

The Evil Demon argument as based on closure plus meta-coherence

Jean Baptiste Guillon¹ 

Received: 18 February 2016 / Accepted: 2 September 2016 / Published online: 16 September 2016
© Springer Science+Business Media Dordrecht 2016

Abstract Descartes’s Evil Demon argument has been the subject of many reconstructions in recent analytic debates. Some have proposed a reconstruction with a principle of Infallibility, others with a principle of Closure of Knowledge, others with more original principles. In this paper, I propose a new reconstruction, which relies on the combination of two principles, namely the Meta-Coherence principle (defended by Huemer) and the principle of Closure of Justification (best defended by Hawthorne). I argue that the argument construed in this way is the best interpretation of what is really at play in the Evil Demon intuition, and also that this argument is dialectically much stronger than previous reconstructions. If this is right, then the “Closure plus Meta-Coherence” argument is what anti-sceptics should really be attacking.

Keywords Evil Demon argument · Scepticism · Closure · Meta-coherence

In the history of philosophical scepticism, Descartes is famous for enriching the repertoire of sceptical arguments with the invention of his notorious Evil Demon argument (EDA). Among the forms of sceptical arguments presented in contemporary classifications, the Ancient Sceptics had versions of the Regress Argument and of the Criterion Argument. But Descartes’ EDA is arguably the ancestor of what is frequently regarded

✉ Jean Baptiste Guillon
guillonjeanbaptiste@gmail.com

¹ Collège de France chaire de Métaphysique et Philosophie de la Connaissance, 11, Place Marcelin Berthelot, 75005 Paris, France

as *the* most important form of sceptical argument, namely the “argument from sceptical possibilities” (SP).¹

Descartes’ argument starts with the construction of a “sceptical hypothesis”, or rather a “sceptical scenario”: his own scenario is one in which I (the epistemic subject reflecting on the argument) am in fact a pure soul, manipulated by an Evil Demon to have various illusory experiences as of an external world (including an illusion of my body).² Descartes also discussed another sceptical scenario, the Dream scenario: in this scenario, I am in fact sleeping, and all my impressions are part of some (very realistic) dream. I will have occasions to discuss some differences between the two kinds of scenarios, but for the time being, we can treat them as equivalent.

The essential step in the argument is to notice that (1) the scenario thus described is, somehow, “a possibility”—it will have to be discussed later whether this is a logical possibility, a metaphysical possibility, an epistemic possibility, or just a kind of conceivability; but at least in very general terms, the argument builds on the intuition that the scenario has some form of possibility. Furthermore, it is evident that (2) the scenario is incompatible with a large amount of what I ordinarily think I know. From these two intuitions, Descartes concludes to the rejection (or doubt) of the corresponding ordinary beliefs, on the ground that they are not justified (or maybe that they are not knowledge).

As (Chisholm 1989, p. 3) has noted, the argument so far reconstructed is invalid: from the fact that the ED scenario is “possible”, it does not logically follow that I am not justified to believe (or don’t know) that I have hands. Some additional premise, or implicit principle, must be at play there in Descartes’ (or the sceptic’s) mind. What is the missing principle? Answering this question has proved to be rather difficult. Of course, if the initial premise were that the sceptical scenario is *true*, then the sceptical conclusion would follow directly from the fact that knowledge is factive: if I am a pure soul, then of course, I don’t “know” that I have hands. But no reasonable argument could start with the *assumption* that the ED scenario is *true*. The premise can only be that the scenario is (somehow) *possible*. And from this premise, it is not so trivial to find the epistemic principle that will raise a threat to my knowing (or justifiably believing) that I have hands.

In contemporary epistemology, several principles have been put forward by different authors, in order to fill in this gap. One principle that has been adduced is the Infallibility principle, which claims that knowledge is logically incompatible with the presence of any error-possibility. But the construal of the sceptical reasoning that has been most discussed in the late XXth century is one based on the Closure principle (or on *some form of* Closure principle). Between the construals based on Infallibility and those

¹ Different authors give this form of argument different names. (Chisholm 1989, p. 2) calls sceptics of that kind the “perhaps-you-are-wrong sceptics”. (Schiffer 2009) calls it “the EPH argument template”, where “H” stands for the sceptical *Hypothesis*. Some use the phrase “Evil Demon Argument” to designate the general form of the argument, even if the scenario is not that of Descartes. Some refer rather to Putnam (1981) alternative scenario and call it the “Brains in vats argument”.

² Putnam (1981) famously proposed an alternative scenario: one in which I am in fact a brain in a vat, manipulated by a scientist to have various illusory sensory inputs as of an external world (including an illusion of my limbs). This alternative scenario is more appropriate to contemporary debates for reasons that will have no essential bearing on the present discussion. So I will stick with the Cartesian scenario.

based on Closure, some less classical (and often more complex) interpretations have been provided, among which Stroud's (1984, p. 29) principle,³ Brueckner's (1994) Underdetermination principle, Pryor's (2000) SPJ principle, etc.

When considering the variety of these construals, and the epistemic principles they are based on, one can start to wonder whether there is such a thing as “the” argument from sceptical possibilities. Wouldn't it be clearer to just acknowledge that there are in fact several completely different arguments, the Infallibility-based argument, the Closure-based argument, the Underdetermination-based argument, etc.? This sounds like a good clarificatory procedure, but it is interesting to notice that many contemporary epistemologists have *not* followed this procedure. Many have maintained the project of “determining just which epistemic principles are required by *the* argument” (Brueckner 1994, p. 827 my emphasis), rejecting such and such a construal (for example the Closure-based construal) because it is not “the most effective formulation of *the* sceptic's reasoning” (Pryor 2000, p. 522 my emphasis). But what is “the most effective formulation” of the sceptic's reasoning? What is the project of those philosophers who privilege this formulation over that formulation? There might be an ambiguity here between two projects: the “strongest argument” project, and the “initial intuition” project.⁴

According to the “strongest argument” project, we should concentrate on the epistemic principle that generates the dialectically strongest argument. Pritchard, for instance, is clearly pursuing this project when he writes: “Given that both sceptical arguments [the infallibility-based and the closure-based] generate the same intellectually devastating sceptical conclusion, there are obvious advantages to opting for the closure-based formulation of the argument that is motivated by a logically weaker epistemic principle.” (Pritchard 2005a, p. 28)

According to the “initial intuition” project, sceptical scenarios like the ED scenario have a natural tendency, when they are consciously entertained, to produce the intuition that our ordinary beliefs don't count as knowledge. This initial intuition must rely on some implicit epistemic principle which is ingrained in our common way of thinking. And the philosopher (sceptic or not) is interested in finding out what this implicit principle is, and in formulating it explicitly. Jim Pryor seems to have this project in mind when he writes: “another problem with the [closure-based] argument is that it does not generalize in the same ways that the sceptic's reasoning intuitively seems to generalize” (Pryor 2000, p. 522).⁵

³ Namely: “if somebody knows something, *p*, he must know the falsity of all those things incompatible with his knowing that *p*”. Stroud constructs this principle after several refinements starting from a common Closure principle.

⁴ I thank Julien Dutant for drawing my attention to this distinction.

⁵ I am not trying to suggest that Pryor is a consistent “initial intuitionist” and Pritchard a consistent “strongest argumentist”. On the contrary, it seems to me very telling that both authors (and others as well) tend to mix the two projects. For instance, (Pritchard 2005a, p. 107) favours the underdetermination-based argument over the closure-based argument because he contests the latter's “importance to capturing the sceptical challenge” or “capturing what is at issue in the sceptical challenge”: this is clearly in the spirit of the “initial intuition” project. And Pryor's (2000, p. 522) first reason to doubt “that the [closure-based] argument is the most effective formulation” is that “some philosophers refuse to allow the sceptic to use claims like ‘I can't know I'm not being deceived’ as *premises* in his reasoning”; this is clearly a “strongest argument” kind of motivation.

In principle, the two projects may come apart. We could imagine, for instance, the following situation: that the initial intuition generated by ED scenarios is based on principle P1 (e.g. Infallibility), and that the argument construed in this way is a weak one, while it is possible to produce a sophisticated argument based on a completely different principle P2 (e.g. Closure), which is a much better argument though no one (perhaps not even Descartes) had initially thought about it.

Now, even though the two projects *may* come apart, we can perhaps understand why they are not always distinguished. The reason has to do with the intention in which such arguments are reviewed. An author who works on the argument(s) from SP is normally trying to assess the rationality or truth of Scepticism in general, and more often than not, such authors are trying to *refute* Scepticism. What does it take to refute Scepticism? Well, on the one hand it seems you have to refute *the* argument on which it is grounded. If sceptical possibilities are often conjured up in support of scepticism, then it seems natural to try and refute *the* argument which is thereby usually used (that is: the natural intuition that usually turns people from the consideration of SP to the sceptical conclusion). On the other hand, if you want to defeat the Sceptic, it seems you'd better refute *all* his arguments, and especially the *best* ones. Does that give you a second and different task? Not so clearly. Suppose we used to consider the initial SP intuition as based on principle P1, and P1 is easily refutable, but there comes another philosopher who manages to devise a much better argument from SP based on principle P2. Should we consider in that case that the “initial intuition” and the “best argument” come apart? Or should we not rather revise our interpretation of what the initial intuition was? Should we not come to believe that, in fact, our initial intuition was probably based on P2? The principle of charity seems to favour the latter solution. Besides charity, I think we also have an independent reason to suppose that the best argument from SP is probably the one that people implicitly have in mind when initially encountering SP. The reason is that a sceptical argument, because it's trying to overturn some very strong (common sense) beliefs *has to* start from premises and principles that are themselves very strong and deeply ingrained in our common thinking (e.g. stronger than the belief that I have hands). Now, it is clear that the principle, whatever it is, that generates the initial intuition that SP are troublesome for our knowledge is deeply ingrained; and if anyone should present us explicitly a deeply ingrained principle which validly goes from SP to scepticism, there is good reason to suppose that this is *precisely* the one we initially had in mind.

