Episteme (2023), **20**, 23–38
doi:10.1017/epi.2020.44

ARTICLE

How to Play the Lottery Safely?

Haicheng Zhao

Department of Philosophy, Xiamen University, Xiamen, Fujian, China
Email: haichengzhao90@gmail.com

(Received 22 May 2020; revised 9 August 2020; accepted 1 September 2020;
first published online 26 October 2020)

Abstract

According to the safety principle, if one knows that p, one's belief that p could not easily have been false. One problem besetting this principle is the lottery problem – that of explaining why one does not seem to know that one will lose the lottery purely based on probabilistic considerations, prior to the announcement of the lottery result. As Greco points out, it is difficult for a safety theorist to solve this problem, without paying a heavy price. In this paper, I first reject three existing safety-based solutions to the lottery problem, due to Pritchard, Sosa, and Broncano-Berrocal. Failure of these accounts reveals that there is something crucial missing in the safety principle. By way of remedying this, I propose to integrate a safety principle with a condition regarding one's own *awareness* of (nearby) error-possibilities. The resulting account, as I argue, enjoys a number of theoretical advantages, including its capacity to handle the lottery problem.

Keywords: Safety; lottery; induction; awareness; testimony

Introduction

According to a rough-and-ready formulation of the safety principle:

> *Safety:* If one knows that p, then one's belief that p is not only true in the actual world, but also true in relevant nearby possible worlds (where the subject forms the belief with the same method as in the actual world.)[1]

Recent discussions have primarily focused on whether or not a principle like *Safety* can be successfully defended as a necessary condition for knowledge.[2] This paper investigates another related issue, i.e. the lottery problem. Very generally speaking, the problem is of explaining why true beliefs based merely on excellent probabilistic evidence don't seem to constitute knowledge. Suppose that someone called Lottie buys a fair lottery ticket. The odds of losing the lottery are massive. Before the result is announced, Lottie forms a true belief that she will lose, based (only) on considerations regarding the very high probability of losing the lottery. Many find it intuitive that Lottie's belief

[1] See Nozick (1981: 179) and Sosa (1999) for why a modal account of knowledge needs to be indexed to the same method. Cf. Zhao (2020*a*).

[2] Recent critics of *Safety* include Baumann (2008), Kelp (2009), Bogardus (2014), Zhao (Forthcoming), among others. Defenders of safety include Sosa (1999), Williamson (2000), and Pritchard (2007, 2008, 2012).

doesn't constitute knowledge, regardless of the very high probability in favor of its truth. This is puzzling because in many other cases where the odds of one's belief being true are lower than in Lottie's case, knowledge does seem to be present. Thus, we would often think that one can know that one loses the lottery purely based on reading the lottery result in a reliable newspaper (cf. Cohen 1998: 292–3). This is so even if the odds of the newspaper misprinting the result may be even higher than the probability of winning the lottery – that is, the odds of forming a true belief conditional on the newspaper result is lower than that conditional on the probability of losing.

This lottery problem has received much attention among epistemologists.[3] It is a desideratum of a theory of knowledge like safety to be able to tackle this perennial problem – in particular, to explain the intuition that Lottie lacks (or seems to lack) knowledge. In section 1, I introduce some background regarding safety and the lottery problem. In section 2 through section 4, I reject three existing safety-based solutions to the problem, due to Pritchard (2007, 2008, 2015), Sosa (2015) and Broncano-Berrocal (2019), respectively. Section 5 presents an account of knowledge that incorporates a safety principle with another condition regarding one's awareness of (nearby) error-possibilities. It is argued that this account enjoys a number of theoretical advantages, including its capacity to tackle the lottery problem. Finally, section 6 addresses some potential objections.

1. Background

Consider the following versions of the safety principle:

> *Safety(weak)*: If one knows that p, then one's belief that p is not only true in the actual world, but also true in *most* nearby possible worlds (where the subject forms the belief with the same method as in the actual world).
> *Safety(strong)*: If one knows that p, then one's belief that p is not only true in the actual world, but also true in *all* nearby possible worlds (where the subject forms the belief with the same method as in the actual world).

For a safety theorist, an obvious way to explain the lack of knowledge in Lottie's case is to appeal to *Safety(strong)*. Notice that although Lottie's belief that she will lose is true in most nearby possible worlds, thus satisfying *Safety(weak)*, given that there is a small number of nearby worlds where she will win, the belief is not true in *all* nearby worlds. Hence, the belief is not *Safety(strong)*-safe, which corresponds to the intuition that Lottie lacks knowledge.

However, Greco (2007) points out that *Safety(strong)* is problematic because it rules out certain cases of inductive knowledge. Consider:

> *Chute Case*
> After dropping a trash bag down the garbage chute of her apartment building, S believes truly that the bag will be in the basement later. Her basis for the belief is inductive: her past experience, background knowledge about the chute, etc. However, it could be the case that her belief is false: it is possible that the bag snags on its way down, though very unlikely. (Adapted from Sosa 1999: 145)

[3]See, e.g. Lewis (1996), Cohen (1998), Hawthorne (2004), McEvoy (2009), Dodd (2012), Mills (2012), Baumann (2016: ch. 4), for relevant discussions.

As per *Safety(strong)*, S's belief is unsafe. In a small range of close worlds, the bag would snag on its way down, rendering S's belief false. Therefore, *Safety(strong)* implies that S does not have inductive knowledge that her bag will be in the basement, which seems quite implausible.

To summarize, safety theorists are caught in a dilemma. On the first horn, by sticking to *Safety(strong)*, one can properly predict the lack of knowledge in Lottie's case, but this form of safety excludes inductive knowledge. On the second horn, *Safety(weak)* helps to preserve inductive knowledge, but this version of safety is too weak to explain the lack of knowledge in Lottie's case.

2. Pritchard's solution

Duncan Pritchard proposes the following safety principle to resolve the dilemma:

> *Safety(weighed)*: If one knows that p, then one's belief that p is not only true in the actual world, but also true in *all very close* possible worlds and *most close* possible worlds. (Pritchard 2007: 292; see also Pritchard 2015: 102)[4]

While preserving *Safety(weak)* as a component, *Safety(weighed)* puts additional weight on the "very close" possible worlds, such that a safe belief is intolerant of errors in any of these worlds. Pritchard thinks that *Safety(weighed)* judges both Lottie's case and *Chute Case* correctly. In the former, not-p worlds are *very close* to the actual world. As he puts it, a "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" (Pritchard 2007: 292). Thus, Lottie's belief that she will lose does not satisfy *Safety(weighed)*, so she lacks knowledge.

