Canadian Journal of Philosophy (2025), 55: 4, 387–404
doi:10.1017/can.2026.10059

ARTICLE

Solving the Lottery Problem: From Modal Accounts to Explanationism

Haicheng Zhao

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

Abstract

Proponents of modal knowledge accounts (safety and sensitivity) concur that one crucial advantage of their accounts is that they solve the so-called lottery problem—the problem of explaining why "lottery beliefs" based merely on statistical evidence do not constitute knowledge. Contra this claim, I argue that epistemic judgments about lottery beliefs do not consistently track what occurs in a specified set of nonactual possibilities. Thus, modal knowledge accounts cannot properly explain beliefs based merely on statistical evidence. Finally, I argue that these beliefs can be better accommodated by a rival theory of modal accounts —namely, explanationism.

Keywords: sensitivity; safety; lottery; statistical evidence; explanationism

1. Background

Modalists in epistemology claim that for one to know that *p*, one's doxastic state regarding *p* must conform to the truth value of *p* in a suitable range of possibilities. Two prominent modal knowledge accounts are safety and sensitivity:

> *Safety*: If S knows that *p* via method M, then S's belief that *p* is true in all close possible worlds in which S continues to believe that *p* via M.

> *Sensitivity*: If S knows that *p* via method M, then in the closest possible worlds in which *p* is false, S does not believe that *p* via M.

Many have argued that these modal accounts hold promise in solving various perennial epistemological problems. Thus, these accounts have gained much popularity in the past decades.[1] Recently, however, the exalted status of modalism has been challenged by *explanationism* (e.g., Bogardus & Perrin, 2022, 2025; Faraci, 2019; Mills, 2012). Explanationists abandon the claim that knowledge essentially requires a modal connection between belief and truth. Rather, they maintain that knowledge demands an explanatory relation between truth and belief in the actual world. In its simplest form, the view is that knowledge is believing that *p* because *p* is true.[2]

[1]For defenses or developments of *Safety*, see, e.g., Beddor and Pavese (2020), Dutant (2016), Greco (2012), Hirvelä (2019), Manley (2007), Pritchard (2007, 2008, 2012), Sainsbury (1997), Sosa (1999), Williamson (2000), Zhao (2024). For defenses or developments of *Sensitivity*, see, e.g., Becker (2007, 2012), Black (2008), Cross (2007), DeRose (1995, 2010), Nozick (1981), Roush (2005).

[2]Earlier advocates of explanationism include Goldman (1984, 1988, 2008) and Jenkins (2006). More recently, the account has been defended by Mills (2012), Faraci (2019), Bogardus and Perrin (2022, 2025).

© The Author(s), 2026. Published by Cambridge University Press on behalf of The Canadian Journal of Philosophy, Inc. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (http://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

In the literature, both modalists and explanationists have maintained that their accounts can solve the so-called lottery problem (Becker, 2012; Bogardus & Perrin, 2022; Cross, 2007; DeRose, 2010; Goldman, 1988; Mills, 2012; Pritchard, 2008; Vogel, 2012). Indeed, sometimes being able to address this problem is regarded as a key advantage of these accounts (cf. Blome-Tillmann, 2017; Mills, 2012; Pritchard, 2008). Generally speaking, this is the problem of explaining why true beliefs based merely on strong statistical evidence do not amount to knowledge.[3] Consider a stereotypical lottery case. Lottie buys a fair lottery ticket with long odds for winning. She has not checked the result yet. However, by reflecting on the extremely low chance of winning, she believes that her ticket is not the winner. Intuitively, Lottie does not know that her ticket does not win, even though probabilistically speaking, her evidence strongly supports her belief.

At first glance, both modalists and explanationists are capable of explaining Lottie's lack of knowledge. Starting with the former, proponents of *Sensitivity* may argue that Lottie does not know because her belief is not sensitive: in the closest worlds where the ticket is the winner, she still falsely believes that it is not the winner based on the same statistical evidence. According to *Safety*, Lottie does not know because her belief could easily have been false. Specifically, although it is very unlikely that she wins, there are close possible worlds where she wins. That is, not much needs to be changed for her to win.[4]

On the other hand, explanationists may point out that Lottie does not know because of the absence of an explanatory link between her belief and the fact that she is not the winner. After all, Lottie has not checked the result of the lottery, so her belief is not explained by the fact that she does not win. Rather, it is explained by her statistical evidence regarding the low chance of winning.

The overarching thesis of this article is that explanationism is better suited than modal accounts in accommodating beliefs based on mere statistical evidence. Here is a roadmap. In the next section, I examine different versions of sensitivity accounts and argue that they fail to either solve the lottery problem or incur significant costs. Sections 3–5 focus on the safety account. Section 3 examines extant objections to the safety-based solutions to the lottery problem and finds them wanting. In light of this, Sections 4 and 5 present new objections, targeting both the traditional safety conditions and the more recent normality-based safety. In Sections 6 and 7, I discuss explanationism vis-à-vis beliefs based on bare statistical evidence and argue that in this regard it fares better than modal accounts.

As a terminological clarification, I use *statistical evidence* throughout this article to refer to evidence that provides probabilistic support for a proposition—a proposition that involves a member (e.g., my lottery ticket) or members that belong to a reference class (e.g., all the tickets in the lottery). The evidential support for the proposition derives solely from general frequencies or odds within that class (e.g., each ticket has a 1 in 100000 chance of winning). Statistical evidence in this sense can justify a high credence but not an extreme credence (1 or 0) (cf. Silva, 2023, 2642).

The question of what evidence consists in is the subject of much controversy. Some argue that false propositions can constitute evidence; others deny this and claim that evidence consists only in true propositions—or even in true propositions one knows (cf. Williamson, 2000). In developing the main line of argument in this article, I will follow the majority of authors writing on statistical evidence in assuming that such evidence may include false propositions. Nevertheless, where relevant, I will also consider how alternative, factive conceptions of evidence would bear on my argument. It turns out that the overall conclusion of this article does not depend on which conception one endorses. Whether evidence is taken to be factive or non-factive, modal accounts

[3]Another epistemological question concerning statistical evidence is whether beliefs based merely on such evidence are justified or rational. See Günther (2024b) for recent discussions of this issue.

[4]This view is defended by Duncan Pritchard, who argues that worlds where one wins the lottery are very close (or similar) to the actual world in the sense that in these worlds "all that need be different is that a few coloured balls fall in a slightly different configuration" (Pritchard, 2007, 292).

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

such as safety and sensitivity fail to accommodate beliefs based solely on statistical evidence, whereas explanationism succeeds.

2. Against Sensitivity

Without further ado, here is a lottery case that counts against *Sensitivity*:

> Kattie buys a lottery ticket with long odds for winning. The game is designed as follows. The tickets are in a lottery vending machine. Each player puts some coins in the machine and gets a random ticket. A series of numbers is printed on each ticket. If one's ticket is the winner, one will immediately know she wins, because the machine delivers the reward money along with the winning ticket. However, if the ticket is not the winner, players will have to wait until the announcement day to be notified if their numbers win. No players know about this. The mechanism is designed to ensure that winners enjoy their pleasure as early as possible and to let losers extend their expectation for winning as long as possible.

Suppose that shortly after Kattie gets her ticket from the machine, she correctly believes that she does not win, based merely on the long odds for winning. Intuitively, just as Lottie, Kattie does not know that her ticket is not the winner. The difference in the winning mechanism does not render Kattie a knower, to the extent that she is not aware of the mechanism at all.

However, *Sensitivity* cannot explain why Kattie does not know. In closest worlds where Kattie's ticket is the winner, she does not believe that her ticket is not the winner based on the same statistical grounds. Rather, in these worlds she gets the reward money and believes that she is the winner. Thus, according to *Sensitivity*, her belief is sensitive.