In any case, whether or not the “initial intuition” project and the “strongest argument” project converge, I will defend that the construal of the EDA proposed in this paper satisfies both projects. That is: I will give reasons to believe that this construal corresponds to our initial intuition about SP, and I will also try to show that this version of the argument is stronger than the other construals proposed in contemporary debates. My general intention, though, in elaborating this version of the argument is to prepare the terrain for an assessment of Scepticism, i.e. to determine clearly what the anti-sceptic has to tackle if he wants to claim victory. So I would rest content if my reader is eventually convinced that the proposed argument is *an interesting and challenging version, which the anti-sceptic has to tackle*, even though this reader were not convinced that it corresponds to our “initial intuition” or that it is “the strongest version” of the argument from SP.

In the first section, I will briefly review the most important shortcomings of the two main construals of the EDA, namely the infallibility-based argument and the closure-based argument. I will also consider briefly the underdetermination-based argument and its relation to the closure-based argument. In the second section, I will follow (Pritchard 2005a) in his diagnosis that the heart of the EDA lies in a kind of “reflective luck”, though this idea of reflective luck will lead me to a principle different from the one used by Pritchard, namely the principle of Meta-Coherence (MC). In the third section, I will give a first formal presentation of my construal, according to which the EDA is based on the combined use of a principle of Meta-Coherence *and* a principle of Closure (for justification). In Sect. 4, I will refine the formulations of Meta-Coherence and Closure, in order to show how deeply ingrained and dialectically strong those principles are. In Sect. 5, I will give a second version of my construal, based on the enhanced versions of Meta-Coherence and Closure. In Sect. 6, I will comment on scenarios like Dream, which differ a little bit from scenarios like Evil Demon or Brains in vats; but I will show that the combination of Meta-Coherence and Closure suffices nonetheless to account for the sceptical potential of such scenarios, be it in a slightly less direct way.

1 The main traditional interpretations of EDA

The two main traditional interpretations of EDA are those based on Infallibilism and on (some form of) epistemic Closure. And many reviews of the EDA start with the former, probably because it is the simpler.⁶

The Infallibility based interpretation starts with the obvious premise that the sceptical scenario (Evil Demon or Brains in vats) is “possible”, and takes this in the undeniable sense of logical, or even metaphysical, possibility. The second premise is the principle of Infallibility, namely that knowledge is incompatible with *any* possibility of error.

The Infallibility-based Interpretation of EDA

(I1) There is a metaphysically possible scenario (ED) in which my belief that I have hands (held on the basis of the same belief-producing process as is in fact its basis) would be erroneous.⁷

(I2) If it is metaphysically possible for a belief to be (held on its actual basis and) erroneous, then it doesn't constitute knowledge.

(C) My belief that I have hands doesn't constitute knowledge.

Even though this argument is very straightforward, most epistemologists nowadays have abandoned it as an interesting interpretation of the EDA. The reason is that the Infallibilist premise (I2) is highly contentious, and in fact rejected by most philoso-

⁶ This is true of (Pritchard 2005a, Chap. 1), but also of (Zalabardo 2012, Sect. 1.3).

⁷ This premise can be considered as the (intermediate) conclusion of the following syllogism: the ED scenario is metaphysically possible, if ED were true, then my belief that I have hands (held on the basis etc.) would be erroneous, therefore (I1). Thanks to an anonymous reviewer for this qualification and other remarks on the best formulation of the infallibility-based argument.

phers.⁸ Even philosophers like (Unger 1975) who accept an infallibilist constraint consider it as a constraint on some limit-case of the meaning of “knowledge”. Now, if all the sceptic could show were that ordinary beliefs lack some extreme form of absolute certainty, it wouldn’t be such a devastating thesis. The real trouble with the EDA is that, intuitively, it seems to threaten our common beliefs, even according to moderate and everyday standards. But the Infallibility-based interpretation is unable to capture this intuition: in our common and everyday practice, we do not consider beliefs as not-knowledge as soon as we discover a possibility of error. In fact, many of our everyday knowledge claims are asserted in full recognition that it is metaphysically possible that the belief be false—which suggests that we do not ordinarily have in mind anything like an infallibility constraint on knowledge. For that reason, the Infallibility-based interpretation can be considered as both dialectically weak, and incapable of capturing the intuitive force of the EDA.

A decisive progress in interpretation has been made when philosophers have discovered that a principle much more convincing than Infallibility could be sufficient to derive the sceptical conclusion from an SP. This principle was the Closure principle. The main idea was that the sceptical scenario was not only metaphysically possible, but it was also something I cannot seem to be able to rule out, or know to be false. And intuitively, since the scenario *entails* that I don’t have hands,⁹ it seems that I must be able to rule it out in order to be in a position to claim that I have hands. More formally, the argument is traditionally put as follows:

The K-Closure-based Interpretation of EDA¹⁰

(K-Closure) For all p and q , if I know that p , and if I know that $p \Rightarrow q$, then I know that q

(1K) I know that Hands $\Rightarrow \sim$ ED

(2K) I don’t know that \sim ED.

(C) Therefore, I don’t know that Hands.

Here, I have formulated the argument in terms of knowledge, but many authors have noticed that a parallel argument can be made about justification:

⁸ (Reed 2002) goes so far as saying that “Fallibilism is endorsed by virtually all contemporary epistemologists.” An anonymous reviewer remarks that Williamson’s (2000) E = K principle could render some version of the infallibilist premise (I2) true, i.e. a version of the premise which would require our belief *as based on the same evidence* to rule out all possibilities of error in order for it to be knowledge. Granted, if evidence just *is* what we know, then evidence is factive and infallible. But then, with such a notion of evidence, it is premise (I1) which would become false, for my belief that I have hands in the ED scenario would *not* be based on the same evidence. In that particular sense, then, we would have to say that my (actual) belief that I have hands (as based on its actual evidence) is in fact “infallible”. What this shows, it seems to me, is not that Williamson is in fact an “infallibilist” in any sense relevant for our present discussion, but rather that his notion of evidence is inappropriate to capture in a straightforward way the very simple idea that our senses are fallible.

⁹ The kind of entailment at play here is strict or metaphysical entailment rather than “logical entailment” strictly speaking (the fact that a pure brain doesn’t have hands doesn’t follow from any *logical constant*). This notion of entailment will be used throughout the paper. Thanks to an anonymous referee for this qualification.

¹⁰ In the semi-formalisations of arguments, I will use the double-arrow “ \Rightarrow ” to symbolize strict implication, and the single-arrow to symbolize the truth-functional conditional (or material implication).

The J-Closure-based Interpretation of EDA

- (J-Closure) For all p and q , if I am justified in believing that p , and that $p \Rightarrow q$, then I have justification for believing that q
- (1J) I am justified in believing that Hands $\Rightarrow \sim$ ED
- (2J) I don't have justification for believing that \sim ED.
- (C) Therefore, I am not justified in believing that p .

Maybe it will be useful to say briefly how the two notions of *knowledge* and *justification* differ. Even though there are important debates as to how to understand both notions, there is fairly large agreement (in this post-Gettier era) that they are not equivalent, or in other words, that knowledge is not just justification added to a true belief. As (Pritchard 2005a, p. 2) notes, if the Gettier problem has taught us anything uncontroversial, it is the “platitude” that in order to constitute knowledge, a (justified) true belief must not be true “as a matter of sheer luck”. Knowledge undeniably involves some “anti-luck” element or condition, which is essentially externalist in nature. (That's why, as Robert Audi (1998, p. 238) puts it, “pure or unrestricted internalism about knowledge” is not a plausible view.) The notion of knowledge I am personally favouring is *restricted* to this externalist “anti-luck” element.¹¹ But I don't know that what I will have to say couldn't in principle be translated with another analysis of knowledge (provided that knowledge remains a notion distinct from justification).

As for justification, I also follow Robert Audi in thinking that “there appears to be at least *some* respect in which justification is internally grounded”, which means that the most defensible views cannot be “purely or unrestricted externalism about justification”. It is certainly uncontroversial that there are certain norms that apply to beliefs which are *purely internalist* norms (e.g. the norm of non-contradiction, or the norm of non-circularity of reasons). Now, externalists about justification may be reluctant to say that the satisfaction of those (purely internalist) norms is sufficient to qualify as *justified* in the proper sense of the term. They will want to add some externalist constraint on justification. I can only say that these authors use a sense of “justification” different from mine, but I think that these authors can understand everything I will say if they replace “justification” by “rationality”, or perhaps by “blamelessly held belief”.¹² It seems to me uncontroversial that there is *some* normative epistemic status (call it justification or rationality or blamelessly held belief) which is purely defined in terms of internalist epistemic norms, i.e. norms having to do with the agent's perspective.

The two Closure-based interpretations of the EDA have been subject to many objections, and many attempts to improve on them. My purpose here is not to give a complete

¹¹ It can be defined thus: S knows that p iff

- (i) S believes that p
- (ii) p is true
- (iii) it is not a matter of sheer luck if S's belief that p is true.

¹² An externalist like Plantinga, for instance, acknowledges that one *could* use a purely internalist notion of justification, which he calls “broad justification” in (Plantinga 1993, p. 27) and “internal rationality” in (Plantinga 2000, p. 108). I tend to think that there is more to (purely internalist) rationality than just blamelessness: following (Pryor 2001), I consider that there are purely internalist aspects of epistemic appraisal that are *not* deontic but strictly epistemic. But the reader may disagree and prefer “blamelessly held belief” as a better translation for what I have to say in the remainder of the paper.

review of these discussions. What I will try to do is to select and put in the right perspective what I take to be the most telling difficulties of both versions. From this brief summary, I will draw (in Sect. 2) the same lesson that Duncan Pritchard (2005a,b) drew in his, namely that the core intuition of the EDA is a form of reflective luck, i.e. a form of luck which can be expressed *neither* in terms of purely externalist knowledge, *nor* in terms of purely internalist justification, but which combines externalist and internalist aspects.