With regard to *Chute Case*, since it is assumed that it is unlikely for the bag to get caught,[5] we can imagine that there is some very minor dysfunction in the lift shaft, but one that is so slight that hardly any bag would snag on it. Under these circumstances, many details would have to be changed for the bag to snag on the chute, and so not-p world is not very close to the actual world. If so, the subject's belief that the bag will be in the basement is safe according to *Safety(weighed),* reflecting the intuition that she knows.[6]

In summary, if Pritchard is correct, it is really the *degree* of modal closeness of error-possibilities that plays a crucial role in determining whether one knows or not. In Lottie's case, not-p world being *very close* to the actual world makes her belief unsafe (as per *Safety(weighed)*), and so she fails to know. By contrast, the subjects in *Chute Case* and the newspaper example can come to know insofar as the error-possibilities are relatively remote.

Now, Pritchard's judgments of modal closeness are made on a primitively intuitive level. These judgements rely on the amount of relevant changes involved in a possible world compared with the actual world. It would be better to specify a more principled way of determining possible world similarity/closeness, so that one can evaluate safety in a more principled manner, thus avoiding the worry of ad hocness (cf. Baumann

[4]Of course, these "very close possible worlds" and "close possible worlds" should also be restricted to those where the subject forms the belief via the same method as in the actual world. See note 1.

[5]Notice that if we assume that a bag frequently gets snagged on its way down, then it's not clear that one has knowledge in this case. See Pritchard (2007).

[6]Pritchard thinks that a similar diagnosis also helps to explain why one knows one loses the lottery by reading the announcement result on a newspaper. See Pritchard (2008).

https://doi.org/10.1017/epi.2020.44 Published online by Cambridge University Press

2008; Kvanvig 2008; McEvoy 2009). Setting this issue aside, however, I want to argue that what epistemically distinguishes Lottie's case from *Chute Case* (as well as the newspaper example) is *not* the degree of modal closeness of error-possibilities. This is so even if we grant Pritchard's methodology of determining modal closeness. Specifically, in what follows I will revise Lottie's case in a way that the error-possibilities are no longer 'very close' to the actual world. Despite this revision, it will be clear that Lottie still lacks knowledge. Therefore, Pritchard's diagnosis of the lottery problem fails. Consider:

> *Inflated Lottery Ball*
> The lottery that Lottie plays works as follows. A player picks three different numbers from 1 to 100. The lottery result is announced via a big lottery machine, which contains 100 balls with different numbers (from 1 to 100) marked on their surfaces. Upon operation, the machine completely randomly picks three balls as the lottery result. Anyone whose ticket numbers match the numbers on the announced three balls will win.
> The day before the announcement, Lottie buys a lottery ticket. One of her chosen numbers is 13. A few minutes later, at t1, by realizing the extremely low probability of winning the lottery, Lottie forms a belief that she will lose. A couple of hours later, at t2, due to some very peculiar chemical effect, ball number 13 in the lottery machine inflates a bit, so that it becomes very difficult for *this* ball to be drawn. That is, the chute in the machine that delivers the drawn balls is too narrow for ball 13 at its current size to successfully go through. Thus, for ball number 13 to be a candidate of drawn balls, quite a few changes need to be made: someone needs to open the machine before the announcement, detect the abnormal size of the ball, and promptly fix the problem, etc.

Now, I take it as quite intuitive that, from t1 to the moment when Lottie learns the lottery result, she *does not know* that she will lose the lottery. After all, her belief is always based on the same probabilistic considerations. She does not gain any new evidence that she will lose, nor is there any new input to her belief-forming process, etc. Sure, the fact that ball 13 inflated at t2 may make it even more unlikely that the numbers Lottie picked will win. Notice, however, that the moral of the lottery example is precisely that high probability in favor of the truth of one's belief does *not* imply that the belief is knowledge. Given that the probability of her losing the lottery is already extremely high, there is no reason to think that any increased probability would turn Lottie into a knower.

However, *Safety(weighed)* delivers a rather strange and implausible verdict here. Given that after t2 quite a few changes need to be made for ball 13 to be drawn, a possible world where Lottie's belief becomes false (i.e. a world in which she wins the lottery) becomes relatively remote. It is no longer the case that for Lottie to win, only a few colored balls need to fall in a slightly different configuration, as Pritchard (2007: 292) conceives in the original Lottie case. Rather, at least after t2, many changes need to be made for Lottie's numbers (especially number 13) to win. This means that, after t2, Lottie's belief that she will lose becomes *safe* according to *Safety(weighed)*, and so, the belief suddenly becomes a good candidate for knowledge. In fact, following Pritchard's analysis, *all* the players whose lottery tickets contain number 13 and who believe that they will lose (purely on the basis of the high probability of losing) may

https://doi.org/10.1017/epi.2020.44 Published online by Cambridge University Press

come to know that they will lose after t2;[7] whereas all the other counterpart players whose numbers do not contain 13 continue to be ignorant. These results are particularly bizarre, since, as we may assume, all the players form the belief based on the same probabilistic reasons, and are not aware that anything happened to the machine, etc.

To clarify, one does not have to be an internalist to find the above results generated by *Safety(weighed)* to be implausible; advocates of various versions of externalism should find them unacceptable as well. Throughout t1 to the announcement of the lottery result, Lottie's belief is caused by the same belief-forming process and is the result of the same cognitive abilities. Therefore, process reliabilism (cf. Goldman 1979) and virtue reliabilism (Greco 2010; Sosa 2015) would *not* suggest that Lottie can suddenly come to know at t2. Also, even at t2, the fact that Lottie's belief is formed by way of properly functioning cognitive faculties, in a congenial epistemic environment, according to a good design plan, remains unchanged. Thus, on proper functionalism (Plantinga 1993), she doesn't suddenly come to know at t2 either.[8] Finally, no matter at t1 or t2, were Lottie to win the lottery, she *would* still believe that she will lose. So, sensitivity (Nozick 1981) also does not lead to the result that she can suddenly know at t2. In sum, all these prominent externalist accounts predict differently than *Safety (weighed)*, which further casts doubts on the plausibility of the latter. It thus turns out that both intuition and theoretical considerations count against *Safety(weighed)*'s verdict on *Inflated Lottery Ball*.[9]

To conclude, *Inflated Lottery Ball* delivers a useful lesson. In particular, it shows that whether Lottie knows or not is not determined by the degree of modal closeness of not-p worlds, contrary to what *Safety(weighed)* suggests. The modal closeness of these worlds can be stipulated, so that they are relatively remote from the actuality, just as not-p worlds seem to be relatively remote in the above *Chute Case* and the newspaper example. Even with such a stipulation, however, the point that Lottie lacks knowledge, whereas the latter examples feature knowledge, remains unchanged. Therefore, what makes Lottie's case epistemically distinctive is not the degree of modal closeness.

