

1-1-2006

Review of Duncan Pritchard, Epistemic Luck

Jason Baehr

Loyola Marymount University, jbaehr@lmu.edu

Repository Citation

Baehr, Jason, "Review of Duncan Pritchard, Epistemic Luck" (2006). *Philosophy Faculty Works*. 30.
http://digitalcommons.lmu.edu/phil_fac/30

Recommended Citation

Baehr, Jason. "Review of Duncan Pritchard, Epistemic Luck." *Metaphilosophy* 37 (2006), 728-36. Print.

This Article - pre-print is brought to you for free and open access by the Philosophy at Digital Commons @ Loyola Marymount University and Loyola Law School. It has been accepted for inclusion in Philosophy Faculty Works by an authorized administrator of Digital Commons@Loyola Marymount University and Loyola Law School. For more information, please contact digitalcommons@lmu.edu.

This is not a final draft. Please cite only the journal version, published in Metaphilosophy 37 (2006).

Epistemic Luck. By Duncan Pritchard. Oxford: Clarendon Press, 2005. Pp. xii + 290.

Duncan Pritchard's *Epistemic Luck* is a wide-ranging, nicely written, and generally masterful treatment of the concept of epistemic luck and its relevance to contemporary epistemology. Its central focus is what Pritchard calls the "epistemic luck platitude," which is the claim that knowledge excludes luck. This claim is widely (if largely uncritically) accepted among epistemologists; and yet an unqualified endorsement of it is problematic, for there are many putative instances of knowledge in which an agent's reaching the truth is in some sense a matter of luck (e.g., the archeologist who "happens" upon an important find or the detective who stumbles across a critical clue). Pritchard's aim is to clarify the sense in which knowledge excludes luck and to consider the resulting implications for the theory of knowledge.

Part I of the book (Chapters 1-4) concerns skepticism. This discussion sets the stage for Pritchard's more direct examination of epistemic luck in Part II.

The primary focus of Chapter 1 is a version of the skeptical challenge that hinges on the so-called "epistemic closure principle," which says that if *S* knows that *p* and *S* knows that *p* entails *q*, then *S* is in a position to know that *q* (27). The argument goes roughly like this: (P1) If I have knowledge of any commonsense empirical propositions (e.g., that I have hands), then via closure I know the denial of any skeptical hypothesis that I know to be incompatible with these propositions (e.g., that I am a brain in a vat, that I am inside the "Matrix," etc.); (P2) But I have no way of knowing the denial of the relevant skeptical hypotheses. (C) Therefore, I cannot have knowledge of any commonsense empirical propositions.¹

Chapters 2 and 3 are primarily a defense of Pritchard's preferred anti-skeptical externalist account of knowledge against two of its main competitors, viz., a "sensitivity-based" and an "attributer contextualist" model of knowledge. According to the former, for *S* to know some contingent proposition *p*, *S*'s belief that *p* must be "sensitive," in the sense that *S* does not believe that *p* in the nearest possible world in which *p* is false (48). Contextualists, on the other hand, hold that the standards for knowledge-ascriptions vary from one conversational context to

¹ Pritchard offers more and less formal versions of this argument on pp. 27-28. My rendering of it incorporates elements of each.

another; more specifically, they claim that in everyday conversational contexts, the standards for knowledge-ascriptions are relatively low, while in skeptical conversational contexts (i.e. those in which skeptical hypotheses are under consideration) they are relatively high (54). While these theories provide a way of blocking the radical skeptical conclusion above (the former via a denial of (P2) and the latter via a qualified denial of (P2) and (C)), Pritchard rejects them both, primarily on the grounds that they are illicitly motivated by certain *internalist* intuitions (53, 59). (Though some contextualists regard themselves as internalists, Pritchard argues on pp. 59-60 that any plausible version of contextualism must embrace externalism.)

Pritchard then turns to a defense of his own anti-skeptical theory, the content and motivation of which he says is thoroughly externalist in nature. He refers to his position as “neo-Mooreanism” on account of its (albeit limited) similarity to G.E. Moore’s famous (infamous?) reply to the skeptic. This account of knowledge incorporates a “safety” (vs. a “sensitivity”) principle, according to which *S* knows a given contingent proposition *p* only if *S*’s belief that *p* is “safe,” meaning that in most nearby possible worlds, *S* only believes that *p* when *p* is true (71). If correct, neo-Mooreanism provides a way of escaping skepticism, for as long as skeptical scenarios are *in fact* modally distant, one’s belief that, say, one is not a brain in a vat will be “safe” and hence, contra (P2), amount to knowledge (75).