So what is wrong (or weak) in the Closure-based arguments? To summarize, we could say that they are both weak, but for exactly opposing reasons: the K-Closure argument is strong in its premise (2K) but weak in its Closure premise, while the J-Closure argument is weak in its (2J) premise but strong in its Closure premise. Let me explain.

In the K-Closure argument, premise (2) is rather strong intuitively (though it *can* be disputed): it is very intuitive to think that my belief that I'm not being deceived cannot *depend* in any way on the *fact* that I'm not being deceived. After all, it is obvious that:

- (i) If I were not being deceived by the ED, I would believe that I'm not being deceived.
- (ii) If I *were* being deceived by the ED, I would *nonetheless* believe that I'm not being deceived.

So my belief that I'm not being deceived seems to be completely independent on the fact. This gives strong (though defeasible) support to the intuition that, even if that belief were true, it would be a matter of sheer luck.¹³ And given that knowledge clearly excludes luck, this means that the belief \sim ED cannot constitute knowledge (2K). Brueckner (1994) thought that this premise was in need of further support, and appealed to the principle of Underdetermination to support it. I will come back to Underdetermination in more details in the discussion of (2J), but as far as (2K) is concerned, the sceptic needs no other principle in support than the anti-luck “platitude” that knowledge excludes luck.¹⁴

What about the Closure premise? Well, as soon as we become aware that knowledge is essentially an externalist notion, involving in particular the kind of counterfactual “truth-tracking” (which is absent in (ii) for instance), we can easily become suspicious about the Closure of Knowledge.¹⁵ This is the lesson that Dretske (1970) and Nozick (1981, pp. 204–211) have famously taught us. Counterfactual tracking can easily be shown *not* to be closed under (known) entailment. Nozick's famous example is that of the proposition “this is a zebra”, which entails “this is not a cleverly disguised mule”: in normal circumstances, my belief that “this is a zebra” counterfactually depends

¹³ I emphasize here that my present point is only one of dialectical or intuitive strength. There are of course ways to resist the intuition that the belief \sim ED is lucky: probably the most important way to do so is to say that “safety” (not “sensitivity”) is what we need to avoid epistemic luck. As a matter of fact, this is a solution I personally endorse; but even if safety in fact preserves the belief \sim ED from being lucky, it should still be acknowledged, dialectically, that there is some initial intuitive tendency to find it lucky.

¹⁴ More on this externalist notion of luck in Sect. 2.

¹⁵ Again, one *may* argue that knowledge is an externalist notion and *yet* is closed under known entailment—this would be the case if one adopts of “safety” understanding of knowledge. My point here is only dialectical: safety (or other notions of closed externalist knowledge), even if I personally hold it to be true, is far from being indisputable, and for that reason, the closure of externalist knowledge is a contentious principle to take as a *premise*.

on the fact that there is a zebra in front of me; and yet, my belief that “this is not a cleverly disguised mule” does *not* counterfactually depend on its content in the same way. Drestke and Nozick were happy to conclude that “the sceptic reasoning” relies on a principle which is demonstrably false (the K-Closure principle). But many authors have put into doubt that they had really captured the essential intuition of the sceptical reasoning: after all, is it not intuitive to think that the original sceptical argument relies on something essentially internalist? *This* aspect is not captured by Dretske’s and Nozick’s strictly externalist discussion. This is, at least, the reason why (Pritchard 2005a) and (Engel 2007), among others, think that the K-Closure interpretation doesn’t capture the gist of the sceptical argument.

Let’s turn then to the J-Closure interpretation which will be better suited to capture the internalist aspect. Turning to justification, the Closure premise becomes much stronger. Indeed, when we’re dealing with internal coherence of a system of beliefs, it seems very plausible to say that an agent is irrational if he doesn’t accept the (acknowledged) logical consequences of what he believes. The intuition of Closure seems to be initially and essentially suited for internalist justification or rationality.

What about premise (2J)? Unfortunately, this premise becomes much more dubious with the notion of justification than it was with the notion of knowledge, as Jim Pryor has clearly expressed it:

(2J*) You can’t be justified in believing you’re not being deceived by an Evil Demon right now.¹⁶

Why on earth should we accept this premise? Even if we can’t know for sure whether or not we’re being deceived by an Evil Demon, isn’t it at least reasonable to assume that we’re not being so deceived, absent any evidence to the contrary? There are important disanalogies between justified belief and knowledge which come into play here, and make (2J*) much less plausible than (2K*).

(Pryor 2000, p. 523)

Pryor doesn’t exclude totally that one might have an *argument* in favour of (2J): what he says is that (2J), unlike (2K), cannot be considered as *intrinsically* plausible, and therefore proposed as a premise. And his diagnosis seems to me to be plainly right.

Therefore, the J-Closure based interpretation, the interpretation which appeals *only* to J-Closure and the direct intuition of (2J), is very weak dialectically, and is probably not a good candidate to capture our initial intuitions about sceptical scenarios: we couldn’t possibly be thinking that “I cannot be *justified* (or rational) in ruling them out”, because it seems so *prima facie* obvious that we *are* justified (or rational) in ruling them out.

The rest of this paper will in fact try to formulate a second argument in favour of proposition (2J), by appeal to the epistemic principle of Meta-Coherence. It should be noticed here that another principle has been appealed to in order to defend (2J),

¹⁶ In this quote, I have adapted the names of the propositions to fit with our present discussion. The proposition which Pryor directly discusses (here (2J*)) is in fact the modal strengthening of our (2J)—hence the * in the name. (Correspondingly, (2K*) stands for the modal strengthening of (2K), which Pryor discusses.) What is important is of course that Pryor’s discussion *commits* him to our (2J), and to the view that (2J) is much more plausible than (2K).

namely Brueckner's principle of Underdetermination.¹⁷ We have seen earlier that the K-Closure-based interpretation did not need to rely on such a principle (because (2K) was in itself sufficiently supported), but the situation is different for the J-Closure based interpretation, since (2J) *is* in need of some argumentative support.¹⁸ How would the Underdetermination Principle (UP) undermine justification for believing that \sim ED? Here is the argument:

- (\sim F) My evidence for believing \sim ED does not favour \sim ED over ED.
- (UP) For all p and q , if my evidence for believing p does not favour p over some incompatible hypothesis q , then I lack justification for believing q .
- (2J) Therefore, I don't have justification for believing that \sim ED.¹⁹

Brueckner has shown that if we accept this principle of Underdetermination in order to complement the J-Closure based argument (by defending its premise (2J)), we have in fact the possibility to make a more direct argument to the sceptical conclusion, one which doesn't appeal to Closure at all:

- The Underdetermination-based Interpretation of EDA
- (\sim F') My evidence for believing Hands does not favour Hands over ED.
- (UP) For all p and q , if my evidence for believing p does not favour p over some incompatible hypothesis q , then I lack justification for believing q .
- (C) Therefore, I don't have justification for believing that Hands.

Brueckner himself thinks that the Underdetermination-based argument and the Closure-based argument are in fact two logically equivalent formulations of the very same argument. I don't think this is right: following (Cohen 1998; Pritchard 2005b; McCain 2013), I accept the diagnosis that (UP) is entailed by (J-Closure) but does not entail it. For that reason, (UP) could provide an independent (and dialectically stronger) route to scepticism... if at least the (\sim F') premise were strongly supported.

This is where I disagree with Pritchard (2005a, pp. 206–211), who seems to think that (\sim F') is indisputable and uncontroversial. Like (Goldman 2007), I consider (\sim F') to be highly disputable. In a nutshell, if we adopt a purely internalist perspective on justification, it seems very plausible to say that my experience as of having hands provides a *prima facie* justification for the propositions Hands and of course no *prima facie* justification for the proposition ED. In *that* sense then (the internalist sense), it seems obvious that my evidence (my experience) “favours” Hands over the hypothesis ED (by providing “more” *prima facie* support for the former than for the latter). If Pritchard thinks it indisputable that my evidence does *not* “favour” Hands over ED, it must be because he means “favour” in some *externalist* sense: my experience would be the same whether Hands or ED be true *as a matter of externalist fact*. In other words, the “non-favouring” intuition is strong only inasmuch it is a way of expressing the intuition of externalist luck which we have already seen as premise (2K). But it is not at all clear (as far as the Underdetermination principle tells us) how this externalist

¹⁷ Thanks to an anonymous reviewer for inviting me to discuss the Underdetermination argument here.

¹⁸ In fact, Brueckner himself presents the problem of Underdetermination as first undermining *justification* and undermining knowledge only in a derivative way, *because* knowledge requires justification.

¹⁹ See Brueckner (1994, pp. 830–832).

problem has effects on internalist *justification*. The connexion between the externalist level and the internalist level seems to be missing in the Underdetermination-based argument.

If we come back to the discussion of the two Closure-based interpretations, the general conclusion is in fact the same: that a connexion seems to be missing between the internalist level and the externalist level. We could summarize the contrastive discussion of the two Closure-based arguments in the following chart:

	With Knowledge (externalist)	With Justification (internalist)
Premise 2	Very strong	Very weak
Closure premise	Very weak	Very strong

In this chart, we can see that there are in fact two strong premises on which the sceptical argument could run:

(J-Closure) For all p and q , if Jp and $J(p \Rightarrow q)$ then Jq
 (2K) $\sim K \sim ED$

The problem is that these two premises use different notions (one internalist, one externalist) and therefore cannot be combined in any trivial way to construct a valid argument. Should we say therefore that our initial intuition that grounds the sceptical argument is based on a confusion between these two different notions, between the externalist and the internalist level? That might be a diagnosis. But in fact, I think that wouldn't be charitable, because it is possible to think that the sceptical argument relies on the legitimate intuition of some "binding" principle between K and J , between the externalist and internalist levels.