---

[7]One may think that since *Safety(weighed)* is only a necessary condition for knowledge, advocates of this account could argue that at t2 Lottie's belief violates some other necessary condition for knowledge, so that she continues to fail to know. But this strategy seems unappealing. For the problem remains – that *Safety(weighed)* explains why Lottie fails to know at t1 but it cannot explain why she continues to fail to know after t2. Such disparity just shows that *Safety(weighed)* is not a good candidate for addressing the lottery problem.

Incidentally, Pritchard (2012) adds to the safety condition an independent "ability/virtue condition," according to which knowledge requires that one's belief that p must be a product of one's relevant cognitive abilities. And this additional condition does not help to explain why Lottie fails to know throughout: her belief indeed satisfies the ability/virtue condition, as it is indeed a product of her relevant cognitive abilities (no matter at time t1 or t2). (See Pritchard 2012: 266–7).

[8]Cf. Plantinga (1993: 166).

[9]There is some other theoretical consideration that counts against *Safety(weighed)*'s verdict. As Roush (2005: 122) points out, one problem with safety is "wrong direction of fit." In particular, knowledge seems to be a matter of one's responsiveness to the way the world is, whereas safety makes a demand in the opposite direction. Roush's own example that illustrates this point is as follows: "imagine a case where S has a fairy godmother, whose special mode of operation is instantaneously to make true anything that S believes. If so, then for this S, for any p, if S believed p it would be true" (Roush 2005: 122). Roush claims that in such a case, although S's beliefs are safe, they can hardly be counted as knowledge. After all, the facts in the world happen to match the subject's relevant beliefs, making them safe, without the subject herself responding to those facts in any way. Similar point applies to Lottie's belief at t2. At this moment, the world suddenly cooperates with Lottie's belief, so that her belief becomes safe according to *Safety (weighed)*, without any sort of responsiveness from Lottie's part. Put colloquially, in both cases, the subject gets safe beliefs too 'cheaply', so much so that the beliefs cannot count as knowledge.

The latter is a red herring; one needs to look somewhere else to handle the lottery problem.

3. Sosa's solution

Let us turn to Sosa. Sosa (2015: 120) thinks that *Safety(weak)* is more plausible than other formulations of safety. Consequently, he is committed to the claim that Lottie's belief is *safe*. Interestingly, however, he thinks that this is a correct result because Lottie does *know* that she will lose. He then offers an error theory to explain why many tend to have the intuition that Lottie does *not* know.

According to the error theory, it is the (weaker) safety condition (i.e. *Safety(weak)*) that is required for knowledge, instead of sensitivity. According to the latter:

> *Sensitivity*: If S knows that p, then S's belief that p, formed via method M, is sensitive, i.e. in the nearest possible worlds where p is false, S does not believe p via M.

Lottie's belief that she will lose is *in*sensitive: If her belief were false – i.e. if she were holding a winning ticket – she would still believe that she will lose (based on the probabilistic information). However, as aforementioned, her belief is (weakly) safe. Moreover, Sosa thinks that safety and sensitivity are easily confused, because it is easy to assume, incorrectly, that the conditional of safety (i.e. $Bp \Box\!\!\rightarrow p$) and that of sensitivity (i.e. $\sim p \Box\!\!\rightarrow \sim Bp$) contrapose (cf. Sosa 1999: note 1). But they are really inequivalent conditions. And it is the (weaker) safety that is a correct necessary condition for knowledge. Therefore, in Lottie's case, people are misled into thinking that she does not know, simply because they have confused the correct (weaker) safety condition with the incorrect sensitivity condition. The former correctly judges that Lottie's belief is safe, whereas the latter misleadingly predicts that the belief is insensitive.

Sosa's explanation is not without difficulties. First, it is an empirical claim that those who have the intuition that Lottie does not know are confusing sensitivity with safety. Such an empirical claim needs to be confirmed by empirical evidence regarding folk psychology; otherwise the claim remains as a speculation.

More importantly, Sosa's above diagnosis leads to some implausible inconsistencies. To illustrate, notice that just as in Lottie's case, in the above descriptions of *Chute Case* (see section 1), the subject's belief is also weakly safe but insensitive. Specifically, if her belief were false, she would still hold the same belief that the bag will be in the basement later (via the same inductive method); hence, the belief is *in*sensitive.[10] At the same time, recall that the belief is true in most relevant nearby worlds, so the belief is safe according to *Safety(weak)*. These results imply that, just as in Lottie's case, it should be expected that in *Chute Case* people tend to incorrectly think that the belief does *not* constitute knowledge, due to confusing safety with sensitivity. However, notice that contrary to Lottie's case, it is almost unanimously agreed that there is inductive knowledge in *Chute Case*. So, it is hard to see how Sosa could consistently deal with both Lottie's case and an example like *Chute Case*. If safety and sensitivity are being confused vis-à-vis Lottie's case, there is no consistent reason to think such confusion does not exist in the other example, given that both cases seem to feature (weakly) safe but insensitive beliefs.

[10]In fact, this is part of Sosa's (1999) objections against sensitivity.

https://doi.org/10.1017/epi.2020.44 Published online by Cambridge University Press

4. Broncano-Berrocal's solution

Broncano-Berrocal (2019) presents an alternative version of safety, which indirectly tracks "appropriate determining conditions," as opposed to directly tracking truths of beliefs:

> *Indirect Safety*: If S knows that p via a method of belief formation M, then the determining conditions in the actual world are appropriate and in nearly all (if not all) close possible worlds in which S continues to believe that p via M, the determining conditions for S's belief that p continue to be appropriate. (Broncano-Berrocal 2019: 48)

Broncano-Berrocal thinks that with *Indirect Safety* (IS), one can properly solve the lottery problem. Before considering his arguments, some clarifications are in order. "Appropriate determining conditions," as Broncano-Berrocal defines the terms, are the kind of conditions that determine whether an agent's doxastic state matches the facts. In perceptual cases, for example, such conditions include proper lighting conditions, distance and the size of an object, etc. Determining conditions are to be differentiated from mere "enabling conditions," such as the existence of oxygen – which enables, without determining – a belief-formation. Furthermore, Broncano-Berrocal (2019) makes another crucial distinction between "conditions that make success likely" and "conditions that make success *in the right way* likely." As he understands it, only the latter refers to conditions that allow a performance to constitute an *achievement*. To illustrate, suppose that all the targets in an archery field are replaced with powerful electromagnets, such that anyone who releases a (metallic) arrow will hit the target easily, regardless of the archer's abilities (Broncano-Berrocal 2019). Here, obviously the conditions that one is in make the success (i.e. hitting the target) extremely likely. In that sense, the conditions are *appropriate* for hitting the target. However, the conditions are *in*appropriate for one's shots to be considered as "successful *in the right way*," or be considered as achievements, simply because the presence of electromagnets are unduly helpful to one's shots.

