KNOWLEDGE, LUCK, AND LOTTERIES

Duncan Pritchard
University of Stirling

0. INTRODUCTION

It is a platitude in epistemology to say that knowledge excludes luck. Indeed, if one can show that an epistemological theory allows 'lucky' knowledge, then that usually suffices to warrant one in straightforwardly 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:

> [...] 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. (Dancy 1985, 134, *my italics*)

Even despite the prevalence of this intuition, however, very few commentators have explored what it 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] It is just such a view that I propose, however, and I hope to give a flavour of what this strategy involves here.[2] In particular, I have two goals in this paper. The first is to outline the general contours of the position and show how such a view can account for the attraction of adducing a safety condition on knowledge, with all the epistemic benefits that this principle holds. Relatedly, I will also explain how an anti-luck epistemology can assist us in determining the best formulation of this principle. The second goal of the paper is to show anti-luck epistemology in action by highlighting how such a view can deal with the various problems posed by lottery-style examples.

1. LUCK

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 his 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 his skill, his 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 the following rough account of luck:

(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. For one thing, it includes that weasel word 'relevant' in the analysis, which is dreadfully vague. Moreover, the account 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). As we will see, however, this rough and partial account of luck will suffice, however, to enable us to get a handle on what an anti-luck epistemology would look like.[3]

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.[4] 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.[5] 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.

With our account of luck, (LE), in mind we can give a rough specification of what would count as a lucky true belief:

(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.[6]

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.[7] 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 word, 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 core constraint on knowledge possession in the recent literature—*viz*., the safety principle, a version of which has been proposed, for example, by Ernest Sosa (e.g., 1999).[8] Very roughly, the safety principle can be characterised as follows:

(SP) *S*'s true 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. Such a principle can, for example, deal with Gettier-style cases, like the stopped clock case that we just considered. Moreover, as I will explain in a moment, this principle also has anti-sceptical appeal in that it appears to show that our beliefs in the denials of sceptical hypotheses—while not sensitive to falsehood or justified by adequate reflectively accessible grounds (as I outline below)—do have some positive epistemic support in virtue of being safe.[9]

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,[10] 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 very similar. If the preceding account of why we should endorse something like (AL) is right, however, 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)—(SP) is, in effect, a condition which simply defines out of existence the problematic form of epistemic luck. Indeed, it is perhaps not surprising that (SP) performs this anti-luck role once one notes that the safety principle is often informally glossed as demanding that the agent's true belief could not have easily been false. 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.

The safety principle is often introduced in contrast to the sensitivity principle, so it is

worth introducing this competing modal epistemic principle in order to get a better handle on what safety demands. Sensitivity can be roughly defined as follows:

(SEN) *S*'s true belief is sensitive *iff* in the nearest possible worlds in which the target proposition is false, *S* no longer believes it on the same basis as in the actual world.

Sensitivity is also able to deal with Gettier-style cases, like the stopped clock case. After all, while the belief may as it happens be true in the actual world, there is a nearest possible world in which what is believed is no longer true—i.e., where the time is different—and yet the agent continues to believe the target proposition regardless. In a sense, then, one might think that both sensitivity and safety can accommodate our anti-luck intuitions.

Interestingly, however, there are cases in which the verdicts issued by sensitivity and safety come apart, but where the anti-luck intuition—encapsulated in (AL)—seems to favour safety. Consider the following case, cited by Sosa:

> 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. (Sosa 2000, 13)

Clearly this belief is not sensitive. After all, in the nearest possible world in which what is believed is false—e.g., the world in which, by some freak accident, the trash bag snags in the chute on the way down—the agent will continue to believe it regardless (and on the same basis as in the actual world). Intuitively, however, Sosa does know that his trash bag is down in the basement. The possibility that the bag snags is, after all, *ex hypothesi* a far-fetched scenario. This intuition is captured by (AL), and thus by (SP). It is not a matter of luck, on this view, that Sosa's belief is true given the way the circumstances are described, since in most worlds like the actual world in which he forms his belief on the same basis as in the actual world his belief continues to be true. His belief is thus safe, and hence on this account knowledge.

