# REVIEW OF DUNCAN PRITCHARD'S EPISTEMIC LUCK

## MARIA LASONEN-AARNIO

Merton College Oxford

**ABSTRACT**

Duncan Pritchard argues that there are two malign varieties of epistemic luck, veritic epistemic luck and reflective epistemic luck. He argues that eliminating veritic luck and meeting sceptical arguments requires an externalist, safety-based theory of knowledge, but that reflective luck is ineliminable. He raises an important challenge for internalists and externalists about knowledge alike, which is to explain how it is that ordinary beliefs can be internalistically justified. I discuss limitations that safety-based accounts have in dealing with various puzzles, and raise a worry about whether Pritchard succeeds in drawing a neat distinction between benign and malign varieties of epistemic luck. I also question Pritchard's sceptical conclusion concerning internalist justification.

**Key words:** knowledge, epistemic luck, safety

That knowledge is incompatible with luck has been an often-made assumption, especially among epistemologists of the post-Gettier era. Many will agree that the beginning of a diagnosis of what goes wrong in Gettier-type cases is recognising that the beliefs of these subjects are true as the result of a bout of good luck. Consequently, eliminating luck has become a central desideratum of contemporary theories of knowledge.

In his book *Epistemic Luck*[1], Duncan Pritchard sets out to subject the 'epistemic luck platitude' to scrutiny and to look at the consequences that attempting to eliminate problematic varieties of luck has for epistemology. In doing so, he offers a subtle diagnosis and discussion of the thrust of sceptical worries about knowledge. Pritchard disentangles two senses in which beliefs can be subject to problematic epistemic luck, what he calls *veritic epistemic luck* and *reflective epistemic luck*. He argues that eliminating veritic luck and meeting closure-based

---

[1] Pritchard 2005.

sceptical arguments requires an externalist, safety-based account of knowledge. But the other malign variety of epistemic luck is ineliminable. On Pritchard's diagnosis, the deep sceptical issue is not whether knowledge is possible but rather, whether almost any ordinary beliefs are justified in the internalistic way we take them to be.

I will focus on the two main claims of Pritchard's book – that a safety-condition for knowledge is necessary and sufficient for answering closure-based sceptical arguments and for eliminating veritic luck, and that since reflective epistemic luck cannot be eliminated, our epistemic position is unhappy. There is much in *Epistemic Luck* that I agree with, and even more that I learned by reading *Epistemic Luck*.

## 1. Veritic luck and safety

Pritchard's general characterisation of luck is modal. A lucky event occurs in the actual world, but 'does not occur in a wide class of the nearest possible worlds where the relevant initial conditions for that event are the same as in the actual world'[2]. Pritchard takes 'wide class' to mean 'approaching half', for, he argues, if an event is likely, then we are reluctant to count its occurrence as lucky.[3] A second general condition is that a lucky event is significant to the agent concerned.[4] There is a *prima facie* worry about whether this condition applies in cases of epistemic luck, for it seems possible to form beliefs on matters that have no significance for one even upon consideration.[5] For instance, a subject might involuntarily form a belief about a distant event that has no significance for her as the result of accidentally overhearing a conversation.

Pritchard distinguishes various ways in which a belief can be affected by luck in benign ways. The proposition believed can be about events that were lucky (*content epistemic luck*).[6] The subject of knowledge can be lucky to have the cognitive capacities by which the belief was formed, as when shortly before forming a visual belief she was lucky to avoid being hit by a branch that would have destroyed her eyesight (*capacity epistemic luck*).[7] Or, she can be lucky to come across a piece of evidence supporting her belief (*evidential epistemic luck*), as in Nozick's example where a bystander is able to recognise an escaping bank robber as Jesse James because his mask happens to slip off.[8] In such cases the subject is also lucky to form a belief in a proposition in the first place (*doxastic*

---

[2]  Pritchard 2005, p. 128.

[3]  Pritchard 2005, p. 130.

[4]  Pritchard 2005, p. 132.

[5]  Pritchard assumes that if a subject forms a belief about an event, then that event must have some impact on her (Pritchard 2005, p. 178, note 2).

[6]  Pritchard 2005, p. 134. This variety of luck is recognised by Unger 1968, p. 159.

[7]  Pritchard 2005, p. 134 – 136. See also Unger 1968, p. 160.

