

DUNCAN PRITCHARD

ANTI-LUCK EPISTEMOLOGY

ABSTRACT. In this paper I do three things. First, I offer an overview of an anti-luck epistemology, as set out in book, *Epistemic Luck* (Oxford UP, 2005). Second, I attempt to meet some of the main criticisms that one might level against the key theses that I propose in this work. And finally, third, I sketch some of the ways in which the strategy of anti-luck epistemology can be developed in new directions.

0. INTRODUCTION

It is a platitude in epistemology to say that knowledge excludes luck. The motivation for this intuition is relatively straightforward. After all, knowledge is an achievement, something for which the agent can take credit, and yet genuine achievements are not attained via luck. That this platitude about epistemic luck is widely held can be seen by considering how we evaluate theories of knowledge in terms of whether they are able to accommodate this claim. If they can't—that is, if they allow lucky knowledge—then this is taken to be a decisive ground for rejecting the view. Nowhere is this more apparent than in the debate regarding Gettier-style counterexamples to the tripartite account of knowledge. Consider, for example, this passage from Jonathan Dancy (1985, p. 134, *my italics*):

[...] justification and knowledge must somehow not depend on coincidence or luck. *This was just the point* of the Gettier counterexamples; nothing in the tripartite definition excluded knowledge by luck.

Gettier's counterexamples are taken to be so devastating to the classical tripartite view of knowledge precisely because they show that such an account is consistent with there being lucky knowledge. Thus, as a commitment to the epistemic luck platitude would dictate, the classical tripartite view is unsustainable.

Interestingly, however, even despite the prevalence of this intuition about epistemic luck, very few commentators have explored what it precisely means to say that knowledge is incompatible with luck. In particular, no commentator, so far as I am aware, has offered an account of what luck is and on this basis identified what it means for a true belief to be non-lucky.¹ Instead, the notion of luck—and thus, *a fortiori*, the notion of epistemic luck—is largely taken as an undefined primitive in the contemporary epistemological debate in such a

fashion that the epistemic luck platitude, while able to play the negative role in evaluating theories of knowledge that was just mentioned, is unable to serve any substantive theoretical role. Intuitively, however, if we could add some flesh to this platitude via an elucidation of the notion of luck, then it ought to be possible to cast light on a number of central debates in epistemology.

It was this observation that prompted me to start thinking about epistemic luck in more detail, and which ultimately resulted in my book on the anti-luck epistemological project, *Epistemic Luck* (2005a). In what follows I aim to do three things. First, to offer an overview of this book, and thus of anti-luck epistemology as I conceive of it. Second, to attempt to meet some of the main criticisms that one might level against the key theses that I propose in this work. And finally, third, to sketch some of the ways in which the strategy of anti-luck epistemology can be developed in new directions.

1. LUCK

Central to the anti-luck epistemological research project is, of course, an account of luck, since without this very little can be usefully said about epistemic luck. What struck me once I began exploring this issue is just how little there is in the philosophical literature regarding this notion, especially since it seems to be centrally relevant to a number of philosophical debates, such as debates about causation, agency, just deserts, free will, and, of course, knowledge.²

Consider a paradigm case of a lucky event, a lottery win. Suppose, for example, that I am jumping up and down right now holding the winning lottery ticket. What is it about this event that makes it lucky? Well, intuitively, the event is lucky because, roughly, this is an event which obtains in the actual world but which does not obtain in a wide class of near-by possible worlds where the relevant initial conditions for this event are the same as in the actual world (where I continue to buy a lottery ticket, the lottery remains free and fair and with long odds, and so forth). Indeed, in *most* near-by possible worlds that meet this description, I am right now tearing up my lottery ticket in disgust.

Contrast this case with a paradigm example of an event which isn't lucky, such as when a skilled archer, in good environmental conditions and in tip-top mental and physical condition, hits the target with her arrow. In this case, not only does the archer hit the target in the actual world, but also in nearly all, if not all, of the near-by possible worlds in which we

keep the relevant initial conditions fixed (such as her skill, her physical and mental condition, the environmental conditions, and so forth). There is no wide class of near-by possible worlds in which the relevant initial conditions for this event are the same as in the actual world and yet the event in question does not obtain.

Cases like these suggest that we should elucidate luck along the following lines:

- (LE) An event is lucky *iff* it obtains in the actual world but does not obtain in a wide class of near-by possible worlds in which the relevant initial conditions for that event are the same as in the actual world.

This is obviously not a complete analysis of luck, and nor is it meant to be (I say an awful lot more about the notion of luck in the book³). For one thing, it includes that weasel word ‘relevant’ in the analysis, which is dreadfully vague. Moreover, as I discuss in the book, (LE) is also incomplete in at least one crucial respect, in that it is only events which are significant to agents in some way that count as lucky. The mere fact that some odd event has obtained which would not obtain in a wide class of relevant near-by possible worlds would not normally be sufficient for us to count it as lucky, since it may be an event that is of no significance to us (think, for example, of something strange that is happening right now on one of Saturn’s moons). This rough and partial elucidation of luck will suffice, however, to enable us to get a handle on what an anti-luck epistemology would look like, which is our goal here.

2. EPISTEMIC LUCK

To begin with we should note that the core worry about luck as regards knowledge possession relates to luck in the *truth* of the relevant belief. We are not, for example, all that worried, from an epistemic point of view at least, about agents who only luckily come across the evidence which supports their knowledge. Suppose, for instance, that it was only a matter of luck that the detective stumbled across the crucial piece of evidence which proves the defendant’s guilt. So long as her resultant true belief in the defendant’s guilt is not lucky, then this poses no problem for the claim that she knows what she believes. In contrast, suppose her belief was only luckily true—suppose, for example, that her belief was based on prejudice rather than evidence, but was true nonetheless—then this would be inconsistent with her possessing knowledge in this regard.

This point is not confined to evidence acquisition. It may, for example, be a matter of

luck that we are even in a position to form true beliefs in the first place (we are lucky to be alive say), but this kind of luck does not by itself impair our ability to acquire knowledge. The point generalises. Luck only undermines knowledge when it impacts directly on the truth of the belief in question; luck in the mere ‘initial conditions’ for knowledge is not by itself knowledge undermining. Part of the initial stages of the epistemic luck project are thus devoted to highlighting various *benign* ways in which one’s knowledge can be lucky—i.e., those ways in which the knowledge in question is lucky but still *bona fide* knowledge nonetheless. This part of the project is important because, as I note in the book, a number of key examples that epistemologists fight over—that is, where the one party claims that because of the epistemic luck involved the agent in question doesn’t know, with the other party claiming that luck isn’t (substantively) involved, and thus the agent does know—in fact trade upon ambiguities in the details of the case in hand, such that it isn’t clear whether it is a benign or non-benign form of epistemic luck that is involved.⁴

With these benign forms of epistemic luck set to one side, we can then get down to the substantive task of outlining the core non-benign variety of epistemic luck, which is luck in truth of one’s beliefs. With (LE) in mind we can give a rough specification of what would count as a lucky true belief as follows:

(LTB) *S*’s true belief is lucky *iff* there is a wide class of near-by possible worlds in which *S* continues to believe the target proposition, and the relevant initial conditions for the formation of that belief are the same as in the actual world, and yet the belief is false.

For obvious reasons, (LTB) won’t apply—at least without being adapted in some way—to necessary propositions (whether the necessity is logical, metaphysical or nomological) since, simply in virtue of the kind of proposition that it is, there are no near-by possible worlds in which the target proposition is false. The initial focus is thus an account of what it is for a true belief in a (fully) contingent proposition to be lucky, with the further question of how to extend the account to true beliefs in propositions which are not (fully) contingent left to a later occasion. For ease of expression, in what follows we will refer to those near-by possible worlds in which the relevant initial conditions for the formation of *S*’s belief are the same as in the actual world, and *S* continues to form a belief in the target proposition, as the ‘relevant’ near-by possible worlds.