As noted above, the further advantage to advancing a safety condition is that it enables us to account for our anti-sceptical knowledge. It is often noted in the literature on scepticism that if the principle of closure is true, then it seems that we cannot have much of the knowledge that we typically ascribe to ourselves without also having lots of anti-sceptical knowledge as well—i.e., knowledge of the denials of sceptical hypotheses. That is, closure holds, roughly, that if one knows one proposition, and knows that this entails a second proposition, then one also knows the second proposition. Accordingly, if one knows that, for

example, one is presently seated, and also knows (as presumably one does) that being seated is incompatible with, say, being a brain-in-a-vat (BIV) who is merely being 'fed' experiences as if one is seated, then it follows that one must know that one is not a BIV. Interestingly, however, there is a widespread intuition that we are unable to know the denials of sceptical hypotheses, such as the BIV hypothesis.

There's certainly something epistemically amiss about our beliefs in these anti-sceptical propositions. For one thing, we do not seem to have any good reflectively accessible grounds for believing these propositions since, *ex hypothesi*, there is nothing reflectively available to one which could indicate that one is not a BIV. Moreover, these beliefs are not sensitive, on most readings of the sensitivity principle at any rate, since in the nearest possible worlds in which what one believes is false—i.e., in which one is a BIV, for example—one will continue to believe that one is not being deceived in this way.[11] This has led proponents of the sensitivity principle—such as Fred Dretske (1970) and Robert Nozick (1981)—to deny closure as a means of blocking the sceptical argument.

As a number of commentators have noted, however, denying closure is a rather tough price to pay for resolving the sceptical problem.[12] Moreover, it is puzzling to hear that our beliefs in the denials of sceptical hypotheses don't count as knowledge even granted that we are able to know everyday propositions—i.e., propositions which would we all take ourselves to know in normal circumstances, such as (in my case) that I am presently seated—which entail such anti-sceptical propositions. After all, if we can know everyday propositions then it follows that sceptical hypotheses must indeed be as modally far-off as we take them to be, since otherwise they would undermine our everyday knowledge. But if they are modally far-off, then there seems no reason to think that our belief that we are not the victim of a sceptical hypothesis is lucky, and so no obvious reason, by the lights of an anti-luck epistemology, for thinking that it isn't knowledge.

We can capture just such an intuition via (AL) and thus (SP). In both cases, it is only near-by possible worlds that count when it comes to determining knowledge, and thus it follows on this view that just so long as sceptical hypotheses are indeed modally far-fetched—such that they only obtain in far-off possible worlds—then our anti-sceptical beliefs are not epistemically lucky, which means that they can be safe and thus count as knowledge. After all, if such worlds are far-off, then it will be the case that not only is one's belief that, say, one is not a BIV is true in the actual world, but also true in all near-by possible worlds in which one continues to believe this proposition. To employ the gloss on safety noted above, it

will be the case that our true belief in this propositions could not have easily been false.

The advocacy of safety, on anti-luck grounds, thus points towards a response to the sceptical problem which denies one of the key premises in the sceptical argument. Note that this feature of the view also indicates the sense in which a safety-based account of knowledge, advocated on anti-luck grounds, is most naturally allied to an externalist epistemology. On standard construals, what makes an account of knowledge internalist is that it insists that to know is to at the very least meet an internal epistemic condition, where what makes an epistemic condition internal is the fact that the facts which determine that this condition has obtained are reflectively accessible to the agent.[13] Clearly, however, the kind of knowledge that is at issue when it comes to anti-sceptical propositions cannot be knowledge which meets the internalist rubric, so conceived, since we have already noted that, intuitively, one has no reflectively accessible grounds for belief in these propositions.