Compare Lottie's belief with Kattie's. It is not as if the latter enjoys any better epistemic status than the former. Both beliefs are based merely on (as we may assume) the same statistical evidence and both fail to be knowledge. *Sensitivity*, however, discriminates these cases and implausibly counts Kattie's belief as sensitive. Thus, it does not succeed in explaining the ignorance in lottery cases. In particular, whether or not one still holds the same belief in closest worlds where the belief is false is inessential to why lottery beliefs are not knowledge. The explanation must lie somewhere else.[5]

Now, it should be noted that *Sensitivity* is not the only way of formulating a sensitivity account of knowledge. Nozick's original account adds the method clause to the antecedent of the sensitivity conditional:

> *Nozick's Sensitivity*: If S knows that p via method M, then were p false and were S to use M to arrive at a belief whether (or not) *p*, S would not believe that p via M (Nozick, 1981, 179).

With *Nozick's Sensitivity*, we need to consider a more restrictive set of possibilities—i.e., closest worlds where not only *p* is false, *but S also uses the same method to arrive at a belief regarding p.*

Applying this account to Kattie's case, we need to consider closest possibilities in which she owns the winning ticket and also uses the same statistical evidence to determine if she is the winner. Imaginably, these are relatively remote possibilities in which although she wins, she does not immediately realize that her ticket is the winner. For instance, maybe the machine is broken so that although her numbers win, the reward money does not come out. In these possible worlds, with the same statistical evidence, she still believes that she does not win. The belief is rendered insensitive.

However, *Nozick's Sensitivity* is not unscathed, for it still does not accommodate all lottery-related beliefs based merely on statistical evidence. To illustrate, note that if in the relevant not-*p*

[5]See also Blome-Tillmann (2017) for a different lottery-based objection against *Sensitivity*.

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

worlds one still uses statistical evidence in favor of *p*, it is guaranteed that one still believes that *p*. Such a method is what Luper-Foy (1984) calls a "one-sided method"—a method that can only result in believing that *p*, instead of not-*p* (cf. Zalabardo, 2012, chapter 3). Therefore, with *Nozick's Sensitivity*, whenever one forms a belief in *p* with statistical evidence that supports *p*, one's belief must be insensitive. But the problem is that strong statistical evidence *can* result in at least some knowledge. Consider a fair lottery case in which there are 100000 tickets and only one of them is the winning ticket. The winner has not been announced yet. Cindy is aware of the chance of winning. She buys a ticket and believes that *p—that she is unlikely to be the winner*—based on the strong statistical evidence that counts in favor of this proposition. Plausibly, Cindy knows that she is unlikely to win.

However, according to *Nozick's Sensitivity*, the belief is insensitive. According to this account, we need to consider closest worlds in which Cindy is *not* unlikely to win *and* she still uses the same statistical evidence in favor of *p* to determine whether she is unlikely to win or not. These may be worlds, for instance, in which the lottery has been rigged in Cindy's favor, so that she is likely to win. In these worlds, by still using the same statistical evidence in favor of *p*, Cindy would still believe that she is unlikely to win. (It could be the case that in these worlds Cindy is not aware of the fact that the lottery has been rigged, so that she still determines whether or not *p* according to the same statistical evidence in favor of *p*.) Again, the method is a one-sided method that only results in the same belief as in the actual world.

One might object that Cindy's case challenges *Nozick's Sensitivity* only if we assume a non-factive conception of evidence. Thus far I have treated Cindy's evidence as consisting in the proposition that the chance of winning is 1 in 100000, and I have assumed that the truth value of this proposition can vary across possible worlds without thereby altering her evidence. Proponents of a factive conception of evidence, however, will reject this assumption. They will argue that it is only in the actual world that Cindy has the evidence which consists in the true proposition that the chance of winning is 1 in 100000. In worlds where she is likely to win, she lacks this evidence, since the proposition is false in those worlds. It follows that, on *Nozick's Sensitivity*, Cindy's belief that *p*—that she is unlikely to win—is trivially sensitive: there are no possible worlds in which this belief is false and in which Cindy employs the same (true) evidence in favor of *p* as in the actual world.[6] The antecedent of *Nozick's Sensitivity* is a counterpossible. This result corresponds to the intuition that Cindy knows that she is unlikely to win. Considered in this light, it appears that a factive conception of evidence rescues *Nozick's Sensitivity*.

However, a slight variation of Cindy's case shows that *Nozick's Sensitivity* remains problematic even when evidence is taken to be factive. Suppose Cindy carefully calculates the odds of winning herself. There are 1000 lottery tickets in total, none of which have been sold yet. Cindy is informed that there will be one winner and wants to figure out each ticket's chance of winning. After counting carefully, she concludes that the chance of each ticket's winning is 1 in 999—she has inadvertently omitted one ticket. On the basis of this premise, she infers that she is unlikely to win were she to buy a ticket.

This is a case of knowledge from falsehood.[7] Intuitively, Cindy knows that she is unlikely to win were she to buy a ticket, even though her inference is based on a false premise. But *Nozick's Sensitivity* fails to explain Cindy's knowledge even when coupled with a factive conception of evidence. On that conception, Cindy's evidence does not consist in the false proposition that each ticket's chance of winning is 1 in 999. At best, her evidence consists in the true qualified proposition that it *seems* to her that the chance of winning is 1 in 999. Given this, her belief is insensitive on *Nozick's Sensitivity*. For there *are* antecedent worlds in which she is likely to win and still employs

[6]I have assumed that Cindy's method is *appealing to statistical evidence in favor of p*. I will consider alternative ways of individuating her method later.

[7]See Warfield (2005) and Arnold (2013) for discussions of this type of case.

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

the same evidence to determine whether she is unlikely to win. (Even in worlds where the lottery is rigged to her favor so that she is likely to win, it could still be true that it *seems* to her that her chance of winning is 1 in 999.) In these worlds, with the same evidence in favor of the proposition that she is unlikely to win, she would still believe that she is unlikely to win. Hence, her belief is insensitive—which is an implausible result. Thus, a factive conception of evidence does not help defenders of *Nozick's Sensitivity* to avoid the objection.

In the above discussions, I have assumed that Cindy's method is *appealing to statistical evidence in favor of p*. One may argue that the method can be individuated otherwise. In particular, it may be more broadly individuated as *statistical evidence in general*. If so, we should consider closest worlds in which Cindy is not unlikely to win and she uses statistical evidence to decide whether or not *p*. In these worlds, her statistical evidence may not indicate that she is unlikely to win, at least not necessarily. For instance, these worlds could be those in which Cindy has purchased a lot more lottery tickets, so that she is not unlikely to win. Crucially, in these worlds her statistical evidence does not lead her to believe that she is unlikely to win. She would rather believe that she is *not* unlikely to win, given that her chance of winning is relatively high. If these are the relevant antecedent worlds, Cindy's belief can be rendered sensitive.

However, there are good reasons not to individuate methods so broadly—at least in the case of inferential methods. To illustrate, consider another case where Cindy* is completely unaware of the chance of winning. However, in the past she played the same lottery a few times (buying one ticket each time) and she did not win. So, this time, after buying a new ticket, she reflects on her very limited past experience and, purely on that basis, she jumps to the conclusion that she is unlikely to win. "After all, in the past I didn't win, so this time it's unlikely that I will win." She thinks.

In this case, it seems that Cindy*'s statistical evidence is insufficient to enable her to know.[8] If statistical evidence could ever generate knowledge, it should be stronger than Cindy*'s. However, if proponents of *Nozick's Sensitivity* individuate methods broadly so that Cindy*'s method is described as statistical evidence, then Cindy*'s belief is also sensitive. Again, the relevant antecedent worlds could be those in which Cindy* has purchased a lot more lottery tickets, so that she is not unlikely to win. In these worlds, her statistical evidence does not lead her to believe that she is unlikely to win. Instead, she would believe that she is *not* unlikely to win. Thus, Cindy*'s belief is sensitive, just as Cindy's. But this is a wrong result. Cindy* is not a knower, even if Cindy is.

