




ARTICLE

Safety, Lotteries, and Failures of the Imagination

Anaid Ochoa 

Department of Philosophy, McGill University, Montreal, QC, Canada

Email: anaid.ochoa@mail.mcgill.ca

(Received 20 November 2023; revised 27 August 2024; accepted 17 October 2024)

Abstract

Safety accounts of knowledge intend to explain why certain true and intuitively justified beliefs fail to be knowledge in terms of such beliefs falling prey to a modal veritic type of luck. In particular, they explain why true and intuitively justified beliefs in “lottery propositions” (highly likely propositions reporting that a particular statistical outcome obtains) are not knowledge. In this paper, I argue that there is a type of case involving lottery propositions that inevitably lies beyond the scope of any reasonable safety account of epistemic luck. I offer counterexamples to accounts of epistemic luck in terms of safety conditions that involve both “locally” and “globally” reliable ways of forming beliefs in nearby worlds. All such counterexamples present a lottery case illustrating the next possibility: the process of selecting the lottery winner might be such that any world in which it delivers a different outcome is extremely far away from the actual world. In addition to being a case of safe ignorance, this type of lottery case shows that, ultimately, either veritic epistemic luck is not unsafe true belief or beliefs in lottery propositions are not epistemically luckily true.

Keywords: Lottery; safety; epistemic luck; safe ignorance

1. Introduction

A large variety of safety conditions have been offered in the literature as complete or partial accounts of the nature of knowledge. Roughly, such conditions state that a subject *S*'s belief in a proposition *p* is *safe* when *S* could not easily have falsely believed *p*.¹ In this paper, I restrict the discussion to safety conditions proposed as complete accounts of *epistemic luck*, which aim to explain why certain true and intuitively justified beliefs fail to be knowledge in terms of such beliefs falling prey to a modal veritic type of luck.² In

¹Ernest Sosa's (1999) and Duncan Pritchard's (2007, 2008, 2012, 2015) accounts of safety as a condition of knowledge are two of the prominent accounts. A third one is Timothy Williamson's (2000), under which safety is a condition of knowledge but not constitutive of it. For a general overview of safety conditions of knowledge, see Danin Rabinowitz (2011).

²Roughly, a belief is epistemically luckily true when, even if justified, it turns out true as a matter of mere accident or coincidence. Reasonably, the fact that a belief turns out true by mere accident prevents it from constituting knowledge. For a general overview of the notion of epistemic luck, see Engel Jr. (2010). The

particular, I assess how well safety conditions can be used to explain why beliefs in lottery propositions – that is, highly likely propositions which are solely supported by statistical evidence – are epistemically luckily true. The general explanation shared by safety accounts is that even though a proposition such as (ticket #34 will lose) is extremely likely given its high odds of being true (in this case, that it is .9999 probable that ticket #34 loses), such a proposition could still easily have been false and thus, whether a belief in it turns out true is simply a matter of luck.

The main rationale for treating lottery beliefs as epistemically luckily true is that there appears to be no other explanation of the fact that they fail to be knowledge despite being true, since they seem justified or rationally acceptable on the basis of purely statistical reasoning, given their extremely high probability (see Hawthorne (2004) and Pritchard (2007)). Additionally, dominant explanations in terms of lottery propositions lacking justification or rational acceptability face serious problems yet to overcome.³ Thus, an account of how lottery propositions are epistemically lucky is called for. Further, it seems natural to understand such luck in terms of safety, since there is large agreement that safety is among the best candidates to account for epistemic luck and it has garnered support by evolving to (presumably) successfully account for our failure to know lottery propositions.⁴

This paper pushes back against the previous considerations by showing that there is a type of case involving lottery propositions that inevitably lies beyond the scope of any reasonable safety account of epistemic luck. To achieve this aim, I first elaborate a counterexample to accounts of epistemic luck in terms of safety conditions that involve “locally” reliable ways of forming beliefs (roughly, ways of belief formation that produce more true beliefs in p than false beliefs in p in close worlds). The counterexample involves a lottery case, dubbed “The Intergalactic Lottery Case,” illustrating the next possibility: the process of selecting the lottery winner might be such that any world in which it delivers a different outcome is extremely far away from the actual world. After defending such counterexample from an objection, I present a variation of it against accounts of epistemic luck in terms of safety conditions that involve “globally” reliable ways of forming beliefs (roughly, ways of belief formation that produce more true than false beliefs in close worlds, not only in p but in any other proposition produced in the

main proponent of the identification of epistemic veritic luck with unsafe true belief is Duncan Pritchard (2005, 2007, 2008a and 2009).