Analogously, Broncano-Berrocal thinks that knowledge must track appropriate determining conditions that allow one to attain true belief in the right way, such that the true belief can be considered as cognitive achievement. This, according to him, is not obtained in the lottery case. For basing one's belief on the mere statistical evidence gives rise to inappropriate conditions in which one's true belief cannot be counted as cognitive achievement. Put simply, mere statistical evidence is unduly helpful to getting truth, just as the presence of electromagnets is unduly helpful to hitting the target. I will let Broncano-Berrocal speak for himself:

> More specifically, the kind of circumstances in which one merely bases one's beliefs on (strong) statistical evidence are analogous to the kind of circumstances in which one shoots metallic arrows in a field with powerful electromagnets. In both cases, the circumstances make success (shooting an arrow accurately) and cognitive success (believing correctly that one will lose the lottery) very likely, almost guaranteed. In other words, in both cases the circumstances are appropriate for *succeeding simpliciter*, cognitively or non-cognitively. However, they are *in*appropriate for *succeeding in the right way*, i.e., for achievement or cognitive achievement (knowledge).
>
> As before, the reason is that, like the effect of the electromagnets, strong statistical evidence is *unduly helpful*. It is *helpful* because it makes one get things

https://doi.org/10.1017/epi.2020.44 Published online by Cambridge University Press

> right in a great number of instances of belief formation, just as the electromagnets make one hit the bull's eye in many instances of shooting. It is *undue* (in the sense of improper) because, just as the electromagnets prevent each particular successful shot from being explained in the right way by the competent exercise of the archer's shooting abilities, *mere* statistical evidence makes the exercise of our cognitive abilities irrelevant in the explanation of why we come to get things right on each particular occasion. (Broncano-Berrocal 2019: 49; italics original)[11]

Thus, Lottie's belief does not track the required appropriate determining conditions, and so it is unsafe according to IS, which explains why Lottie does not know.

However, I find Broncano-Berrocal's diagnosis of the lottery problem still problematic. In particular, the claim that the conditions (or circumstances) in the lottery case are such that one's true belief is not a cognitive achievement seems dubious. To illustrate, notice that in the standard lottery case, there are two cognitive steps required for forming the belief that one will lose the lottery:

(1) Via a certain method, one comes to believe that the *odds* of her ticket losing the lottery are massive. (Call this belief(probability).)
(2) Based on belief(probability), one believes that she will *lose* the lottery. (Call the latter belief(losing).)

Discussions of the lottery problem frequently (and narrowly) focus on (2). With more attention being paid to (1), however, we will find Broncano-Berrocal's diagnosis implausible. Consider:

> *Amy and Bill*
> Sally gives Amy and Bill each a lottery ticket, without telling them any information about the odds involved for winning. Amy is epistemically careful. She asks herself: "Couldn't the probability of losing not be so high?" Back at home, she decides to use her brilliant mathematical abilities to calculate the odds. Half an hour later, she gets the correct answer: 1 out of 1 million chance for winning. She then realizes that the chance of losing is too high, and thus believes that she will lose. Bill, with his pessimistic character, does not do any calculation, investigation, or the like. He just tells himself: "I doubt I could be the lucky guy. The odds of losing must be pretty high, anyway …" It turns out that the odds of their tickets' losing are the same, and neither of them won.

This story shows that one may form belief(probability) in different ways, either haphazardly or carefully. Such a difference, as it appears, does not affect whether one *knows* she will lose or not. To the extent that both Amy and Bill base their belief(losing) merely on the probability involved, it seems neither of them know that they will lose.[12] That

[11]Broncano-Berrocal (2019) further illustrates this last point by referring to an example by Adler (2005): Suppose that one already knows that 999 out of 1000 apples in a barrel are good. Then, judging that each apple in the barrel is good merely because most of the apples in the barrel are good would be epistemically improper, in the sense that such judgment does not constitute cognitive achievement. By contrast, one's correct judgment that a particular apple is good would be cognitive achievement, if such judgment is a result of one's careful and skilled examination of the apple. In the latter, the correct judgment is attributable to one's apple-sorting abilities.

[12]Could one plausibly insist that Amy's belief constitutes knowledge? Not without embracing the following bizarre result: If Amy knows, then a lottery player just needs to sit down and meticulously calculate the odds, in order to figure out whether she wins, without the need to wait for the announcement at all. This is

https://doi.org/10.1017/epi.2020.44 Published online by Cambridge University Press

said, it is clear that Amy's, but not Bill's, belief(losing) constitutes cognitive *achievement*. Put colloquially, it is Amy who *earned* her true belief(losing) with her distinctive cognitive abilities. She deserves the credit.

Hold on, though. To the extent that Amy forms her belief(losing) based purely on the probability, aren't the conditions still "unduly helpful," rendering the belief not cognitive achievement? I think such a response misses the point. Amy's (true) belief(losing) can be counted as cognitive achievement because she did a brilliant job in the first step – figuring out what the probability is. And such etiology matters. To illustrate, consider a swimmer who wins a gold medal in an Olympic 100 m freestyle. Suppose that she performed very well in the first half of the race (let's say her average speed in this period far exceeds the world record.) Then, her winning the gold medal certainly constitutes achievement, even if her performance in the latter half of the race is not as impressive as the first half. What matters is that her excellent performance during the first half is causally/explanatorily responsible for her winning the race. Analogously, just because Amy bases her belief(losing) purely on the probability – a step that Broncano-Berrocal deems as 'unduly helpful' – does not mask the fact that she did a wonderful job in the previous step of calculating the probability. To the extent that the latter is explanatorily/causally responsible for Amy's formation of (true) belief (losing), this belief does constitute cognitive achievement.

Thus, contra Broncano-Berrocal, at least sometimes one's belief(losing) *could be* cognitive achievement. Or, in his terms, the conditions *do not have to be* such that they are inappropriate for the belief to constitute achievement. Therefore, his solution to the lottery problem fails. And such failure leaves us with an interesting question: If one's belief(losing) could be cognitive achievement, then how to explain the intuition that the belief is not knowledge? To answer this question, we may compare Amy and Bill again. Notice that, although they have taken different routes to come to believe that the chance of losing is very high, their commonality is that they are both aware of the possibility of winning the lottery (i.e. the possibility in which their belief is false.) As I shall argue in the next section, it is precisely such awareness that serves to explain why subjects like Amy and Bill both fail to know, despite their differences in terms of achievement.