Thus, some statistical evidence is strong but some is not. Needless to say, only the former could generate knowledge. However, if proponents of *Nozick's Sensitivity* individuate inferential methods broadly, so that in cases of inference based on statistical evidence whether one's evidence is strong or not is not considered as a constitutive component of the method, then these theorists will be unable to differentiate between cases of knowledge based on strong statistical evidence and cases of ignorance based on insufficient statistical evidence.

In sum, *Nozick's Sensitivity* is not better suited than *Sensitivity* in accommodating beliefs based on pure statistical evidence. On the one hand, if proponents of *Nozick's Sensitivity* individuate methods specifically so that the evidential strength of statistical evidence is taken into account, then they could explain Kattie's ignorance but not Cindy's knowledge. On the other hand, if they individuate methods more generally so that the evidential strength of statistical evidence is not considered as part of the method, they could better accommodate Cindy's knowledge but not Cindy*'s ignorance.[9]

[8]If you think Cindy* has a background belief that it is *in general* unlikely to win a lottery and it is this belief that enables her to know that she is unlikely to win, then we may stipulate the case so that this background belief is unjustified. For instance, suppose that there are various lottery games that are on the market in Cindy*'s town and some of them have a relatively high chance of winning (though the prize money is not much at all).

[9]Günther (2024a) defends an interesting and novel form of sensitivity condition which is dubbed "epistemic sensitivity," in the context of legal proof. The central idea is that a defendant should be found liable for *p* only if the fact finder believes that her

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

3. Against Safety I: Some Extant Counterexamples

Different from sensitivity theorists, safety theorists appeal to the closeness of error in explaining beliefs based merely on statistical evidence. This is how Pritchard deals with ignorance cases where one believes that her particular ticket is not the winner based merely on statistical evidence. As Pritchard argues: "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, Pritchard thinks that possible worlds where one's particular ticket is the winner are rather close to the actual world. So, one's belief that one does not win is unsafe, which explains why one does not know one is not the winner.[10]

Notably, Broncano-Berrocal (2019) and Paterson (2022) have argued against safety-based explanations of lottery cases. They present variants of lottery counterexamples and argue that a safety account is unable to deliver the correct result in such cases.[11] In what follows, I will point out the limitations of these counterexamples. And then, in the next section, I will present a more successful objection. Consider first Broncano-Berrocal's counterexample:

> Eaten Ticket: Lisa bought a ticket for her twin sister Lottie and put it in her drawer to keep it there until the lottery draw. Lottie comes to know that Lisa bought a ticket for her and, based on the odds, comes to believe that she won't win. As a matter of fact, the winning number is Lottie's ticket number. Unbeknownst to both, when cleaning the drawer, the household help dropped the ticket on the floor where it was swiftly eaten by the dog (weeks ago), so Lottie no longer owns the ticket. (Broncano-Berrocal, 2019, 42)[12]

---

evidence e is sensitive to p. An agent believes e is sensitive to p iff she believes two conditionals: $p > e$ and $\sim p > \sim e$. These conditionals are interpreted as follows: an agent believes $p > e$ iff upon supposing all the most likely scenarios where p is true, she would still believe e; an agent believes $\sim p > \sim e$ iff upon supposing all the most likely scenarios where p is false, she would believe $\sim e$ (Ibid., 1356). When applied to the analysis of knowledge, this yields an epistemic sensitivity account according to which S knows that p based on evidence e only if S believes that $p > e$ and $\sim p > \sim e$.

One worry with this account is that it is too demanding to accommodate lottery beliefs based purely on statistical evidence. Consider again Cindy's case: she believes that p—that she is unlikely to win the lottery—based on the strong statistical evidence in favor of losing. According to epistemic sensitivity, whether she knows that p depends on what she would believe upon supposing the most likely scenarios in which p and $\sim p$ hold. The most likely $\sim p$ scenarios in which Cindy is likely to win may include ones where e still holds. In those scenarios, the announced odds of winning remain the same, but some unrevealed factor (such as fewer actual tickets or a secret rigging) makes Cindy's chance of winning relatively high. It is therefore possible (and we may also stipulate) that were Cindy to consider the most likely $\sim p$ scenarios, she would still believe e. Thus, she does not believe $\sim p > \sim e$. Epistemic sensitivity thus fails to explain why Cindy knows that she is unlikely to win.

Epistemic sensitivity also appears too permissive. Consider Amy, who has good reasons to believe that Bob, a loyal employee of a lottery company, possesses inside information about which tickets will win or lose. She buys a ticket and asks Bob whether it will win. Unbeknownst to Amy, Bob has no such information; he merely responds at random, saying: "Yes, it will!" Amy then believes that her ticket will win. Intuitively, Amy does not know that her ticket will win, even if it in fact does. Nevertheless, her belief satisfies epistemic sensitivity. Upon supposing the most likely $\sim p$ scenarios—where her ticket loses—Amy would believe that Bob would not have said the same thing; hence, she believes that $\sim p > \sim e$ holds. Upon supposing the most likely p scenarios—where her ticket wins—she would believe that Bob would have said the same thing; hence, she believes that $p > e$ holds as well. Her belief therefore implausibly satisfies the epistemic sensitivity condition. (See also Zhao, 2026 for critical discussions of Günther's 2024a view.)

[10]What is assumed here is a naïve similarity conception of closeness of possible worlds. On this view, if a world is very similar to the actual world in the relevant respects, it is counted as close. (For further discussions, see Bogardus, 2014, 294–296.) It is the safety-based solution to the lottery problem relying on this conception of closeness that is the target of the arguments in this section and Section 4. It is worth noting that this is not the only way of reading the safety condition. On a reading that aligns more closely with our ordinary language intuitions of closeness, one might argue that the belief that one will not win the lottery is indeed safe. After all, ordinarily we think that it's not easy to win the lottery (and that's why so few people play it). In *this* sense, the belief that one will not win the lottery could not easily have been false. I thank an anonymous referee for raising this point.

[11]See also Mills (2012), Zhao (2023), Zhao and Baumann (2021) for relevant discussions of safety and lottery.

[12]Broncano-Berrocal (2019, 42) presents another counterexample (Address Lottery case), which shares the same structure as Eaten Ticket. My following argument applies to Address Lottery case as well.

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

As Broncano-Berrocal points out, this is a Gettierized lottery case. It involves a coincidence that is composed of two unrelated events: (1) Lottie's ticket number happens to be the winning number and (2) Lottie's ticket is eaten by the dog. Broncano-Berrocal thinks that as long as it is stipulated that event (2) is bound to occur—that is, that it occurs in all close possible worlds—Lottie's belief that she won't win would also be true in all the close possible worlds (Ibid., 42–43). That is, assuming that in all the close possible worlds the ticket is eaten by the dog so that Lottie no longer owns a ticket, she won't win in those worlds either. Hence, her belief is safe according to *Safety*. However, intuitively Lottie does not know she won't win.

Though interesting, this case has only limited power of challenging safety-based explanations of lottery cases. Given that in all the close possible worlds Lottie won't win, the example resembles cases of beliefs in necessary propositions or beliefs in contingent but modally robust propositions. In the literature, a standard way of accommodating such beliefs is to globalize safety. In particular, a globalized safety condition demands that for one's belief that p to be knowledge, the method that produces the belief that p could not easily have produced false beliefs in relevantly similar propositions. What counts as "relevantly similar propositions"? Hirvelä (2019) offers an answer: relevant propositions are those that belong to the same subject matter of inquiry. In lottery cases, the subject's inquiry is to determine whether or not she will win. Then, all the propositions that belong to this inquiry will be counted as relevantly similar. A number of proponents of safety have argued that the globalized safety is superior to *Safety*, which requires only that beliefs in a single proposition p are true in close worlds (Blome-Tillmann, 2017; Hirvelä, 2017; 2019; Pritchard, 2012; Williamson, 2000; Zhao, 2024).