Consider how (LTB) captures some of the standard cases that are cited of lucky true belief. Take a belief formed as a result of a lucky guess, for example.⁵ My true belief that the horse ‘Lucky Lass’ will win the 4.20pm race—based entirely on a guess—is not knowledge

because, aside from anything else that might be epistemically wrong with it, it is only luckily true. In terms of (LTB), this means that although my belief is true in the actual world, in a wide class of relevant near-by possible worlds (i.e., worlds in which I'm still forming this belief via a guess, for example), my belief is false. After all, there will be a wide class of relevant near-by possible worlds in which Lucky Lass does not win the race but where I continue to believe that she will because this is what my guess leads me to believe.

Contrast this case with that of someone who forms her belief that Lucky Lass will win based on the fact that she has personally fixed the race (she's drugged the other horses, say). The truth of this agent's belief is not a matter of luck, and this is reflected in the fact that her belief is not only true in the actual world, but also true in most relevant near-by possible worlds as well. For example, those relevant near-by possible worlds in which Lucky Lass does not win the race will tend to be worlds in which the race wasn't fixed, and since our protagonist is personally involved in this race-fixing we would expect her not to form a belief that Lucky Lass will win the race in these cases. There is thus nothing on this score to prevent this agent from being counted as possessing knowledge.

Next consider a Gettier-style case, such as the 'stopped clock' example. Here we have an agent who forms her justified true belief about what the time is by looking at what is, unbeknownst to her, a stopped clock. Intuitively this is not knowledge because the agent's belief is only luckily true. (LTB) captures this fact. Although the agent happened to form a true belief in the actual world by looking at this stopped clock, in a wide class of relevant near-by possible worlds (i.e., worlds in which she is still forming her belief by looking at the stopped clock for example), her belief will be false (think, for example, of those near-by worlds in which the clock stopped a couple of minutes earlier or a couple of minutes later).

In contrast, had the agent in this example formed her true belief by looking at a working clock then we would not have counted her true belief as lucky, and this is borne out by (LTB). For not only would the agent have formed a true belief in the actual world, but in most relevant near-by possible worlds her belief would continue to be true. Had the time been slightly different, for example, then this would have been reflected in the time shown by the clock, and thus the agent would not have continued to form the same belief about the time as she did in the actual world.

With (LTB) in mind, then, coupled with our initial intuition that knowledge excludes luck, we can give a rough formulation of a necessary condition for knowledge—a condition that is not met in the 'lucky guess' and 'stopped clock' cases, but which is met in their counterpart cases where no luck is involved (and where, pending further details about the

examples at least, we are inclined to attribute knowledge):

(AL) *S*'s true belief is non-lucky *iff* there is *no* wide class of near-by possible worlds in which *S* continues to believe the target proposition, and the relevant initial conditions for the formation of that belief are the same as in the actual world, and yet the belief is false.

By offering a rough general analysis of luck, (LE), and applying that to the case of true belief, (LTB), we are thus able to give a rough formulation of an anti-luck condition on knowledge, (AL).

3. EPISTEMIC LUCK AND THE SAFETY PRINCIPLE

The observant reader will have already spotted that (AL) bears a striking similarity to an epistemic principle that has been defended as a necessary condition on knowledge in the recent literature—*viz.*, the safety principle, a version of which has been proposed, for example, by Ernest Sosa (e.g., 1999).⁶ Very roughly, the safety principle can be characterised as follows:

(SP) *S*'s belief is safe *iff* in most near-by possible worlds in which *S* continues to form her belief about the target proposition in the same way as in the actual world the belief continues to be true.

As has been widely noted, there are a number of advantages to adducing such a principle on knowledge, the chief among them being the anti-sceptical appeal of using safety to show that our beliefs in the denials of sceptical hypotheses—while not sensitive to falsehood or justified by adequate reflectively accessible grounds—do have some positive epistemic support in virtue of being safe. The advocacy of this principle thus lies at the heart of an anti-sceptical view—which I have elsewhere termed “neo-Mooreanism”—whose primary anti-sceptical thesis is that we can, *contra* the sceptic, know the denials of sceptical hypotheses after all.⁷

Since we might reasonably suppose that the ‘relevant initial conditions for the formation of the belief’ roughly corresponds to the ‘way’ in which the belief was formed,⁸ and that ‘in no wide class of near-by possible worlds not-*X*’ is roughly equivalent to ‘in most near-by possible worlds *X*’, then (SP) and (AL) are essentially the same. If the preceding account of why we should endorse something like (AL) is right, therefore, then it seems that the underlying motivation for (SP) comes from our anti-luck intuition about knowledge, expressed in (AL). That is, we find the safety principle (SP) intuitive *because* it captures our

anti-luck intuition about knowledge which is explicitly expressed in (AL). As with (AL), in what follows in our discussion of (SP) we will refer to those near-by possible worlds in which the relevant initial conditions for the formation of *S*'s belief are the same as in the actual world, and *S* continues to form a belief in the target proposition, as the 'relevant' near-by possible worlds.

With (AL) in mind as the underlying motivation for (SP), we can make headway with some of the problems that have been raised for safety-based theories of knowledge. In particular, one such problem concerns how robust the principle needs to be if it is to deal with an appropriate range of cases. Consider, for instance, the problem that the lottery example poses for (SP). The lottery example concerns one's belief, prior to the draw being announced, that one owns a losing ticket (where the lottery is free and fair, and with long odds). Intuitively, this is not a case of knowledge even despite the excellent statistical grounds one has in favour of one's belief because, even when the belief is true, it is nevertheless a matter of luck that this belief is true. After all, there are relevant near-by possible worlds in which one's belief is false (i.e., those worlds in which one owns the winning lottery ticket). If this is right, however, then it seems that demanding that the agent's belief be true in 'most' near-by possible worlds, as (SP) demands, will not suffice, since, intuitively, this condition is met in the lottery case where there are only very few relevant near-by possible worlds in which one forms a false belief in the target proposition. (SP) thus predicts knowledge in this case, even though our intuitions, guided by the anti-luck platitude, dictate otherwise.

Of course, (SP) is only meant to be a necessary condition on knowledge, and so one could argue that there are further conditions on knowledge available to deal with this case, but since the intuition here is that knowledge is lacking *because* of the presence of luck, and since, as we have noted, (SP) turns out to be a natural way of understanding the anti-luck condition, this dialectical move is not very appealing. Accordingly, it seems that one must instead understand safety along much stronger lines as demanding that the agent's belief be true not just in most of the relevant near-by worlds, but in nearly all (if not all) of them. We thus get (SP*):

(SP*) *S*'s belief is safe *iff* in nearly all (if not all) near-by possible worlds in which *S* continues to form her belief about the target proposition in the same way as in the actual world the belief continues to be true.

On the face of it, this kind of refinement to safety could well be in the spirit of (AL), in that given the kind of substantial cognitive achievement at issue in the possession of knowledge it is not that implausible to suppose that any trace of luck must be eliminated

before one can have knowledge, and thus that there should be *no* class of relevant near-by possible worlds in which the agent's belief is false, not just that there be no wide class. As a number of commentators have noted, however, tightening-up (SP) in this way would seem to make it unable to deal with more mundane cases of knowledge possession.⁹ Consider, for example, the 'rubbish chute' case described by Sosa (2000, p. 13):

On my way down to the elevator I release a trash bag down the chute from my high rise condo. Presumably, I know my bag will soon be in the basement. But what if, having been released, it still (incredibly) were not to arrive there? That presumably would be because it had been snagged somehow in the chute on the way down (an incredibly rare occurrence), or some such happenstance.