While Pritchard favors this neo-Moorean reply to skepticism over its externalist competitors, Chapter 4 is devoted to examining a major shortcoming of it (and any other version of externalism). Pritchard argues that when properly understood, the skeptical challenge is fundamentally an *internalist* one; it concerns the quality of the *grounds* or *evidence* we are capable of having for our beliefs. More precisely, it concerns the “underdetermination principle,” which says that *S* has good evidence for or is internalistically justified in believing that *p* only if the evidence *S* has relative to *p* favors *p* over any other hypothesis *q* which *S* knows to be incompatible with *p* (108). The problem is that *none* of our evidence for commonsense empirical claims seems to satisfy this condition, for the evidence in question consists in sensory *experience*, and *ex hypothesi* our experience is precisely what it would be if we were victims of a skeptical scenario. The apparent result is that ultimately none of our commonsense empirical beliefs are well-grounded.

As Pritchard himself acknowledges, this line of argument deals a major blow to the neo-Moorean and any other externalist “solution” to skepticism, for while these views may rescue a

kind of “brute” externalist knowledge from the skeptical bog, they do nothing to mitigate the argument’s radical and unsettling conclusion. Pritchard concludes: “[r]esponding to scepticism with externalist knowledge alone as the neo-Moorean ultimately does is thus beside the point” (119); and “the sceptic is still alive and kicking, despite the many well-publicized claims [by externalists] to the contrary” (118).

There is much to be said in favor of this first part of Pritchard’s book; however, in the interest of space, I will limit myself to just two points. First, these chapters provide a lucid and informative survey of a variety of cutting-edge epistemological issues: e.g., skepticism, epistemic closure and its denial, infallibilism vs. fallibilism, internalism vs. externalism, “safety” vs. “sensitivity,” contextualism, underdetermination, and more. I would recommend a careful reading of these chapters (perhaps together with Chapter 6) to anyone with at least a minimal philosophical background who is looking to get “up to speed” on the contemporary scene in epistemology.

Second, Pritchard’s analysis of the dialectic between internalism, externalism, and skepticism is insightful and refreshingly balanced. Internalists sometimes charge externalist replies to skepticism with “changing the subject” or even with capitulating to the skeptic. While Pritchard’s sympathies lie with externalism, he does an excellent job of elucidating the basis of such complaints. On the other hand, he does not (as some internalists do) undersell the attraction of externalism: he makes clear how an externalist can rescue at least a certain, not altogether insignificant, concept of knowledge from the clutches of the skeptic.

I turn now to Part II of the book, in which Pritchard addresses the topic of epistemic luck head on, and in a way that sheds further light on the dialectic between the internalist, externalist, and skeptic. I will focus on his discussion in Chapters 6 and 8, since the remaining chapters mainly serve either to set up or to reinforce the central claims of these chapters.

After developing a general, counterfactual account of luck in Chapter 5, Pritchard turns in Chapter 6 to a discussion of two main varieties of epistemic luck: viz., “veritic luck” and “reflective luck.” A person’s belief is *veritically* lucky if it is a matter of luck that the belief is true (146). Pritchard illustrates this notion with the Gettier case of “Gullible John” who, in his characteristically gullible way, accepts his friends’ fabricated claim that his house is on fire. Unknown to his jesting friends, however, John’s house *has* just caught fire, which he promptly

discovers when he returns home to extinguish the blaze. Clearly John's belief that his house is on fire (at least prior to his actually seeing it) does not amount to knowledge (146).

Pritchard explains how an externalist, safety-based account of knowledge is well-suited to exclude veritic luck (and thus to overcome the Gettier problem). Recall that according to neo-Mooreanism, John knows his house is on fire only if, in most nearby possible worlds, John believes that his house is on fire only when it is on fire. But of course there are many possible worlds in which John, being gullible, accepts his friends' testimony even when it is false. Thus on Pritchard's view, knowledge is incompatible with the relevant kind of epistemic luck (146). Pritchard goes on to argue that, when properly amended (more on this amendment below), neo-Mooreanism can also handle a host of other problematic examples widely discussed in the epistemological literature, all of which point to an incompatibility between knowledge and veritic luck (e.g., the lottery paradox and barn façade cases) (161-173).