Crucially, with the globalized safety, one *can* deliver the result that Lottie's belief in Eaten Ticket is unsafe. After all, based on the same statistical evidence, Lottie could easily have formed false beliefs in other propositions that belong to the same subject matter of inquiry, such as *the numbers on the ticket are not the winning numbers* or *the ticket that Lisa bought is not a winning ticket.* Even if it is granted that Lottie's belief that *she won't win* is stably true in close worlds because it is a modally stable fact that the ticket is eaten by the dog, these other beliefs in relevantly similar propositions are not. The truth value of these other propositions is unaffected by whether or not the ticket is being eaten. Whether or not it is being eaten, in some close possible worlds (including the actual world) these propositions are false. So, Lottie's belief is unsafe according to the globalized safety: she could easily have formed false beliefs in relevantly similar propositions. As such, the case fails to be a counterexample against the globalized safety.

Next, consider the following intriguing counterexample from Paterson (2022):

> Suppressor Lottery: Lottie has bought a ticket in a fair lottery. The numbers have been drawn, but Lottie is ignorant of this. Given the enormous number of tickets, Lottie concludes that her ticket is a loser. She is correct: Lottie has lost the lottery. Moreover, whilst the lottery is 'fair', the owners of the lottery have an interest that the prize is not collected. Consequently, they have ensured that anybody that checks the results who possesses a winning ticket will activate a Doxablock^{TM} and will consequently be unable to form the belief that they have won the lottery and collect their prize. (Paterson, 2022, 1094)

Paterson argues that in this case Lottie could not easily have formed a false belief that she has lost the lottery. For even if she had won the lottery, she would not have been able to form the belief that she had won because of the activation of Doxablock^{TM}. In spite of this, given that Lottie's belief is based purely on statistical grounds, she does not know her ticket is a loser (Ibid).

Paterson intends this to be a counterexample against safety—in particular, the globalized safety. However, the main flaw of the above argument is that it neglects the method relativization of safety accounts. The case is such that *were Lottie to check the result*, she would not be able to form the belief that she wins. But checking the result is clearly not the same method as merely reflecting on one's chances of winning. And only the latter is the method Lottie uses in the actual world. Thus, possible

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

worlds where Lottie checks the result is irrelevant anyway, regardless of whether or not she is able to form a belief in those worlds. Relevant possibilities are worlds in which, either after or before the numbers have been drawn, Lottie employs the same method (reflecting on the odds) to determine whether she wins. Plausibly, in at least some such worlds, her numbers are (or will be) the winning numbers, and so she ends up with a false belief that her ticket is a loser. Considered this way, her belief is indeed unsafe.

4. Against Safety II: A New Variant of the Lottery Case

In this section, I will present a novel variant of the lottery-style counterexample against safety accounts, one that can overcome the limitations associated with the above objections. Before proceeding, recall Pritchard's aforementioned comment on lottery. He has in mind a quite specific lottery mechanism, where the winner is determined by the placement of lottery balls. In such cases, it seems true that not much need be changed for one to win (only a few balls need to fall in a slightly different configuration in order for one to win). So, there are close error possibilities where the target belief that one is not the winner is false. However, the problem is that, even if the mechanism is changed, so that there are no such close error possibilities, the intuition that one does not know she is not the winner remains. Consider:

> Jimmy buys a lottery ticket. The result is not announced yet, but Jimmy forms a true belief that he will not win, based merely on statistical evidence. The game is such that the lottery company releases 100000 tickets and, as the company advertises, only one of them is the winner. Unbeknownst to every player, however, the company's budget has run out. So they have decided to save money by not delivering any reward. The plan is to announce that one of the unsold tickets is the winner, so that they need not offer reward to any players.[13]

Intuitively, Jimmy does not know that he will not win the lottery. The company's untoward trick notwithstanding, this is not something that turns Jimmy into a knower. To the extent that he is not aware of anything going on inside the company, he does not know he will lose.

However, according to *Safety*, Jimmy's belief is safe: there are no close worlds in which he wins the lottery. For him to win, quite a few changes need to be made: the company's financial situation gets much better, they change their mind about picking an unsold ticket as the winner, Jimmy's number happens to be the winning number, etc. Presumably, worlds where such events all occur are not close to the actual world. Thus, given the actual circumstances, even if technically speaking Jimmy still could win (in the sense that it is not metaphysically impossible that he wins), at best he wins in relatively remote possible worlds. Of course, closeness is a vague notion and it is hard to pin down precisely how far away the worlds where Jimmy wins are. Be that as it may, it suffices for our purposes that worlds where he wins are farther out from the actuality, compared to worlds where one wins the lottery by a few balls being placed in a slightly different configuration.

As noted earlier, the lottery problem, broadly speaking, is the problem of explaining why a belief based merely on statistical evidence does not amount to knowledge (even when, probabilistically, the evidence strongly supports the belief). The safety theorist's explanation is that such beliefs are unsafe—in particular, statistical evidence leaves the possibility of error too close to the actual world.

[13]This case shares similarities with Kelp's (2018, 113) rigged-lottery case. In Kelp's version, the lottery is rigged so that no ticket wins; Jimmy's case involves a more detailed story of how the lottery is rigged: the financially distressed company ensures that no sold ticket wins by announcing an unsold ticket as the winner. Kelp deploys his case mainly in discussion of virtue-reliabilist views, with brief remarks about Pritchard's safety, whereas Jimmy's case is used here to test a wide range of safety conditions, including normality-based accounts.

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

Jimmy's case, however, precisely challenges this diagnosis. The case shows that even if the possibility of error is remote from the actuality, our intuition about ignorance in a lottery case remains. Thus, closeness of error is not an essential factor in explaining why lottery beliefs based on statistical evidence alone do not constitute knowledge (see also Zhao, 2023). A more adequate solution to the lottery problem must therefore be sought elsewhere.

Also, note that Jimmy's case differs from Paterson's Suppressor Lottery. In the latter, there *are* close possible worlds where Lottie's ticket wins. In these worlds, by using the same method of statistical evidence, she forms a false belief that she does not win. This is why her belief is unsafe. But in the present case, there are no close possible worlds in which Jimmy's ticket wins, regardless of what Jimmy's method is. So, Jimmy's belief is safe.

Furthermore, different from Broncano-Berrocal's case, Jimmy's scenario not only casts doubt on *Safety* but also challenges the globalized safety. Granted, Jimmy has a false belief that his chance of winning is 1 in 100000. But on a plausible description of the case, this probabilistic belief should be considered as an input belief that is constitutive of his method, rather than an output belief that is issued by his method. More specifically, we may stipulate that it is *based on* his probabilistic belief that Jimmy *infers* the actual belief that he will not win the lottery. Then, Jimmy's method, though it involves a false belief, could not easily have produced false beliefs in similar propositions. The method could easily have produced beliefs such as *the chance of winning is very low* or *I am unlikely to win the lottery and collect the prize*, etc. But these beliefs are not false, strictly speaking. (There is still a chance of winning, although the chance is lower than 1/100000 given what is going on inside the company.) So, even on the globalized safety condition, Jimmy's belief is still implausibly safe.[14]

Those who deny that evidence may consist in false propositions (e.g., Williamson, 2000) might argue that Jimmy's statistical evidence does not include the proposition that his chance of winning is 1 in 100000, since this proposition is false. They might instead maintain that, at best, Jimmy's evidence consists in a qualified true proposition such as *it seems that the chance of winning is 1 in 100000*. Moreover, suppose we describe Jimmy's method of belief formation as inferring on the basis of his evidence (i.e., the qualified proposition). On this description, Jimmy could easily have formed relevantly similar false beliefs via the same method. In particular, he could easily have believed falsely that his chance of winning is 1 in 100000 on the basis of his belief that it seems so. Hence, Jimmy's belief that he will not win is unsafe according to the globalized safety.[15]

In reply, note first that even if it is possible for Jimmy to believe—and know—that it seems that his chance of winning is 1 in 100000, it is far less plausible that he must hold such a belief. If he is highly confident about the odds of winning, it is more natural to think that he only believes the unqualified proposition that the chance of winning *is* 1 in 100000, rather than the qualified proposition that *it seems so*.