5. Safety-awareness account

In this section, I will propose an alternative way to solve the lottery problem on behalf of a safety theorist. In fact, as will be shown, the benefits of my account are not restricted to resolving the lottery problem; it has a broader appeal.

First off, some preliminaries. Recall that Pritchard adopts *Safety(weighed)* as opposed to *Safety(weak)* because he thinks that the latter cannot properly solve the lottery problem. But abandoning *Safety(weak)* straightaway solely on the basis of its inability to solving this problem seems a bit too quick. For one thing, setting aside the lottery cases, *Safety(weak)* seems remarkably more attractive than the other formulations of safety introduced in sections 1 and 2. Compared to *Safety(weighed)*, *Safety(weak)* does not rely on the murky distinction between "*close* possible worlds" and "*very close* possible worlds." As critics (e.g. Baumann 2008; Xu Forthcoming) have pointed out, giving a plausible and systematic account of "closeness relation" is not an easy task. *Safety(weighed)*, however, seems to exacerbate the difficulty here by introducing the additional distinction between "close" and "very close" worlds. (If giving an account of "closeness relation" is already challenging, how could one get a good grip on the yet

clearly untrue, otherwise the announcement largely becomes a redundancy, designed only for people who lack abilities in calculation, or who are lazy in calculating.

https://doi.org/10.1017/epi.2020.44 Published online by Cambridge University Press

more complicated distinction here?) In a word, although "closeness relation" is a crucial component in a safety principle, it would be preferable for a safety theorist to let such relation play a relatively minimal role, without adding further complexities to it.[13] In this regard, *Safety(weak)* seems better than *Safety(weighed)*.

Also, compared to *Safety(strong)*, *Safety(weak)* moderately requires truths of beliefs in the majority of nearby worlds – as opposed to all of them – and so the latter nicely accommodates inductive knowledge.

Perhaps more importantly, the problem with *Safety(weak)* vis-à-vis the lottery cases is that it fails as a sufficient condition, rather than as a necessary condition. That is, Lottie's belief that she will lose is *Safety(weak)*-safe but does not amount to knowledge, which points out that this form of safety is insufficient for knowledge. Hence, a safety theorist has the option of retaining *Safety(weak)* and handling the lottery problem with other theoretical resources.[14]

So, what else could be added to *Safety(weak)* to solve the lottery problem? Or, what is lacking in this account that prevents it from tackling the problem? To answer these questions, observe first that safety in general is an externalist modal condition, such that whether a belief is safe or not depends on the facts in the *world* – in particular, the facts which may secure modal stability of one's belief. Thus, safety leaves out factors regarding one's own *perspective* on the modal stability of her beliefs. More specifically, safety leaves out whether or not a *subject herself* is aware of any error-possibilities of her beliefs. I think once this "subject's part" is taken into account, a solution to the lottery problem on behalf of a safety theorist will naturally emerge.[15]

Thus, notice that in lottery cases, since we assume that the subject forms her belief that she will lose on the basis of the probabilistic considerations alone, she is de facto aware of the (nearby) possibility of winning the lottery, where "awareness" can be simply construed as "belief." That is, the subject *believes* (either implicitly or explicitly) that there is some close error-possibility in which she wins the lottery and thus her belief that she will lose is false. This is true even in *Inflated Lottery Ball*. Since Lottie does

[13]It is not my ambition here to fully address the issue of determining closeness relation, but it is worth making the following clarification. Following most discussants of safety, I understand "closeness relation" roughly along the line of Lewisian possible world semantics. Closeness of a possible world is thus determined by the similarity of that world compared with the actual world (i.e. the more similar a world is to the actual world, the closer it is). In particular, we need to consider the kind of *relevant* changes (/differences) involved in a possible world compared with the actual world (i.e. the kind of changes that carry more 'weight'), in order to determine whether that world is close or not. In the context of safety, such relevant changes must be *epistemic* ones. Specifically, I think these changes must be factors that *make the target belief either true or false*, as opposed to factors that are merely *related to* the truth or falsity of a belief. Without bearing in mind such a distinction, one may easily get confused by Pritchard's claim that the possible world in which Lottie wins is a close one (assuming that this is the 'normal' lottery case without an inflated lottery ball or anything like that). After all, the possible world in which Lottie wins the lottery certainly involves many conceivable changes compared with the actual world, if we imagine that Lottie would become very wealthy and thereby lead a quite different life. But her being wealthy is certainly not what makes her belief true or false, although it is related to the falsity of her belief. What contributes to the truth/falsity of her belief is just how the lottery mechanism works (whether the lottery machine is working properly, whether the lottery is being rigged, etc.) And only these factors are what we need to focus on in determining closeness. (But of course, putting the lottery case aside, it is not always clear what kind of factors make a belief true or false, thus making it difficult to decide what factors determine closeness relation in a principled manner.) Thanks to an anonymous referee for pressing me on this issue.

[14]Most safety theorists indeed defend safety as necessary, as opposed to a sufficient condition, for knowledge. See, e.g., Williamson (2000), Pritchard (2012), Greco (2016).

[15]Cf. Zhao (2020*b*: 854–5), Zhao and Baumann (Forthcoming).

https://doi.org/10.1017/epi.2020.44 Published online by Cambridge University Press

not know anything happened with ball 13, she is misled to believe that there is some close possibility in which she wins the lottery.

Some clarifications must be considered. First, it is not intended here that Lottie's awareness/belief must explicitly involve such terms as "close possibility" or "nearby possible world." It is unlikely that ordinary people who play the lottery would form beliefs that precisely involve such concepts. The point is rather that what Lottie believes can be *analyzed* or *interpreted* in terms of "close possibility" or "nearby possible world." For example, it is plausible that an ordinary player like Lottie would believe that there need not be many changes for her to win the lottery ("I just need to pick slightly different numbers!"), and such a belief can then be interpreted as belief regarding *nearby/close error-possibilities*.

Second, understanding "awareness" as belief means that my technical usage of the term "awareness" is, to some extent, detached from how this word is used in ordinary life. Ordinarily, "awareness" is factive – e.g., one cannot be said to be aware that it is raining outside if it is not really raining outside (Silva 2019: 725). However, since awareness in the current context is understood as belief, just as belief is non-factive, the sense of "awareness" here also inherits the feature of being non-factive. This is not itself a problem against the account that I will present below. After all, the purpose of my theorizing is not conceptual analysis of the term "awareness." (If the reader finds my usage of "awareness" unpalatable, simply substitute the phrase "awareness of nearby error-possibility" with "belief that there is nearby error-possibility.") Indeed, as will be shown momentarily, the non-factive feature of my "awareness" allows the following account to accommodate the phenomenon of misleading defeat.