This point is further confirmed once one reflects on those cases around which the epistemic internalist/externalist debate tends to focus. Consider, for example, the famous 'chicken sexer' example, which concerns an agent who has a highly reliable ability to distinguish between male and female chicks, but who lacks any good reflectively accessible grounds to back-up her beliefs so formed (on the usual telling of the example, the agent has false beliefs about how she is forming her beliefs in this regard, supposing herself to be touching or seeing something distinctive when there is in fact nothing distinctive to see or touch, the faculty of smell being the one that is operative).[14] Since no internal epistemic condition is met in this case, the epistemic internalist will typically argue that the agent lacks knowledge. It is open to the externalist, in contrast, to accord knowledge in this case just so long as an external epistemic condition is adequately met, such as a reliability condition. Interestingly, the external epistemic condition that we have identified here, the anti-luck condition encapsulated in (SP), is met in this case, in that the chicken sexer's belief is safe. That is, given that the agent really does have the sexing-sexing ability in question, then her true belief as regards the sex of the chick will be such that in most near-by possible worlds in which she continues to form her belief in the same way as in the actual world her belief will continue to be true. Anti-luck epistemology, encapsulated here in terms of the safety principle, thus goes hand-in-hand with epistemic externalism.

4. REFINING THE SAFETY PRINCIPLE: LUCK AND LOTTERIES

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 hearing the result of the draw and based on the odds involved—that one owns a losing ticket for a free and fair lottery 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 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.

It is worth noting that sensitivity, which we saw above to be inferior to safety when it comes to dealing with certain cases, appears to hold the advantage over safety when it comes to this example. After all, one's belief that one has lost the lottery when it is based on the long odds involved will tend to be insensitive, in that in the nearest possible world in which it is no longer true—i.e., the world in which one wins the lottery—one will continue to form a belief in this proposition on the same basis. A sensitivity-based theory of knowledge would thus generate the right result in this case.

Moreover, note that it is fairly essential to any anti-luck epistemology that it is able to deal with the lottery case. After all, the moral of these cases seems to be that what is primarily important to knowledge possession is not the quality of one's evidence in favour of one's belief, but more specifically whether or not one's belief is epistemically lucky. After all, there is, intuitively at least, no problem with gaining one's knowledge that one has lost the lottery by reading the results in a reliable newspaper, even though the odds that the newspaper has printed the right results are almost certainly going to be a lot lower than the odds involved in the lottery draw itself. Thus, given that the strength of probabilities here must surely reflect the strength of one's evidential support for one's belief that one has lost the lottery, whether or not one knows is not a function of the strength of one's evidence. Moreover, sensitivity can

again accommodate this fact, since the key difference between forming one's belief that one has lost the lottery by looking up the results in a reliable newspaper rather than simply reflecting on the odds involved is that the former belief *does* meet the sensitivity requirement—in the nearest possible worlds in which the target belief is false (where one has won the lottery), one will no longer believe that one has lost via this method, but believe instead that one has won (because the lottery numbers printed will be different). How is the proponent of an anti-luck epistemology, and thus (SP), to deal with this problem?

One sort of response might be to supplement a safety-based theory of knowledge with further epistemic conditions that are able, in conjunction with safety, to deal with this case. The trouble with this proposal, however, is that since the guiding 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. Essentially, one would be accounting for the epistemic luck in this case not by appealing to one's supposed anti-luck condition, but rather to other epistemic conditions instead, and that obviously undermines the theoretical attraction of the view.

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 true 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 safety in this way would seem to make it unable to deal with more mundane cases of knowledge possession.[15] Consider again, for example, the 'rubbish chute' case described by Sosa which we quoted above. Intuitively, we want to say that Sosa does know that his trash bag is in the basement. The trouble is, however, if safety is understood along very robust lines as (SP*) then, it has been claimed, this knowledge seems to be ruled-out since surely there are quite a few near-by possible worlds—i.e., 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-cut 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 may well satisfy the demands laid down by (SP*) after all, since the possibility that the bag snags is in fact 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 very secure, 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 there is a wide class of relevant near-by possible worlds in which Sosa 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 clearly is a matter of luck that his belief is true in this regard given the nearness of the relevant error-possibility. On this reading, then, it seems that Sosa's belief fails to meet the conditions laid down by (SP*), but since there is now no intuition that Sosa has knowledge in this case—because of the nearness of the relevant error-possibilities—this is not troubling.

Everything therefore 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 is not to say, however, that we should simply understand safety as (SP*) and leave the matter at that, since further reflection on the anti-luck motivation of the safety principle highlights a sense in which we can avoid construing it too austerely while still accommodating both the lottery and rubbish chute cases. And, naturally, if a weaker rendering of safety is available that can deal with both cases then this is

to be welcomed.