Whatever its success in connection with the problem of veritic luck, Pritchard acknowledges that neo-Mooreanism, or any other version of externalism, is incapable of excluding *reflective* luck. A belief is reflectively lucky, says Pritchard, if *given what the agent has access to on the basis of reflection alone*, it is a matter of luck that this belief is true (175). The problem for Pritchard (and externalists in general) is that the beliefs of, say, an unwitting but reliable clairvoyant might be reflectively lucky and yet perfectly "safe" in Pritchard's sense (or otherwise externalistically up to par). Pritchard stops short of repudiating externalism on this account, however, for he maintains that while the exclusion of reflective luck is no doubt epistemically *desirable*, it may not be *necessary* for knowledge (184).

In Chapter 8, Pritchard argues that externalism is not alone in its inability to overcome the problem posed by reflective luck, for reflective luck, he says, is in principle *ineliminable*, and hence is a problem for the internalist as well as the externalist. Here again Pritchard's focus is the problem of underdetermination, which suggests that we can never have *any* good reasons for thinking that *any* of our commonsense empirical beliefs are true. Given the connection between reflective luck and the possession of good epistemic grounds, this amounts to the claim that our commonsense empirical beliefs are ineluctably reflectively lucky (206).

There is also a great deal to be said in favor of Pritchard's discussion in Part II of the book; but again I shall limit myself to just two quick points. First, Pritchard's treatment of luck in general (Chapter 5), of moral luck (Chapters 5, 10), and especially of epistemic luck (Chapters

5, 6, and 8) constitute substantial contributions to the literature on these topics. Regarding the latter, Pritchard's analysis is more comprehensive and subtle than any other treatment of epistemic luck in the literature and will undoubtedly (and rightly) form the basis of much future work on this topic.

Second, I think Pritchard makes good on (what is apparently – see below) the central claim of the book: namely, that by getting a better handle on the nature and varieties of epistemic luck, we are able to better understand some of the principal differences between internalist and externalist accounts of knowledge and how they fare with respect to the skeptical challenge. He shows that epistemologists would benefit from making the notion of epistemic luck – and especially the distinction between veritic and reflective luck – a permanent and more central part of their conceptual repertoire.

Having noted some of the many virtues of *Epistemic Luck*, I want now to raise some criticisms. Doing so will, I hope, facilitate a better understanding of Pritchard's project and its subject matter.

The first criticism has to do with a notable gap in Pritchard's treatment of the dialectic between internalism and skepticism. We saw above that a dominant theme of Pritchard's discussion is that, owing to the problem of underdetermination, all of our empirical beliefs are ultimately internalistically ungrounded. This is because there is apparently no way to adjudicate between what we might call a commonsense hypothesis (CH) or explanation of our experience and any number of skeptical hypotheses (SH). The result is that experience apparently fails to justify our beliefs about the external world.

In assessing the internalist's ability to deal with this problem, Pritchard neglects a general line of argument with notable historical and contemporary credentials, and which, if successful, would have significant implications for his discussion. Something like this argument seems clearly to be driving Locke's discussion of knowledge of the external world in his *Essay Concerning Human Understanding* (see, e.g., Bk. IV, Sec. XI); more recently, it has been defended by Laurence Bonjour and Jonathan Vogel (among others). While the details of the argument cannot be explored here, the basic idea is that CH is rationally preferable to SH on the grounds (1) that CH is *simpler* than SH (e.g., because SH is *parasitic* on CH) and (2) that as a result of its greater simplicity, CH is *more likely to be true* than SH.

Pritchard does make reference to an argument of this general sort, but only (to my knowledge) in two footnotes. And in each case, he says merely that “it is far from clear ... just how this line of thought can help the internalist ... because it is not obvious that this use of inference to the best explanation presents us with *epistemic* grounds for our antisceptical beliefs as opposed to merely *pragmatic* grounds ...” (63, nt. 11; cf. 112, nt. 11). The problem is that while some internalists have not said very much in favor of (2) above (thus making Pritchard’s remark understandable to an extent), others have. Bonjour, for instance, has argued at length and in several places for the claim that CH’s greater simplicity gives it an *epistemic* advantage over SH.²

This is not an insignificant oversight on Pritchard’s part, for his claim that there are no good internally accessible reasons for rejecting skepticism looms extremely large throughout the book. Indeed, if the sort of argument just noted were to succeed, it would, among other things, undercut Pritchard’s case against contextualism (59) and lessen the motivation for his own safety-based account of knowledge (68). More importantly, it would alter the overall tone and conclusion of the book, for as it stands *Epistemic Luck* is largely a *defense* of skepticism the cogency of which depends on the failure of just this kind of argument. In light of its dialectical significance, such an argument (regardless of its ultimate plausibility) should have been given serious consideration.