Moreover, although it is possible that Jimmy's method involves inferring from the true qualified proposition, it is implausible that the method must be characterized in this way. Surely a method can involve inference from false propositions. Indeed, even Williamson—the foremost defender of the non-factive conception of evidence—maintains that, in assessing whether a belief is safe, we should consider whether one could easily have formed false beliefs on a similar *basis*. For Williamson, a "basis" is akin to a "process." In response to Alvin Goldman, he writes "…by 'bases', I did not mean grounds; I meant something more like processes than Goldman appreciates. I had in mind a

[14]In view of the present objection, proponents of safety may further strengthen the globalized safety condition. In particular, they may demand that not only the method could not easily have produced false beliefs in any relevantly similar propositions but also *the method itself must not involve any false beliefs*. But this requirement is too strong, as it excludes cases of knowledge from falsehood. For instance, suppose Tom's watch is set 1 minute earlier but he is unaware of this. When he looks at it, he sees that it reads 2:30 pm. So he falsely believes that it is 2:30 pm. Based on this belief, he infers that it is past 2 pm (or that it is afternoon). Plausibly, Tom can come to know that it is past 2 pm (or that it is afternoon), even though his method involves a false belief (cf. Warfield, 2005).

[15]I am grateful to the editors for raising this concern.

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

very liberal conception, on which the basis of a belief includes the specific causal process leading to it and the relevant causal background" (Williamson, 2009, 307). On this conception, a basis may involve a false belief, since a false belief can be part of the causal process that produces another belief.

Accordingly, in the lottery case one's method (or Williamson's "basis") need not be inference from true propositions. We might instead stipulate that Jimmy's belief that he will not win is consciously inferred from (and partly caused by) his belief in the unqualified false proposition that the chance of winning is 1 in 100000. On this description, the false belief that the chance of winning is 1 in 100000 is not a belief that he could easily have formed *via* the same actual method; rather, it is constitutive of that very method. According to the globalized safety, insofar as he could not easily have formed other false beliefs via this method (as I have argued he could not), his belief counts as safe—even though the method itself involves a false belief.[16]

Finally, a safety theorist might point out that their account specifies only a necessary condition for knowledge, not a sufficient one. Since the above objection merely shows that safety is not sufficient for knowledge, there remains the option of adding further conditions to explain Jimmy's ignorance. However, even if such additional conditions are available, the challenge against the safety account persists. To see this, note that although Jimmy's case differs from the paradigmatic lottery case (i.e., Lottie's case in Section 1) in terms of the modal profile of the subject's belief, the two cases are not wholly unrelated. In both, the subject believes that they will not win the lottery based on mere statistical evidence. If the safety theorist explains Lottie's ignorance by appealing to the closeness of error possibilities, it is puzzling why an appeal to the modal profile cannot be applied to explain Jimmy's ignorance. The safety theorist thus faces the challenge of explaining this discrepancy, regardless of whether an additional condition would yield the correct verdict in Jimmy's case.

More broadly, this issue points to a limitation in safety's explanatory power: given that the above safety accounts struggle with Jimmy's case, safety theorists have to concede that their accounts do not fully capture what is epistemically problematic about beliefs grounded solely in statistical evidence. This flaw in explanatory power should be troubling to the safety theorist, since the ability to solve the lottery problem is often cited as a crucial advantage of the safety account.[17]

5. Against Safety III: Normality

A potentially more promising option, on behalf of a safety theorist, is to envisage relevant possibilities in a different fashion. The traditional safety accounts like *Safety* (or the globalized safety) focus on "close" possibilities. Recently, a number of theorists argue that it is better to relativize safety to "normal" possible worlds:

> *Normality Safety*: If one knows that p, then one's belief that p is not only true in the actual world, but in all normal possible worlds where one continues to believe that p via the same

[16]See also note 14.

[17]The dialectic here is reminiscent of the Gettier problem for the traditional JTB account of knowledge. Gettier-style counterexamples show that the JTB account is *insufficient* for knowledge, not that it is unnecessary. Nevertheless, many take these counterexamples to constitute a serious challenge to the account. As I see it, the concern here primarily concerns explanatory power: Gettier-style cases reveal that the JTB account cannot adequately explain why justified true beliefs formed by luck do not amount to knowledge. A similar point applies to lottery-style counterexamples against safety. Even if such cases do not demonstrate that safety is unnecessary for knowledge, they still pose a challenge for the safety theorist, since they indicate that the account is not explanatorily powerful enough to accommodate beliefs based on mere statistical evidence.

Furthermore, in the Gettier literature, epistemologists tend to regard a theory of knowledge as unsatisfying if it only addresses some Gettier-style cases (such as those proposed by Gettier himself) but cannot account for other Gettier-style cases that share the same key features. (Indeed, this is often how debate in the literature proceeds, see, e.g., Mortini, 2024, Miracchi, 2015.) Analogously, a safety theory is unsatisfying if it only accounts for some lottery cases (such as Lottie's case) but not others (such as Jimmy's case).

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

> method as in the actual world, the belief continues to be true. (cf. Beddor & Pavese, 2020, Dutant, 2010, Smith, 2016, Goodman & Salow, 2023, Littlejohn & Dutant, 2020)

Here, normal possible worlds are those in which conditions are at least as normal as those obtained in the actual world. Could *Normality Safety* better accommodate lottery beliefs based solely on statistical evidence? Much depends on how one understands the concept of normality. In what follows, I will proceed with an intuitive notion of normality and apply theoretical accounts where needed.

Consider the stereotypical lottery case in which Lottie believes that she will not win, based on the excellent statistical evidence against winning. Given that the lottery is fair and not rigged, winning the lottery—though highly unlikely—does not appear to be an abnormal event. In Smith's (2016) terminology, in this case winning is on an "explanatory par" with losing—it does not require any special explanation. So, according to *Normality Safety*, winning worlds should be included for the purpose of determining whether or not the belief is safe. Therefore, Lottie's belief that she will not win is rendered unsafe, which is a desirable result (see also Beddor & Pavese, 2020, 71).

Unfortunately, Jimmy's case—which is a counterexample against standard formulations of safety (*Safety* and the globalized safety)—constitutes an objection to *Normality Safety* as well. Recall that in this case, the lottery company is in a poor financial situation and decides not to announce a winner. Given this, it appears that possible worlds where Jimmy wins the lottery, *despite* the company's current financial situation and the decision, are even more abnormal. So, such worlds should be ruled out. This verdict is also supported by Smith's (2016) account of normality. Possible worlds in which Jimmy wins require more explanation than the actual world in which he does not win. Accordingly, those worlds are correspondingly more abnormal. Hence, on *Normality Safety*, such worlds are ruled out.

One may argue that, since Jimmy is unaware of what is going on inside the company and his evidence consists only of his belief that his chance of winning is 1 in 100000, he *himself* may not regard the possibility of winning as abnormal. However, evaluations of normality should not be determined solely by the subject's own perspective (cf. Zhao & Baumann, 2021). A subject's evidence could be limited so that her judgments about the normality of a possibility depart from a more objective standard of normality. In my view, at least for the purposes of epistemic evaluation, it is the objective assessment of normality that is more relevant and worth theorization. A subject's personal sense of what is normal may vary from person to person, though the epistemic statuses of their beliefs remain the same. Thus, were normality to be evaluated solely in terms of the subject's own perspective, this would lead to an implausible form of epistemic relativism.

To illustrate, suppose Lottie participates in a fair lottery and believes that she will not win, based solely on statistical evidence in favor of losing. She also thinks that (as Smith, 2016 argues) the possibility of her winning, though extremely unlikely, is not an abnormal one. Now compare Lottie with Tommy, who participates in the same lottery and believes that he will not win. The reason that Tommy believes so is that he is misled by someone else into believing that the lottery has been rigged in favor of another player, making it virtually impossible for him to win. As a result, Tommy regards the possibility of his winning as highly abnormal.