The motivation for this 'intermediate' formulation of safety comes from the observation that the problem posed by the lottery case isn't so much that there is a small class of relevant near-by possible worlds in which the belief is false, 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 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 true 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.

Moreover, note that this formulation of safety is also in accordance with the intuition noted above that while one lacks knowledge in the lottery case where one forms one's belief by considering the odds involved, one can gain knowledge that one has lost the lottery by reading the result in a reliable newspaper, even though the odds in favour of the truth of one's belief may be substantially lower. After all, the possible world in which one forms a false belief by using the former method is much closer than the possible world in which one forms a false belief using the latter method (indeed, given the level of proof-reading that goes into the production of a reliable newspaper, one would typically not expect there to be any near-by possible world in which one forms a false belief via the latter method).

This highlights an important feature of anti-luck epistemology, and indeed modal epistemology in general. The rubric against which worlds are ordered is a function of the similarity of the worlds relative to the actual world, and not a function of the probabilities of the events involved. This is why low probability events, such as the winning of a lottery, can occur in near-by possible worlds. What counts when we are trying to see if an event is lucky, however—and thus what counts when we are trying to determine whether a true belief is only luckily true—is the *similarity* of the possible world in which one uses one's actual method to form a false belief in the target proposition. Anti-luck epistemology is thus able to deal with the oddity of the result that knowledge is not a function of the strength of evidential support where that is in turn understood along probabilistic lines, since it is the *nearness* of the relevant possible worlds that counts on this view, not the probability of the events that obtain in those worlds.[16]

5. ANTI-LUCK EPISTEMOLOGY AND THE LOTTERY PUZZLE

Anti-luck epistemology can also enable us to deal with another key problem in epistemology—known as the 'lottery puzzle'—which emerges out of the lottery case that we have just considered. The claim is that if we are unable to know that we are in possession of the losing lottery ticket, prior to hearing the results of the draw, then it seems that there are an awful lot of other propositions which we fail to know as well. This puzzle takes centre-stage in John Hawthorne's (2004) recent and highly influential book. Here is how he 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? (Hawthorne 2004, 1-2)

Hawthorne goes on to formulate the puzzle in terms of the following 'lottery argument':

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

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. This is never a pleasant situation to be in. For example, suppose one responded to this argument by conceding that (L1) is, contrary to intuition, simply false. The problem is that once one allows this, scepticism cannot be too far behind, for if I don't know *this* proposition, then how can I know anything of substance? Accepting the conclusion, (LC), would be another revisionary possibility, as would following Dretske (1970) and Robert (1981) in rejecting the closure principle (which (L2) appears to presuppose), but for obvious reasons neither move is particularly appealing.

A more subtle approach to the problem, though one that is still revisionary in a fairly radical way, would be to respond to this puzzle with some form of contextualism. David 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 the attributer contextualists (though like all contextualist views, his position makes some awkward claims as well).[17]

What I will be suggesting is that rather than accepting that there is a puzzle here which needs to be responded to, and therefore considering which of the various revisionary epistemic proposals we should endorse to deal with this problem, we should instead look a little closer at just what the puzzle is supposed to be. Indeed, I will claim that a closer

inspection of how this problem is characterised reveals that it isn't clear that there is a genuine puzzle here at all, and that an anti-luck epistemology can bring this fact into clear view. 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.

Let's start off with the example as it is described by Hawthorne above. Take the claim that the subject knows that he won't have enough money to go on safari this year. This isn't quite as obvious as it may at first appear. I think we would all grant that the agent knows that he's not going to have enough money to go on safari next week, or even next month, but when we move to suitably distant periods of time the intuition weakens. Suppose, for example, we ask the question of whether the agent knows that he will not have enough money to go on a safari in the next *ten* years. If this really is an agent who would go on a safari if he could, who yearns to one day go on a safari, then I think we would be cautious about saying that he knows that he will *not* have enough money to go on a safari sometime in the next ten years. After all, a lot can happen in ten years; a lot of small sums of money could be saved, for example. Can a lot happen in one year? Possibly. Here our intuitions (under reflection) aren't so clear.