[8]  Pritchard 2005, pp. 136 – 137. The example is in Nozick 1981, p. 193.

*epistemic luck*).[9] Indeed, Pritchard argues that evidential and doxastic luck tend to accompany one another.[10]

Pritchard argues that nearby worlds in which a subject forms a false belief in a proposition only present a problem for knowledge of that proposition when the belief is formed *in the same way* as in the actual world. On Pritchard's diagnosis, this is precisely what happens in Gettier-cases.[11] For instance, in Russell's 'stopped clock case' the subject forms a true belief by looking at a clock that happened to stop exactly 12 hours before.[12] Because the clock could easily have stopped at a slightly different time, looking at it could easily have resulted in a false belief. Such cases are instances of *veritic epistemic luck*[13]. The safety-requirement for knowledge as defended by Pritchard is equivalent to a thesis stating that knowing is incompatible with veritic luck.

Given Pritchard's initial characterisation of luck, we should expect veritic luck to arise only if a subject forms a false belief by the same means as in the actual world in *most* nearby worlds in which she forms a belief by those means. However, this isn't true of the stopped clock as described by Pritchard. Unless we characterise the way in which the subject forms a belief as 'looking at clock *C* while *C* has stopped' rather than just as 'looking at clock *C*'[14], the subject does not hold, as Pritchard[15] claims, a false belief in *most* of the nearest worlds in which she forms a belief about the time in the same way as in the actual world. The clock is assumed to be generally reliable, and hence, in most nearby worlds it won't have stopped at the time the subject forms a belief by looking at it.[16] It looks as though a more stringent criterion of luck is needed to count the stopped clock case as one involving veritic luck.

Pritchard recognises a pressure to revise his account, though one created not by the stopped clock case, but by the intuition that a subject cannot know, prior to the draw, that a ticket in a fair lottery will lose.[17] Because such beliefs can be true and, it seems, justified by strong probabilistic evidence, they seem to constitute a class of Gettier-

---

[9] Pritchard 2005, p. 137 – 138.

[10] Pritchard 2005, p. 140. There seem to be less controversial cases than those Pritchard discusses that I cannot describe here for reasons of space.

[11] Pritchard 2005, p. 148.

[12] Pritchard 2005, pp. 148–149.

[13] Pritchard 2005, p. 146 – 149.

[14] Building the fact that the clock is stopped into the initial conditions is problematic, for if a subject in fact looks at a clock that is *not* stopped, why should this not similarly be part of the way in which her belief is formed? Similarly, why should not the fact that Henry in 'barn façade country' is looking at a real and not a fake barn be part of the way in which his belief is formed? If this was the case, we would have to give the counterintuitive verdict that Henry *does* know that there is a barn before him.

[15] Pritchard 2005, p. 148; p. 156.

[16] Pritchard 2005, p. 149: 'in most nearby possible worlds it won't be broken right now'.

[17] Pritchard 2005, p. 162 – 163.

cases. At least in such cases, ruling out veritic luck will demand avoiding error in *all* of the relevant nearby worlds. Pritchard's final formulation of safety is the following:

*Safety III*

For all agents, S, if an agent knows a contingent proposition φ, then, in nearly all (if not all) nearby possible worlds in which she forms her belief about φ in the same way as she forms her belief in the actual world, that agent only believes that φ when φ is true. (Pritchard 2005, p. 163)

Ultimately, Pritchard settles with the stronger reading of Safety III on which knowledge requires avoiding false belief in *all* of the relevant nearby worlds.[18]

Demanding that a subject avoid error in all relevant nearby worlds draws into question whether all of the Gettier-type cases discussed by Pritchard can correctly be said to involve luck in the ordinary sense of the term. In discussing general intuitions about luck, Pritchard notes that if the odds are in favour of an event, we are less inclined to describe it as being lucky.[19] For instance, suppose that I win a prize if I pick a black ball out of a hat containing 999 black balls and one red ball. If I then go on to pick a black ball, it seems odd to say that my winning the prize was a lucky event. Why, then, if I form the belief that I will win the prize based on the odds, is it lucky that I form a true belief? After all, the conditions in which I do so are identical with those in which I win the prize. The difference cannot be explained by appeal to stakes, for we can assume that both winning the prize and holding a true belief about the matter are of very little importance to me. This worry is perhaps not conclusive: if anything, it shows that knowledge excludes luck in a more specific, technical sense.