A second worry concerns the amended and final version of Pritchard’s “safety-based” account of knowledge put forth in Chapter 6. Recall that Pritchard originally claims that *S* knows that *p* only if, in *most* nearby possible worlds, *S* believes that *p* only when *p* is true (71). This view runs into difficulty in connection with the so-called lottery paradox. Consider a scenario in which I purchase a ticket in a fair lottery, but subsequently form the belief that it is not the winning ticket. Intuitively, I do not *know* that I will lose the lottery; yet my belief seems to satisfy the safety principle above, since in most nearby possible worlds I do in fact lose the lottery. In response to this problem, Pritchard concludes that “when it comes to knowledge the agent’s belief must not only be true in *most* of the nearby possible worlds where she forms her belief in the same way as in the actual world, but in *nearly all* (if not all) of them” (162). Thus amended, Pritchard’s account says (roughly) that I know I will lose the lottery only if in nearly all (if not all) nearby possible worlds, I believe that I will lose only when I do lose. But since

² Bonjour 1985, 179-88; Bonjour and Sosa 2003, 87-96; and elsewhere.

there is a nearby possible world (or “small cluster” of possible worlds, as Pritchard says) in which I believe I will lose when in fact I will win, Pritchard can explain the intuition that my belief does not qualify as knowledge (163).

Setting aside the apparent equivocation between “nearly all” and “all,” the problem is that Pritchard now seems to have set the requirements for knowledge too high, for he is essentially claiming that *S* fails to know that *p* if there are any – or nearly any – possible worlds in which *S* believes that *p* when *p* is false. Pritchard is well aware of this problem, and he develops a strategy to try to overcome it (164-65). It seems to me, however, that this strategy is ineffective with respect to a variety of fairly obvious and mundane cases. Suppose, for instance, that I am driving on the freeway and that the sign before me indicates that my exit is a mere quarter of a mile away. I form the (true) belief that I will shortly be exiting the freeway. Suppose, moreover, that I have been making this commute for years and have never been this close to my exit and had any trouble making it off the freeway in a reasonably short amount of time. It seems clear, on the one hand, that my belief does not satisfy the revised, more demanding version of Pritchard’s safety principle. After all, the driver next to me could spill his coffee and lose control of his car; the truck ahead of me could drop its load; a devastating earthquake could occur; etc. These represent a few of several “nearby” possible worlds (i.e., worlds that are not very different from the actual world) in which I mistakenly believe that I will shortly be making my exit. And yet, in light of my accumulated experience on the matter (or alternatively, in light of the comparatively very large number of nearby possible worlds in which my belief is true), there is little reason to doubt that my true belief amounts to knowledge. It appears, then, that in modifying his view to deal with the lottery problem, Pritchard has made the requirements for knowledge too demanding.

A third point concerns Pritchard’s conception of reflective luck. We saw above that Pritchard initially offers a purely *epistemic* formulation of this notion, according to which a belief is reflectively lucky if “[g]iven only what the agent is able to know by reflection alone, it is a matter of luck that her belief is true” (175). In subsequent parts of the book, however, Pritchard characterizes the notion of reflective luck in a variety of more overtly *normative* terms. He indicates, for instance, that *S*’s belief that *p* is reflectively lucky just in case: (a) *S* fails to be “cognitively responsible” (in the standard deontological sense) in how she forms or maintains her belief that *p* (184); (b) *S* is unable to “take cognitive responsibility for” the truth of her belief

that *p* (181, 184, 202); (c) *S* is unable to take “credit” for the truth of her belief that *p* (184-85); (d) *S* cannot “claim knowledge” of *p* (182, 185); and (e) *S*’s belief that *p* does not amount to a personal (vs. a “sub-personal”) “cognitive achievement” of *p* (184, 190, 195).

Clearly (a) – (e) are closely related; but they are not identical. Nor would they seem to be coextensive. While I cannot explore the details of the relevant cases here, it seems, for instance, that certain items of testimonial knowledge might be lucky in the sense of (c) but not (a): that is, one might be cognitively *responsible* in, say, accepting the testimony of a reliable expert on some topic about which one otherwise knows very little even though one cannot take *credit* for the truth of the relevant claim or claims (only the expert is in a position to do that). Similarly, many of our cognitively spontaneous sensory beliefs are not lucky in the sense of (a), (b), or (d), but are lucky in the sense of (e), in that they do not amount to personal (vs. sub-personal) cognitive “achievements.”