³The main views in the literature that explain (or suggest an explanation of) why lottery beliefs are not rationally acceptable propose either what we will call a “formal” property of rationally acceptable beliefs or a “non-formal” property, where a formal property is (roughly) a property reducible to logical and probabilistic notions (for examples, see Sharon Ryan (1996), Igor Douven (2002), and Lin & Kelly (2012)). Douven and Williamson (2006) have argued that formal views are untenable by providing a proof to show that if there is a property of rationally acceptable propositions that lottery propositions lack, it cannot be a formal property (although if rational acceptability is relativized to certain formal contexts, some formal properties of rationally-acceptable beliefs so relativized could escape their criticism. For such alternative, see Hannes Leitgeb’s (2014)). Salient views that propose non-formal properties in a knowledge-first vein, such as Kelp’s (2017) virtue epistemology and Ichikawa’s (2014) knowledge-first views of justification, are non-informative with respect to why we fail to know lottery propositions. Another salient view, Douven’s pragmatic view (2012), relies on problematic claims about the individuation of the content of beliefs (see Ochoa, 2025). A promising property may be captured by Martin Smith’s (2016) normic support condition (although see Michael Blome-Tillmann (2020b)).

⁴Safety remained an exception among the various post-Gettier accounts of epistemic luck insofar as it continuously developed into versions that avoided shortcomings arisen for its earlier versions. For example, a version of safety evolved from requiring the target belief to be true in “most close worlds” into a version requiring such belief to be true *also* in “all very close” worlds (Pritchard, 2008b), driven by the need to explain lottery beliefs as unsafe beliefs.

same way). Inspired in safety conditions that involve globally reliable ways of belief formation, I propose an account of epistemic luck in terms of safety that avoids the second counterexample. I argue, however, that with enough imagination, a counterexample to such account may be found. The Intergalactic Lottery Case, in all of its proposed variants, is not only intended to be a case of safe ignorance, but it also shows that either veritic epistemic luck is not unsafe true belief or beliefs in lottery propositions are not true by mere epistemic luck.

2. Safety, epistemic luck, and lottery propositions

To cover all the safety conditions on offer in the literature without examining their various details and peculiarities, we can follow Di Yang (2019) in identifying two types of sub-conditions of any safety condition:

- i) a “modal distance” condition, capturing a degree of a relevant type of dissimilarity⁵ among the actual world⁶ and other possible worlds, in terms of which a special class C of close worlds is defined, and
- ii) a “modal frequency” condition, capturing the frequency in which a given event e (proposition, fact) obtains in the C-worlds.⁷

A lucky event is then conceived by combining (i) and (ii) as follows:

An event e is lucky if, and only if, e obtains but the frequency of C-worlds in which e obtains is not sufficiently high.

According to this modal conception of luck, beliefs that are epistemically luckily true are dealt with by treating them as unsafe, that is:

S’s belief in p is epistemically luckily true if, and only if, S’s belief in p is true but *unsafe*, that is: S’s belief in p is true but the frequency of C-worlds in which a belief

⁵The relevant type of similarity in safety theories is the relation of overall similarity (similarity considering all respects, rather than similarity in one respect). An exception among safety accounts is Blome-Tillmann’s (2020a) “non-reductive safety,” which takes modal distance to capture a degree of *epistemic* similarity, where similarity is epistemically restricted to similarity “in those respects that are relevant for knowledge.” It is worth exploring whether non-reductive safety can account for why we fail to know the modally stable lottery propositions presented in this paper as part of the counterexamples to safety. However, to cover most safety theories, I will here restrict my analysis to those involving the notion of overall similarity. This could be problematic for the counterexamples here presented if safety conditions are better cashed out with a more restricted notion of similarity rendering such counterexamples ineffective (while counting intuitive cases of knowledge as safe beliefs). Such a restricted notion of similarity, however, would need to be specified before it poses a real problem to the proposed counterexamples. Thanks to Michael Blome-Tillmann and Ulf Hlobil for discussion on this point.

⁶Throughout the paper, I use “the actual world” to designate the world of evaluation of the relevant safety-claims corresponding to a given scenario (even if the scenario is counterfactual, such as the Intergalactic Lottery Case), rather than our actual world. This sense of “actuality” can be roughly understood in the following way: a world w is the actual world in a context c if, and only if, it is being evaluated in c whether a given modal claim (e.g., a safety claim) is true in w .

⁷These sub-conditions are not independent from each other, since the special class C in (ii) is a function of there being a determinate threshold for the dissimilarity relation in (i). Also notice that the class C is defined by the dissimilarity relation *with respect to the actual world*, so C might be a different class when different worlds are taken as the actual world, that is, as the world of evaluation.

of S in p is formed in the same way as in the actual world and p is true is not sufficiently high.