However, I am not convinced that Safety III can live up to all of the demands that Pritchard sets for it. In particular, I doubt i) that it resolves what Pritchard calls the '"lottery" puzzle'; ii) whether it can rule out cases where a subject's belief is formed in what seems like an epistemically irresponsible way; and iii) whether it is sufficient for ruling out Gettier-cases in which a subject seems to have an internalistically justified, true belief, but to lack knowledge. I will take these issues up in turn.

(i) What Pritchard calls the '"lottery" puzzle' is that even though a subject cannot know that her ticket will lose a lottery based on the odds, she *can* know that her ticket has lost once she has read this in the newspaper, even though it is likelier that there is a misprint in the newspaper than that her ticket is a winner. I take Pritchard's solution to be that though there is a nearby world in which the subject wins the lottery, there is no nearby world in which the newspaper contains a misprint. But the problem was created by the assumption that the probability that the newspaper contains a misprint

[18] Pritchard 2005, p. 165

[19] Pritchard 2005, p. 130.

is higher than the probability that one will win the lottery. At least if the sense of 'probability' employed here is uniform, is there any non-arbitrary way of ruling a misprint in a newspaper as a more distant occurrence than a lottery-win?

The problem is sharpened by noticing that the outcomes of many events we take ourselves to have knowledge about depend on processes that could be described in ways that show them to be the outcomes of events very similar to lottery-draws.[20] If the world is indeterministic, very many future events will be the results of chancy cosmic lotteries. For instance, current interpretations of quantum mechanics tell us that there is a non-zero chance that my pen will 'tunnel' through the table when I drop it. How, then, can I know that if I drop my pen it will land on the table?[21] Or, suppose that when performing a mathematical calculation, there is a non-zero chance that owing to a quantum-event in my brain, I make a mistake. How, then, can I know the outcome of my calculation? The 'lottery puzzle' is not easily dismissed.

(ii) Pritchard argues that Safety III excludes problematic cases in which a subject's belief 'concerns an event which is fairly stable across nearby worlds and which is matched by a stable true belief, albeit one that is formed in the wrong kind of way'[22]. Pritchard considers the following case. A mother believes truly that her son is innocent of a murder solely based on her love for him, despite having good reasons to believe the contrary. Intuitively, because the mother ignores the evidence, she lacks knowledge. But in fact, the son could not very easily have committed the murder and hence, it looks as though in the nearest relevant worlds the mother's belief continues to be true. Pritchard responds to this case by insisting that if there really are good reasons for thinking that the son committed the crime – if a reputable law enforcement agency charges him with murder, etc. – then it must be the case that there is at least some nearby world in which he commits the murder.[23]

It is difficult to see how this response could generalise to all cases in which a subject forms a belief 'in the wrong kind of way'. All we need to do is to alter the case described by Pritchard by supposing that someone cunningly framed the son of murder, and that there is very strong though misleading evidence pointing to his guilt. Suppose, for instance, that the mother does not know her son well at all, but believes that he is innocent simply because he is her son. There is no nearby world in which the mother falsely believes in the innocence of her son on this basis and yet, her belief seems problematic in the same way as in Pritchard's original case. And not all problematic cases involve misleading evidence. Suppose that the mother has not seen her son since he was born, and 60 years later believes, as the result of wishful thinking, that he has never been in

---

[20] For such arguments, see Vogel 1990 and Hawthorne 2004.

[21] Cf. Hawthorne 2004, pp. 4–5.

[22] Pritchard 2005, p. 166.

[23] Pritchard 2005, pp. 166 – 167.

an accident of any sort. In fact, the son is neurotically cautious and never steps out of his house, and as a result, could not easily have been in an accident. As in the original case, the mother's belief is formed in the wrong kind of way, but continues to be true in nearby worlds.

One way to deal with such cases is to construe safety not as demanding that a subject must avoid falsely believing the very proposition *P* she actually believes in any of the relevant nearby worlds (as Safety III does), but as demanding that the way in which she comes to believe *P* in the actual world does not lead her to form a false belief in *any* proposition in nearby worlds.