With these points in mind, let us consider the following proposal:

Safety-Awareness Account (*SAA*)
S knows that p, *only if*

(1) *Safety(weak)*: S's belief that p is not only true in the actual world, but also true in *most* nearby possible worlds (where the subject forms the belief with the same method as in the actual world).
(2) *Lack-of-Awareness*: S is not aware of nearby error-possibilities, i.e. nearby possible worlds in which her belief that p is false.

(Notice that the phrase "*Lack-of-Awareness*" is not to be construed literally in terms of its face value. In many contexts, if one lacks awareness of something, it may sound as if the person is being negligent – e.g. "He lacks awareness of how serious the current situation is." In this way, "lack of awareness" may sound undesirable and negative.[16] But it certainly seems odd that something undesirable and negative is necessary for knowledge. Indeed, in my usage of the term, "*Lack-of-Awareness*" essentially means "no-defeater," which I will explain below.)

Now, *SAA* handles the aforementioned problematic cases nicely.[17] In Lottie's cases, again, her belief that she will lose satisfies *Safety(weak)*, but since she is aware of some nearby error-possibilities of winning the lottery, her belief violates *Lack-of-Awareness*. So, she fails to know that she will lose. Turning to *Chute Case* and the newspaper example, since it is assumed that these examples feature knowledge, presumably the

[16] I address this particular sense of "*Lack-of-Awareness*" in note 20.

[17] Zhao (2020*b*), Zhao and Baumann (Forthcoming) argue that an account like *SAA* also addresses some other objections to safety.

https://doi.org/10.1017/epi.2020.44 Published online by Cambridge University Press

chute is functioning properly and the newspaper is a reliable one (otherwise there is no reason to think that the subjects know in these cases.) Yet, given these assumptions, there appears no special reason to think that these subjects are aware of any nearby error-possibilities under these ordinary circumstances.[18] So, the subjects' beliefs plausibly satisfy *Lack-of-Awareness*. This, coupled with the point that the subjects' beliefs are *Safety(weak)*-safe, implies that the current proposal preserves inductive knowledge and renders it possible to know one loses the lottery by reading a reliable newspaper.

Furthermore, adding *Lack-of-Awareness* to *Safety(weak)* is not an ad-hoc maneuver. For as mentioned earlier, the former can be independently defended as a "no-defeater condition." That is, when S is aware of a nearby error-possibility where her belief that p is false, it seems quite plausible that she simultaneously acquires a (doxastic) defeater against her belief that p. Hence, her belief that p does not constitute knowledge. In this way, the notion of "awareness of nearby error-possibility" helps the safety theorist to accommodate the important phenomenon of knowledge-defeat, as any theory of knowledge should do. To illustrate, consider a paradigmatic example of knowledge-defeat from Chisholm (1966: 48). (See also Lasonen-Aarnio 2010). At a time t1 Suzy comes to know that a certain object is red based on perception (she has normal perceptual abilities, the lighting is normal, etc.). At a slightly later time t2, she believes based on someone's false testimony that the object is illuminated by peculiar red lighting, lighting that would make objects of any color look red. In our terminology, it can be said that at t2 Suzy loses her knowledge that the object is red, because she is misled to believe(/aware) that some nearby error-possibilities exist – the possibilities in which the object is illuminated by the red lighting, without itself being red.

Indeed, understanding *Lack-of-Awareness* as a no-defeater condition also delivers fresh insights in solving the lottery problem. Although there are various attempts of addressing this problem in the literature, few have attempted to consider the case as one that involves the phenomenon of defeat. However, such an attempt seems to be both justified and interesting. In particular, we may claim that, the reason Lottie does not seem to know is because she has a doxastic, modal defeater – a defeater that essentially involves such contents as "I *could* be the winner," "My belief of losing the lottery *could* turn out to be false," etc. Furthermore, notice that "could" here – as explained above – refers to a relatively close error-possibility, at least from the subject's own perspective. And plausibly, such a defeater goes some way in undermining the epistemic status of Lottie's belief that she will lose. Specifically, due to the defeater's *modal* character, Lottie's mere probabilistic evidence becomes insufficient for adequately supporting (epistemically) her belief. It is insufficient in the sense that she must gain additional resources to eliminate the error-possibility of winning (say, by learning from the lottery announcement that she is indeed not the winner). And lacking such resources, she fails to know that she will lose. Put another way, the modal defeater functions in a way that it places a particular kind of epistemic demand on the subject – i.e. a demand for eliminating the close error-possibility.

[18]Perhaps there are some peculiar scenarios in which the subjects hold their beliefs that p (that the bag will be in the basement/that my ticket numbers didn't win), but they are misled to believe some nearby error-possibilities exist. For example, the subjects could be paranoid so that they falsely believe that the bag could easily have been snagged or that the newspaper could easily have misprinted the result. However, I think it is unclear that these paranoid subjects really know that p. Or, at least, the intuition here seems quite unclear. Thus, not predicting knowledge in such scenarios is not an objection against *SAA*.

6. Some lingering worries

6.1. Is SAA too permissive?

The set-up of the standard lottery cases – specifically, the point that Lottie forms her belief that she will lose based purely on *statistical/probabilistic* evidence – guarantees that Lottie is aware of the possibility of winning. In fact, even if a lottery player's belief is not formed on such a basis, it is hard for one who has minimal background information regarding a lottery to not be aware of the possibility of winning (unless one already knows the lottery is rigged). Being aware of such possibility is just part of the game, so to speak. Furthermore, I have argued that such awareness can be plausibly analyzed or interpreted in terms of awareness of *nearby error-possibility*, thus rendering Lottie's belief a violation of *Lack-of-Awareness*. That said, we may still conceive cases where one does not have such awareness. And such cases may put some pressure on *SAA*. The most reasonable counterexample I could think of goes as follows.

Imagine a five-year-old named Tom. Tom has no background information whatsoever about how lottery works. One day, his dad gives him a lottery ticket. His mom, who believes that the ticket is not a winner (based purely on the probabilistic considerations), tells Tom that he will lose the lottery. Tom trusts his mom, as usual, and comes to believe that he will lose the lottery.