Accordingly, if we are to straightforwardly regard the first premise of the lottery argument as intuitive then we are either going to need to know a lot more detail about the example or we are going to have to shorten the time-period in question. For the sake of simplicity, I will opt for the latter alternative and reformulate the lottery argument as follows:

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

A further advantage of this reformulation is that it also tightens-up our intuitions regarding the conclusion of the lottery argument. If we simply talk about whether or not the subject knows that he won't win a major lottery prize in the next year this leaves it open just what his lottery playing habits are. Perhaps he only plays intermittently, and may not play at all in the next year. Clearly, however, if he never buys a ticket then he has no chance of winning, and we'd have no problem with ascribing knowledge to him that he won't win in *that* case. With the new formulation, however, we can just stipulate that he has bought a ticket for a lottery, the draw for which is next week. On this construal, it is clearly intuitive that he doesn't know that he won't win a major lottery prize in the relevant time period.

We will also add one further *caveat* in order to simplify our discussion of the lottery argument, which is to stipulate that it is true that the agent won't have sufficient funds to go on a safari next week, and thus that it is also true that he will not win a major lottery prize in the next week. Remember that the issue is whether the agent *knows* these propositions, not whether he truly believes them. Thus, there can be no harm in stipulating that he does in fact truly believe them.

With (L1*) in mind, note that part of what drives our intuition regarding whether or not the subject will have enough money to go on a safari next week is the implicit suggestion that there is no specific reason to expect him to come into money next week. If we knew, for example, that his rich great uncle was about to expire, and that he was planning to bequeath our subject a substantial body of funds, then our intuition that (L1*) is true would subside accordingly. After all, there would now be a wide class of relevant near-by possible worlds in which the agent forms the target belief and yet forms a false belief as a result (because he has inherited a fortune). The belief is thus only luckily true, and thus it is not safe (as we are now understanding that notion) and so not a case of knowledge.

Notice, however, that insofar as we are swayed by considerations like this, then we ought to be similarly suspicious of (L1*) once it is stipulated—as it needs to be, if (LC*) is to be regarded as intuitively false—that the agent is in possession of a lottery ticket, the draw for which is next week. After all, given this further detail it will again be true that there is a wide class of relevant very close near-by possible worlds in which the agent forms the target belief and yet forms a false belief as a result (because he has won the lottery). Again, the belief is now only luckily true, and thus it is not safe and so not a case of knowledge.

Here is the crux: If we ask the question of whether the agent will have enough money to go on a safari next week *independently* of specifying that he is the owner of a lottery ticket—(or, for that matter, that he has a rich, but seriously ill, close relative)—then I think the intuition is clear that we should answer positively. If we ask that same question in the light of the further detail about the lottery ticket (or the rich dying relative), however, then the intuition changes accordingly. Crucially, of course, it is essential to the lottery puzzle that this further detail is stipulated, since it is only if the agent is indeed in possession of a lottery ticket that the truth of (LC*) is at all counterintuitive.

So provided that we fill-in the details to this example properly, then there is no good reason to think that our intuitions about the first premise of the lottery argument and its conclusion are in tension. Accordingly, it isn't at all clear that there is even a *prima facie* puzzle that needs to be responded to here. Note that the underlying problem here is a failure to

pay due attention to how filling-in the details of the situation will alter the degree to which we will regard the true belief of the agent as lucky, and thus the degree to which we will regard the belief as safe. Provided that we consistently understand the facts of the situation in such a way that (L1) is intuitive—such that the agent's belief is therefore safe—then our intuition that (LC) is false subsides, since the agent's belief in that case will be safe as well. Conversely, provided that we consistently understand the facts of the situation in such a way that (LC) is counterintuitive—such that the agent's belief is not safe—then we will no longer find (L1) intuitive either, since the agent's belief at issue in (L1) won't be safe on this reading of the facts. 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 in the lottery puzzle, and thus there is no puzzle that needs to be responded to. Anti-luck epistemology is thus able to resolve a key problem in contemporary epistemology, and thereby undermine the motivation that many have seen in this problem for a form of epistemological revisionism.