2 The central intuition of the EDA: reflective luck

Duncan Pritchard (2005a, chap. 6) has put forward an illuminating distinction between two kinds of epistemic luck, in order to pin down exactly what was problematic about sceptical possibilities. The first kind of luck, which he calls "veritic epistemic luck" is the one we already have encountered: a belief is veritically lucky when it has no external dependence whatsoever on the truth of its content. This kind of (purely externalist) luck is typical of Gettier cases. As is clear from the Gettier discussion, this kind of luck prevents beliefs from constituting knowledge, but it doesn't prevent them from being justified (or rational).

But Pritchard draws our attention to another kind of epistemic luck which has a bearing on the justification of the belief:

Even with veritic epistemic luck eliminated, another form of epistemic luck remains that is potentially just as epistemologically significant, if not more so. This type of epistemic luck concerns the manner in which, *from that agent's reflective position*, it is a matter of luck that her belief is true.

(Pritchard 2005a, p. 174)

Pritchard calls this second kind of luck “reflective epistemic luck”. The simplest and most famous case of reflective luck is Foley’s (1987) chicken-sexer scenario: the chicken-sexer has the ability to discern reliably between male and female chicken. The beliefs he forms on such matters are *not* veritically lucky (they track the truth). But the chicken-sexer has no reflective awareness as to how he forms such beliefs, and in particular, he isn’t aware that those beliefs are reliable. In this scenario, we can say that “*from the chicken-sexer’s own reflective position*, it is a matter of luck that his belief is true”, even though from an externalist perspective it is *not* a matter of (veritic) luck that his belief is true.

To put it in other terms, we could say the chicken-sexer’s belief constitutes (unreflective or externalist) *knowledge*, but that he doesn’t have an (internalist) *justification* (or a rational support) to believe that it is knowledge. We could formalize the situation thus:

$$Kp \& \sim JKp$$

According to Pritchard, this notion of reflective luck is at the heart of the sceptical reasoning in EDA, at least in its most convincing construal.²⁰ The sceptical scenario may *not* render our beliefs veritically lucky: maybe our common perceptual beliefs are tracking the truth, and maybe even the beliefs in the negation of SP are reliable in the sense that they couldn’t easily have been wrong. But what the sceptical scenario still seems to show is that our beliefs are *reflectively* lucky: even if our beliefs are reliable (as a matter of external fact), it seems to remain (epistemically) possible, from our internal perspective, that they are unreliable. And *this* fact (reflective luck) seems to be troublesome for the internal justification of our beliefs. In order to be fully justified in having such and such a belief, it seems that we should have a justification (or rational support) to consider them as reliable (or knowledge).

The reasoning of the last sentence relies, in fact, on an epistemic principle which has been recently studied and defended by Huemer (2011), who calls it the “principle of Meta-Coherence”. We will see different versions of this principle later, but let us start with the simpler one:

Principle of Meta-Coherence (MC)

For all p , if I am justified in believing that p , then I must be justified in believing that I know that p .

For all p , $Jp \rightarrow JKp$

(or conversely)

For all p , if I am not justified in believing that I know that p , then I am not justified in believing that p .

²⁰ This is why he considers as insufficient even the Neo-Moorean response to scepticism, which tries to show that our beliefs in the negation of SP can be “reliable” after all, if not in the sense of *sensitivity* (or tracking), at least in the sense of *safety* (i.e. the belief could not easily have been wrong, it is true in all or most *nearby* possible worlds). This is insufficient because *safety* remains a purely externalist property: a belief could be safe while I had no idea whatsoever that it is; and in such a case, my belief would remain *reflectively* lucky.

In Pritchard’s terminology, we could say that MC is the epistemic constraint stating that reflective luck must be ruled out for a belief to be (internally) justified. And in my view, this principle is the best way to capture Pritchard’s problem of reflective luck.²¹

Now, how could we derive a sceptical argument from this problem of epistemic luck? Well, Pritchard seems to say that our ordinary beliefs (that I have hands) are themselves rendered reflectively lucky by the sceptical possibility. If this is true, then we could construct an interpretation of the sceptical argument which uses only the principle MC.

The Meta-Coherence-based interpretation of EDA

(MC) For all p , if I am justified in believing that p , then I am justified in believing that I know that p

(MC2) I am not justified in believing that I *know* I have hands (that this belief is reliable).

(C) Therefore, I am not justified in believing that I have hands.

or in formal terms:

(MC) For all p , $Jp \rightarrow JKp$

(MC2) $\sim JKh$

(C) $\sim Jh$

What should we think of this MC-based interpretation? Does it capture the intuitive force of the EDA? And is it cogent? The answer to both questions is probably no. The main problem here is premise (MC2): I am not justified in believing that my belief of having hands constitutes knowledge or is reliable. Taken as a premise, this proposition is exceedingly weak: there is no need to be a committed Moorean to acknowledge that *prima facie* and ordinarily we naturally tend to consider our perceptual beliefs as *knowledge*, and we seem to be justified in doing so. So if anyone wants to build on a proposition like $\sim JKh$, he should provide a strong *argument* for it, and not take it as a premise. Now, of course, the implicit idea in our reconstruction was that *the consideration of sceptical possibilities* leads us to consider that we cannot (anymore) consider our perceptual beliefs as reliable or as knowledge. But then the argument would have to involve a second epistemic principle explaining us why and when perceptual beliefs become reflectively lucky. We could try and write down candidates for such a principle, but I don’t think this route is very promising; another route is much more promising in interpreting how reflective luck can be involved in the EDA.

If it is far from clear that ordinary beliefs are reflectively lucky, there is however another belief in the neighbourhood of EDA which seems *prima facie* and very intuitively reflectively lucky: it is the belief in the negation of ED. It seems very probable that:

²¹ Even though Pritchard himself develops his idea with another epistemic principle, namely Underdetermination. In defense of the MC interpretation of reflective luck, we can remark that the second-order form of the MC principle (not KKp but JKp) immediately corresponds to the “reflective” aspect of reflective luck; it also captures the internalist aspect of this notion. The principle of Underdetermination is not so clearly “second-order” or reflective, and as I noted earlier, it doesn’t show clearly the relation between the externalist level and the internalist level.

(1) I am not justified in believing that my belief \sim ED is reliable or constitutes knowledge
 \sim JK \sim ED

Why is this proposition very cogent intuitively? We can understand this easily if we remember the intuitive force of premise (2K) in Sect. 1:

(2K) I don't know that \sim ED
 \sim K \sim ED

We've said that this premise was very intuitive because it lies in the nature of sceptical possibilities, by definition, to be "undetectable", and therefore it appears that a belief denying those possibilities can be at best (veritically) lucky even if it is true. This intuition (that \sim ED cannot be *known*) is certainly an *a priori* intuition which is accessible to any competent reflective thinker. But this entails that any competent reflective thinker will be in a position to have *justification* for believing \sim K \sim ED. (The rule of inference here applied could be called the "rule of competent *a priori* justification".) In other words, for any competent reflective thinker, it seems that we have:

(1*) J \sim K \sim ED

Now, this is not exactly premise (1) above, because of the scope of the negation. But one thing seems clear: if it is problematic for my belief \sim ED that I *lack* justification to consider it as knowledge (which is what (1) states), it will be all the more problematic that I *have positive justification* to think that it is *not knowledge* (which is what (1*) states).²² If I have positive justification to consider a certain belief as unreliable (not knowledge), then it seems obvious that my justification for this belief is defeated. Here, we are in fact encountering a second (and weaker) version of the principle of meta-coherence²³:

The principle of Negative Meta-Coherence (MC-)

For all p , if I am justified in believing that I don't know that p , then I am not justified in believing that p .

(or conversely)

For all p , if I am justified in believing that p , then I am not justified in believing that I don't know that p .

For all p , $Jp \rightarrow \sim J \sim Kp$

With premise (1*)—which is in fact derived from (2K) as we've seen—together with the principle of negative meta-coherence, we now have the resources to get the argument off the ground. This is what I will do in the next section.

²² In fact, as an anonymous reviewer suggests, it seems possible to show that J \sim K \sim ED is strictly stronger than (entails) \sim JK \sim ED if we assume that our justified beliefs are consistent. This assumption amounts to accepting axiom D for justification: for all p , if Jp then $\sim J \sim p$

²³ I will say more in 4 about the two versions of the principle.

3 The interpretation of EDA based on closure plus meta-coherence

In the preceding sections, we have encountered a few principles and premises, which we have considered very intuitive, but that have not yet been put together to build a valid sceptical argument. Typically, the cogent premises were encountered in various arguments in which they were combined with far less cogent ones. Let me remind the three main premises.

(2K) My belief that \sim ED doesn't constitute knowledge.
 \sim K \sim ED

And from the fact that premise (2K) is *a priori* (together with the “rule of competent *a priori* justification”) we have seen that we could derive for any competent reflective thinker another proposition:

(1*) I am justified in believing that my belief that \sim ED doesn't constitute knowledge.
 J \sim K \sim ED

Then we have seen the intuitive force of the principle of Closure *for justification*:

(J-Closure) For all p and q , if I am justified in believing that p , and that $p \Rightarrow q$, then I have justification for believing that q
 For all p and q , $(Jp \& J(p \Rightarrow q)) \rightarrow Jq$

Finally, we have seen the intuitive force of the principle of (negative) Meta-Coherence

(MC-) For all p , if I am justified in believing that p , then I am not justified in believing that I don't know that p .
 For all p , $Jp \rightarrow \sim J \sim Kp$

Now, when we combine all those principles, taken in various arguments, it happens that we *can* construct a valid sceptical argument. We can see this in two steps.