Do we want to say that Sosa has knowledge in this case? Intuitively we do, but if safety is understood along very robust lines as (SP*) then, it has been claimed, this seems to be ruled-out since surely there are quite a few near-by possible worlds—those in which the bag snags—in which Sosa continues to form his belief on the same basis and yet forms a false belief as a result. We thus seem to be stuck between two opposing intuitions. On the one hand, that safety must be understood rather robustly, as (SP*), in order to eliminate the luck in play in the lottery case; on the other, that safety must be understood relatively weakly, as (SP), in order to accommodate more mundane cases of knowledge, like that in the case of the rubbish chute example.

Closer inspection of the rubbish chute example reveals, however, that the challenge it poses to (SP*) is not nearly as clear as many have thought. To begin with, it is worthwhile looking again at how Sosa describes the example, especially where he says that the snagging of the bag would be “an incredibly rare occurrence”. If this is right then it seems that Sosa's belief *isn't* lucky since the possibility that the bag snags is actually quite remote. I take it that what we are supposed to be imagining here is that, say, there is an imperfection in the lift shaft, but one so slight that it would be hard for a bag to snag on it, and situated in a position in the shaft where hardly any bag would be likely to make contact with the snag anyway. On this reading of the example the intuition that Sosa knows that his bag is in the basement is not under threat, but now it isn't so obvious that there are very many (if any) relevant near-by possible worlds in which his belief is false—the worlds in which his bag snags on the chute seem relatively far-off.

In contrast, we can stipulate the details of the case such that it is clear that Sosa's belief *is* lucky by demanding that there is a wide class of relevant near-by possible worlds in which the agent forms a false belief because the bag has snagged in the chute. In order to make this supposition plausible we would have to imagine, for example, that there is a snag in

the chute that the bag is *almost* snagging on each time. If that's right, however, then it would be odd to think that Sosa does know that his bag is down in the basement, since it is clearly a matter of luck that his belief is true in this regard given the nearness of the relevant error-possibility.

Everything thus depends on how we are understanding the details of the case. Filling-in the detail here thus highlights that formulating safety in a fairly strict fashion as (SP*) so as to deal with the lottery case may not result in our denying knowledge in cases like the rubbish chute example after all. This, in any case, is the formulation of safety offered in the book. I consider some problems one might pose for this formulation below, and also examine a further refinement of safety which is also in the spirit of anti-luck epistemology.

4. NEO-MOOREANISM AND REFLECTIVE EPISTEMIC LUCK

As noted above, one of the key advantages of safety—and thus of the anti-luck epistemology that motivates the adoption of safety—is that it offers a very straightforward way of dealing with the sceptical problem as it is usually understood in the contemporary literature. This construal of the sceptical challenge can be roughly expressed as follows, where 'E' is an 'everyday' proposition which we would all typically take the subject to know, and 'SH' is a radical sceptical hypothesis—such as the brain-in-a-vat (BIV) sceptical hypothesis—which is known by the subject to be inconsistent with the target E-type proposition:

The Template Radical Sceptical Argument

- (S1) S does not know \neg SH
- (S2) If S does not know \neg SH, then S does not know E.
- (SC) S does not know E.

Both of the premises here are plausible. The first is plausible because such hypotheses are *defined* such that there is no detectable difference between ordinary circumstances and counterpart sceptical circumstances. The second premise is plausible because it rests on the highly intuitive closure principle that, roughly, if one knows one proposition, and one knows that this proposition entails a second proposition, then one knows the entailed proposition. Since the argument is also valid, however, then a compelling and intuitive argument for a radical sceptical conclusion is in the offing. After all, since one can vary the sceptical hypothesis at issue here in order to undermine knowledge of just about any E-type proposition, it follows that we are unable to know just about any of the propositions which we

would ordinarily take ourselves to know.

Adopting safety as characterised above gives us a way of denying the first premise of this sceptical argument outright, and thereby enables us to meet the problem without resorting to one or other of the main forms of epistemological revisionism presently on offer in the literature, such as, principally, non-closure (Dretske 1970; Nozick 1981), contextualism (DeRose 1995; Lewis 1996; Cohen 2000), and contrastivism (Schaffer 2005). As noted above, I have termed this direct response to the sceptical problem, “neo-Mooreanism”.

It is important to recognise the crucial role that an anti-luck epistemology plays here in motivating safety, since it is only with this motivation in place that we can account for why we can, contrary to intuition, know the denials of sceptical hypotheses. The point is that so long as these sceptical hypotheses do indeed only obtain in far-off possible worlds, then our true belief that we are not the victim of a sceptical hypothesis will not be lucky and thus will not be unsafe either. In all near-by possible worlds we will believe this proposition, and in all of those worlds this belief will be true. Of course, we lack an internalist justification for our belief in this respect, since our internalist justification—which will concern such facts about our experiences as that it does not seem to one as if one is a BIV—will fail to discriminate between sceptical and non-sceptical scenarios, and thus will give us no reason to believe that we are not presently in a sceptical world (or, for that matter, in a world that could very easily have been a sceptical world). Nevertheless, this fact alone does not mean that our true belief is lucky, and thus should not disqualify it for knowledge. On this view, our knowledge is to some degree hostage to the way the world is, but that claim is, it transpires, not equivalent to the sceptical claim that knowledge of the world is impossible.

What this point illustrates is that anti-luck epistemology naturally leads to an externalist epistemology—that is, an epistemology which does not demand of knowers that they are able to offer an internalist justification for their beliefs. In this respect, then, our anti-sceptical knowledge is akin to the kind of knowledge possessed by the notorious chicken-sexer, an agent who has a highly reliable ability to distinguish between male and female chicks but who lacks any good internalist grounds in support of her beliefs so formed. Internalists have typically balked at the idea of ascribing knowledge to such agents, and this is because they regard meeting an internalist justification condition as being essential to knowledge possession. Externalists, in contrast, have been more open to knowledge ascription in this case, and we can now offer a rationale for this apparent epistemic permissiveness. After all, if the chicken sexer really does have the reliable cognitive ability in question, and is forming her beliefs about the sexes of the chicks on the basis of this ability, then the fact that

those beliefs are not internalistically justified does not mean that the true beliefs she thereby forms are lucky. Instead, we would expect those beliefs to meet our anti-luck condition—and thus the safety condition—on the grounds that such a cognitive ability ought to be able to enable the chicken sexer to form true beliefs about the target proposition in near-by possible worlds as well as the actual world. Contrary to first appearances, then, and despite the protestations of some internalists, externalists are not sanguine about allowing lucky knowledge.

Nevertheless, one should, I think, grant that there is something correct underlying the very natural complaint that the internalist makes about knowledge of this brute sort. As one might put the point, although it is not an *objective* matter of luck that the agent's belief is true in the chicken sexer case—and the same goes for one's belief that one is not the victim of a sceptical hypothesis provided that the sceptical worlds are indeed far-off—it is a *subjective* matter of luck that this agent's belief is true since, *ex hypothesi*, the agent in question lacks good reasons for thinking that her belief is true. I think the externalist, and thus the neo-Moorean, can concede this point without undermining her commitment to externalism. Indeed, I think conceding this point actually *strengthens* the case for externalism.

The way this is achieved is by granting the existence of a second variety of epistemic luck, one that is epistemically problematic in the sense that it is epistemically desirable to eliminate this kind of luck in one's true beliefs, but a form of epistemic luck which is nonetheless compatible with knowledge possession. I refer to such luck as *reflective epistemic luck*, and contrast it with the *veritic epistemic luck* described above. Essentially, the difference between reflective and veritic epistemic luck lies in how the ordering of the possible worlds is determined purely in terms of what is known by reflection alone by the agent concerned, thereby capturing the idea that an agent's belief can be subjectively lucky even if it is not objectively lucky.