Some such fix seems to be needed anyway to deal with a variant of the stopped clock case. In this case, the subject forms the accidentally true belief that it is 8.22 am based on looking at the stopped clock. The clock could easily have stopped, for instance, 10 minutes before it actually did, and in that case the subject would have come to falsely believe that it is 8.12. This precludes the subject from knowing that it is 8.22, despite the fact that (let us suppose), she could not easily have falsely believed *this very proposition*.[24] This revision might help with the mother-and-son cases, in so far as the mother could easily have formed false beliefs about her son out of love for him, or as the result of wishful thinking.

But the suggested solution isn't altogether unproblematic. Suppose that I am walking on the street, and need to know both what time it is and where the nearest tube station is. I decide to consult the nearest passer-by. He is in fact John, who is meticulously reliable about the time, but likes to send tourists like me out on little detours around the city looking for non-existent sights and stations. I randomly decide to only ask him what time it is. It seems that in this case I can come to know the time, despite the fact that by consulting John I could very easily have formed a false belief in a different proposition. This problem can be solved by gerrymandering the way in which I form my belief by characterising it as 'relying on John's testimony for the time'. This raises an issue about whether there is any general way of individuating methods, for this way of individuating methods seems to give the wrong results in other cases. For instance, in the mother-and-son cases described, by describing the mother's way of forming beliefs in a suitably fine-grained way, such as 'forming beliefs about her son on matters of safety based on wishful thinking', we can get the beliefs to come out as knowledge.

These difficulties give reasons to doubt whether there is any way to formulate safety in a way that would rule out all cases of believing a proposition 'in the wrong kind of way'. But is there really any need to rule out such cases? One of the main conclusions of *Epistemic Luck* is that subjects who merely meet a safety condition have mere externalist

---

[24] All we need to do is describe the case so that the subject could not easily have looked at the clock at any other time than at 8.22.

knowledge that they cannot properly claim or take cognitive responsibility for. And one diagnosis of what is wrong with beliefs formed in epistemically irresponsible ways is that they are not internalistically justified. In this connection it is good to bear in mind that it was never the purpose of the safety requirement to guarantee that subjects have an internalistic justification for their beliefs.

(iii) A further worry is that there seem to be Gettier-type cases in which a subject has a justified, true belief that we would be very reluctant to classify as knowledge, but in which the sort of luck that is involved would be classified by Pritchard as being evidential, not veritic. Recall the well-known thought-experiment of Henry in a country replete with cleverly constructed papier-mâché façades of barns.[25] Henry happens to look at the only real barn in the country, and forms the belief that the object before him is a barn. The case involves veritic epistemic luck, as Henry could easily have formed a false belief by looking at the barn-like objects surrounding him. In this case, the way in which Henry forms his belief can plausibly be described as one that could easily have given rise to a false belief.

However, the case can be altered so that this will no longer hold. Pritchard is inclined to individuate perception as a way of forming beliefs distinct, for instance, from being fed experiences by a neuroscientist.[26] Then perception and hallucination should count as different ways of forming beliefs. Now imagine a case in which Henry is in 'barn hallucination country', where there are peculiar hallucination-inducing currents in the air. As Henry walks around, he seamlessly switches from hallucinating barns to perceiving them. Henry in fact perceives a barn, and forms the belief that there is a barn before him, but he could very easily have falsely believed this based on hallucinating a barn. If Henry doesn't have knowledge in the original barn façade case, it is difficult to see how he could have knowledge in the hallucinated barn case. It is easy to construct further cases – for instance, suppose that a subject is walking around in a country full of concealed loudspeakers emitting recorded bird-calls, but happens to come across the only real (but presently silent) bird, forming the visually based belief that there is a bird in her vicinity.

Besides threatening the neat distinction between problematic and unproblematic varieties of epistemic luck, such cases illustrate that whether or not a case is classified as evidential or veridic depends on how the way in which a belief is formed, or the method employed, is individuated. The worry is that if the distinction between evidential and veritic luck is to be retained, there will be no general guidelines for how this is to be done but rather, judgements about which method a subject employs will be guided by judgements about whether or not the subject has knowledge. Pritchard mentions the

---

[25] This case is described in Goldman 1976, pp. 772–773.

[26] Pritchard 2005, p. 160.

'Generality Problem' for process reliabilism,[27] but doesn't discuss the way in which it arises for his own safety-account of knowledge.

Let me now turn to reflective luck.