The formulation of this principle can be made more precise in multiple respects, e.g., by specifying what class C is (the class of *nearby* worlds, of *very nearby* worlds), by specifying what counts as a “way” in which a belief is formed (e.g., a belief-forming method, a belief-forming process, a basis for believing), or by specifying how high the frequency of C -worlds in which a believed proposition is true should be in order for such proposition to count as safe (e.g., it must be true in all C -worlds, most C -worlds, at least one C -world). The lack of specificity of this principle is useful to cover a large range of safety conditions differing in those respects. Acceptable specifications are constrained by whether the resulting version of the principle covers all and only cases of epistemic luck. Those specifications that count very far off worlds from the actual one as C -worlds, and tolerate only a small number of C -worlds in which a belief is false for it to count as safe, risk counting non-luckily true beliefs as unsafe. For example, if a belief is taken to be safe iff it is true in *all* worlds similar to the actual one in *some* respect, most ordinary beliefs would be unsafe, since they are false in some sceptical world which at least share some features with the actual world. Those specifications that only admit very close worlds to the actual one as C -worlds and tolerate a large number of C -worlds in which a belief is false for it to count as safe, risk counting luckily true beliefs as safe (since very close worlds might fail to differ enough from the actual world to make a luckily true belief false). For example, if a belief is taken to be safe iff it is true in some world similar to the actual one in all respects, any true belief would count as safe, including the luckily true ones. In particular, a specification identifying safety with truth in a large number of close worlds clashes with the verdict that beliefs in lottery propositions are unsafe, since they are typically true in a large class of close worlds (given the low chances for lottery propositions to be false, they would be true in the great majority of close worlds). For that reason, throughout this paper I assume that the correct frequency for a belief in a lottery proposition to be unsafe is for such proposition to be false in at least one C -world in which a belief in it is formed in the same way as in the actual world, that is:

2.1. Safety

S 's belief in p is epistemically luckily true if, and only if, S 's belief in p is true but *unsafe*, that is: S 's belief in p is true but there is some C -world in which a belief of S in p is formed in the same way as in the actual world, and p is false.

In this paper, I take safety as adequately capturing most safety conditions intended to explain lottery cases as involving epistemic luck. Safety, however, is still inadequate in that it only encompasses safety conditions whose application to a belief in a given proposition depends on the truth value (with respect to the C -worlds) of that very same proposition. These safety conditions are said to involve a *local* type of reliability⁸ because if a belief in a proposition p is safe in the sense defined by them, the way in which such belief was formed is supposed to “track” (throughout the C -worlds) the truth of p but not the truth of any other proposition (even if such a proposition is the object of a belief formed in the same way). This contrasts with safety conditions involving a *global* type of reliability, where if a belief in a proposition p is safe in the sense defined by them, the way

⁸See, for example, Dani Rabinowitz' The Safety Condition for Knowledge (2021).

in which the belief was formed is supposed to “track” (with respect to C-worlds) the truth of *any* proposition that is the object of a belief formed in the same way. This gives rise to the next variation of Safety:

2.2. Global safety

S’s belief in *p* is epistemically luckily true if, and only if, S’s belief in *p* is *globally unsafe*, that is: S’s belief in *p* is true but there is a C-world in which a belief of S in a proposition *q* is formed in the same way in which S’s belief in *p* was formed in the actual world, and *q* is false (where *q* might be different a proposition than *p*).

For the moment, let us leave Global Safety aside. Whichever is the best way to formulate Safety, its application to deal with lottery propositions is very simple. No matter how unlikely it is that a lottery proposition is false, only a small difference in the actual world is required for that to happen. Just a small difference in the lottery winner selection process – e.g., a couple of different numbered balls coming out from a tombola – is needed to yield a different lottery result. Given this, for every lottery proposition, there is a C-world in which a false belief in it is formed in the same way as in the actual world, so the belief is unsafe. Thus, Safety provides a simple explanation of why beliefs in true lottery propositions are epistemically luckily true.

Only two things are needed for Safety to account for the case of lottery propositions: that the modal distance condition does not track the likelihood of the target proposition (so it is not true that the more likely a proposition is, the closer a world in which it is true is to the actual world, and vice versa), and that the differences needed from the actual world to make the target proposition false are small enough to stay below the threshold of modal distance for being a C-world. Unfortunately, the simplicity with which Safety explains lottery cases turns against it, or so I argue next.