Here, it seems far-fetched to say that little Tom is aware of the nearby possibility of winning the lottery. Indeed, we could even stipulate that he doesn't quite understand what "winning the lottery" really means – he just trusts his mom and literally believes that he will lose. So, Tom's belief satisfies *Lack-of-Awareness*. Assuming the belief is also *Safety(weak)*-safe, it seems that *SAA* implies that Tom can come to *know* that he will lose the lottery. But isn't this implausibly *permissive*? After all, even Tom's testifier – his mom – doesn't know that the ticket is not a winner. Then, how could a child who knows even less in general about the lottery than his mom have the fortune of *knowing* that he will lose? It thus seems that whether one is aware of nearby error-possibility or not doesn't matter.

In reply, I think that the reason why Tom doesn't know is fundamentally different from why we tend to think that the subject in the standard lottery case – or, for that matter, someone like Tom's mom – doesn't know. Put simply, these cases warrant different theoretical treatments. In the following, I will offer two resolutions, leaving it to the reader to decide which one is better.

First, since the above alleged counterexample only puts pressure on the sufficiency of *SAA*, we may appeal to some other account to explain it away. Consider transmission theory about testimony, the idea that if a hearer knows that p based on a speaker's testimony, the speaker must also know that p.[19] On this view, Tom doesn't know simply because Tom's testifier – his mom – doesn't know. Transmission theory is compatible with *SAA*, to the extent that *SAA* is not aimed at giving sufficient conditions for knowledge.

Second, and perhaps more ambitiously, a *SAA*-theorist may avail herself with a relatively more refined version of *Safety(weak)*. In his criticisms of Pritchard's early "anti-luck epistemology," Greco (2007: 302) argues that "safety must have its seat in S's cognitive abilities (or virtue) …." What he means by "seat" here can be understood as one's "method" (broadly construed) that gives rise to the belief. In other words, not just any method can give rise to a safe belief. For a belief to be properly safe, in Greco's sense, it must be produced by a *virtuous method* that manifests the subject's cognitive abilities. (See also Greco 2016.)

[19]For some critical discussions of transmission theory, see Lackey (1999).

https://doi.org/10.1017/epi.2020.44 Published online by Cambridge University Press

Indeed, I find Greco's arguments for the above point plausible: Sometimes one's belief can be "safe" just by default, without manifesting the subject's cognitive abilities at all. A wishful thinker's belief that p is "safe," if the world happens to be arranged in a way that p is stably true across nearby possible worlds (see Greco 2007: 302). Plausibly, restricting safety to "virtuous method" avoids counting beliefs of this kind as knowledge.

Now, what counts as a "virtuous method" is a tricky issue that I cannot fully address here (cf. Hirvelä 2019, 2020). But it seems plausible to take Tom's belief as *not* being formed via a virtuous method. This is because a virtuously formed testimonial belief at least requires an adequate understanding of the content of the testimony. Given the set-up of the example, however, Tom does not have relevant background information about how a lottery works. He (almost blindly) trusts his mom that he will lose, without proper linguistic-interpretive understanding of what the testimony really means. (Notice that if he does understand properly, he should also be aware of the possibility of winning, in which case the imagined scenario would not be a counterexample at all.) Hence, Tom's belief is not properly safe, according to our more sophisticated safety principle. And this explains why he doesn't know he will lose. Compare this with another scenario where Tom's mom says to him: "Your granny has come to visit you – she is downstairs!" Assuming what mom says is true and that Tom accepts her testimony, we may take him to *know* that his granny has come to visit him. For even a five-year-old may have the linguistic-interpretive abilities to understand the content of the testimony here: he knows which individual is his grandma, he understands what is meant by "visiting him," etc. In this way, Tom's testimonial belief that his granny is visiting him is virtuously formed and is properly safe.[20]

6.2. Is SAA too demanding?

Another worry takes a different direction. Notice that most (if not all) human cognitive faculties are fallible. Thus, for a subject S and a belief p that she holds, S may (at least implicitly) believe that p is possibly false. But then, isn't such awareness an awareness of "nearby error-possibility"? If so, skepticism seems to be a result of *Lack-of-Awareness*. And so, the latter cannot be what is required for knowledge.

We may call the fact that human cognitive faculties are fallible a kind of "generic error-possibility." Merely being aware of such a fact is not a violation of *Lack-of-Awareness*. For such a generic error-possibility cannot be evaluated as either nearby or remote – it is just a brute fact. One needs to look at a particular scenario that one is aware of – a scenario which *instantiates* the generic error-possibility – in order to determine whether what one is aware of is remote or nearby. And crucially, in the kind of good cases where knowledge is intuitively present, the scenario one is aware of tends to be remote. Take memory as an example. Right now, I am (episodically)

[20]One may think that *SAA* is too permissive for another reason – i.e., because of "normative defeaters." A normative defeater is a proposition that one does not actually believe to be true, but given the circumstances that one is in, one *should* believe it to be true. See Lackey (1999). Cases involving a normative defeater typically feature a subject who does not perform well epistemically (e.g., being careless, narrow-minded, dogmatic etc.), so that there is some proposition related to her certain belief that she inappropriately does not believe to be true. *SAA*, as it stands, does not take into account normative defeaters. A quick fix would be to add another clause: there are no such nearby error-possibilities that one should be aware of but is not actually aware of. In addition, invoking "virtuous method" for our safety principle also goes some way towards accommodating normative defeaters. Again, a subject who has a normative defeater against her belief is not epistemically performing well. In that sense, her belief is arguably not virtuously formed, all things considered.

remembering that I met my colleague Joe *a moment ago*. Indeed, I *just* had a conversation with him in the hallway of our department building. My memory of meeting him is remarkably vivid and clear. And I, let's assume, have normally functioning, though not infallible, memory systems. Granted, even under this benign circumstance, there is still the possibility of me misremembering – a possibility that I could be aware of. But what would such a possibility look like? Presumably, we could only conceive of rather remote ones: Perhaps a few hours ago someone managed to put a hallucination drug in my coffee mug, so that now I am hallucinating that I just met Joe? Or perhaps a demon just controlled my mind, etc. Assuming that the actual world is just the way we think it is, these possibilities are clearly remote ones – they involve too many changes from the actuality. Fortunately, being aware of such remote error-possibilities doesn't violate *Lack-of-Awareness*, as the latter only concerns nearby ones. Thus, to the extent that in the good cases error-possibilities one could be aware of do tend to be remote, embracing *Lack-of-Awareness* does not commit one to skepticism.