## 2. Reflective luck

Pritchard defines internalism about justification as the claim that an agent's belief in a proposition is justified if and only if 'the agent is able to know the facts which determine that justification by reflection alone', where reflection includes 'a priori reasoning, introspective awareness of one's own mental states, and one's memory of knowledge that has been gained in either of these ways'.[28] I will understand this as the claim that in order to have an internalistically justified belief, a subject must be in a position to know by reflection the facts that her belief is (justificatorily) based on.

Pritchard motivates his discussion of the second malign variety of epistemic luck, *reflective luck*, by describing the case of a 'naïve chicken-sexer' who 'has a natural and highly reliable ability to distinguish male and female chicks, but who has no idea how she is doing this nor even aware that she is reliable in this respect'[29]. As a result, the naïve chicken-sexer is said to lack an internalist justification for the beliefs she forms based on her ability. Compared to her enlightened counterpart, the epistemic position of the naïve chicken-sexer seems problematic and impoverished.

What the example of the naïve chicken-sexer demonstrates is that a subject can in fact have a belief that is not true by luck without having any reasons of an internalistic kind for thinking that this is the case. A subject's belief can be safe and yet, 'from that agent's reflective position, it is still a matter of luck that her belief is true'[30]. The demand seems to be that to avoid reflective luck, a subject must have an internalist justification for believing that her belief is safe – or that it meets certain conditions entailing that the belief is safe. In the case of antisceptical propositions, this would mean having an internalist justification for believing that sceptical possible worlds are far-off and hence, having an internalist justification for believing that the actual world is not one of them.[31]

Pritchard argues that the problematic nature of reflective luck provides a strong motivation for thinking that meeting an internalist justification-condition is necessary for

[27] Pritchard 2005, p. 187.

[28] Pritchard 2005, p. 42.

[29] Pritchard 2005, pp. 43–44; see also especially pp. 174–177.

[30] Pritchard 2005, p. 174.

[31] Pritchard 2005, p. 248.

knowledge and hence, that *internalism about knowledge*[32] is correct.[33] However, meeting an internalist justification-condition can rule out reflective epistemic luck only if being internalistically justified in believing a proposition guarantees having reflectively accessible grounds for thinking that one's belief in it is safe (or not subject to veritic luck). This doesn't follow from Pritchard's characterisation of justification internalism, since it is one thing to have reflectively accessible reasons for believing a proposition *P*, and another to have reflectively accessible reasons for believing that one's belief in *P* is not subject to veritic luck, or that it is reliably formed, or safe. As there is no *prima facie* reason for why these justifications shouldn't sometimes come apart, Pritchard seems to move a bit too swiftly in stating the connection between reflective luck and internalist justification.

Granted that having an internalistic justification for a belief is essential for avoiding reflective luck, Pritchard's argument for the conclusion that reflective luck is ineliminable is completed by appeal to what he calls the 'underdetermination principle'. Before looking at this principle, I briefly discuss Pritchard's claim that it is impossible to be internalistically justified in believing antisceptical propositions.

## 3. Scepticism and the underdetermination principle

The claim that internalism about knowledge entails the impossibility of knowing antisceptical propositions plays a pivotal role in *Epistemic Luck*, especially in Part I. This is why, according to Pritchard, internalism about knowledge is unable to deal with the closure-based sceptical argument save by denying closure[34]; why sensitivity and contextualist theories of knowledge must be externalist[35] and hence, are motivated by conflicting intuitions in agreeing that at least in some contexts agents cannot know the denials of sceptical hypotheses[36]; at least in part why Moore's response to the sceptic fails[37]; and why asserting the falsity of a sceptical hypothesis is always conversationally inappropriate[38].

Pritchard argues that a subject in an ordinary scenario and her counterpart in a sceptical scenario both lack a justification for believing themselves not to be in the sceptical scenario:

---

[32] Pritchard 2005, p. 44.

[33] Pritchard 2005, p. 177.

[34] Pritchard 2005, p. 44.

[35] Pritchard 2005, p. 59

[36] Pritchard 2005, pp. 53, 60.

[37] Pritchard 2005, ch. 3.

[38] Pritchard 2005, p. 85.