In a similar fashion, anti-luck epistemology also highlights how one should go about responding to other quasi-lottery puzzles that are often expounded in the literature. These are arguments which do not concern lotteries at all but which are meant to share the essential features of the lottery puzzle and therefore pose the same sort of challenge to contemporary epistemology. Consider the following case, also given by Hawthorne:

> […] I am inclined to think that I know that I will be living in Syracuse for part of this summer. But once the question arises, I am not inclined to think that I know whether or not I will be one of the unlucky people who, despite being apparently healthy, suffer a fatal heart attack in the next week. […] Indeed, I am just as unwilling to count myself as knowing about the heart attack as I am to count myself as knowing about the lottery. […] Just as I have excellent grounds for supposing that any given lottery ticket will lose, I have excellent statistical grounds for supposing that a given apparently healthy person will not have a heart attack very soon. (Hawthorne 2004, 3)[18]

The parallel argument to the lottery argument in this case is thus supposed to be:

(H1) *S* knows that *S* will spend part of his summer in Syracuse.
(H2) If *S* knows that *S* will spend part of his summer in Syracuse, then *S* is in a position to know that *S* will not die of a heart attack next week.
(HC) *S* is in a position to know that *S* will not die of a heart attack next week.

As before, in order to keep matters simple, we will simply stipulate that it is true that the subject will spend part of his summer in Syracuse, and thus that it is also true that he doesn't die of a heart attack next week.

Although there are obvious analogies between this 'heart attack' argument and the

corresponding lottery argument, the two cases are disanalogous in a crucial respect. The issue concerns the 'statistical' parallel that Hawthorne notes between the heart attack case and the lottery case, the point being, presumably, that there are low odds of the target event (heart attack/lottery win) occurring in both cases. While this is undoubtedly true, the nature of these low probabilities is importantly different and this has ramifications for whether we regard the related belief as being epistemically lucky, thus affecting our epistemic intuitions in this respect.

Consider the claim that there is a low probability that an apparently healthy middle-aged man will die of a heart attack next week. Although this is surely true, note that it does not follow that for each *particular* man who meets the relevant description there will, in fact, be a low probability that he will die of a heart attack next week. After all, this statistic represents an *average*, and therefore glosses over the fact that there are some unfortunate souls who, due to, say, their lifestyle or their genetic make-up, are middle-aged and apparently healthy but for whom it is quite likely that they will die of a heart attack next week. From the way Hawthorne describes the example we are clearly meant to suppose that we have no way of knowing whether he falls into the 'at risk' category, which is, of course, the usual situation anyway in these cases. Notice, however, that once we fill-in the details in this regard then our intuitions about whether Hawthorne knows that he will be alive next week vary accordingly.

On a natural reading of the example, such that we don't regard Hawthorne as falling into the 'at risk' category (after all, we've been given no specific reason for thinking that he does fall into this category), then our intuition is, I take it, that he *does* know that he will be alive next week. The reason for this is (at least in part) because his true belief is not epistemically lucky. On this reading of the example there are *no* near-by possible worlds in which he dies from a heart attack in the next week because his health is perfectly fine (his arteries aren't clogged etc.,), and thus his true belief in this regard is not lucky but safe.

In contrast, if we stipulate that Hawthorne is one of the people in the 'at risk' category, then we will tend to revise our intuitions accordingly and hold that he *doesn't* know that he will be alive next week. This change in our judgements reflects the fact that his true belief that he will be alive next week is lucky and thus unsafe. This is because there is a class of near-by possible worlds where the relevant initial conditions for the formation of his belief are the same as in the actual world, and yet his belief is false (i.e., those near-by possible worlds in which he *does* expire from a heart attack in the next week).

The difference between the lottery case and the heart attack case is thus that while it is

straightforwardly true that those who play lotteries will win a major prize in some near-by possible worlds, it is not straightforwardly true that those who fall into a category which is subject to a low probability of death from a heart attack in the next week will die from a heart attack in the next week in some near-by possible worlds. It will all depend on the details of the case in hand. The crucial point is that if the possibility is near, then it will undermine subject's knowledge in this regard (since it will make the belief unsafe); while if it is not then it won't (because the belief will not be unsafe). But just knowing the probability by itself in this case will not tell us whether the relevant worlds are near or not, and thus whether or not the true belief is safe.