3. The Intergalactic Lottery Case: a counterexample to safety

Given the simplicity with which Safety deals with beliefs in lottery propositions, the only thing needed to show that it fails is to provide a lottery case in which the process of selecting the winner of the lottery delivers different results from the actual one only in *very far away* possible worlds (given a reasonable specification of the modal distance condition), while maintaining the intuition that the corresponding lottery proposition is not known. Here I provide one such case, having the following structural features:

- A) The process by means of which the winning lottery ticket is selected is so spatially large, causally complex, and having a long duration, that a different result could not easily have obtained. Because of the nature of this process, the more a particular result deviates from the actual result, the harder it is that it could have obtained.
- B) If (A), then there is no C-world in which a different winning ticket is selected, so any true belief in a relevant lottery proposition is safe (and, therefore, not epistemically luckily true, according to Safety). This is especially true with respect to results differing the most from the actual result.

A scenario satisfying both conditions (A) and (B), while maintaining the intuition that the relevant lottery proposition is not known, constitutes a counterexample to

Safety, since the conjunction of both conditions is incompatible with such principle. One such scenario is provided by the following story:

Intergalactic Lottery Case: Every five thousand years, the Intergalactic Federation of Planets (IFP) runs an Intergalactic Lottery in which one (and only one) citizen of the IFP is prized with the unique opportunity to live a very long and happy life on Planet Paradise IFP (excluding those citizens already living in Planet Paradise). The lottery is run this way because it takes five thousand years for an inhabitant of Planet Paradise to peacefully die, and only then a vacancy for a new inhabitant opens.

The process for selecting the winning ticket is as follows. First, the number of each lottery ticket has as many digits as there are planets in the IFP, and given that the number of planets in the IFP is astronomically large, the number of digits in the number of each lottery ticket is astronomically large as well. Second, each planet of the IFP gets to select one digit of the number of the winning ticket by means of some local subprocess with the following characteristics:

- a) it is spatially very large, relative to the size of the planet in which it occurs,
- b) it is causally very complex, and
- c) it has a very long duration (nearly five thousand years).

Given that each of the subprocesses is gigantic at a planetary scale and there is an astronomical number of them, the process for selecting the number of the winning ticket is astronomically larger than each of its gigantic planetary-scale subprocesses.

To illustrate the nature of the subprocesses, consider the following instance (let us keep in mind that, as previously stipulated, any digit-selecting process is comparable in scale, causal complexity and duration to the following one): Planet Fungus creates multiple gigantic enclosed platforms (each one of 16,000 km²) in which fungus is freely allowed to grow at a very slow rate, with the expectation that a gigantic digit-like pattern naturally forms in one of them. If a platform is filled with fungus and no digit-like pattern forms, the process is restarted in that platform. The first recognizable digit-like pattern to form is the selected digit. Once the digit is obtained, it is securely registered, codified, and sent to the organizers of the Intergalactic Lottery.

In this fanciful scenario, an Earth citizen of the IFP, Annie, reasons in the following usual way: “The Intergalactic Lottery has an astronomical number of tickets, so the chances of my ticket being the winner are close to zero. Ergo, if I can be sure of something, it’s that my ticket is not the winner.” As the result of this reasoning, Annie forms the belief that her ticket is not the winner. And Annie is right, her ticket is not the winner. Thus, Annie ends up forming a true belief in a lottery proposition.

Intuitively, Annie is justified in thinking that her lottery ticket is not the winner (on the basis of purely statistical reasoning), and the proposition believed by her is true, yet it is also intuitive that she does not know such proposition. With respect to these intuitions, Annie’s belief is no different from any other typical belief in a lottery proposition (except

for the lottery proposition that is about the actual winning ticket). However, besides its science-fictional features, the Intergalactic Lottery stands apart from normal lotteries in a crucial aspect: the *astronomical scale and complexity* of the winner-selecting process. To understand why this is crucial, let us momentarily restrict our attention to the digit-selection subprocesses. Given their nature, the counterfactual variations in the outcome of *any* of them need a significant dissimilarity from the actual world, significantly larger than is required for counterfactual variations in the outcome of any normal digit-selection process (e.g., one of the numbered balls coming out from a tombola is a 6 rather than a 2). Thus, for a *single* digit of the number of the winning ticket to have been different, a significant degree of dissimilarity from the actual world is necessary – e.g., significantly different causal chains of fungus growth during a very long period of time, forming a very different gigantic digit-like pattern than the one produced in the actual world. Thus, we obtain the following intuitive result:

3.1. Distant difference

The number of the Intergalactic Lottery's winning ticket could only have differed in one of its digits in far-off worlds

Let us say that p is an intergalactic proposition if, and only if, p truly reports that a particular ticket of the Intergalactic Lottery is a loser. Accepting Distant Difference entails that, for every intergalactic proposition p , p is only false in a far-off world, so (contrary to Safety) every typical belief in p is safe. Yet, it is intuitive that Annie does not know that her ticket is a loser, despite her belief being true, justified and (according to Distant Difference and Safety) safe. Thus, we have a class of counterexamples to Safety.