Proposition (1*), together with the principle of meta-coherence, directly entails the intermediate conclusion that I am not justified in believing the negation of ED:

First step: Meta-Coherence

(1*) I am justified in believing that my belief that \sim ED doesn't constitute knowledge.

(MC-) For all p , if I am justified in believing that my belief that p doesn't constitute knowledge, then I am not justified in believing p .

(C1) I am not justified in believing \sim ED.

(1*) J \sim K \sim ED

(MC-) For all p , $Jp \rightarrow \sim J \sim Kp$

(C1) $\sim J \sim$ ED

Now it is interesting to notice that conclusion (C1) is in fact the proposition (2J), which the argument from J-Closure wanted to use as a *premise*. We've seen that, as a premise, it was very weak, but now that we have an *argument* for it (from powerful premises), we can use J-Closure as a second step to derive the sceptical conclusion:

Second step: J-Closure

(C1) I am not justified in believing \sim ED.

(2) I am justified in believing that Hands entails \sim ED.²⁴

(J-Closure) For all p and q , if I am justified in believing that p , and that $p \Rightarrow q$, then I have justification for believing that q

(C2) I am not justified in believing Hands.

(C1) \sim J \sim ED

(2) J ($h \Rightarrow \sim$ ED)

(J-Closure) For all p and q , ($Jp \& J(p \Rightarrow q)$) $\rightarrow Jq$

(C2) \sim Jh

We could summarize the argument with the following (simplified) version:

The interpretation of EDA based on Closure plus Meta-Coherence²⁵

(1*) J \sim K \sim ED

(MC-) J \sim ED $\rightarrow \sim$ J \sim K \sim ED

\therefore (C1) \sim J \sim ED (by *modus tollens*)

(J-Clos.) Jh \rightarrow J \sim ED

\therefore (C2) \sim Jh (by *modus tollens*)

In this argument, it is important to understand the nature of the two steps that allow us to move from a second order claim about an anti-sceptical proposition (J \sim K \sim ED) to a first order claim about an everyday proposition (\sim Jh). The first step (meta-coherence) is a step from the second order to the first order; the second step (closure) is a step from the anti-sceptical content (\sim ED) to the everyday content (h).

This distinction of the two steps can help us see what is unsatisfactory in the traditional externalist reconstruction and response to the EDA (by Dretske and Nozick), namely the diagnosis according to which the EDA wrongly presupposes Closure of Knowledge. Dretske and Nozick's interpretation of the EDA consists in applying the move from the anti-sceptical proposition to the everyday proposition directly at the level of knowledge: the first conclusion that the sceptic draws (in this interpretation) is therefore that we lack knowledge of everyday propositions: \sim Kh. This conclusion could then be used (though Dretske and Nozick do not apply this move explicitly) as a threat on the *justification* (or rationality) of everyday perceptual beliefs (\sim Jh), which seems worse than just saying that they are not "knowledge". So, in a sense, we could also use K-Closure in order to produce a two-step sceptical argument based on Closure and Meta-Coherence:

²⁴ This premise, which is obviously necessary for the validity of the argument, is another application of the "rule of competent *a priori* justification": since it is an *a priori* truth that Hands entails \sim ED, it is a truth for which any competent thinker will have justification.

²⁵ For simplification, I have conjoined premises (2) and (J-Closure) as a single premise in this version. An instantiation of (J-Closure) is ($Jh \& J(h \Rightarrow \sim$ ED)) \rightarrow J \sim ED which is logically equivalent to $J(h \Rightarrow \sim$ ED) \rightarrow (Jh \rightarrow J \sim ED). And premise (2) is just the antecedent of this conditional. Therefore, by *modus ponens*, premises (J-Closure) and (2) entail our premise (J-Clos.) Jh \rightarrow J \sim ED.

The interpretation of EDA based on K-Closure plus Meta-Coherence²⁶

(2K) $\sim K \sim ED$

(K-Clos) $Kh \rightarrow K \sim ED$

$\therefore (C1') \sim Kh$ (by *modus tollens*)

$\therefore (C1'') J \sim Kh$ (because (C1') is *a priori*)

(MC-) $Jh \rightarrow \sim J \sim Kh$

$\therefore (C2') \sim Jh$ (by *modus tollens*)

This two-steps interpretation differs from the previous one only in the *order* in which it applies the two necessary steps: it applies the “anti-sceptical to mundane proposition” step *first*, and the “second order to first order” step in a second time, while the previous argument applied the two steps exactly in the opposite order.

In such an argument, Dretske and Nozick are happy to draw attention to the fact that K-Closure is highly controversial,²⁷ and they believe this is sufficient to block any worry concerning the knowledge (and rationality) of everyday beliefs. But the comparison with the previous two-steps interpretation shows clearly that this is a non-starter; it is a non-starter because when you have blocked *this* route to $\sim Jh$ (i.e. K-Closure and then Meta-Coherence), there remains another route to the same conclusion (i.e. Meta-Coherence and then J-Closure); this second route is much stronger because J-Closure is much stronger (dialectically) than K-Closure, and Dretske and Nozick’s classical response doesn’t offer the beginning of a response to this kind of sceptical reasoning. For that reason, the traditional externalist reconstruction and response to Evil Demon scepticism misses the real and strongest threat to everyday perceptual beliefs.

At this stage, it may seem that the argument is too complex or refined to be a plausible interpretation of what is really at stake in the sceptic’s reasoning. Or other readers may have doubts as to whether the principle of Meta-Coherence or of J-Closure are as so cogent as I say they are; which would mean that maybe the argument is not as *strong* as I say it is. This is partly because I gave the argument a simplified form, which helps seeing it’s general structure, but does not present the principles in their strongest version. In the next section, I will elaborate on both principles in order to show that (under a refined version) they are (i) very much ingrained in our unreflective / everyday treatment of beliefs, and (ii) very hard to deny philosophically speaking.

4 Why closure and meta-coherence are so intuitive and cogent

The principles of Closure and Meta-Coherence, as I have stated them in the preceding sections, are objectionable and need some major qualification in order to be considered as really cogent principles. In short, the needed qualification amounts to formulating the principles as principles of *epistemic commitment*. This has been done by Hawthorne (2005) for Closure, and by Michael Huemer for Meta-Coherence, but

²⁶ For premise (K-Clos), see the previous footnote: this premise isn’t an instantiation of K-Closure properly speaking but it is derived from an instantiation of K-Closure together with the premise $K (h \Rightarrow \sim ED)$.

²⁷ I don’t want to affirm here that K-Closure is false. I am far from convinced that it is. What I want to say is just that, dialectically, K-Closure is a controversial position that requires argument, while J-Closure is immediately a very cogent principle, especially in the form that we will see in the next section.

the two discussions have been kept separate. In this section, I will present the needed improvement of both principles as strictly parallel in spirit.

The main objection that can be made against J-Closure as I have presented it so far ($Jp \rightarrow Jq$ for all p and q such that $p \Rightarrow q$), is that it seems impossible to have justification for *absolutely all* the logical consequences of our beliefs: that would require us to have an infinite number of justifications present to our minds at the same time. Similarly, against the principle of Meta-Coherence (for all p , $Jp \rightarrow JKp$), it seems impossible to have a justification for the meta-proposition Kp for *absolutely all* proposition p that one may believe; that would obviously generate an infinite regress (because I would therefore need justification for KKp , for $KKKp$, etc.), and it doesn't seem that we could possibly have in mind justifications for all those levels of propositions at the same time.

The idea of *epistemic commitment* solves this problem because it avoids the requisite of an infinity of justifications. In the case of Closure, the idea goes like this:

(J-Closure Commitment) For all p and q , if I am justified in believing p , and *if I become consciously aware that I can deduce q from p* , then I am committed to believing q justifiably, otherwise my initial justification for p is defeated.²⁸

In order to satisfy this principle, I do *not* need to have, at a given time t , all the justifications for all the logical implications of my belief; but *if* the question arises explicitly about one particular implication, then I am *committed* to believe that implication; and if I am not in a position to believe it justifiably, then I should (epistemically) refrain from believing it, and *that* (together with the present awareness of the entailment relation) would constitute a justification defeater for my belief in the antecedent.

For Meta-Coherence, a move exactly parallel can be made:

(MC Commitment) For all p , if I am justified in believing p , and *if I consciously entertain the meta-epistemic question whether this belief constitutes knowledge*, then I am committed to believing that I know that p , otherwise my initial justification for p is defeated.²⁹

In fact, we have seen earlier, that we could hesitate between a negative and a positive form of Meta-Coherence. The difference between the two has to do with *what* exactly would constitute a defeater for my initial justification for p . When the issue whether the belief that p constitutes knowledge (whether Kp) explicitly arises, there are three possible scenarios:

²⁸ The formulation of (Hawthorne 2005, p. 29) is not completely equivalent, but relies on relevantly close intuitions: "If one knows P and competently deduces Q from P , thereby coming to believe Q , while retaining one's knowledge that P , one comes to know that Q ." The main difference is that Hawthorne is elaborating on a principle of K -Closure, not J -Closure. But the elaboration he gives is very close to the one I propose to J -Closure; namely, it concerns *only* what happens for implications that are explicitly entertained, and it allows for the possibility that, in considering the implication, I may thereby *lose* the justification (or knowledge) I had before considering it ("*while retaining one's knowledge that P* ").

²⁹ Here is Huemer (2011, p. 2) exact formulation: "Categorically believing that P commits one, on reflection, to the view that one knows that P ".

- (i) I may be in a position where I have justification to believe Kp
- (ii) I may be in a position where I have justification to believe $\sim Kp$
- (iii) I may be in a position where what I'm justified to do is only to remain in doubt as to whether the belief that p constitutes knowledge or not.