Setting the problem of scepticism aside for a moment (the reason for this *caveat* will become apparent in a moment), consider how this distinction applies to the chicken sexer case by considering a counterpart to the chicken sexer that we described above who is “enlightened” in the sense that not only has she the ability in question but also has good reflectively accessible grounds in support of the beliefs so formed. There is clearly something epistemically advantageous about being an enlightened chicken sexer rather than a naïve chicken sexer, and the distinction just drawn between two different types of luck can account for this intuition. After all, although neither of these chicken sexers will be forming their

beliefs about the sex of the chicks in a veritically lucky fashion—since they both genuinely have the cognitive trait in question—it will be true that the naïve chicken sexer’s true beliefs will be reflectively lucky. This is because if we order the possible worlds at issue in our assessment of whether or not this true belief is lucky not in terms of the facts that obtain in the actual world but purely by considering what the agent can know by reflection alone, then it will follow that just about *any* possible world can count as a near-by world when it comes to the naïve chicken sexer’s belief, including worlds in which she does not possess the cognitive trait in question and is in fact forming her beliefs in a random way. The same is not true of the enlightened chicken sexer, since the good reflectively accessible grounds she has in favour of her belief ensures that the subjective ordering of worlds will be much the same as the objective one—she has, for example, good grounds for thinking that she has the cognitive trait in question, and this will mean that the near-by possible worlds on this ordering (as with the objective ordering) will tend to be worlds in which she retains this trait.

The distinction between veritic and reflective epistemic luck can thus account for our intuition that there is something lucky about an agent’s beliefs in those cases of pure externalist knowledge, like chicken sexer cases, without also thereby granting that such knowledge is lucky *per se*. We can thus capture the internalist intuition without thereby abandoning externalism. Indeed, since one can capture this intuition without thereby conceding that pure externalist knowledge is lucky knowledge, it seems that the externalist is now on dialectally strong ground. After all, the externalist can argue that the internalist, by failing to make this distinction between two different types of luck, has confused a type of luck which it is *desirable* to eliminate from one’s true beliefs from a type of luck which it is *essential* to eliminate from one’s true beliefs if those beliefs are to count as knowledge. The externalist can thus give expression to the commonly held claim that internalists are setting the standard for knowledge too high, and diagnose this mistake on the internalists’ part in terms of a failure to distinguish between two types of epistemic luck.

I argue in the book that this distinction can do a lot of work in casting light on a number of crucial debates in epistemology. I have already mentioned the externalism/internalism debate in this regard. A related debate is that between those virtue epistemologists who take their lead from reliabilism—such as Sosa (1991), Alvin Plantinga (1993), and John Greco (1999)—and those virtue epistemologists—such as Lorraine Code (1987) and Linda Zagzebski (1995)—who offer instead what is known as a “responsibilist” or neo-Aristotelian virtue epistemology, one that stresses the roles of the intellectual virtues in

the production of knowledge over that of the (mere) cognitive faculties. This is, in effect, one manifestation of the externalism/internalism debate within a specific area of epistemology. Moreover, as I show in the book, like the externalism/internalism debate more generally, one can diagnose what is going on in the arguments proffered by these two camps in terms of a debate about whether merely eliminating veritic luck is sufficient for knowledge, or whether one should also eliminate reflective luck as well. So construed, however, a lot of the impetus for this debate starts to diminish, since the agent reliabilists can recognise the value of eliminating reflective epistemic luck—and thus the importance of the intellectual virtues—without thereby conceding that knowledge that is purely the product of a reliable cognitive faculty, such as a chicken-sexing faculty, is not *bona fide* knowledge nonetheless.

This is not to say that I wholeheartedly side with the agent reliabilist in this regard, since I also argue that properly understood anti-luck epistemology calls into doubt the very project of virtue epistemology. In essence, my claim is two-fold. First, that virtue epistemologists are going to have to cite an anti-luck condition in order to deal with the Gettier cases and thus offer a fully-fledged theory of knowledge. Second, that insofar as the virtue epistemologist concedes this much to anti-luck epistemology, then the case for the view is lost. This is because an anti-luck condition can, I argue, deal with all the cases the agent reliabilist wants to cover, but can do so without making any essential mention of cognitive faculties or intellectual virtues. This is not to deny, of course, the importance of the virtue-theoretic story in accounting for how creatures such as us come to have knowledge; the point is rather that this virtue theoretic story comes *after* an anti-luck account of knowledge, and need form no essential part of that account. Anti-luck epistemology thus not only casts light on a key debate in contemporary epistemology, but also serves to undermine the defining claim of one the major movements in the recent literature.¹⁰

The advantages of offering the distinction between reflective and veritic luck do not end there. For example, I also claim that one can make sense of what is at issue in the parallel debate in ethics regarding moral luck using this distinction. Here I focus on the two classic papers on this topic by Thomas Nagel (1976) and Bernard Williams (1976), arguing that they motivate an illusory problem regarding moral luck by implicit appeal to a genuine problem regarding reflectively epistemic luck. The anti-luck epistemology proposed in the book thus has ramifications which extend beyond the purely epistemological realm.

5. REFLECTIVE EPISTEMIC LUCK AND EPISTEMIC *ANGST*

The tone of the book is not, however, wholly upbeat, since I also maintain that highlighting the distinction between veritic and reflective epistemic luck points to where the real moral of the sceptical argument lies. Recall that we have seen how an anti-luck epistemology which regards the elimination of veritic epistemic luck as essential to knowledge can deal, in a neo-Moorean fashion, with the sceptical problem outlined above in a fairly immediate way. Many feel, however, that responses to the sceptical problem of this sort somehow by-pass the real source of the sceptical difficulty. I argue in the book that there is something essentially correct about this observation, but that it cannot be properly expressed without recourse to the distinction between veritic and reflective epistemic luck. In particular, I claim that the real sceptical problem is not a problem about possessing non-lucky knowledge in the veritic sense, but is in fact targeted at the kind of epistemic standing that is at issue when one can eliminate reflective epistemic luck.

Consider, for example, those—such as Barry Stroud (1984) and Richard Fumerton (1995)—who advance a so-called “metaepistemological” scepticism against externalist responses to the sceptical problem and claim that the sort of knowledge rescued from the sceptic’s grasp by these anti-sceptical theories is such as to make the anti-scepticism at issue inevitably intellectually dissatisfying. One could characterise this form of objection against externalist anti-sceptical accounts as being little more than an expression of a prior commitment to epistemic externalism, and thus nothing for the externalist to unduly concern herself with. By factoring the veritic/reflective epistemic luck distinction into the sceptical debate, however, we can identify the source of the complaint here. Essentially, it boils down to the claim that the kind of knowledge that is rescued from the sceptic in externalist anti-sceptical theories puts knowledge on a par with the brute knowledge possessed by the naïve chicken-sexer where, recall, such knowledge, while not veritically lucky, is indeed reflectively lucky.

The externalist will in all probability concede this point, of course, and simply maintain that it only applies to one’s knowledge of the denials of sceptical hypotheses—where it is granted that we perforce lack a good internalist justification in favour of the target belief—and not also to one’s knowledge of everyday propositions. It is far from clear, however, how such a ‘containment’ strategy is to work. How could it be that we have knowledge of everyday propositions which is not reflectively lucky, know that such

knowledge entails the denials of anti-sceptical hypotheses, and yet fail to have knowledge of the denials sceptical hypotheses which is also not reflectively lucky? If this is right, then the reflective luck inherent to our knowledge of the denials of sceptical hypotheses spreads out to infect our everyday knowledge. *All* our knowledge now becomes brute knowledge on a par with that possessed by the naïve chicken sexer. If this is right, then it is little wonder that externalist responses to the sceptical problem—and neo-Mooreanism in particular—while formally able to block the sceptical argument, are not such as to fully intellectually satisfy us.