It might be argued that Distant Difference is not a decisive threat to Safety. For all we have shown, it might be argued, Distant Difference is compatible with there being a safety condition with a modal distance threshold that counts worlds in which intergalactic propositions are false as C-worlds, according to such safety condition. This is not a rebuttal of the counterexample unless the safety condition in question is specified. However, in the absence of an argument showing that no such safety condition exists, its mere possibility renders the counterexample inconclusive.

Granting for the moment the previous response (there are good reasons to think that it should not be granted, which I discuss shortly), there still are intergalactic propositions for which it does not work. Consider the number of a ticket, k , that differs from the number of the actual winner in *all its digits*. And consider an intergalactic proposition p_k reporting that the ticket $\#k$ is not the winner. A world in which p_k is false (that is, in which ticket $\#k$ is the winner) is *astronomically far-off* from the actual world. This is because i) if a number differs in all of its digits from the number of the actual winner, the outcomes of *all the subprocesses* of the Intergalactic Lottery would have to be different for that ticket to be the winner and ii) a world that differs from the actual world in the outcomes of all such subprocesses is *astronomically far-off* from it (since it involves an astronomical number of planetary-scale differences from the actual world). No specification of the class of C-worlds seems able to count this kind of (astronomical) difference among worlds as below the threshold of a reasonable modal distance condition for Safety, since we expect of such condition to count significantly smaller differences among worlds as above that threshold. To see this, consider the next three variations of the Intergalactic Lottery Case:

S₁) Annie has the Intergalactic Lottery ticket with # k . Given the astronomical number of digits of each number of an Intergalactic Lottery ticket, people need to rely on an extremely powerful and reliable software designed by the IFP (used in extremely powerful and reliable computers) to check and compare the numbers of lottery tickets. Annie forms the belief in the proposition, p , that her ticket is not the winner (expressed to herself by the sentence “My ticket is not the winner”) in the way specified above.⁹ Based on that belief, she infers that the ticket with *that* number (that is, k) is not the winner, whichever is the number of her ticket. Assuming Annie’s inference allowed her to form a belief, Annie ends up believing that the ticket # k is not the winner, that is: she ends up believing p_k (which, as we have seen, is true because k differs from the number of the winning ticket in each one of its digits).

Based on her previous beliefs, Annie decides to run a scam. Being a very skillful programmer, she creates multiple copies of her lottery ticket that are indistinguishable from the original for the ticket-reading software, although the copies would easily be revealed as fake if they were examined by agents of the IFP with a more sophisticated software. Her plan is to sell each copy to a number of her acquaintances who do not know each other and who live far away from one another. Annie thinks that the scam is likely to go unnoticed if her ticket does not win, so the risk of being discovered is insignificant. After making all the copies for the scam, Annie puts them in a box, together with a list of the people she intends to scam, inside a secret drawer under the desk in her office, and the real lottery ticket next to that box. Annie ends up justifiably and truly believing the proposition, q , that her ticket is next to the box.

S₂) Everything is like S₁, except for the following: Annie’s boyfriend, Philipp, discovers Annie’s secret drawer by accident, and all its contents. Philipp discovers that he is one of the people that Annie intends to scam. Believing that she is going to win the lottery and to teach her a lesson, Philipp replaces the original ticket with one of the copies and leaves it next to the box, so q ends up being false.

S₃) Everything is like in S₁, except for all the facts about the winning ticket selection process that need to be different for p and p_k to be false.

Intuitively, the belief in q is knowledge in S₁, while it is not knowledge in S₂ (given that q is false in S₂). Also, intuitively, the beliefs in p and p_k are not knowledge in any of the previous scenarios, even though they are true and (presumably) justified in S₁ and S₂. Let us compare, however, the modal distance between all such scenarios to assess whether these intuitions can be captured by a safety condition.

⁹I am assuming that Annie formed her belief in p , and her belief in p_k , after the result of each subprocess was settled (maybe by years, decades or centuries), so if the counterfactual capturing the relevant safety condition requires that everything before the target belief is formed is kept fixed throughout the possible scenarios required to evaluate such counterfactual, that by itself guarantees that there is no nearby world in which Annie formed the belief in p (or in p_k) and her ticket was the winner of the Intergalactic Lottery.

It is also interesting to notice that if a Millian view on indexical expressions such as ‘my ticket’ (that is, the claim that the only semantic content of those expressions is their designatum) is a correct view, ‘My ticket is not the winner’ expresses the same proposition in the present context as ‘Ticket # k is not the winner’, regardless of whether Annie is aware that ‘my ticket’ and ‘ticket # k ’ are co-designative terms in such context. In that case, p and p_k are the very same proposition, so the inference form p to p_k in S₁ is unnecessary.

