

ANTI-LUCK EPISTEMOLOGY

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.¹ 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.² In outline at least, this is what I propose to do here. As we will see, this approach casts light on a number of issues in epistemology, and I focus on two issues in particular. First, I will argue that an anti-luck epistemology explains the attraction of adducing a safety condition on knowledge, and also assists us in determining the best formulation of this principle. Second, I will claim that an anti-luck epistemology can show us how best to respond to the co-called ‘lottery puzzle’.

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). 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.

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.

With our account of luck 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.³

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.⁴ 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).⁵ 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.⁶

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,⁷ 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). 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.⁸ Consider, for example, the 'rubbish chute' case described 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)

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 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 austere whilst 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 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.

4. EPISTEMIC LUCK AND THE LOTTERY PUZZLE

Anti-luck epistemology can also enable us to deal with other key problems in epistemology. Consider the so-called 'lottery puzzle'. It is commonly held these days that our apparent inability to know, prior to the draw being announced, that our lottery ticket is not a winner poses a fundamental problem for epistemology. A prominent example of this is John

Hawthorne's (2004) recent and highly influential book, which takes this puzzle as its starting-point. Here is how Hawthorne 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)

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

- (L1) *S* knows that he won't have enough money to go on safari this week.
- (L2) If *S* knows that he won't have enough money to go on safari this week, then he is in a position to know that he won't win a major lottery prize this week.
- (LC) *S* is in a position to know that he won't win a major lottery prize this week.⁹

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.¹⁰ 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 he 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 him to come into money next week. If we knew, for example, that his 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 he 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 his belief is false

(i.e., those worlds in which his uncle dies in the next week leaving him 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 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 he 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 his belief that he 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 he 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 he knows that he 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 he 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. 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.¹¹

5. CONCLUDING REMARKS

By taking the anti-luck platitude seriously, and thereby formulating an anti-luck condition on knowledge, we can thus cast light on (i) the motivation for, and correct construal of, the safety principle; and (ii) the lottery puzzle. 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.

REFERENCES

- 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.
- Greco, J. (2003). 'Virtue and Luck, Epistemic and Otherwise', *Metaphilosophy* 34, 353-66;
- Hawthorne, J. (2004). *Knowledge and Lotteries*, Oxford: Clarendon Press.
- Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy* 74, 549-67.
- Nozick, R. (1981). *Philosophical Explanations*, Oxford: Oxford University Press.
- 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.
- Sosa, E. (2000). 'Skepticism and Contextualism', *Philosophical Issues* 10, 1-18.
- Unger, P. (1968). 'An Analysis of Factual Knowledge', *Journal of Philosophy* 65, 157-70.
- Williamson, T. (2000). *Knowledge and Its Limits*, Oxford: Oxford University Press.

NOTES

¹ 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 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*)

² 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.

³ For obvious reasons, (LTB) won’t apply—at least without being adapted in some way—to logically 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 focus here is thus an account of what it is for a true belief in a (fully) contingent proposition to be lucky. I suggest some ways in which one might adapt (LTB) so that it applies to necessary propositions in Author [Book].

⁴ I shall ignore the issue of whether a belief could ever be the product of what is, self-consciously, a guess.

⁵ 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 Author [Article].

⁶ The safety principle thus lies at the heart of a broadly ‘Moorean’ response to the sceptic. See Sosa (1999) and Author (Article).

⁷ 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.

⁸ For a recent discussion along these lines, see Greco (2003).

⁹ This is roughly as Hawthorne (2004, 2) formulates it, 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 he truly believes them. Thus, there can be no harm in stipulating that he does in fact truly believe them.

¹⁰ 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).

¹¹ 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 he will be living in Syracuse next summer and the intuition that he does not know that he 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.