Relatedly, it's worth noting that the present observation also helps to address some other problematic cases in epistemology.[21] Take Greco's (2007: 300) typing monkey case as an example: "A monkey sits down at my computer and starts banging away at the keyboard. I believe that he won't type out a perfect copy of *War and Peace*." (Vogel 1999: 165 presents a similar example.) Here, intuitively, my belief constitutes knowledge, even if I am aware of the possibility of the monkey typing out *War and Peace*. *SAA* can explain the presence of knowledge here: the possible world in which the monkey types out the book is fairly far away from the actual world.[22] After all, given the length of *War and Peace*, in the relevant error-possibility the monkey would have to type numerous keys differently compared with the actual world (assuming that in the actual world the monkey is not typing out anything close to *War and Peace*.) Thus, the error-possibility involves many relevant changes compared with the actual world, so that even if I am aware of it, my awareness does not lead to a violation of *Lack-of-Awareness*. In addition, plausibly my belief is true in most nearby possible worlds where I continue to hold the same belief. In this way, the belief is also *Safety (weak)*-safe. *SAA* thus correctly predicts knowledge in the present example.[23]

References

Adler J. (2005). 'Reliabilist Justification (or Knowledge) as a Good Truth-Ratio.' *Pacific Philosophical Quarterly* **86**, 445–58.

Baumann P. (2008). Is Knowledge Safe? *American Philosophical Quarterly* **45**, 19–30.

Baumann P. (2016). *Epistemic Contextualism: A Defense*. New York, NY: Oxford University Press.

Bogardus T. (2014). 'Knowledge Under Threat.' *Philosophy and Phenomenological Research* **86**(1), 289–313.

Broncano-Berrocal F. (2019). 'Knowledge, Safety, and Gettierized Lottery Cases: Why Mere Statistical Evidence is Not a Source of Knowledge.' *Philosophical Issues* **29**(1), 37–52.

Chisholm R. (1966). *Theory of Knowledge*. Englewood Cliffs, NJ: Prentice-Hall.

Cohen S. (1998). 'Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery.' *Australasian Journal of Philosophy* **76**(2), 289–306.

Dodd D. (2012). 'Safety, Skepticism, and Lotteries.' *Erkenntnis* **77**(1), 95–120.

[21]Thanks to an anonymous referee.

[22]Here I assume the way of interpreting "closeness relation" that I mentioned in note 13.

[23]Earlier drafts of this paper were presented at 2020 Philosophy in Progress Conference at the University of Nottingham and a workshop at Saint Louis University. Thanks to Peter Baumann, John Greco, Mark Jago, Sahar Joakim, Luke Kallberg, Francesco Praolini, Joe Salerno, Katherine Sweet, Yiling Zhou for helpful comments and discussions. I'd also like to thank an anonymous referee, whose insightful comments made me aware of a number of ways to improve the previous draft.

Goldman A. (1979). 'What is Justified Belief.' In G. Pappas (ed.), *Justification and Knowledge*, pp. 1–25. Boston, MA: D. Reidel.

Greco J. (2007). 'Worries about Pritchard's Safety.' *Synthese* **158**, 299–302.

Greco J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.

Greco J. (2016). 'Knowledge, Virtue and Safety.' In M.Á. Fernández (ed.), *Performance Epistemology*, pp. 51–61. New York, NY: Oxford University Press.

Hawthorne J. (2004). *Knowledge and Lotteries*. New York, NY: Oxford University Press.

Hirvelä J. (2019). 'Global Safety: How to Deal with Necessary Truths.' *Synthese* **196**(3), 1167–86.

Hirvelä J. (2020). 'No Safe Haven for the Virtuous.' *Episteme* **17**(1), 48–63.

Kelp C. (2009). 'Knowledge and Safety.' *Journal of Philosophical Research* **34**, 21–31.

Kvanvig J. (2008). 'Epistemic Luck by Duncan Pritchard.' *Philosophy and Phenomenological Research* **77**, 272–81.

Lackey J. (1999). 'Testimonial Knowledge and Transmission.' *Philosophical Quarterly* **49**(197), 471–90.

Lasonen-Aarnio M. (2010). 'Unreasonable Knowledge.' *Philosophical Perspectives* **24**(1), 1–21.

Lewis D. (1996). 'Elusive Knowledge.' *Australasian Journal of Philosophy* **74**(4), 549–67.

McEvoy M. (2009). 'The Lottery Puzzle and Pritchard's Safety Analysis of Knowledge.' *Journal of Philosophical Research* **34**, 7–20.

Mills E. (2012). 'Lotteries, Quasi-Lotteries, and Scepticism.' *Australasian Journal of Philosophy* **90**(2), 335–52.

Nozick R. (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.

Plantinga A. (1993). *Warrant and Proper Function*. New York, NY: Oxford University Press.

Pritchard D. (2007). 'Anti-luck Epistemology.' *Synthese* **158**(3), 277–97.

Pritchard D. (2008). 'Knowledge, Luck and Lotteries.' In V. Hendricks (ed.), *New Waves in Epistemology*. Basingstoke: Palgrave-Macmillan.

Pritchard D. (2012). 'Anti-Luck Virtue Epistemology.' *Journal of Philosophy* **109**(3), 247–79.

Pritchard D. (2015). 'Anti-luck Epistemology and the Gettier Problem.' *Philosophical Studies* **172**(1), 93–111.

Roush S. (2005). *Tracking Truth: Knowledge, Evidence, and Science*. New York, NY: Oxford University Press.

Silva P. (2019). 'Beliefless Knowing.' *Pacific Philosophical Quarterly* **100**(3), 723–46.

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

Sosa E. (2015). *Judgment & Agency*. New York, NY: Oxford University Press.

Vogel J. (1999). 'The New Relevant Alternatives Theory.' *Philosophical Perspectives* **13**, 155–80.

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

Xu Y. (Forthcoming). 'A Probabilistic Approach to Epistemic Safety from the Perspective of Ascribers.' *Episteme*. doi: 10.1017/epi.2020.7.

Zhao H. (2020*a*). 'Knowledge Without Safety.' *Synthese* **197**(8), 3261–78.

Zhao H. (2020*b*). 'Revisiting Amodal Completion and Knowledge.' *Philosophia* **48**(2), 847–56.

Zhao H. (Forthcoming). 'Better Virtuous than Safe.' *Synthese*. doi: doi.org/10.1007/s11229-019-02501-5.

Zhao H. and Baumann P. (Forthcoming). 'Inductive Knowledge and Lotteries: Could One Explain Both 'Safely'?' *Ratio*.

Haicheng Zhao works at Xiamen University, China. His research interests include epistemology (safety, sensitivity, virtue epistemology, etc.) and philosophy of mind (e.g. mental causation). He completed his PhD in philosophy at Saint Louis University in June 2020.

Cite this article: Zhao H (2023). How to Play the Lottery Safely? *Episteme* **20**, 23–38. https://doi.org/10.1017/epi.2020.44