On a subjectivist account of normality, *Normality Safety* produces the counterintuitive verdict that Tommy's belief that he will not win is safe, while Lottie's belief is not. But this is implausible: to the extent that Tommy's belief is based on misleading testimony, he does not know he will not win either. This problem can be avoided if proponents of *Normality Safety* adopt an objective standard of normality. From an objective standpoint, the possibility of Tommy's winning is no more abnormal than Lottie's winning, given that they participate in the same fair lottery.

Another possible objection is that, even if we grant that the possible worlds where Jimmy wins the lottery (despite the company's dishonest decision) are abnormal possibilities, defenders of *Normality Safety* could still deliver the verdict that Jimmy fails to know by incorporating other, more normal possibilities. These are possibilities in which Jimmy participates in a non-rigged

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

lottery conducted by an honest company. In at least some of these intuitively more normal worlds, Jimmy wins the lottery. So his belief is unsafe.[18]

However, the actual world can be redescribed so as to avoid incorporating these possibilities. Suppose that Jimmy is a dedicated follower of statistical evidence and normally acts only in accordance with what his statistical evidence recommends. As such, he has never played the lottery before, since he considers it a foolish game. He purchases a ticket on this occasion only because his best friend, Joe, is a loyal employee of the (dishonest) company. Joe persistently urges Jimmy to try their lottery. That's why Jimmy buys the ticket—just to stop his friend's nagging about the lottery.

Given the addition of this twist, the more normal possible worlds are not ones in which Jimmy participates in a different lottery, but rather ones in which he does not participate in any lottery at all. In Smith's (2016) terminology, given how Jimmy thinks about lotteries in general, worlds in which he participates in a different lottery require special explanation. Perhaps in those worlds his outlook on lottery has changed, or perhaps his friend Joe has switched to a more honest lottery company and persuades Jimmy to buy one of their tickets, etc. No matter how we fill in the details, such worlds appear to require special explanation. By contrast, worlds in which Jimmy does not play any lottery require no special explanation. On Smith's account, it is these worlds that should be taken into account. But then, proponents of *Normality Safety* fail to explain Jimmy's ignorance. In the normal possible worlds, he simply forms no beliefs about lottery at all. His actual belief is trivially safe.

In general, the lesson is that what counts as normal for a subject is relative to her interests, background beliefs, behavioral dispositions, etc. These subjective factors can be easily stipulated so that the subject's holding the belief that p constitutes an abnormal departure from her normal doxastic and behavioral patterns. As such, the more normal possible worlds are ones in which she does not form any relevant belief concerning p at all (Valaris, 2022, 399–400). If so, her true belief that p becomes trivially safe, since in no normal possible worlds does she believe that p.

Finally, a safety theorist may attempt to develop an alternative notion of normality—one that can better handle lottery cases. However, the above objections against *Normality Safety* are based on a conception of normality that closely aligns with our pre-theoretical intuitions. Given this, an alternative account risks becoming detached from our pre-theoretical notion of normality. That does not seem a promising path to pursue.

6. Modal Accounts versus Explanationism

I have argued that sensitivity and safety accounts cannot properly accommodate beliefs based solely on statistical evidence. More specifically, it turns out that our epistemic judgments about such beliefs do not consistently track what occurs in a specified set of nonactual possibilities, be they normal possibilities, close possibilities, or possibilities in which the believed proposition is false.

The failure of these modal knowledge accounts delivers the following useful lesson: it seems preferable to pay more attention to the *actual* profile of lottery beliefs when explaining their epistemic defectiveness, as opposed to what occurs in counterfactual situations. In this regard, one notable candidate is *explanationism*, which has recently been defended as a main rival of modal accounts.[19] In the remainder of this article, I will examine the explanatory power of explanationism. My purpose is not to launch a full-fledged defense of this account. I shall be content with this conclusion: explanationism is better suited than modal accounts in accommodating beliefs based solely on statistical evidence. So, considerations of statistical evidence count in favor of explanationism, as opposed to its modal rivals.

[18]I assume that in these worlds Jimmy forms the belief via the same method as in the actual world.

[19]For debates between explanationism and modal accounts, see, e.g., Bogardus and Perrin (2022, 2025), Faraci (2019), Mills (2012), Mortini (2024).

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

Whereas proponents of modal accounts demand that knowledge requires a suitable modal connection between belief and truth in certain nonactual possibilities, explanationists jettison this claim. In their view, truth must feature as an important explanation of why the *actual* belief is being held. Thus, Goldman (1984, 101), in his early defense of explanationism, claims that "a belief counts as knowledge when appeal to the truth of the belief enters prominently into the best explanation for its being held." Similarly, in their recent work, Bogardus and Perrin (2022, 179, italics original) define their explanationism as follows: "knowledge requires only that truth play *a crucial role* in the explanation of your belief."[20] In general, explanationists hold that there must be an explanatory connection between the fact that renders one's believed proposition true and one's belief. In its simplest form, the connection is such that one believes that p because p is true. That is, truth constitutes a prominent explanation of why the belief is being held.

Before proceeding, some clarifications are in order. One may worry that if the explanatory relation between p and believing that p entails a modal relation between them, then explanationism is not a genuine alternative to modal accounts. For instance, it may be tempting to think that if p is a prominent explanation of one's belief that p—as explanationists demand—then were p not to obtain, one would not believe that p either. But this is just a simple version of the sensitivity account.

However, explanationists like Bogardus and Perrin (2022, note 15) explicitly reject the idea that an explanatory relation entails a modal relation. They think this inference faces obvious counterexamples from the causation literature. Consider a preemption case. Suppose the first shooter fires at the victim's heart, causing the victim's death. A moment later, a second shooter fires at the victim's heart as well. Clearly, the victim's death is explained by the first shooter's shot. But the aforementioned modal relation does not hold: in the closest worlds where the first shooter does not fire, or where the shot fails to kill the victim, the victim would still have died due to the second shooter's shot. (Moreover, as I will argue later, the lottery-style counterexamples that afflict sensitivity and safety accounts do not pose a threat to explanationism. This further corroborates the point that explanationism does not imply a modal account.)

The notion of "explanation" in explanationism is worth further unpacking. Explanationists tend to understand explanation as an answer to a *why*-question. When S believes that p, we may ask: why does S believe that p? Explanationism requires that the truth of p should figure prominently in the answer to this question. That is, S's believing that p is made *intelligible* by appeal to the truth of p. Put another way, appealing to the truth of p should provide a clear and illuminating account of why S believes that p (cf. Bogardus & Perrin, 2022; Jenkins, 2006).

More specifically, explanationists typically understand the explanatory relation in objective terms. On Jenkins' (2006, 139) account of explanationism, the truth of p should be a good explanation of S's believing that p, where "good explanation" is understood from the perspective of a "rational outsider"—that is, someone not acquainted with the particular details of S's situation. Bogardus and Perrin (2022, note 15) claim that "An explanation may succeed, render intelligible a phenomenon, and demystify it, even if someone disagrees or fails to appreciate this." These views suggest that explanationism is an *externalist* account of knowledge: whether S's believing that p is explained by the truth of p does not depend on whether S herself has good reasons for thinking that the explanatory relation holds (or something in the vicinity). Rather, it depends on whether the belief is *in fact* explained by the truth.

To further illustrate, consider a classic Gettier case. Suppose Catherine sees an object in clear view that looks like a sheep. She therefore believes that p—there is a sheep in the field. Unbeknownst to her, the object she sees is a sheepdog whose appearance resembles a sheep. Coincidentally, a real sheep is hidden behind the sheepdog. Catherine thus forms a justified true belief that there is a sheep in the field (Chisholm, 1966, 23). Given that Catherine's belief is justified, she may well have good reasons to think that her belief is explained by the truth of p. But to the extent that the belief is not

[20]See also Bogardus and Perrin (2025), Faraci (2019), Goldman (1988, 2008), Jenkins (2006), Mills (2012).

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

actually explained by the truth of *p*, Catherine's belief does not satisfy explanationism. That is, it is not the fact that there exists a real sheep in the field (a fact that makes Catherine's belief true) that explains why Catherine holds the belief that there is a sheep in the field. Rather, she believes so because she sees the sheepdog. Thus, from an objective standpoint, her belief is not explained by the truth of *p*. Accordingly, explanationism delivers the correct verdict that Catherine does not know there is a sheep in the field.