One can put this point more succinctly by saying that there is a sense in which reflective epistemic luck is ineliminable. This claim might sound odd since we seemed to have highlighted a case in which reflective epistemic luck was eliminated above, when we contrasted the predicament of the enlightened chicken sexer and her naïve counterpart. Recall, however, that we were only able to motivate this distinction because we bracketed the sceptical problem. If we had not done this, then the distinction between the enlightened and naïve chicken sexer would radically diminish. After all, since the enlightened chicken sexer lacks good reflectively accessible grounds for thinking that she is not, say, a BIV, it follows that the BIV-world can count as a near-by possible world on the subjective ordering at issue in reflective epistemic luck, and the same goes for other possible worlds in which sceptical hypotheses are true. With this in mind, there will be a wide class of near-by possible worlds on the subjective ordering in which even the enlightened chicken sexer continues to form a belief in the target proposition, but where the belief formed is false. Thus, although there is a difference of degree between the enlightened and the naïve chicken sexer, it is not a difference which ensures that the enlightened chicken sexer's beliefs—or any other beliefs about the world that are internalistically justified for that matter—are not reflectively epistemically lucky provided we do not bracket the sceptical problem in advance. I call our implicit recognition of the eliminability of reflective epistemic luck, "*epistemic angst*."¹¹

I further illustrate this point by considering a number of what I claim are related issues. The first is the underdetermination-based sceptical argument that focuses on evidence rather than knowledge. I claim that this argument poses a more dramatic sceptical challenge because by considering the evidential basis of our knowledge it draws out the more fundamental problem that the sceptic is seizing upon—not the possibility of knowledge *simpliciter*, but rather the possibility of knowledge that is grounded in a way that is immune to reflective epistemic luck.¹²

Second, I relate this claim to classical Pyrrhonian scepticism which, I argue, can be

understood as explicitly focussed on a certain kind of cognitive responsibility which is associated with the elimination of reflective epistemic luck rather than simply on the possession of knowledge. I connect this thesis to a more general issue about the propriety of our claims to know in sceptical contexts. In the first part of the book I motivated the neo-Moorean stance by showing how contextualism is best understood as a pragmatic thesis about the context-sensitivity of the propriety conditions for claims to know (and ascriptions of knowledge more generally) rather than as a semantic thesis about the context-sensitivity of ‘knows’. Moreover, I argued that one can use this account to explain why it would be so problematic to explicitly advance the neo-Moorean anti-sceptical argument in the light of the sceptical challenge, even though the claims made in such an argument could nevertheless be true. As I show in the second part of the book, however, the anti-sceptical import of this apparently positive neo-Moorean thesis is in fact double-edged, in that it incorporates a significant concession to the sceptic by allowing that in the sceptical context the only legitimate response to the sceptical challenge is a mute one, thereby limiting the degree of cognitive responsibility that we can take for our beliefs. I claim that Pyrrhonian sceptical challenges, by being primarily directed at claims to know, are specifically targeted at cognitive responsibility of this sort.

Finally, third, I develop my quasi-sceptical thesis by considering the very different approaches to scepticism found in the work of John McDowell (e.g., 1995) and Wittgenstein (1969).¹³ In particular, I contrast the anti-sceptical appeal to factive empirical reasons which is a central plank of the McDowellian view with the account of “hinge” propositions offered by Wittgenstein, and which offers limited anti-sceptical comfort. As I show, the McDowellian thesis, far from neutralising the problem about epistemic luck that lies at the heart of the sceptical challenge, in fact just relocates the luck elsewhere in the epistemic architecture, and I highlight this point via a consideration of the role of “hinge” propositions in our everyday practices of offering grounds in favour of our beliefs.

6. VERITIC EPISTEMIC LUCK AND SAFETY

Perhaps the most controversial aspect of the book, and certainly that part of the book which commentators have most seized upon, is the idea that knowledge is, at root, non-veritically lucky true belief, where this is in turn cashed-out in terms of a formulation of the safety principle. Many balk at this suggestion because they think that there are obvious cases of

true beliefs that meet this principle but which are clearly not cases of knowledge. (See, for example, the excellent article by Avram Hiller and Ram Neta (200X) in this volume). Relatedly, others—such as Greco (200X) in his incisive contribution to this volume—have argued that there is no consistent way in which one can read the safety principle such that it deals with the full range of cases that one would expect it to deal with if it were indeed capturing what is central to knowledge. If these objections are well-founded, then the project as a whole is certainly in serious trouble, as are the particular claims that I make which arise out of the project, such as that an anti-luck epistemology highlights why we shouldn't define knowledge in terms of the cognitive faculties and intellectual virtues.

Before getting into the nitty-gritty of this issue, however, there are a few preliminaries that I want to mention in order to set the scene for this discussion. The first is to recall that I am only offering an elucidation of knowledge as it applies to fully contingent propositions, and so the most obvious counterexamples—those that, for example, appeal to nomic necessities and stubborn, but groundless, beliefs in these necessities in order to ensure that the target beliefs are safe—will not be applicable here.

The second is to point out that I am not trying to offer necessary and sufficient conditions for knowledge, where this is understood as ruling-out the possibility of there being, for example, penumbral cases about which intuitions might diverge. I am, instead, trying to get to the heart of the matter by giving an account of what is most central to knowledge, and this more modest aim is compatible with the existence of penumbral cases. Indeed, I would counsel analytical philosophers to remember that not all counterexamples to a position are on a par. Gettier cases, for example, have the import that they do not simply because they are counterexamples to the tripartite position—in that they are cases in which the conditions imposed by this position are met but knowledge is not, intuitively, possessed—but because they do something far more devastating in showing how a certain view of knowledge is, on closer inspection, entirely implausible. Not all counterexamples have this effect. Some merely highlight the vagueness inherent in the extension of philosophically interesting terms (and which philosophically interesting terms do not have vague extensions?), while others simply show that there are recalcitrant intuitions that need to be dealt with by the theory (and which philosophical theory does not have to contend with recalcitrant intuitions at some juncture?).

With these two *caveats* in mind, we can return to the issue at hand. Interestingly, we can by-pass the discussion of each counterexample on a case-by-case basis because of the fact that it is explicitly part of the view to ascribe knowledge in at least one type of case which has

all the hallmarks of the core counterexamples offered to safety in this respect—*viz.*, those cases which involve belief in the denials of sceptical hypotheses. The problematic feature shared by the paradigm counterexamples to safety and our putative knowledge of the denials of sceptical hypotheses is that we have a supposed case of knowledge where the knowledge is purely in virtue of a stubborn belief in the truth of a proposition which is, *ex hypothesi*, only false in far-off worlds. This is in contrast to standard cases in which one has knowledge in virtue of meeting the safety condition, in that typically there is a near-by possible world in which what one believes is not true and one does not believe that it is true on that basis. There is no parallel for the case of our knowledge of the denials of sceptical hypotheses because insofar as knowledge is possessed in this case then there are no near-by possible worlds in which the target proposition is false. So while safety will typically capture a sense in which a true belief is responsive both to counterfactual falsity as well as to counterfactual truth, cases like the sceptical case are exceptions, in that there is no relevant falsity to be responsive to. Without such responsive to falsity, however, the question arises as to whether we should treat the true belief in question as an instance of knowledge at all.

It should be clear that cases with this feature will necessarily be rare, since most propositions which are not false in near-by possible worlds are at least nomically necessary, and so not covered by the account of knowledge on offer here. Nevertheless, as the sceptical case illustrates, there can be propositions which are not false in near-by possible worlds which are nonetheless not necessary in any sense.

