literature. But it can also cast light on problems that arise in a range of domains beyond epistemology, such as law (e.g., disputes about legal evidence) and aesthetics. I now have a bunch of papers on the modal account of risk and its applications, and I suspect there are many more applications that I simply haven't foreseen yet.

15. *Any other preview of forthcoming projects?*

DP: Right now I'm working on a backlog of commissioned papers that built-up when I was trying to finish *Epistemic Angst*. Once the backlog is cleared, then I'm not sure where I will go next. In fact, this is kind of excit-

ing, as it is the first time in my career where I've finished a monograph and not already been under contract for the next monograph! At one point I did have a rather nutty idea of trying to write a book that weaves all of my work in epistemology together within a single volume. The idea was to start with meta-epistemological issues (e.g., epistemological methodology, radical scepticism, etc), then turn to normative epistemology (e.g., theory of knowledge, epistemic axiology, etc), and then close with applied epistemology that put the ideas previously set out into action within a specific domain (e.g., epistemology of law, epistemology of education, etc). Half in jest, this had a working title of "Epistemic Panorama". I now think that such a book would be absurdly ambitious, not least because it would be monumentally long. In terms of other projects, I would like to write something that brings together my ideas on risk and luck in one place (this would be my first monograph that doesn't have a primarily epistemic focus). I probably also have a book in me on the epistemology of religious belief too, a topic that I keep returning to at regular intervals. I'd also like to write something on applied epistemology. Another idea I've been mulling over for some time is to do something on the meaning of life (and in particular the role of philosophical paradoxes in this regard). So I've lots of vague ideas of where I might go next, but no clear path as yet, which suits me just fine right now.

16. *Here are my three final and straightforward curiosities. The first is a brutal thought experiment. I hope you don't mind playing this game for a second. Suppose that today you can select within your philosophical production one work that would become a milestone for future generations of philosophers and philosophy students, whereas the rest of your production will be lost. (In your case, the problem is not very serious, as your books are already extremely popular and discussed all over the world.) But suppose for a moment that you can save only one article or book. Which one would you pick?*

DP: I guess it would have to be *Epistemic Angst*, since that's the book I'm most proud of, and it does constitute a philosophical labour of love that takes me right from my first engagement with philosophy to my current work. (It would be a pain for my work on anti-luck/risk virtue epistemology – and the wider work that supports this view – to be lost though, as that's the other philosophical contribution that I'm particularly proud of).

17. *Second curiosity, I bet a far easier one. Could you share with the readers of APhEx the titles of the three books that have most radically influenced*

your philosophy? And could you tell us which philosophers, mentors or scholars – among the ones you've worked with – have been crucial for your research?

DP: My writings on radical scepticism have been very much informed by two books: Barry Stroud's *The Significance of Philosophical Skepticism*, and Michael Williams's *Unnatural Doubts*. My work in epistemology in general owes an enormous debt to the incredible body of work produced by Ernie Sosa (I wouldn't want to pick a particular book, since it is the body of work as a whole that has exerted the influence). Indeed, Ernie isn't just an intellectual influence on me, due to his writing, but also a personal influence, as I think he is an exemplar for others to follow in the profession, with his modesty, his intellectual generosity, and his general virtue, intellectual and otherwise. One doesn't have to be wise to write important works on the nature of wisdom, but Ernie has the rare distinction of both being wise and making a distinctive contribution to the debate about wisdom. If only more of the leading philosophers in our profession could be more like Ernie!

18. *Finally. What would you suggest to young scholars in philosophy, who often have to balance their passion for philosophical questions with the difficulty of starting a degree in a discipline that is undervalued in many countries?*

DP: As I mentioned earlier, I count my lucky stars that I somehow made it as a professional philosopher. Even quite late in the day I thought that it couldn't happen, and I didn't really have a plan B in my back pocket if things didn't work out (except that I knew that under no circumstances would I return to Wolverhampton in that case; I think emigration would have likely been on the cards). My hurdle was class, but there are many other hurdles (disability, race, gender, ageism and so on, and of course many people face several hurdles of this kind). So it's not going to be easy for many people, if anyone. That said, for what it is worth, I think the current pessimism about careers in philosophy is not particularly evidence-based. Many philosophy departments in the UK, for example, have massively expanded in the last couple of decades, as Universities have opened their eyes to how cost-effective it can be to have a world-class philosophy department (no need for expensive labs for one thing!). I also think there are academic jobs out there that simply didn't exist before, such as posts at the intersection of academia and administration which explore such things as developing impact of research, or promoting outreach. My guess is that such posts

will proliferate in the coming years, and that philosophers will be well-placed to occupy them. (As noted above, I also think that working on impact of research is perfectly intellectually respectable thing to do. It's not for everyone, of course, but that's consistent with it being a worthwhile thing for some philosophers to engage in). Moreover, philosophers are increasingly engaging in interdisciplinary research. While obviously this doesn't suit all philosophical inquiries, when one is able to conduct projects along these lines then it obviously opens up further possibilities for employment post-PhD. Finally, one thing I tell my PhD students is that what will ultimately ensure that you make it in academia is publications in very good journals (you don't need lots; indeed, a couple in the very best journals is much better than half a dozen all in weaker journals). If you can publish well, then you have a decent shot at an academic career, and given that being paid to do philosophy is a marvellous thing, a decent shot is probably enough to give it a go, if that's what you really want to do.

References

Carter J. A., Kallestrup J., Palermos S. O., Pritchard D., 2014, «Varieties of Externalism», *Philosophical Issues* 24, 1, pp. 63-109.

Carter J. A., Clark A., Kallestrup J., Palermos S. O., Pritchard D. (eds.), forthcoming, *Extended Epistemology*, Oxford, Oxford University Press.

Carter J. A., Clark A., Kallestrup J., Palermos S. O., Pritchard D. (eds.), forthcoming, *Socially Extended Epistemology*, Oxford, Oxford University Press.

Haddock A., Millar A., Pritchard D., 2010, *The Nature and Value of Knowledge*, Oxford, Oxford University Press.

Kallestrup J., Pritchard D., 2014, «Virtue Epistemology and Epistemic Twin Earth», *European Journal of Philosophy*, 22, 3, pp. 335-357.

Pritchard D., 2005a, *Epistemic Luck*, Oxford, Oxford University Press.

Pritchard D., 2005b, «The Structure of Sceptical Arguments», *Philosophical Quarterly*, 55, pp. 37-52.

Pritchard D., 2010a, «Relevant Alternatives, Perceptual Knowledge and Discrimination», *Noûs* 44, pp. 245-268.

Pritchard D., 2010b, «Cognitive Ability and the Extended Cognition Thesis», *Synthese* 175, pp. 133-151.

Pritchard D., 2012, *Epistemological Disjunctivism*, Oxford: Oxford University Press.

Pritchard D., 2015, *Epistemic Angst: Radical Skepticism and the Groundlessness of Our Believing*, Princeton (NJ), Princeton University Press.

Pritchard D., forthcoming, «Epistemic Risk», *Journal of Philosophy*.

Stroud B., 1984, *The Significance of Philosophical Scepticism*, Oxford, Oxford University Press.

Williams M., 1991, *Unnatural Doubts*, Princeton (NJ), Princeton University Press.

Wittgenstein L., 1969, *On Certainty*, (eds.) G. E. M. Anscombe, R. Rhees, (tr.) G. E. M. Anscombe, Oxford, Blackwell.