> If there is, *ex hypothesi*, no phenomenological difference available to the subject which could indicate to her that she is not a victim of this [sceptical] scenario, then it follows that there is not going to be anything reflectively available to the subject that could suffice to indicate to her that her belief in this antisceptical proposition is true. In this sense then, she cannot be internalistically justified in believing this proposition, even if her belief is true. (Pritchard 2005, p. 44)

Presumably, 'suffice to indicate' means minimally that evidence that is reflectively available to the subject must support or favour the antisceptical proposition over its negation (a sceptical proposition), by making the antisceptical proposition more likely to be true. But, given that there is no difference in the reflectively accessible evidence between the ordinary and sceptical scenarios, Pritchard thinks that this evidence cannot favour the belief that the scenario does not obtain.

It is false that 'everyone agrees that we lack internalist justification for believing in the denials of sceptical hypotheses'[39]. There are internalists who argue that Moore-type reasoning is a perfectly acceptable way of gaining an internalist justification for antisceptical propositions.[40] Pritchard agrees with Wright and Davies that Moore-type arguments exhibit failures of warrant transmission, since the evidence or grounds one has for believing ordinary propositions don't constitute (sufficient) grounds for believing antisceptical propositions.[41] Moreover, Pritchard seems reluctant to regard an a priori warrant or entitlement to assume antisceptical propositions as constituting an internalist justification for believing them.

However, the failure of internalistic evidence for ordinary propositions to transmit to antisceptical propositions doesn't show Moore-style reasoning to be illegitimate, for it still leaves open the sort of position defended by Peter Klein[42]. It looks as though a subject's reflectively accessible evidence *can* justify a belief in an ordinary proposition, such as the proposition that one is wearing black shoes. But then, why couldn't the proposition that one is wearing black shoes provide an internalist justification for believing that one is not a brain in a vat, for surely wearing black shoes makes it more likely that one is *not* a brain in a vat? Some further argument is needed to block this response.[43]

Pritchard presents with approval the idea he attributes to Wittgenstein that the reasons offered in favour of a proposition need to be more certain than that proposition itself, but because nothing can more certain than an antisceptical proposition, one cannot

---

[39] Pritchard 2005, p. 211.

[40] See, for instance, Klein 1981, 1995 and Pryor 2004.

[41] Pritchard 2005 p. 94. See Wright 2000, 2003; Davies 2003.

[42] Klein 1981, 1995.

[43] Pritchard rejects Klein's view on the grounds that it is impossible to be internalistically justified in believing an antisceptical proposition (*EL*, p. 45), but in the present context this reply is question-begging.

give reasons for it.[44] There is much to unpack in this line of reasoning. Luckily, the conclusion that one cannot give reasons for believing antisceptical propositions is also supported by what Pritchard calls the 'underdetermination principle'. This principle, Pritchard argues, is also what the real sceptical argument rests on:

## *The underdetermination principle*

> For all S, $\varphi$, $\psi$, if S's evidence for believing $\varphi$ does not favour $\varphi$ over some hypothesis $\psi$ that S knows to be incompatible with $\varphi$, then S is not internalistically justified in believing $\varphi$.[45] (Pritchard 2005, p. 108)

'Evidence' is here to be understood internalistically[46], as being something reflectively available and, Pritchard assumes, constant across the ordinary and relevant sceptical scenarios. By the underdetermination principle, a subject cannot be justified in believing an ordinary proposition *O* unless her (reflectively accessible) evidence *E* favours *O* over a known to be incompatible sceptical hypothesis *SH*. Because an ordinary proposition *O* entails the antisceptical proposition *~SH*, this means that *E* must also favour *~SH* over *SH*, by making the antisceptical proposition more likely to be true. Klein's reasoning is blocked, since it doesn't look as though a subject's reflectively accessible evidence for the proposition that she is wearing black shoes, for instance, makes it more likely that she is not a brain in a vat than that she is a brain in a vat.

This, Pritchard argues, is the real gist of scepticism, as even externalists would like to think that many ordinary beliefs are internalistically justified.[47] Because reflective epistemic luck can be eliminated only if a subject can meet the demands for internalistic justification set by the underdetermination principle, reflective luck is ineliminable.[48] It turns out that even the beliefs of the enlightened chicken-sexer are subject to reflective luck, since her reflectively accessible evidence doesn't favour her beliefs over their sceptical alternatives.[49]

Given the essential role the underdetermination principle plays in *Epistemic Luck*, we should expect Pritchard to do more than merely state that it is intuitive. Why, for instance, can the relevant intuition not be captured by an alternative principle stating that a subject's evidence must favour a proposition over non-distant, or relevant, known to be incompatible sceptical alternatives? For instance, Pritchard argues that contextualist

[44] Pritchard 2005, pp. 86 – 89; pp.  226 – 227.