When it comes to this sort of objection to the view, I'm inclined to bite the bullet that is offered here, for two reasons. The first is that it is open to me to account for the epistemic lack in such cases without conceding anything substantive to the sceptic (or internalist) about knowledge by highlighting how such beliefs, while not veritically lucky, are indeed reflectively lucky. Thus, I can, at least in part, accommodate some of the counterintuitions being appealed to here without acceding to them.

The second reason why I'm inclined to stick to my guns on this score is that if this is a problem, then it is a problem shared by *any* anti-sceptical theory that allows that we can know the denials of sceptical hypotheses, such as contextualist and virtue-theoretic views. Since the sceptical point is that we fail to have *any* substantive degree of epistemic support for our beliefs in the denials of sceptical hypotheses, why should the mere lowering of the relevant epistemic standards ensure that an ascription of knowledge expresses a truth, as the contextualist claims? Similarly, what is epistemically virtuous about believing such anti-sceptical propositions? If one does not opt for an anti-luck account of such knowledge, then it

is unclear what sort of defence one could mount as regards our knowledge in this case. Moreover, as I show in the book, those anti-sceptical theories that do not take the route of allowing that we can know the denials of sceptical hypotheses face even more severe difficulties, so there is no comfort from this direction.

Nevertheless, I'm conscious of the fact that a number of commentators have been troubled by a worry of this sort, so it would be useful to say something further about this issue. One thing to note in this regard is that one can separate out the project of anti-luck epistemology as I outline it in the book with the specific working-through of this project that I offer. That is, one could be persuaded by my claim that anti-luck epistemology is a fruitful way of approaching the core problems in the theory of knowledge without thereby thinking that the particular anti-luck theory I advocate is correct. With this in mind, one possibility open to those concerned by this sort of objection, but who wish to retain an anti-luck epistemology, would be to try to tweak the formulation of the safety principle so that it could deal with this problem while still staying within the general spirit of anti-luck epistemology.¹⁴

Since I have no need of this further thesis, I will not explore the merits of this possibility further here. That said, I do think that there is one modification that should be made to our understanding of the safety principle as it is set out in the book which can help us to deal with a different sort of problem that is often raised against safety-based views of knowledge. I noted earlier that there are those, such as Greco, who are sceptical about the idea that there is a reading of safety which can accommodate both cases of 'ordinary' knowledge, such as the rubbish chute example, while also accounting for lottery cases. I explored, and discounted, this objection in the book, and I set out the general contours of the line of reasoning I employed there above. Even though I still think that this approach is essentially correct, I do think that there are some further comments I can usefully make on this topic which supplement that reasoning, and which might help it to be more persuasive.

Essentially, the idea is that one can, entirely in the spirit of anti-luck epistemology, offer a formulation of safety which gives greater weight to near-by possible worlds as opposed to far-off possible worlds, and in doing so accommodate some of the intuitions that lie behind the sort of objection pushed by Greco. For notice that what is distinctive about the lottery case isn't so much that there is a small class of relevant near-by possible worlds in which the target proposition is false (but believed nonetheless), but rather that the worlds in question are actually *very close* to the actual world. After all, the possible world in which I win the lottery is a world just like this one, where all that need be different is that a few coloured balls fall in a slightly different configuration.

Crucially, however, the *nearness* of the relevant possible worlds has an impact on our judgements about the presence of luck. One can see this point by considering cases of lucky events more generally. For example, that you narrowly avoided being hit by a bullet that was fired directly at you at close range by a competent marksman is clearly luckier than avoiding being hit by a bullet that was fired from further away by the same marksman, and which missed you by a couple of feet. What makes the first event luckier than the second is the fact that the world in which the bullet hits you is much closer in the first case than in the second.

If this is right, then we can account for what is going on in the lottery case without thereby accepting, as I do in the book (and thus above), that a non-lucky true belief is one that remains true in all, or nearly all, relevant near-by possible worlds, since the point about the lottery case is that the relevant near-by possible worlds in which your belief is false are worlds that are *very close* to the actual world. Accordingly, rather than opting for the strict formulation we can stick with the weaker (SP) but simply add the further stipulation that the belief must be true in all of the *very close* relevant near-by possible worlds. We thus get (SP**):

(SP**) *S*'s belief is safe *iff* in most near-by possible worlds in which *S* continues to form her belief about the target proposition in the same way as in the actual world, and in all very close near-by possible worlds in which *S* continues to form her belief about the target proposition in the same way as in the actual world, the belief continues to be true.

With (SP**) in mind, consider again the formulation of the rubbish chute example in which it clearly isn't lucky that Sosa's belief that his rubbish is in the basement is true (the formulation in which the possibility of the bag snagging is actually quite remote). This formulation was such that although one might plausibly suppose that there were some relevant near-by possible worlds in which Sosa formed this belief and yet his belief was false, it was certainly the case that (unlike in the lottery example) none of the *very close* possible worlds fitted this description. Paying close attention to the way in which our judgements about luck are affected by the closeness of the relevant near-by possible worlds thus highlights a relatively minor modification that we need to make to our formulation of the safety principle in order to meet this particular problem.

7. PROSPECTS FOR AN EPISTEMIC LUCK RESEARCH PROGRAMME

I want to close by considering some of the various avenues along which one might develop the anti-luck approach to epistemology that I set-out, and to say a little about one particular application that I think the view has. Some potential developments of the view are fairly immediate. One would like to know, for example, just how (if at all) one could adapt anti-luck epistemology so that it can deal with non-fully contingent propositions. Moreover, I think it is crucial to say more about just why having non-lucky true beliefs is so important to knowledge, and this extension of the project will inevitably lead into issues about epistemic value which are currently at the forefront of discussion.

Two further ways of extending the project are less obvious, but have a great deal of potential on further inspection. The first is suggested by Sandy Goldberg's (200X) thought-provoking contribution to this volume, in which he explores the prospect of applying the anti-luck machinery to the epistemology of testimony. As Sandy shows, there is real promise in this programme. Moreover, I have a personal interest in such a programme being fruitful, since it would allow me to connect my work on testimony (e.g., Pritchard 2004), which at present makes no mention of epistemic luck, with my more general work on epistemic luck.

I have a similar motivation for wishing to see if the project can be explored in a different direction which takes in recent work on the epistemological ramifications of a disjunctivist account of perceptual experience, especially as this is set out by McDowell. Although, as explained above, I hold that McDowell's anti-sceptical view does not achieve what it advertises, I am nevertheless also persuaded that there may be a form of disjunctivism which is also epistemically internalist, and which therefore allows that agents can have reflective access to factive empirical reasons.¹⁵ If this is right, then this would complicate the story I tell about reflective epistemic luck, since one could no longer take it as given—as I do, following convention on this score—that what is reflectively accessible to the agent is thereby not empirically factive. It could be, of course, that the complications are fatal, but my suspicion is that one would simply need to tell a more complicated story about how the two different types of epistemic luck relate to the externalism/internalism distinction granted this non-standard form of epistemic internalism.

As a sort of promissory note as to the attractions of an anti-luck epistemology, I will end by outlining one application of the proposal to a contemporary discussion. What I have in mind here is the so-called 'lottery puzzle', as described, for example, by John Hawthorne (2004) in his recent and highly influential book, which takes this puzzle as its starting-point.

Here is how Hawthorne (2004, pp. 1-2) states the problem:

Suppose someone of modest means announces that he knows that he will not have enough money to go on an Africa safari this year. We are inclined to treat such a judgement as true [...] However, were that person to announce that he knew that he would not win a major prize in a lottery this year, we would be far less inclined to accept his judgment as true. We do not suppose that people know in advance of a lottery drawing whether they will win or lose. But what is going on here? The proposition that the person will not have enough money to go on an African safari this year entails that he will not win a major prize in a lottery. If the person knows the former, then isn't he at least in a position to know the latter by a simple deduction?