Of course, if I find myself in the first scenario, there is no problem: my initial justification for p is retained, and even confirmed by the meta-justification for Kp . If I find myself in the second scenario, that is if I justifiably come to believe that my belief that p *doesn't* constitute knowledge, most epistemologists would agree that this constitutes an “undercutting” or an “undermining” defeater for my justification to believe that p .³⁰ This diagnosis about case (ii) is all that is needed to deliver negative J-Closure:

(MC-Defeater) For all p , if I am justified in believing p , if I consciously entertain the meta-epistemic question whether this belief constitutes knowledge, and *if I justifiably come to believe that it does not*, then my initial justification for p is defeated.

In order to move from this negative principle to the principle of a positive commitment, we have to consider what happens in case (iii), i.e. if I justifiably remain in *consciously entertained doubt* about whether p constitutes knowledge or not. The defender of Meta-Coherence—(Huermer 2011, p. 7) but also (McGrath 2013b, Sect. 6.2) though he makes the point in a different lingo—considers that a consciously entertained doubt about whether p would *also* constitute an undermining defeater for the belief that p . This point is not as uncontroversial as the preceding one (that positive anti-knowledge belief constitutes an underminer). In a sense, this is not so important for the present discussion, because, as we've seen in the previous section, the two-steps reconstruction of the EDA only needs negative Meta-Coherence, not positive Meta-Coherence. But I will nonetheless provide a defence of positive Meta-Coherence, for two reasons: first, because in Sect. 6 we will see that the Dream argument may need positive Meta-Coherence; second, because the defence of positive Meta-Coherence, i.e. the Meta-Coherence *commitment*, is (I think) intuitively stronger when it is made in parallel with the defence of the Closure commitment. So this is what I am going to do now: I will present the main argument for both commitments in a parallel fashion.

The main argument in favour of a commitment of Closure has been very nicely put by John Hawthorne. His strategy is to show what an epistemic agent would look like who would deny his being committed to Closure. According to Hawthorne, this would generate an intuitively unacceptable form of “inconsistency”. I will call this the “problem of the inconsistent interlocutor”.

[Problem of the inconsistent interlocutor]:

Suppose Q is a “heavyweight” consequence of P and S knows P and also that P entails Q. I ask S whether she agrees that P. She asserts that she does: “Yes,” she says. I then ask S whether she realizes that Q follows from P. “Yes,” she says. I then ask her whether she agrees that Q. “I'm not agreeing to that,” she

³⁰ The idea of “undermining defeaters” has been introduced by (Pollock and Cruz 1999, p. 196), and has later been used by many epistemologists, sometimes with refinements, like in (Bergmann 2006, p. 156).

says. I ask her whether she now wishes to retract her earlier claims. “Oh no,” she says. “I’m sticking by my claim that P and my claim that P entails Q. I’m just not willing to claim that Q.” Our interlocutor now resembles perfectly Lewis Carroll’s Tortoise, that familiar object of ridicule who was perfectly willing to accept the premises of a *modus ponens* argument but was unwilling to accept the conclusion.

(Hawthorne 2005, p. 32)

In this little dialogue, the inconsistent interlocutor seems clearly irrational from an internalist point of view: if she is aware of the logical implication, either she is justified in accepting the consequent, or she should retract her earlier claims. This thought-experiment by Hawthorne gives in my view the strongest (and hardly deniable) intuition of a commitment of Closure.

An exactly parallel argument can be given for the commitment of Meta-Coherence, which I will call the “problem of the incoherent interlocutor”.

Problem of the incoherent interlocutor

Suppose I ask S whether she agrees with *p*. Her response is affirmative: “Yes,” she says. I then ask S whether she considers her belief that *p* as knowledge (whether she is willing to say “I know that *p*”). “I’m not agreeing with that,” she says. I ask her whether she now wishes to retract her earlier claims. “Oh no,” she says. “I’m sticking by my claim that *p*. I’m just not willing to claim that I know that *p*. What I’m saying is that *p* is true, but that this is not something I know, or at least it is completely possible that I don’t know it; and there is no logical contradiction here”.

The incoherent interlocutor is right to point out that there is no *logical contradiction* in claiming both “*p*” and “I don’t know that *p*” (or rather “my belief that *p* isn’t knowledge”). But in spite of the absence of logical contradiction, a conjunction of the form “*p* but I don’t know *p*” is infelicitous, for the same kind of reasons that other “Moore-paradoxes” (like “*p* but I don’t believe that *p*”) are infelicitous. Moore himself, though his examples are most often about “belief”, has given an explicit version of the “knowledge” version of the paradox:

Dogs bark, but I don’t know that they do. (Moore 1962, p. 277)

In order to get a *positive* constraint of Meta-Coherence, we also have to consider as Moore-paradoxical beliefs of the following form:

MP *p*, but it is completely possible (epistemically) that I don’t know that *p*.

The idea here is that a conjunction of this form, though it is not logically inconsistent, cannot be rationally believed for reasons having to do with internal constraints on rationality.³¹

³¹ It is important to notice that the interpretation of what is problematic in Moore-paradoxes is not uncontroversial. There are in fact two main traditions, one which supports Meta-Coherence as a fundamental principle of justification and one not so clearly. The latter seems to have been that of Moore himself, and is now aptly defended by (Williamson 2000, Sect. 11.3); we could call it the “norm of assertion” interpretation. According to this interpretation, Moore-paradoxes are problematic “because by asserting *p* positively you

If anyone has doubts whether a conjunction of beliefs like MP is really irrational or unjustified, maybe it will be of some help to recall the definition of externalist (anti-luck) knowledge. As I noticed in Sect. 1, post-Gettier epistemology seems to have reached the consensus that a belief cannot count as knowledge if it is true “merely as matter of sheer luck”. This is an uncontroversial *necessary* condition on knowledge. It is also possible to consider it as a *sufficient* condition (together with truth and belief) at least for some weak sense of knowledge.³² The definition of knowledge would then look like this:

- S knows that p iff
- (i) S believes that p
 - (ii) p is true
 - (iii) it is not a matter of sheer luck if S’s belief that p is true.

Now, suppose that someone believes a conjunction of the form MP, “ p but it might just as well be that this belief doesn’t constitute knowledge”. For what reason might this belief fail to be knowledge (in the view of the believer)? Is the believer believing that p is true while being in complete doubt about condition (ii) (that p is true)? That would be a direct contradiction. Is the believer believing p while being in complete doubt whether (i) (whether he believes that p)? That would be obviously irrational: any rational person who consciously entertains the belief that p cannot fail to be reflectively aware that she believes that p . So it seems the best chance of the “incoherent interlocutor”, in order to avoid irrationality, would be to say that she doubts whether she knows that p *because* she doubts whether condition (iii), the anti-luck condition, is satisfied or not.

The question of Meta-Coherence then comes down to our intuition about the rationality of the following conjunction:

- (PC) p but though my belief that p is true, still given the way in which I formed this belief it would be a matter of sheer luck if it comes out true.

Footnote 31 continued

imply, though you don’t assert, that you know that p ” (Moore 1962, p. 277). Therefore, a contradiction appears for conversational reasons between the asserted content ($\sim Kp$) and a conversational implicature (Kp , as an implicature of asserting p). Why should an assertion that p involve the implicature that Kp ? If we follow Grice’s theory (Grice 1975), implicatures appear according to *pragmatic norms* that preside over assertions. In the present case, the idea would be that “knowledge is the norm of assertion” (I should assert that p only if I know that p , and therefore, if I assert that p , my interlocutors have reason to suppose that I have respected the norm, i.e. that I know that p). This Moore-Williamson interpretation, which relies on a norm of *assertions*, has one major drawback: it only explains the paradoxical nature of an *assertion* of the form MP; therefore, it says nothing about the rationality of an agent who would merely believe a proposition of that form, without asserting it. As (Almeida 2001, p. 33) puts it: “an adequate explanation of the nature of Moorean absurdity is, first and foremost, an explanation of the oddity of (the objects of) certain *beliefs*”. Indeed, it seems evident that beliefs of the form MP would be as infelicitous as assertions of the same form, and if we account for the more fundamental infelicity of beliefs, we could also explain (derivatively) the infelicity of assertions (since belief is undoubtedly a norm of assertion). As (Shoemaker 1995, p. 76) says: “what can be (coherently) believed constrains what can be (coherently) asserted”. That’s why a second tradition of interpretation of Moore-paradoxes situates the problem in more fundamental rational norms of belief, and not in norms of assertion (Shoemaker 1995; Almeida 2001; Huemer 2011, Sect. 3). This is the tradition I follow here.

³² Sosa’s (1991, p. 125) notion of “animal knowledge” is very close to this idea.

It is, I think, very hard to deny the intuition that such a conjunction of beliefs is irrational. Maybe an example from applied epistemology will help understanding the kind of obvious irrationality I have in mind. In the epistemology of religious belief, Montaigne defended that one should always consider as true the religion of the province in which we were born, and (he adds) we should thank Heavens for being born in the province that has the true religion. This conjunction is (willingly, I suppose) paradoxical and provocatively irrational. For if you adopt a religion *only because* it is that of the place where chance had you be born, then obviously it can only be a matter of pure chance if it comes out true.

Of course, there are notions of “knowledge” that are stronger than “non-lucky true belief”. For example, there is the notion of “absolutely certain non-lucky true belief”. Someone who takes knowledge in this strong sense could (perhaps) rationally assert and believe “*p*, but my belief doesn’t constitute *knowledge*”, if all he means in so saying is “I believe that *p*, but my belief is not *absolutely certain*”. Indeed, there seems to be real life cases of this kind in which we maintain some belief for which we are not prepared to apply the word “knowledge”, because of some strong conception of knowledge. My defence of Meta-Coherence does *not* rule out these possibilities. My defence is really about the irrationality of consciously considering one’s beliefs as epistemically *lucky*. It seems to me very hard to refuse that someone who maintains a belief that *p* while also believing that this belief is (or might be) lucky, someone of that kind displays an “incoherence” which is just as problematic as Hawthorne’s “inconsistent interlocutor”.