Let d_1 be the degree of modal distance between S_1 and S_2 (that is, the degree of dissimilarity between S_1 and S_2), and let d_2 be the degree of modal distance between S_1 and S_3 .¹⁰ Now, let m be the modal distance threshold of an arbitrary safety condition s , that is: a world w belongs to the class C of worlds in terms of which s is cashed out iff the degree of modal distance between w and the actual world (in the present case, S_1) is less than m . To obtain the expected intuitive results, d_1 should be greater than m (so S_2 does not belong to C when S_1 is taken as the actual world, and the belief in q is safe in S_1). Now, whichever degree of modal distance d_2 is, it is clear that d_2 is greater than d_1 (astronomically so!), so d_2 should also be greater than m . Therefore, a reasonable modal threshold allowing for the belief in q to count as knowledge in S_1 would count the belief in p and the belief in p_k as safe in S_1 as well.

The previous result is bad enough, since it shows that Safety cannot deal correctly with beliefs in *some* intergalactic propositions, that is, those that involve tickets whose number differs from the winning one in a large number of their digits. But we can show that the problem generalizes to the rest of the intergalactic propositions. This is because the response we temporarily granted with respect to counterfactually different outcomes of isolated subprocesses is not very good either, since we can think of variations of the Intergalactic Lottery Case in which each subprocess has a significantly larger scale than the one specified in the original scenario, and there seems to be no limit in how much that scale can be increased to make the subprocesses large enough to have different outcomes only in far-off worlds, according to any reasonable modal distance threshold.

4. Does the Intergalactic Lottery Case work against global safety?

So far, we have seen that the Intergalactic Lottery Case functions as a counterexample to Safety. Global Safety, however, seems to provide an easy way to block it. According to such principle, a belief an agent has is (globally) safe iff some false belief could be formed by the same agent, in the same way, in some C -world. Let $n_{@}$ be the actual winning ticket's number. Consider now a world w in which the number of Annie's ticket is $n_{@}$ instead of k , and everything else in w is the same as in the actual world. In w , Annie forms a belief in the same way in which she actually formed her belief in the target proposition (that is, that her ticket is not a winner), but the lottery proposition she believes in w is different from the target proposition (the former proposition is true in a given world iff $n_{@}$ is not the winning ticket's number in such world, while the target proposition is true in a given world iff k is not the winning ticket's number in that world).¹¹

Global Safety theorists can now proceed as follows: Intuitively, w is very close to the actual world, since there are only slight differences needed for Annie's ticket to have a different number (that is, $n_{@}$) than the number of her actual ticket (that is, k). If so, there is a very close world – plausibly, belonging to the relevant class of C -worlds – in which Annie forms a false belief in the same way in which she actually formed her belief in the target proposition that her ticket (that is, ticket $\#k$) is not the winner of the lottery.

¹⁰As indicated in note 4, I am assuming Safety is cashed out in terms of the notion of overall similarity.

¹¹The claim that both propositions are different is based on the assumptions a) that the demonstrative term 'my ticket' is a rigid designator and b) that the definite description 'the ticket $\#k$ ' is not a rigid designator. In my view, (a) is sufficiently supported by the pre-theoretical intuition about the rigidity of demonstratives, while (b) is supported by the intuition that a ticket (understood as the concrete object playing a ticket-role in a lottery) has the number it has contingently. This second assumption might be wrong, if a ticket is an abstract object that is essentially tied to its number. In that case, the scenario proposed is not a possible one.

Consider now a specification of Global Safety requiring that there be no C-worlds in which S forms a false belief in the same way in which S actually formed her belief in the target proposition in order for the latter belief to be safe. Given such specification of Global Safety, the target belief is unsafe (thus, being epistemically luckily true), allowing to block the counterexamples proposed so far.

Fortunately, the Intergalactic Lottery Case can be modified to obtain a counterexample to Global Safety. To do so, let us begin by noticing that the process of selecting a winning ticket is, in the scenario of the Intergalactic Lottery Case, a *modally stable* process, that is, a process that could not have easily had a different outcome than the actual one. It is because such a process is modally stable that the target proposition, p , could not have easily been false. In contrast, the correlation between the next elements is not modally stable: i) the way in which S's belief in p was actually formed and ii) that the belief of S has p as its object (it is not modally stable since S could have easily formed a belief in a different proposition in the same way in which S actually formed the belief in p). Once we notice this contrast, it is easy to see that all we need to do to obtain a counterexample to Global Safety is to add further details to the description of the way in which S's belief that p is actually formed, in such a way that the correlation between (i) and (ii) is modally stable. We can do this in as follows:

Consider the Intergalactic Lottery case depicted earlier. Suppose additionally that hundred of thousands of years ago when the Intergalactic Lottery was first instituted, a flawless technology was created to analyze the genetic code of the citizens of the IFP existing at that time, to obtain (by means of a secret, unchanging, and nonrandom algorithm) a numeral with an astronomical number of digits serving as the lottery ticket of each citizen (thus, each lottery ticket was essentially its number rather than a concrete object). All the tickets thus generated were significantly different from one another (even if the genetic codes of the ancient citizens did not vary significantly, the algorithm nonetheless generated a radically different ticket for each one). After a citizen died, her ticket was inherited to her direct descendants, if she had any (if she had many, her ticket was inherited to one of them, and the algorithm provided her other descendants a significantly different ticket). This hereditary practice continued for thousands of years up to the present. Given this process for generating tickets, Annie could not have easily had a significantly different ticket. Since the winning ticket could not easily have been different, and since Annie's ticket varied significantly from the winning ticket, Annie could not have easily formed a belief in a false proposition in the same way in which she actually formed her belief that her ticket is a loser.

With these modifications, the Intergalactic Lottery case now functions as a counterexample to both Safety and Global Safety. However, the next revision of Global Safety seems to be sufficient for blocking the present counterexample:

4.1. *Global safety

S's belief in p is epistemically luckily true if, and only if, S's belief in p is **globally unsafe*, that is: there is some C-world in which a belief of an agent S^* (possibly different from S) in a proposition q is formed in the same way in which S's belief in p was formed in the actual world, and q is false.

The rationale for using *Global Safety to deal with the previous version of the Intergalactic Lottery Case is captured in the next line of reasoning:

Any epistemic agent in the Intergalactic Lottery Case could have formed a belief in the lottery proposition about her own ticket in the same way in which Annie formed her belief in the target proposition. This applies to the citizen of the IFP, named Zhee-Vra, who actually has the winning ticket. Had Zhee-Vra formed a belief in the lottery proposition that her ticket is not the winner in the same way Annie's belief in the target proposition was actually formed (all else remaining as in the actual world), she would have ended up forming a false belief. Plausibly, Zhee-Vra could have easily formed her belief in her corresponding lottery proposition in the same way Annie actually formed hers. Thus, some agent (that is, Zhee-Vra!) could have easily formed a false belief in the same way Annie actually formed the target belief. Therefore, the target belief is unsafe according to *Global Safety.

This reasoning makes plausible that *Global Safety can block the version of the Intergalactic Lottery Case presented in this section. Nevertheless, that same reasoning indicates a way to further modify the scenario in question so that it also works as a counterexample to *Global Safety. We only need to suppose that Zhee-Vra does not have the ability to carry out the same reasoning that gave rise to Annie's belief in the target proposition (maybe Zhee-Vra is an infant, lacks the proper conceptual skills to perform the reasoning in question, or lacks the linguistic competence to make self-reference by using a demonstrative term), and that such deficiency is modally stable (perhaps, it is due to some genetic deficiency recurring in all of Zhee-Vra's ancestors). Whatever detail is chosen here to produce the desired version of the Intergalactic Lottery Case, it seems plausible that such details can be worked out into a counterexample to *Global Safety.

Is there any amendment to *Global Safety that avoids the counterexample? Suppose that some such amendment were obtained. Plausibly, it would come from identifying some modal instability in either:

- a) the state of affairs making the target proposition true, or
- b) the correlation between the way in which the agent's belief in such proposition is formed and other possible beliefs that such agent could have formed in the same way.

However, with sufficient imagination, it always seems possible to make further stipulations to the scenario to increase the modal stability of both (a) and (b), obtaining a counterexample for the amended version of *Global Safety.

Finally, it might be objected that the Intergalactic Lottery Case in its different variations is just too far-fetched and designed too intricately to merit serious consideration. But while it involves plenty of detailed science-fictional content, it is conceptually possible, and it does not require something bizarre or outrageous about the cognitive life of the epistemic agents in it. Quite the opposite. From an epistemological point of view, the Intergalactic Lottery is just like any normal lottery, except for the astronomical number of tickets and the astronomical number of their digits. So, despite its fancifulness and intricate design, the Intergalactic Lottery Case shows something important about the conceptual relation (or lack of it) between the epistemic status of

lottery propositions and their condition of being safe or unsafe: we can be ignorant of modally-stable lottery propositions.¹²

Moreover, the intuition that an epistemic agent like Annie does not know that her ticket is a loser remains strong despite the detailed and far-fetched fictional content of the scenario. Given that Annie's belief is true and (intuitively) justified, the only reason left for why it fails to be knowledge is that it is epistemically luckily true. In that regard, the counterexample might only seem to be a case of safe ignorance. However, it shows the next two claims to be incompatible: i) epistemically luckily true beliefs are unsafe beliefs (that is, safety or one of its variants is true) and ii) (typical) beliefs in lottery propositions are epistemically luckily true. If beliefs in lottery propositions fail to be knowledge for a reason other than being epistemically luckily true (for example, for being unjustified), the Intergalactic Lottery Case is no counterexample to Safety. If, on the other hand, Safety is false, the Intergalactic Lottery Case is compatible with beliefs in lottery propositions failing to be knowledge because they are epistemically luckily true. The main moral to be drawn from the above cases is that conceptually tying epistemic notions to non-epistemic modal notions, such as safety, is not a safe theoretical move to make.¹³

Competing interests. The author declares none.