So provided that the details of the case are stipulated and kept fixed for both of the target propositions, then the 'heart attack' argument outlined above poses no problem. If Hawthorne falls into the 'at risk' members of the class such that there are near-by possible worlds in which he dies of a heart attack next week, then it is *not* intuitive that he knows that he will be alive next week, and neither is it now intuitive that he will spend part of his summer in Syracuse. Both of these true beliefs will be lucky, and thus unsafe, and hence (H1) and (HC) will be both false. Alternatively, if Hawthorne is not one of the 'at risk' members of the class so that there are no near-by possible worlds in which he dies of a heart attack next week, then it *is* intuitive that he knows that he will be alive next week and also that he knows he will be in Syracuse for part of the summer. Neither of these true beliefs will be lucky, and thus neither will be unsafe, and hence there is every reason to regard (H1) and (HC) as true.[19]

6. CONCLUDING REMARKS

By taking the anti-luck platitude seriously, and thereby formulating an anti-luck condition on knowledge, we can thus cast light on the motivation for, and correct construal of, the safety principle. Moreover, we can also deal with the various problems posed by lottery cases. Indeed, I would argue that the merits of this approach do not end there, since an anti-luck epistemology can contribute to the discussion of a number of other central epistemological issues, such as the problem of scepticism and the externalism/internalism dispute. These further applications are, however, topics for another occasion.[20,21]

REFERENCES

Black, T. (2002). 'A Moorean Response to Brain-In-A-Vat Scepticism', *Australasian Journal of Philosophy* 80, 148-63.

Brandom, R. (1998). 'Insights and Blindspots of Reliabilism', *The Monist* 81, 371-92.

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.

——— (2005*a*). 'The Case Against Closure', in *Contemporary Debates in Epistemology*, (eds.) M. Steup & E. Sosa, 13-25, Oxford: Blackwell.

——— (2005*b*). 'Reply to Hawthorne', in *Contemporary Debates in Epistemology*, (eds.) M. Steup & E. Sosa, 43-6, Oxford: Blackwell.

Foley, R. (1987). *A Theory of Epistemic Rationality*, Cambridge, Massachusetts, Harvard University Press.

Greco, J. (2003). 'Virtue and Luck, Epistemic and Otherwise', *Metaphilosophy* 34, 353-66;

Hawthorne, J. (2002). 'Lewis, the Lottery, and the Preface', *Analysis* 62, 242-51.

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

——— (2005). 'The Case For Closure', in *Contemporary Debates in Epistemology*, (eds.) M. Steup & E. Sosa, 26-42, Oxford: Blackwell.

Kornblith, H. (ed.) (2001). *Epistemology: Internalism and Externalism*, Oxford: Blackwell.

Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy* 74, 549-67.

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

Pritchard, D. H. (2002). 'Resurrecting the Moorean Response to the Sceptic', *International Journal of Philosophical Studies* 10, 283-307.

——— (2005). *Epistemic Luck*, Oxford: Oxford University Press.

——— (*forthcoming*). 'How to be a Neo-Moorean', *Internalism and Externalism in Semantics and Epistemology*, (ed.) S. Goldberg, Oxford: Oxford University Press.

Pritchard, D. H., & Smith, M. (2004). 'The Psychology and Philosophy of Luck', *New Ideas in Psychology* 22, 1-28.

Sainsbury, R. M. (1997). 'Easy Possibilities', *Philosophy and Phenomenological Research* 57, 907-19.

Sosa, E. (1999). 'How to Defeat Opposition to Moore', *Philosophical Perspectives* 13, 141-54.

——— (2000). 'Skepticism and Contextualism', *Philosophical Issues* 10, 1-18.

Unger, P. (1968). 'An Analysis of Factual Knowledge', *Journal of Philosophy* 65, 157-70.