More importantly, there is a notable gap between these ways of characterizing reflective luck and the more central, epistemic characterization, according to which *S*’s belief that *p* is reflectively lucky if, relative to *S*’s *grounds* or *evidence* regarding *p*, the truth of *p* is a matter of epistemic luck. For instance, in situations of extreme “epistemic poverty,” a person might ultimately lack good evidence for one of his beliefs and yet not be flouting any epistemic *duties* or behaving in an epistemically *irresponsible* way. Or again, one might have good evidence for a particular testimonial belief and yet not be in a position to take *credit* for its truth; or one might have good experiential grounds for accepting certain cognitively spontaneous beliefs about, say, the appearance of one’s surroundings even though these beliefs would not qualify as personal (vs. sub-personal) cognitive *achievements*.

This equivocation is not without consequence. At certain points, it obscures what is epistemologically most central vis-à-vis the notion of reflective luck. For instance, because of the strong connection he perceives between reflective luck and cognitive responsibility, Pritchard devotes an entire chapter (Chapter 7) to virtue-theoretic accounts of knowledge. His reasoning is that since these theories focus on the epistemic *agent* (the very *locus* of cognitive responsibility), they may be suited to overcome the problem of reflective luck. But one of the two main views he discusses in this chapter (viz., a reliabilist version of virtue epistemology) is a version of epistemological *externalism*. This is problematic because in the immediately prior chapter Pritchard argues at length for the claim that externalist accounts of knowledge are *incapable* of

dealing with the problem of reflective luck. Therefore, while the discussion in question does shed some interesting light on the terrain of virtue epistemology, it is apparently irrelevant to the problem of reflective luck, at least as this problem is originally (and, I think, correctly) formulated.

A fourth and final worry concerns how exactly Pritchard intends for us to understand the overall significance of his project. The text suggests both a weaker and a stronger interpretation. According to the former (which is the interpretation made above), Pritchard's inquiry into the nature and varieties of epistemic luck is important primarily because it helps illuminate certain preexisting epistemological positions and debates, especially the debate between internalism, externalism, and skepticism. On this reading, the concept of epistemic luck provides a kind of interpretive lens through which to view and gain a deeper understanding of certain issues that are themselves logically distinct from the concept of epistemic luck. According to a stronger interpretation, however, the concept of epistemic luck plays an epistemologically or conceptually more basic role such that one cannot adequately understand or even properly formulate the relevant issues and debates absent an appeal to epistemic luck.

At certain points, Pritchard seems merely to be defending the former, weaker reading of his central claim (e.g., 221); and as noted above, I think he succeeds in this defense, and that the result is an important contribution to epistemology. At other times, however, his discussion seems to favor the stronger reading (e.g., 7-9, 95, 119). But I see no reason to accept the relevant stronger claim about the dialectical importance of the notion of epistemic luck, and I certainly do not find any argument in Pritchard's discussion to support it. Indeed, I think Pritchard himself does an excellent job (e.g., in Chapter 4) of showing how the central epistemological issues with which he is concerned can be formulated and understood in an entirely adequate way with little or no reference to epistemic luck. Furthermore, while usually helpful and illuminating, the application of the concept of epistemic luck to other issues in epistemology at times comes across as forced or "tacked on." This is the case, in my estimation, with respect to much of the discussion of epistemic luck in Part I (e.g., 26, 29, 30, 39).

These shortcomings notwithstanding, *Epistemic Luck* is a rich, engaging, and groundbreaking work. It is a fine example of the kind of original and exciting work being done at the frontier of epistemology today.

Jason Baehr

Department of Philosophy

Loyola Marymount University

One LMU Drive

Los Angeles, CA 90045

jbaehr@lmu.edu

References

- BonJour, Laurence and Ernest Sosa. 2003. *Epistemic Justification* (Oxford: Blackwell Publishing).
- BonJour, Laurence. 1985. *The Structure of Empirical Knowledge* (Cambridge, MA: Harvard UP).
- Locke, John. 1996. *An Essay Concerning Human Understanding*, ed. Kenneth Winkler (Indianapolis, IN: Hackett Publishing).