References

- Blome-Tillmann M.** (2020a). 'Non-Reductive Safety.' *Belgrade Philosophical Annual* **33**(33), 25–38.
- Blome-Tillmann M.** (2020b). 'Statistical Evidence, Normalcy, and the Gatecrasher Paradox.' *Mind* **129**(514), 563–578.
- Douven I.** (2002). 'A New Solution to the Paradoxes of Rational Acceptability.' *British Journal for the Philosophy of Science* **53**(3), 391–410.
- Douven I. and Williamson T.** (2006). 'Generalizing the Lottery Paradox.' *The British Journal for the Philosophy of Science* **57**(4), 755–779.
- Douven I.** (2012). 'The Lottery Paradox and the Pragmatics of Belief.' *Dialectica* **66**(3), 351–373.
- Engel Jr M.** (2010). 'Epistemic Luck.' In J. Dancy, E. Sosa and M. Steup (eds), *A Companion to Epistemology, Second Edition*, pp. 336–340. Hoboken, NJ: Blackwell.
- Hawthorne J.** (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Ichikawa J.J.** (2014). 'Justification is Potential Knowledge.' *Canadian Journal of Philosophy* **44**(2), 184–206.
- Kelly K.T. and Lin H.** (2012). 'A Geo-logical Solution to the Lottery Paradox, with Applications to Nonmonotonic Logic.' *Synthese* **186**(2), 531–575.
- Kelp C.** (2017). 'Lotteries and Justification.' *Synthese* **194**(4), 1233–1244.
- Leitgeb H.** (2014). 'The Stability Theory of Belief.' *Philosophical Review* **123**(2), 131–171.
- Ochoa A.** (2025). *Lotteries and the Roads to Knowledge Failure*. Montreal, QC: McGill University.
- Pritchard D.** (2005). *Epistemic Luck*. Oxford: Clarendon Press.
- Pritchard D.** (2007). 'Anti-luck Epistemology.' *Synthese* **158**(3), 277–297.
- Pritchard D.** (2008a). 'Sensitivity, Safety, and Anti-Luck Epistemology.' In *The Oxford Handbook of Skepticism*, pp. 437–455. Oxford: Oxford University Press.
- Pritchard D.** (2008b). '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), 247–279.
- Pritchard D.** (2015). 'Anti-luck Epistemology and the Gettier Problem.' *Philosophical Studies* **172**(1), 93–111.

¹²Special thanks to an anonymous referee for a version of this objection.

¹³Thanks to the audiences at McGill University and the participants of the Epistemology seminar at Universidad Nacional Autónoma de México (UNAM) for their feedback. Special thanks to Michael Blome-Tillmann, Melahuac Hernández, Peter Klein, Stephanie Leary, and Lourdes Valdivia for their helpful discussion on earlier drafts.

- Rabinowitz D.D.** (2011). 'The Safety Condition for Knowledge.' *Internet Encyclopedia of Philosophy*.
- Ryan S.** (1996). 'The Epistemic Virtues of Consistency.' *Synthese* **109**(2), 121–141.
- Smith M.** (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford: Oxford University Press.
- Sosa E.** (1999). 'How to Defeat Opposition to Moore.' *Philosophical Perspectives* **13**, 137–149.
- Williamson T.** (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Yang D.** (2019). 'What's Wrong with Modal Conceptions of Luck and Risk.' *Erkenntnis* **86**(4), 773–787.

Anaid Ochoa recently earned her PhD from McGill University (2025) and received her MA from Universidad Nacional Autónoma de México. Her primary areas of research are traditional and social epistemology, as well as their intersections with metaethics. She also discusses issues in metaphysics, philosophy of mind, and moral psychology.

Anaid's current research focuses on the epistemic status of lottery propositions and how lotteries illuminate discussions in both applied epistemology (especially those centered on legal standards of proof) and more traditional epistemology, particularly with respect to the notions of good/bad reason, justification, ignorance, and knowledge. Anaid is also increasingly interested in discussions surrounding epistemic attributions to members of oppressed social groups.