Vogel, J. (1990). 'Are There Counterexamples to the Closure Principle?', in *Doubting: Contemporary Perspectives on Scepticism*, (eds.) M. Roth & G. Ross, Dordrecht: Kluwer.

Williams, M. (1991). *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*, Oxford: Blackwell.

Williamson, T. (2000). *Knowledge and Its Limits*, Oxford: Oxford University Press.

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' true belief, though the notion of 'accident' here is treated as an undefined primitive.

[2] I develop the view in detail in Pritchard (2005).

[3] I offer an extended discussion of the notion of luck in Pritchard (2005, chapter 5). See also Pritchard & Smith (2004) for a survey of the psychological and philosophical literature on luck.

[4] In Pritchard (2005, chapter 5), I call such a benign form of epistemic luck, "evidential luck". Although this form of epistemic luck clearly is compatible with knowledge possession, I also show how divergent intuitions about certain examples in epistemology are often the result of failing to spell-out the details of the case to a sufficient degree to make it clear whether it is merely evidential luck at issue or a non-benign form of epistemic luck.

[5] In Pritchard (2005, chapter 5), I refer to this benign form of epistemic luck as "capacity luck".

[6] 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 in these cases there will be no near-by possible worlds in which the target propositions are false. The focus here is thus an account of what it is for a true belief in a (fully) contingent proposition to be lucky.

[7] I shall ignore the issue of whether a belief could ever be the product of what is, self-consciously, a guess.

[8] 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 defended a version of this principle myself—see Pritchard (2002).

[9] The safety principle thus lies at the heart of a broadly 'Moorean' response to the sceptic. See Sosa (1999) and Pritchard (2002; *forthcoming*).

[10] 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.

[11] Not everyone agrees that sensitivity has this consequence. As Williams (1991, chapter 9) and Black (2002) have argued, provided that one understands the 'same basis' clause in the right kind of way, then the agent is not forming her belief in the BIV world on the same basis as in the actual world, and thus it does not follow from the fact that if the agent were a BIV she would continue to believe that she isn't a BIV that she lacks a sensitive belief in this anti-sceptical proposition.

[12] For the most recent survey of the problems facing the denial of closure, see Hawthorne (2005; cf. Dretske 2005*a*), to which Dretske (2005*b*) responds.

[13] For more on the epistemic externalism/internalism distinction, see Kornblith (2001).

[14] For further discussion of the chicken-sexer example, see Foley (1987, 168-9), Lewis (1996), Zagzebski (1996, §2.1 & §4.1), and Brandom (1998).

[15] For a recent discussion along these lines, see Greco (2003).

[16] I think this also explains why we would not regard the purchase of a lottery ticket, with long odds, as being irrational in the same way that we would regard the placing of a bet, of the same monetary value, on an event occurring that has the same long odds (say, that Wales will win the next five finals of the Baseball World Series). This is because although the two target events have, *ex hypothesi*, the same chances of occurring, it is only the lottery win that is an event that occurs in worlds like the actual world (think, for example, of all the features of the actual world that would have to be different for the Welsh baseball team to achieve this feat). Our judgements about the rationality of placing certain wagers are thus, I would suggest, responsive to the modal nearness of the events in question, rather than just their probability.

[17] 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] Hawthorne (2002, 244) makes essentially the same point, as does Vogel (1990), who Hawthorne credits as being the first to notice the parallel.

[19] The same applies to the other 'statistical' examples that are often cited in the literature as being supposedly analogous to the lottery case. For example, given the small probability that my car has just been stolen, do I know where it is parked right now? Hawthorne (2004, 4) is one of the many commentators who thinks not, and assimilates this case to the lottery puzzle. Again, however, it depends on the details of the situation, such as whether there really are car thieves operating in the area and what sort of vehicle they're looking for (as in the 'heart attack' case, the issue is whether one's car is one of those that is in the 'at risk' category). In short, it depends on whether there are near-by possible worlds in which my car is stolen, and one can't tell that simply by considering the mere fact that there is a low likelihood that, say, a locked Ford Ka will be stolen from outside a campus university in broad daylight on a working weekday.

[20] For more on the application of anti-luck epistemology to these issues, see Pritchard (2005).

[21] Acknowledgements.