Because of these two thought-experiments, therefore, it seems to me that the commitment of J-Closure and the commitment of Meta-Coherence are very intuitive and cogent principles, perhaps principles of “common sense”. In the next section, I will use these new versions of the principles in order to give an enhanced formulation of the two-step interpretation of EDA.

5 The double challenge formulation of EDA

I defend that, fundamentally, the intuition of EDA (and what makes it a very cogent argument) is the combination of two epistemic commitments that anyone is able to recognize for any belief one consciously entertains. The first commitment is the commitment to accept the logical implications of one’s beliefs, when one notices these implications, and on pains of being defeated. The second commitment is the commitment to consider that one’s beliefs are not true by sheer luck, when the question is explicitly raised, on pains of being defeated.

We could write down a complete version of the two-steps argument with the qualifications contained in the commitment-versions of MC and J-Closure. But this would be a bit tedious, and I think that the commitment principles allow for a much more intuitive presentation, namely a presentation in the form of dialogue between the believer and an interlocutor (maybe the believer herself) who makes her aware of her commitments by presenting “challenges”.

The double-challenge presentation of the EDA

- believer: I have hands.
 challenge: But do you accept that you are not a pure soul manipulated by an Evil Demon (surely you notice that it's entailed by your belief)?
 believer: Well, yes, I accept that I'm not a pure soul manipulated.
 challenge: but do you *know* what you're now believing, i.e. that you're not being manipulated (wouldn't it be sheer luck even if it's true)?
 believer: ...

In this short dialogue, the first challenge draws attention on the commitment of J-Closure, the second challenge draws attention on the commitment of Meta-Coherence. The first challenge is not so hard to meet: after all, any sane person would be ready to believe that she is not a pure soul manipulated by an Evil Demon. Similarly, if the interlocutor decided to challenge the believer by using *only* the commitment of meta-coherence, the interlocutor would be in a position to meet it quite naturally: a sane person would be ready to consider her belief that she has hands as non-lucky or as knowledge.

The problem arises only because the *two* challenges are combined. The strength of the EDA is that it is not clear at all how a sane person could meet the second challenge. One strong intuition³³ seems to push her in the direction of having to respond:

- believer: well, in fact, no, *that* cannot be known.

And that response would lead her to retract her second belief (that she is not a pure soul manipulated) because of the MC commitment, and then also her first belief (that she has hands) because of the Closure commitment.

Now, maybe the intuition that anti-sceptical beliefs cannot constitute knowledge (or can be at best *true by sheer luck*) is wrong. Maybe we don't have positively this intuition but rather we find ourselves confused when confronted with the second challenge. The reaction would be something like this:

- believer: well, I don't quite know what to say, it's not clear to me whether the belief \sim ED can be non-lucky in *some* way, or whether perhaps it doesn't have to be non-lucky, or maybe there are other possibilities I'm not presently seeing, it's just that I can't feel too confident to positively call this belief knowledge.

Notice that this reaction, though it is much weaker than the previous one, will not suffice to satisfy the *commitment* of *positive* Meta-Coherence. If positive Meta-Coherence is accepted, the in order to meet the challenge, the believer *has* to be prepared to say:

- believer: Yes, my belief that I am not a pure soul manipulated constitutes knowledge and is not lucky.

And this reaction doesn't seem so easy to have when one considers the obvious fact that "even if I *were* a manipulated soul, I would form all the same perceptual beliefs in just the same way".

³³ The intuition here is of course that there is no counterfactual dependence whatsoever between the belief and its content.

I am *not* claiming that the argument is successful, and that the challenge cannot be met. In fact, I think the challenge *can* be met and that the most important elements of response are already present in the epistemological literature.³⁴ What I have been trying to do in this paper is to propose a novel reconstruction of the problem *to which* they are elements of response. I am not pretending that the task of the anti-sceptic is an impossible task; I have only tried to describe *what* his task is, with all its exigencies.

6 The dream argument

In this last section, I want to consider an objection to my reconstruction of the EDA, one inspired by Stroud (1984, p. 14) insistence that sceptical scenarios like the Dream scenario are in fact more worrisome and more complex than the Evil Demon scenario proper.

Stroud's objection against traditional closure interpretation of scepticism is that there are sceptical scenarios which have no relation of logical implication with our ordinary beliefs. Granted, if I am a pure soul manipulated by the Evil Demon, *then* (logically) I don't have hands, there are no chairs, etc. But if I am dreaming that I am speaking in the House of Lords (as did the Duke of Devonshire), it does not follow that I am not (in fact, the Duke woke up to find that he *was* speaking in the house of Lords). The principle of Closure is able to account for the fact that the ED scenario is "epistemically bad" for my perceptual beliefs, but it is unable to account for the fact that the Dream scenario is just as bad (or perhaps worse). So it doesn't seem to be a proper interpretation of the threat raised by sceptical scenarios in general.³⁵

To solve this problem, Stroud elaborates a complex epistemic principle according to which the knowledge that *p* requires the knowledge of any *q* which is entailed not only by *p* but also by the *fact that one knows that p*.

Stroud's principle

if somebody knows something, *p*, he must know the falsity of all those things incompatible with his knowing that *p*. (Stroud 1984, p. 29)

This principle, I'm afraid, seems fairly *ad hoc* in its complex formulation. It doesn't have *prima facie* the ring of self-evident truth. The only reason we could have to accept it is that it seems to fit a number of cases for which we *do* have clear intuitions, but since there are also cases for which we *do not* have so clear an intuition (namely the controversial cases of anti-sceptical beliefs), I don't see why we should accept it as a universal truth which entails problems for the (controversial) anti-sceptical beliefs.

(Pryor 2000, n. 12) notes that Stroud's principle could be derived from closure *plus* the luminosity of knowledge (or KK-principle). That would be a much less *ad hoc* way of defending the principle, but it wouldn't make for a very cogent argument either, because the KK-principle itself is highly controversial.

³⁴ I have in mind the literature on the "safety" interpretation of the anti-luck condition for knowledge, which is, to my understanding, the best response to my own interpretation of the EDA, but I won't pursue here the project of responding to the argument.

³⁵ (Pryor 2000, p. 523) makes just the same objection to closure interpretations.

What I propose here, is that we can solve the problem Stroud is trying to solve not by combining the controversial principle of K-Closure with the highly controversial KK-principle, but rather by combining in the appropriate way the highly plausible commitment of J-Closure together with the highly plausible commitment of Meta-Coherence (which is, so to speak, a JK principle instead of a KK principle).

If we try to apply both principles in the way we have applied them so far (in the two-steps formulation), that will not do. Remember that our reconstruction *started* with a challenge of Closure: where the believer asserted “I have hands”, the interlocutor asked her whether she was ready to accept the *logical implication*, which was the negation of the sceptical scenario. But now, with the *Dream* scenario, the negation of the scenario is *not* a logical implication of the initial claim *Hands*: the belief that I have hands (or that I am in the House of Lords) doesn’t imply that I am not dreaming; I might very well dream that I have hands and really have hands (or dream that I am in the House of Lords and really be in the House of Lords). Fortunately, there is another way to construct a problematic challenge with J-Closure and MC.

The three-step challenge of the Dream scenario

- | | |
|--------------------|--|
| believer: | I have hands. |
| MC challenge: | But are you ready to consider this first belief as <i>knowledge</i> ? |
| believer: | Well, yes! That I have hands is certainly something that I <i>know</i> . |
| Closure-challenge: | But surely you are aware that you wouldn’t <i>know</i> that you have hands, if you were merely <i>dreaming</i> to have hands. So are you willing to accept the implication of your second claim, namely that you are not dreaming? |
| believer: | Well, yes. I certainly believe that I am not dreaming. |
| MC challenge: | But are you ready to consider this new belief as <i>knowledge</i> (in spite of the fact that you would still have that belief even if it were false)? |

In this reconstruction, we use the commitment of MC twice, but we still make use of no other principle but J-Closure and MC.³⁶

Perhaps some readers will find the formal presentation (with the imprecise principles $Jp \rightarrow Jq$ and $Jp \rightarrow JKp$) more telling. (Notice that this presentation is “in reverse order” so to speak relative to the commitment-presentation.)

³⁶ It is interesting to notice that (Huemer 2011, p. 10) has thought of the *first* use of MC in this scenario, and proposed to combine this use of MC with an application of Closure. What he has not seen, apparently, is the *second* use of MC, i.e. using MC *after* Closure in order to challenge the belief in the *negation of ED or Dream*. This second use is more important in that it is common to all sceptical scenarios and because it is only due to this move that we can avoid the problems of K-Closure.

The Meta-Coherence plus J-Closure Interpretation of the Dream Argument³⁷

- (1) $J \sim K \sim D$
 (MC-) $J \sim D \rightarrow \sim J \sim K \sim D$
 $\therefore (C1) \sim J \sim D$ (by *modus tollens*)
 (J-Clos.) $JKh \rightarrow J \sim D$
 $\therefore (C2) \sim JKh$ (by *modus tollens*)
 (MC+) $Jh \rightarrow JKh$
 $\therefore (C3) \sim Jh$ (by *modus tollens*)

From this presentation of the Dream argument, we can make two important remarks.

First, this argument makes no use at all of the famous KK principle. As said before, Pryor (2000, n. 12) interprets Stroud's version of the Dream argument as relying implicitly on a combination of a K-Closure plus the KK principle. And it is true that this combination of principles would allow us to derive (C2) from premise (1), in the following way:

The K-Closure plus KK principle interpretation of the Dream Argument³⁸

- (1') $\sim K \sim D$
 (K-Clos) $KKh \rightarrow K \sim D$ (because $Kh \Rightarrow \sim D$)
 $\therefore (C1') \sim KKh$ (by *modus tollens*)
 (KK) $Kh \rightarrow KKh$
 $\therefore (C2') \sim Kh$ (by *modus tollens*)
 $\therefore (C2) \sim JKh$ (because (C2') is *a priori*)

But this way to derive (C2) from (1) suffers from the dialectical weaknesses of both K-Closure and the KK principle. If Pryor is right in thinking that Stroud implicitly relies on this combination of principles, then the interpretation we are giving here provides a much more solid argument, one which, instead of using K-Closure and then KK principle, uses Meta-Coherence (or the "JK principle") and then J-Closure—both principles being, as we have seen, very strong ones. As in our discussion of Dretske and Nozick in Sect. 3, it seems that Stroud's argument (or Pryor's interpretation thereof) applies the two steps (from second order to first order, and from Dream to $\sim K$) in just the wrong order.