[45] The principle is a slightly modified version of one formulated in Brueckner 1994.

[46] See Pritchard 2005, p. 117.

[47] Pritchard 2005, p. 44, p. 206, p. 221.

[48] Pritchard 2005, p. 207.

[49] Pritchard 2005, p. 208.

strategies cannot deal with sceptical arguments that rest on the underdetermination principle[50], but he doesn't discuss the possibility of construing the underdetermination principle in a contextualised way, as demanding that a subject's evidence must favour a proposition over known to be incompatible *contextually salient* alternatives.

If the underdetermination principle is defended only by its intuitiveness, then the dialectical situation in the end of Pritchard's book threatens to be a standoff. Reflective luck, Pritchard argues, is problematic because even externalists will want to think that we are in a position similar to that of the enlightened, rather than the naïve, chicken-sexer with respect to a vast number of our ordinary beliefs. Hence, Pritchard's argument rests crucially on the intuition that the naïve chicken-sexer *could* improve her epistemic position in a significant way. But this intuition could be turned against the underdetermination principle by arguing that the demands set by the principle for the elimination of reflective luck are simply too strong – the position of the enlightened chicken-sexer really is good enough.

## Concluding remarks

*Epistemic Luck* makes an important contribution to epistemology. It offers an in-depth treatment of the varieties of luck relevant for knowledge, and is well-written and accessible. Of particular significance is its attempt to understand both the merits and limitations of externalism about knowledge. I am inclined to agree with Pritchard that safety-based accounts are the most plausible externalist theories, though I argued that their ability to deal readily with certain puzzles and intuitions is more limited than Pritchard is willing to admit. I also questioned the neat distinction between evidential and veritic epistemic luck.

Even if one doesn't agree with the demands Pritchard sets for eliminating reflective luck, *Epistemic Luck* convincingly presents a challenge that even externalists about knowledge must meet: either explain how it is that ordinary beliefs can be internalistically justified, or deal with the rather severe sceptical consequences of admitting that it is impossible to have internalistically justified beliefs.

[50] Pritchard 2005 , pp. 113–114.

REFERENCES

Brueckner A. (1994), "The Structure of the Sceptical Argument", *Philosophy and Phenomenological Research*, 54, pp. 827–835

Davies M. (2003), "The Problem of Armchair Knowledge", in Nuccettelli 2003, pp. 23–57

Goldman A. (1976), "Discrimination and Perceptual Knowledge", *Journal of Philosophy*, 20, pp. 771–791

Hawthorne J. (2004), *Knowledge and Lotteries*, Oxford: Clarendon Press

Klein P. (1981), *Certainty: A Refutation of Skepticism*, Minneapolis, MN: University of Minnesota Press

Klein P. (1995), "Skepticism and Closure: Why the Evil Genius Argument Fails", *Philosophical Topics*, 23, pp. 213–36

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

Nuccetelli S. ed. (2003), *New Essays on Semantic Externalism and Self-Knowledge*, Cambridge, MA. & London, England: MIT Press

Pritchard D. (2005), *Epistemic Luck*, Oxford: Clarendon Press

Pryor J. (2004), "What's Wrong with Moore's Argument?", Philosophical Issues, 14, pp. 349–378

Roth M. & G. Ross ed. (2000), *Doubting: Contemporary Perspectives on Scepticism*, Dordrecht: Kluwer

Unger P. (1978), "An Analysis of Factual Knowledge", *Journal of Philosophy*, 6, pp. 157 – 170

Vogel J. (1990), "Are There Counterexamples to the Closure Principle?" in Roth & Ross 2000, pp. 13-27

Wright C. (2000), "Cogency and Question-Begging", *Philosophical Issues*, 10, pp. 140–163

Wright C. (2003), "Some Reflections on the Acquisition of Warrant by Inference", in Nuccettelli 2003, pp. 79–97

Merton College Oxford
Oxford OX1 4JD, UK
maria.lasonen@philosophy.ox.ac.uk