We can formulate the puzzle in terms of the following 'lottery argument':

- (L1) *S* knows that she won't have enough money to go on safari this week.
- (L2) If *S* knows that she won't have enough money to go on safari this week, then she is in a position to know that she won't win a major lottery prize this week.
- (LC) *S* is in a position to know that she won't win a major lottery prize this week.¹⁶

Since this is, *ex hypothesi*, a valid argument with intuitively true premises and a counterintuitive conclusion, we seem forced to respond to it with some form of epistemological revisionism, and this is never a pleasant situation to be in.¹⁷ As we will see, however, clarifying the role of luck in this example will show us why there is, in fact, no puzzle here at all. If this is right, then the importance of the lottery 'puzzle' to contemporary epistemology has been radically overestimated, and the motivation to advance a form of epistemological revisionism in order to deal with this problem is undermined.

To begin with, note that part of what drives our intuition that *S* knows that she won't have enough money to go on safari next week—i.e., that (L1) is true—is the implicit suggestion that there is no specific reason to expect her to come into money next week. If we knew, for example, that her rich uncle was about to expire, and that this uncle was planning to bequeath our hero a substantial body of funds, then our intuition that (L1) is true would subside accordingly. And note the reason *why* this is relevant. This further information affects our intuitions because with this new information added it is now the case that there are relevant near-by possible worlds in which *S* has enough money to go on safari, and given our anti-luck condition set-out above—expressed now in terms of safety, (SP**)—this is clearly important to the question of whether or not *S* knows. Indeed, if this further detail is added, then (pending further details about the example at any rate) *S*'s true belief that she won't have enough money to go on a safari next week is, it seems, lucky (and so not a case of knowledge) because there is a wide class of relevant near-by possible worlds in which her belief is false (i.e., those worlds in which her uncle dies in the next week leaving her a large sum of money).

This point is important, because if the issue of whether *S* has a rich dying uncle is

relevant then so too should be the issue of whether *S* owns a lottery ticket (the draw for which is in the next week). After all, were it to be true that *S* owns a lottery ticket, then it would also be the case that there are very close relevant near-by possible worlds in which *S* has enough money to go on safari next week—i.e., those very close relevant near-by possible worlds in which she enjoys a large lottery win. Similarly, then, we should regard *S*'s belief in this case to be only luckily true, since there will be a class of very close relevant near-by possible worlds where her belief that she won't have enough money to go on safari next week is false. So on this reading of the example the claim that *S* knows that she won't have enough money to go on safari next week is *not* intuitive, and thus we ought not to assent to (L1).

If we are to get the intuition that (L1) is true off the ground it is thus important that we stipulate that *S* is *not* the owner of a lottery ticket. The problem is that adding this detail to the example has the effect of undermining our intuition that (LC) is false. After all, if *S* does not own a lottery ticket, then it is hardly counterintuitive to suppose that she knows that she won't win a major lottery prize in the next week, for there is *no* near-by possible world in which someone who does not own a lottery ticket wins a major lottery prize. In order to ensure that (LC) is counterintuitive it is thus essential that we suppose that *S* owns a lottery ticket, since it is only on this reading of the example that it is plausible that there is a class of very close relevant near-by possible worlds in which *S* wins a major lottery prize and yet continues to believe that she hasn't regardless.

The lottery argument thus plays on an ambiguity in how we are reading (L1) and (LC). Provided that we consistently understand the facts of the situation in such a way that (L1) is intuitive, then our intuition that (LC) is false subsides. Conversely, provided that we consistently understand the facts of the situation in such a way that (LC) is counterintuitive, then we no longer find (L1) intuitive. Taking anti-luck epistemology seriously thus enjoins us to be crystal clear about what the facts of the situation are in cases like this and to keep them fixed throughout. Once we do that, however, there is no tension between the two key claims, and thus there is no lottery problem that needs to be responded to.¹⁸ Anti-luck epistemology is thus able to resolve a key puzzle in contemporary epistemology, and thereby undermine the motivation that many have seen in this puzzle for a form of epistemological revisionism.¹⁹

REFERENCES

- Axtell, G. (2003). '*Felix Culpa: Luck in Ethics and Epistemology*', *Metaphilosophy* **34**, 331-52.
 Axtell, G. (200X). '*Felix Culpa Revisited: Two For the Show*', *Synthese (this volume)*.

- Becker, K. (forthcoming), 'Reliabilism and Safety', *Metaphilosophy*.
- Code, L. (1987). *Epistemic Responsibility*, Hanover, New Hampshire: University Press of New England.
- Cohen, S. (2000). 'Contextualism and Skepticism', *Philosophical Issues* **10**, 94-107.
- Coffman, E. J. (200X). 'Thinking about Luck', *Synthese (this volume)*.
- Dancy, J. (1985). *Introduction to Contemporary Epistemology*, Oxford: Blackwell.
- DeRose, K. (1995). 'Solving the Skeptical Problem', *Philosophical Review* **104**, 1-52.
- Dretske, F. (1970). 'Epistemic Operators', *The Journal of Philosophy* **67**, 1007-23.
- Fumerton, R. (1995). *Metaepistemology and Skepticism*, Lanham, Maryland: Rowman & Littlefield.
- Goldberg, S. (2007). 'How Lucky Can You Get?', *Synthese (this volume)*.
- Greco, J. (1999). 'Agent Reliabilism', *Philosophical Perspectives* **13**, 273-96.
- Greco, J. (2003). 'Virtue and Luck, Epistemic and Otherwise', *Metaphilosophy* **34**, 353-66.
- Greco, J. (200X). 'Worries about Pritchard's Safety', *Synthese (this volume)*.
- Harman, G. (1973). *Thought*, Princeton University Press, Princeton.
- Hawthorne, J. (2004). *Knowledge and Lotteries*, Oxford: Clarendon Press.
- Hiller, A., & Neta, R. (2007). 'Safety and Epistemic Luck', *Synthese (this volume)*.
- Lackey, J. (2005). 'What Luck is Not', *manuscript*.
- Lackey, J. (200X). 'Why We Don't Deserve Credit for Everything We Know', *Synthese (this volume)*.
- Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy* **74**, 549-67.
- McDowell, J. (1995). 'Knowledge and the Internal', *Philosophy and Phenomenological Research* **55**, 877-93.
- Nagel, T. (1976). 'Moral Luck', *Proceedings of the Aristotelian Society* (supp. vol.) **76**, 136-50.
- Neta, R., & Pritchard, D. H. (forthcoming). 'McDowell and the New Evil Genius', *Philosophy and Phenomenological Research*.
- Nozick, R. (1981). *Philosophical Explanations*, Oxford: Oxford University Press.
- Plantinga, A. (1993). *Warrant and Proper Function*, New York: Oxford University Press.
- Pritchard, D. H. (2002a). 'Recent Work on Radical Skepticism', *American Philosophical Quarterly* **39**, 215-57.
- Pritchard, D. H. (2002b). 'Resurrecting the Moorean Response to the Sceptic', *International Journal of Philosophical Studies* **10**, 283-307.
- Pritchard, D. H. (2003a). 'McDowell on Reasons, Externalism and Scepticism', *European Journal of Philosophy* **11**, 273-94.
- Pritchard, D. H. (2003b). 'Virtue Epistemology and Epistemic Luck', *Metaphilosophy* **34** (2003), 106-30.
- Pritchard, D. H. (2004). 'The Epistemology of Testimony', *Philosophical Issues* **14**, 326-48.
- Pritchard, D. H. (2005a). *Epistemic Luck*, Oxford: Oxford University Press.
- Pritchard, D. H. (2005b). 'Scepticism, Epistemic Luck and Epistemic Angst', *Australasian Journal of Philosophy* **83**, 185-206.
- Pritchard, D. H. (2005c). 'The Structure of Sceptical Arguments', *The Philosophical Quarterly* **55**, 37-52.
- Pritchard, D. H. (2006a). 'How to be a Neo-Moorean', *Internalism and Externalism in Semantics and Epistemology*, (ed.) S. Goldberg, Oxford: Oxford University Press.
- Pritchard, D. H. (2006b). 'Knowledge, Luck, and Lotteries', *New Waves in Epistemology*, (eds.) V. F. Hendricks & D. H. Pritchard, Aldershot: Ashgate.
- Pritchard, D. H. (2006c). 'McDowellian Neo-Mooreanism', *New Essays on Disjunctivism*, (eds.) A. Haddock & F. Macpherson, Oxford: Oxford University Press.
- Pritchard, D. H., & Smith, M. (2004). 'The Psychology and Philosophy of Luck', *New Ideas in Psychology* **22**, 1-28.
- Riggs, W. (200X). 'Why Epistemologists Are So Down On Their Luck', *Synthese (this volume)*.
- Sainsbury, R. M. (1997). 'Easy Possibilities', *Philosophy and Phenomenological Research* **57**, 907-19.
- Schaffer, J. (2005). 'Contrastive Knowledge', *Oxford Studies in Epistemology*, (eds.) T. Gendler & J. Hawthorne, Oxford: Oxford University Press.
- Sosa, E. (1991). 'Intellectual Virtue in Perspective', *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge, UK: Cambridge University Press.
- Sosa, E. (1999). 'How to Defeat Opposition to Moore', *Philosophical Perspectives* **13**, 141-54.