With these preliminaries out of the way, let us now turn to the lottery cases. Note that the ignorance scenarios discussed above all share the following feature: the subjects believe that they will not win, based solely on the low chance of winning, *without checking the result of the lottery*. Why do beliefs formed in such a way fail to constitute knowledge, if not because of some modal failure? Explanationists may answer as follows: these beliefs are formed in a way that has nothing to do with the fact that the subjects will not win. After all, they have not checked the result of the lottery, so that they simply have no contact with the fact that their ticket is a loser. Hence, the truth of their belief does not figure in the explanation of why they hold the belief.[21] Rather, the belief is explained by their knowledge (or belief) about the chance of winning, as well as their abilities to infer on that basis (cf. Goldman, 1988, 52–53). By contrast, consider the case in which Cindy, also based merely on statistical evidence, comes to know that *she is unlikely to win*. Here, Cindy is a knower precisely because it is the fact that she is unlikely to win—given the mechanism of the lottery—that explains her corresponding belief. The truth of her belief thus plays a crucial role in explaining why she holds it. Or so an explanationist may argue (cf. Bogardus & Perrin, 2022, note 25).

In addition, explanationists can resolve a related lottery puzzle with ease. Although one cannot know one will not win the lottery based merely on the low probability of winning, one can know this by reading the lottery result on a reliable newspaper. This is puzzling because there remains a chance that the newspaper misprints the result. And the probability of this occurring may be no lower than the probability of one's winning (Pritchard, 2009, 35).

Pritchard explains the discrepancy here by arguing that, although worlds where one wins the lottery are very close to the actual world, worlds in which a reliable newspaper misprints the result are not. According to him, quite a few changes need to occur for an event like that to happen (Ibid., p.36). But not everyone is persuaded by this diagnosis. McEvoy (2009) insists that worlds in which the newspaper makes an error can be rather close—indeed, these errors even occur in the actual world.

I shall not adjudicate whose intuition regarding the closeness of error worlds is more reliable. It suffices to note that appealing to modal closeness provides a rather meager means of resolving the puzzle, given the vagueness of the similarity metric. By contrast, explanationists can offer a more straightforward account. Reading the result in a reliable newspaper can yield knowledge because the subject's belief is explained by the result as reported, which in turn is explained by the actual draw of the lottery. So, the fact that one's ticket is not the winner (according to the actual draw) does constitute a crucial explanation of why the subject believes that her ticket does not win.

The above discussions already indicate that explanationism enjoys considerable initial plausibility with respect to lottery cases. In the next section, I will further defend explanationism by addressing some potential objections.

7. Objections and Replies

7.1. Deviant causal chain

To start with, one may worry that even explanationism faces lottery-style counterexamples. More generally, accounts like explanationism may struggle with cases that involve deviant causal chains.

[21]Broncano-Berrocal's case deserves a slightly different explanation. In that case, Lottie's ticket *is* the winning ticket but it is eaten by the dog, so her belief that she will not win is true. Here, the fact that she will not win since the ticket is being eaten does not figure into the explanation of her belief at all. What explains Lottie's belief is her belief regarding the low chance of winning and her inference based on this. So, her belief does not meet the requirement of explanationism.

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

And it is not hard to construct such a scenario that involves a lottery. It could be the case that one forms a belief that one's ticket is a loser based solely on statistical evidence, so that one fails to know, but what explains one's believing on such a basis is the fact that the ticket is a loser. Here is a concrete case. Suppose that little Tom asks his father whether his father's lottery ticket is a winner. The father checks the result and comes to know that it is a loser. However, in order to not disappoint Tom (or for whatever other reasons), he only tells Tom how low the chance of winning is, without telling him that the ticket is not a winner. Based on the very low chance of winning, Tom then believes that the ticket is a loser. Here, it may be argued that it is because the ticket is not a winner that Tom's father tells him that the chance of winning is very low, which in turn explains why Tom believes that the ticket is not a winner. Thus, there appears to be an explanatory connection between Tom's belief and its truth.

In reply, putative counterexamples like this can be addressed once it is emphasized that explanationists require the truth to constitute a *prominent explanation* of or play a *crucial role* in explaining the belief. Tom's belief does not satisfy this requirement. The most prominent and immediate explanation of Tom's belief is his father's testimony that the chance of winning is very low (plus Tom's inference on that basis). Also, even if we ask why his father offers this testimony, the explanation is not simply the fact that makes Tom's belief true (i.e., that the ticket is not a winner). Rather, it is largely due to the father's intention and decision of not disappointing Tom. In this sense, the fact of losing is at best a partial explanation of Tom's belief.

Contrast this with the newspaper case, where the immediate explanation of the subject's belief is the newspaper report. When seeking further explanation for that report, we need look no further than the very fact that makes the subject's belief true. That is, it is precisely the drawing of the lottery (which indicates that the subject's ticket is not the winner) that explains the information appearing in the newspaper, which in turn explains the subject's belief. Here, truth remains a salient explanation of the belief. But Tom's case is not like this: the father's intention and decision muddy the explanatory link and render the truth a less prominent factor in explaining Tom's belief.

7.2. Generalization issues

Another worry is that the explanationist's solution to the lottery problem cannot be generalized to accommodate structurally similar cases of inductive knowledge. On its face, it is curious how an explanationist could explain the possibility of inductive knowledge, especially inductive knowledge about the future. After all, it seems implausible to say that my belief that the sun will rise tomorrow is explained by the fact that the sun will rise tomorrow.

In reply, let us first consider how explanationists have attempted to explain inductive knowledge (Bogardus & Perrin, 2022; Jenkins, 2006). In cases where one acquires inductive knowledge about particulars, such as the sun will rise *tomorrow*, one also holds a corresponding belief that involves a generalization, such as that the sun rises *every day*.[22] One believes the sun rises every day because one has observed the sun rising many times in the past. Plausibly, the fact that the sun rises every day explains why one has observed the sun rising many times in the past. As such, it is the fact that the sun rises every day that explains one's belief that it rises every day. According to explanationism, one thus knows that the sun rises every day.

Now, the proposition that the sun rises every day entails that the sun will rise tomorrow. So, when one infers that the sun will rise tomorrow based on one's knowledge that the sun rises every day, the fact that the sun will rise tomorrow enters into the explanation of one's belief that it will rise

[22]Of course, given that the sun has its own lifespan, "every day" here does not mean "eternally." It roughly means "over a very long period of time (billions of years or so)."

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

tomorrow (Bogardus & Perrin, 2022, 191). This is how an explanationist may accommodate inductive knowledge.[23]

Note that the pattern of explanation here is not restricted to this one example. It applies to other cases of inductive knowledge as well. Consider Sosa's much-discussed trash chute case.

"On my way to the elevator I release a trash bag down the chute from my high rise condo. Presumably I know my bag will soon be in the basement. But what if, having been released, it still (incredibly) were not to arrive there? That presumably would be because it had been snagged somehow in the chute on the way down (an incredibly rare occurrence), or some such happenstance." (Sosa, 1999, 145)

Here, it seems that I know the bag will arrive in the basement despite the possibility that it might be snagged on its way down. Explanationists could argue as follows. In this case as well, I have an implicit belief in a generalization to the effect that *under normal circumstances, when one drops a bag into the chute, the bag arrives in the basement* (call this *p*). Plausibly, given my inductive evidence regarding the reliability of the chute, I know *p*. Explanationists can explain my knowledge that *p* as follows. My inductive evidence in this case consists of my past experiences with the chute—namely, I have dropped trash bags into it many times before, and those bags have consistently arrived in the basement without incident. Such evidence is best explained by the truth of *p*. Thus, my belief in *p*, formed on the basis of the inductive evidence, is also explained by the truth of *p*.