Secondly, our reconstruction shows another important peculiarity of the Dream argument: its third step (from (C2) to (C3)) can work *only* if we accept *positive* MC, and not just negative MC. This is because the *Dream* scenario leads us to *doubt* whether our perceptual experiences are reliable, but it doesn't lead us to *positively believe* that they are unreliable (in order to bring us to believe that, the Dream scenario would have to be assumed as *true* and not just *possible*). This is, I think, a good reason to suppose that, in general, the sceptical intuition based on sceptical scenario uses the positive MC principle and not the (dialectically stronger) negative principle. Does it

³⁷ For premise (J-Clos.), see footnote 25: (J-Clos.) is not an instantiation of J-Closure properly speaking, but is derived from an instantiation of J-Closure, namely $(JKh \& J(Kh \Rightarrow \sim D)) \rightarrow J \sim D$ together with the following premise $J(Kh \Rightarrow \sim D)$ which comes from the fact that the dream scenario is metaphysically incompatible with knowledge of h (and that this incompatibility can easily be grasped *a priori*, see footnote 24).

³⁸ For premise (K-Clos), see footnotes 26 and 37.

make the Dream argument dialectically weaker than the EDA, since the EDA can be formulated by using only negative MC?³⁹ This is not completely obvious to me: the overall comparison of the strengths of the two arguments will depend on the relative strengths of other premises as well. The Closure premise will make no difference, but the first premise of both arguments (the “luck” premise) could be argued to be dialectically *stronger* for the Dream argument than it is for the EDA.⁴⁰ Now, if the “luck” premise of the Dream argument is *much* stronger than that of the EDA, while its MC premise is only *slightly* weaker, then perhaps the Dream argument is somewhat stronger after all, even though it requires strong Meta-Coherence.

7 Conclusion

In this paper, I have defended that a combination of two very cogent and intuitive principles, namely the J-Closure commitment (“accept the implications of what you believe”) and the MC commitment (“do not hold on to beliefs that might be utterly lucky”) allow us to construct a version of the EDA which is both (i) very cogent (perhaps the strongest version of the EDA, dialectically speaking), and (ii) probably the original intuition that anyone has when grasping that the ED scenario “threatens” our beliefs.

The two principles are cogent, because it is very hard to deny the intuition that the agents who explicitly violate them (the “inconsistent interlocutor” and the “incoherent interlocutor”) display some highly problematic form of irrationality. In fact, I think that the two principles cannot be plausibly denied at all, and that an anti-sceptic who wants to resist the argument based on MC plus J-Closure should rather attack the very first premise of the sceptical argument, namely that we can’t know the denial of sceptical scenarios. The anti-sceptic will have, therefore, to provide an account of epistemic luck according to which anti-sceptical beliefs are *not* lucky, even though they are not counterfactually tracking truth.⁴¹ Though I haven’t tried to give such a response in the limits of this paper, it seems to me to be the only satisfactory answer to the challenge raised by the EDA.

The principles are intuitive as well, and probably principles of common sense, because it is very natural and common for any rational agent to *avoid* being explicitly forced to tortoise-like inconsistencies, and it seems also very natural and common to *avoid* representing one’s belief as matters of sheer luck. If those practices are natural and deeply ingrained in us, we can also presume that we can apply them, and even *combine* them, quite rapidly. That is, if someone presents me the ED scenario, I am able to grasp in a second (i) that I should deny the scenario (if I want to retain most of my beliefs) and (ii) that I should represent myself as *knowing* the negation of the scenario.

³⁹ Thanks to an anonymous reviewer for drawing my attention to this question.

⁴⁰ This is because the anti-sceptic can reply to the EDA that the ED scenario is a pure philosophical scenario, with no basis whatsoever in actual reality, while the Dream possibility is certainly closer to our actual experience. In that sense then, it seems easier to claim epistemic “safety” for the belief in \sim ED than for the belief in \sim Dream.

⁴¹ Again, I have in mind here the “safety” branch of the anti-luck tradition, the notion of safety being famously defended in (Williamson 2000) and (Pritchard 2005a).

If I can grasp those two commitments very rapidly, I will therefore be confronted with the question whether it's possible to *know* the negation of ED, and *this* is the question which is hard to answer and provokes the original uneasiness, or the original challenge of the EDA. For those reasons, it is quite likely that the combination of these principles is really what drives the “initial sceptical intuition” that one tends to have when first encountering the ED scenario.

The motivation for reconstructing “the original” interpretation of EDA, and also (arguably) the strongest version of it, was in fact to determine what the anti-sceptical has to tackle, if he wants to be in a position to confidently say that he has refuted EDA scepticism. Of course, there are *other* sceptical arguments than the EDA (the regress argument, the criterion argument, etc.). But the EDA has proved one of the most important, and it is certainly an essential part of any anti-sceptical philosophy that it could confidently consider itself as having satisfactorily addressed “what is really at play” in this particular argument. I have argued that versions of the argument based on K-Closure and/or on the KK-principle were unduly presupposing some controversial principles that are not needed in order to account for the essential intuitions of the argument.

If my reconstruction is correct, we can consider that recent epistemological principles have shed decisive light on Descartes's original argument, which sets us in a much better position to respond satisfactorily to it. But the response will have to wait for another occasion.

References

- Almeida, C. (2001). What Moore's Paradox is about. *Philosophy and Phenomenological Research*, 62(1), 33–58.
- Audi, R. (1998). *Epistemology: A contemporary introduction to the theory of knowledge*. London: Routledge.
- Bergmann, M. (2006). *Justification without awareness: A defense of epistemic externalism*. Oxford; New York: Clarendon Press; Oxford University Press.
- Bruce, A. (1994). The structure of the skeptical argument. *Philosophy and Phenomenological Research*, 54(4), 827–835.
- Chisholm, R. M. (1989). *Theory of knowledge*. Englewood Cliffs, NJ: Prentice-Hall.
- Cohen, S. (1998). Two Kinds of Skeptical Argument. *Philosophy and Phenomenological Research*, 58(1), 143–159.
- Dretske, F. (1970). Epistemic Operators. *The Journal of Philosophy*, 1007–1023.
- Engel, P. (2007). *Va Savoir!: De La Connaissance En Général*. Paris: Hermann.
- Foley, R. (1987). *The theory of epistemic rationality*. Cambridge, MA: Harvard University Press.
- Goldman, A. H. (2007). The underdetermination argument for brain-in-the-vat scepticism. *Analysis*, 67(1), 32–36.
- Grice, H. P. (1975). Logic and conversation. In *Studies in the way of words* (pp. 22–40). Cambridge, MA: Harvard University Press.
- Hawthorne, J. (2005). The case for closure. In M. Steup & E. Sosa (Eds.), *Contemporary readings in epistemology* (pp. 26–43). Malden, MA: Blackwell Pub.
- Huemer, M. (2011). The puzzle of metacoherence. *Philosophy and phenomenological research*, 82(1), 1–21.
- McCain, K. (2013). Two skeptical arguments or only one? *Philosophical Studies*, 164(2), 289–300.
- McGrath, M. (2013). Dogmatism, underminers and skepticism. *Philosophy and Phenomenological Research*, 86(3), 533–562.
- Moore, G. E. (1962). *Commonplace book, 1919–1953*. London; New York: Allen & Unwin; Macmillan.
- Nozick, R. (1981). *Philosophical explanations*. Cambridge, MA: Harvard University Press.

- Plantinga, A. (1993). *Warrant and proper function*. New York: Oxford University Press.
- Plantinga, A. (2000). *Warranted Christian belief*. New York: Oxford University Press.
- Pollock, J. L., & Cruz, J. (1999). *Contemporary theories of knowledge*. Lanham, MD: Rowman & Littlefield Publishers.
- Pritchard, D. (2005a). *Epistemic luck*. New York: Oxford University Press.
- Pritchard, D. (2005b). The structure of sceptical arguments. *Philosophical Quarterly*, 55(218), 37–52.
- Pryor, J. (2000). The skeptic and the dogmatist. *Nous*, 34(4), 517–549.
- Pryor, J. (2001). Highlights of recent epistemology. *The British Journal for the Philosophy of Science*, 52(1), 95.
- Putnam, H. (1981). *Reason, truth, and history*. Cambridge University Press.
- Reed, B. (2002). How to think about fallibilism. *Philosophical Studies*, 107(2), 143–157.
- Schiffer, S. (2009). Williamson on knowledge. In G. Patrick & P. Duncan (Eds.), *Evidence = knowledge: Williamson's solution to skepticism?* (pp. 183–202). Oxford University Press.
- Shoemaker, S. (1995). Moore's paradox and self-knowledge. *Philosophical Studies*, 77(2–3), 211–228.
- Sosa, E. (1991). *Knowledge in perspective: Selected essays in epistemology*. Cambridge [England]; New York: Cambridge University Press.
- Stroud, B. (1984). *The significance of philosophical scepticism*. Oxford [Oxfordshire]; New York: Clarendon Press; Oxford University Press.
- Unger, P. (1975). *Ignorance: A case for scepticism*. Oxford [Eng.]: Clarendon Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford; New York: Oxford University Press.
- Zalabardo, J. L. (2012). *Scepticism and reliable belief*. OUP.