- Sosa, E. (2000). 'Skepticism and Contextualism', *Philosophical Issues* **10**, 1-18.
- Stroud, B. (1984). *The Significance of Philosophical Scepticism*, Oxford: Clarendon Press.
- Unger, P. (1968). 'An Analysis of Factual Knowledge', *Journal of Philosophy* **65**, 157-70.
- Williams, B. (1976). 'Moral Luck', *Proceedings of the Aristotelian Society* (supp. vol.) **76**, 115-35.
- Williamson, T. (2000). *Knowledge and Its Limits*, Oxford: Oxford University Press.
- Wittgenstein, L. (1969). *On Certainty*, (eds.) G. E. M. Anscombe & G. H. von Wright, (tr.) D. Paul & G. E. M. Anscombe, Oxford: Blackwell.
- Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge, UK: Cambridge University Press.

NOTES

- ¹ Some commentators come close to offering an account of knowledge along these lines. Unger (1968), for example, defines knowledge in terms of 'non-accidental' belief, though the notion of 'accident' here is treated as an undefined primitive.
- ² Interestingly, there is quite a lot of material in the psychological literature on luck. For a survey of some of this material which sets it in its philosophical context, see Pritchard & Smith (2004).
- ³ Though, on reflection, not enough I think. At the time I was, in effect, working in a vacuum, since there was hardly anything else in the literature on the philosophy of luck. It is gratifying to note that one of the effects of the book has been to stimulate debate on this topic. See, for example, Lackey (2005), Coffman (200X), and Riggs (200X).
- ⁴ I claim, for example, that Harman's (1973, 142-54) 'assassination' case is like this. See Pritchard (2005a, 168-73).
- ⁵ I shall ignore the issue of whether a belief could ever be the product of what is, self-consciously, a guess.
- ⁶ Sosa is not the only person to offer a safety-type principle. Similar principles can be found in recent work by, for example, DeRose (1995), Sainsbury (1997) and Williamson (2000, chapter 8), and I have previously defended a version of this principle myself. See Pritchard (2002b).
- ⁷ I coined the term in Pritchard (2002a, §8). For more on neo-Mooreanism, see Sosa (1999) and Pritchard (2002b; 2005a, chapter 3; 2006a).
- ⁸ Just how closely related these two clauses are depends to a large degree on the extent to which one allows the 'way' in which a belief is formed to be determined by facts in the agent's environment. This is a complex issue, however, and I will not be exploring it further here.
- ⁹ For two recent discussions along these lines, see Greco (2003; 200X). I consider this issue further below.
- ¹⁰ For more on this issue, see the symposium on virtue epistemology and epistemic luck which features Axtell (2003), Greco (2003), and Pritchard (2003b). See also Axtell (200X) and Becker (*forthcoming*).
- ¹¹ For more on epistemic *angst*, see Pritchard (2005b).
- ¹² I also claim that the underdetermination principle on which this form of scepticism turns is logically weaker than the closure principle on which the standard sceptical argument turns. For a fuller development of this point, and for a thorough discussion of the two types of sceptical argument, see Pritchard (2005c).
- ¹³ See also Pritchard (2003a).
- ¹⁴ I explore one variation on the safety principle—what I term "super-safety"—in Pritchard (2002b).
- ¹⁵ I sympathetically discuss such a position in Neta & Pritchard (*forthcoming*) and Pritchard (2006a; 2006c).
- ¹⁶ This is roughly as Hawthorne (2004, 2) formulates it (aside from the change in gender), though note that I've tightened-up the example so it talks of 'next week' rather than 'next year'. The reason for this is that (L1) strikes me as *prima facie* problematic on the latter reading. After all, a *lot* can happen in a year. In what follows I will simply take it as granted that *S*'s beliefs in this regard are true. Remember

that the issue is whether *S* *knows* these propositions, not whether she truly believes them. Thus, there can be no harm in stipulating that she does in fact truly believe them.

¹⁷ For example, one could, following Dretske (1970) and Nozick (1981), respond to the argument by rejecting the principle that knowledge is closed under known entailment (which (L2) appears to presuppose). Another revisionary line could be to employ some form of contextualism. Lewis (1996), for example, offers an attributer contextualist treatment of ‘knows’ which captures a sense in which both (L1) and (LC) are true and also a sense in which they are both false, thereby preserving closure while also paying due attention to our conflicting intuitions in this regard. Hawthorne’s own response to the problem is of this general contextualist sort, although on his view it is the *subject’s* context that is important rather than the attributer’s context, and this enables him to avoid endorsing some of the more awkward claims made by attributer contextualists (though, like all contextualist views, his position makes some awkward claims as well). (Note that Hawthorne doesn’t refer to his view as a contextualist position, but this is because he uses the term ‘contextualism’ to exclusively refer to *attributer* contextualist theses).

¹⁸ This diagnosis of where the lottery argument goes wrong also has application to other lottery-style arguments that are offered in the literature, such as, to use one of Hawthorne’s examples (Hawthorne 2004, 3), the supposed tension between the intuition that a subject knows that she will be living in Syracuse next summer and the intuition that she does not know that she won’t die of a heart attack in the next week. I leave the details of how an anti-luck epistemology would deal with this case as an exercise for the reader. I consider the issue of how the anti-luck programme should be applied to puzzles like the lottery puzzle in more detail in Pritchard (2006b).

¹⁹ This paper arose out of a Pacific APA author-meets-critics symposium on my book, *Epistemic Luck*, which was organised by Richard Greene and held in San Francisco in March 2005. My thanks go to the audience that day, and also my commentators—Guy Axtell, Sandy Goldberg, John Greco, and Ram Neta—and Richard Greene. I am also grateful to a Mellon Foundation funded epistemology discussion group on the book which was organised by John Turri and Ernie Sosa at Brown University in November 2005. My thanks go to John Turri and Ernie Sosa and the rest of the members of this discussion group. Finally, thanks also to Kelly Becker, E. J. Coffman, Vincent Hendricks, Jennifer Lackey, Wayne Riggs, and John Symons.