Furthermore, *p* entails that the bag will arrive in the basement on this occasion (which is a normal circumstance). Thus, when I infer that the bag will arrive in the basement on this occasion from *p*, the fact that the bag will arrive in the basement on this occasion enters into the explanation of my belief that it will.

Indeed, the above reasoning delivers a promising way to distinguish lottery cases from cases of inductive knowledge. In lottery cases where one believes that one will lose based merely on statistical evidence, one holds a corresponding belief concerning the odds of winning, such as *the chance of winning is very low*. Crucially, this proposition does not entail that one's particular ticket will lose. Despite the high probability of losing, it is still possible that one wins. So, when one infers that one's ticket will lose from the probabilistic belief, the fact that one's particular ticket will lose does not enter into the explanation of one's belief that it will, as this fact is not entailed by the premise. This differs from cases of inductive knowledge where the proposition involving the generalization does entail the truth of the believed proposition concerning the particular. As such, in cases of inductive knowledge—but not in lottery cases—the subject's belief satisfies explanationism.

8. Conclusion

To summarize, I have established two conclusions in this article. First, on the negative side, I have argued that, contra sensitivity and safety theorists, their accounts face lottery-style counterexamples. As a result, these modal theories fail to properly accommodate beliefs based merely on statistical evidence. Second, on the more positive side, I have shown that where modalists stumble is where explanationists rise. The latter can better address the problematic lottery cases that modalists struggle with. As such, explanationism offers a more promising framework for accommodating beliefs based solely on statistical evidence.

Acknowledgments. I would like to thank the editors and anonymous reviewers at the *Canadian Journal of Philosophy* for their very helpful feedback on multiple earlier drafts of this article.

Funding statement. This research was supported by the National Social Science Fund of China (Grant Number 25CZX037).

[23]On this view, one acquires inductive knowledge that the sun will rise tomorrow based on a deductive inference from "the sun rises every day" to "the sun will rise tomorrow." Be that as it may, the knowledge that the sun will rise tomorrow is inductive in the sense that the epistemic status of this belief depends on one's evidence for the premise (that the sun rises every day). And *that* evidence is inductive—it mainly consists in one's past observations of the sun.

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

Haicheng Zhao is a faculty member in the Department of Philosophy at Xiamen University. His main research interests cover epistemology and normative ethics. He is currently working on a nationally funded project exploring paradoxes of statistical evidence.

References

Arnold, A. (2013). Some evidence is false. *Australasian Journal of Philosophy*, 91(1), 165–172.
Becker, K. (2007). *Epistemology modalized*. Routledge.
Becker, K. (2012). Methods and how to individuate them. In K. Becker & T. Black (Eds.), *The sensitivity principle in epistemology* (pp. 81–97). Cambridge University Press.
Beddor, B., & Pavese, C. (2020). Modal virtue epistemology. *Philosophy and Phenomenological Research*, 101(1), 61–79.
Black, T. (2008). Defending a sensitive neo-Moorean invariantism. In V. Hendricks & D. Pritchard (Eds.), *New waves in epistemology* (pp. 8–27). Palgrave-Macmillan.
Blome-Tillmann, M. (2017). Sensitivity actually. *Philosophy and Phenomenological Research*, 94(3), 606–625.
Bogardus, T. (2014). Knowledge under threat. *Philosophy and Phenomenological Research*, 88(2), 289–313.
Bogardus, T., & Perrin, W. (2022). Knowledge is believing something because it's true. *Episteme*, 19(2), 178–196.
Bogardus, T., & Perrin, W. (2025). A defense of explanationism against recent objections. *Episteme*, 22(1), 35–46.
Broncano-Berrocal, F. (2019). Knowledge, safety, and Gettierized lottery cases: Why mere statistical evidence is not a source of knowledge. *Philosophical Issues*, 29, 37–52.
Chisholm, R. (1966). *Theory of knowledge* (1st ed.). Englewood Cliffs, NJ: Prentice Hall.
Cross, T. (2007). Comments on Vogel. *Philosophical Studies*, 134(1), 89–98.
DeRose, K. (1995). Solving the skeptical problem. *Philosophical Review*, 104(1), 1–52.
DeRose, K. (2010). Insensitivity is back, baby! *Philosophical Perspectives*, 24(1), 161–187.
Dutant, J. (2010). Two notions of safety. *Swiss Philosophical Preprints*, 87, 1–19.
Dutant, J. (2016). How to be an infallibilist. *Philosophical Issues*, 26(1), 148–171.
Faraci, D. (2019). Groundwork for an explanationist account of epistemic coincidence. *Philosopher's Imprint*, 19(4), 1–26.
Goldman, A. H. (1984). An explanatory analysis of knowledge. *American Philosophical Quarterly*, 21(1), 101–108.
Goldman, A. H. (1988). *Empirical knowledge*. Berkeley, CA: University of California Press.
Goldman, A. H. (2008). Knowledge, explanation, and lotteries. *Noûs*, 42(3), 466–481.
Goodman, J., & Salow, B. (2023). Epistemology normalized. *Philosophical Review*, 132(1), 89–145.
Greco, J. (2012). Better safe than sensitive. In K. Becker & T. Black (Eds.), *The sensitivity principle in epistemology* (pp. 193–206). Cambridge University Press.
Günther, M. (2024a). Epistemic sensitivity and evidence. *Inquiry*, 67(6), 1348–1366.
Günther, M. (2024b). Legal proof should be justified belief of guilt. *Legal Theory*, 30(3), 129–141.
Hirvelä, J. (2017). Is it safe to disagree? *Ratio*, 30(3), 305–321.
Hirvelä, J. (2019). Global safety: How to deal with necessary truths. *Synthese*, 196(3), 1167–1186.
Jenkins, C. S. (2006). Knowledge and explanation. *Canadian Journal of Philosophy*, 36(2), 137–164.
Kelp, C. (2018). *Good thinking: A knowledge first virtue epistemology*. Routledge.
Littlejohn, C., & Dutant, J. (2020). Justification, knowledge, and normality. *Philosophical Studies*, 177(6), 1593–1609.
Luper-Foy, S. (1984). The epistemic predicament: Knowledge, Nozickian tracking, and scepticism. *Australasian Journal of Philosophy*, 62(1), 26–49.
Manley, D. (2007). Safety, content, apriority, self-knowledge. *Journal of Philosophy*, 104(8), 403–423.
McEvoy, M. (2009). Safety, the lottery puzzle, and misprinted lottery results. *Journal of Philosophical Research*, 34, 47–49.
Mills, E. (2012). Lotteries, quasi-lotteries, and scepticism. *Australasian Journal of Philosophy*, 90(2), 335–352.
Miracchi, L. (2015). Competence to know. *Philosophical Studies*, 172(1), 29–56.
Mortini, D. (2024). The explanationist and the modalist. *Episteme*, 21(2), 371–386.
Nozick, R. (1981). *Philosophical explanations*. Harvard University Press.
Paterson, N. (2022). Safety and necessity. *Erkenntnis*, 87(3), 1081–1097.
Pritchard, D. (2007). Anti-luck epistemology. *Synthese*, 158(3), 277–297.
Pritchard, D. (2008). Knowledge, luck and lotteries. In V. Hendricks (Ed.), *New waves in epistemology* (pp. 28–51). Palgrave-Macmillan.
Pritchard, D. (2009). Safety-based epistemology: Wither now? *Journal of Philosophical Research*, 34, 33–45.
Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy*, 109(3), 247–279.
Roush, S. (2005). *Tracking truth: Knowledge, evidence, and science*. Oxford University Press.
Sainsbury, R. M. (1997). Easy possibilities. *Philosophy and Phenomenological Research*, 57(4), 907–919.
Silva, P. (2023). Merely statistical evidence: When and why it justifies belief. *Philosophical Studies*, 180(9), 2639–2664.
Smith, M. (2016). *Between probability and certainty: What justifies belief*. Oxford University Press.
Sosa, E. (1999). How to defeat opposition to Moore. *Philosophical Perspectives*, 13, 141–153.
Valaris, M. (2022). Normality, safety and knowledge. *Philosophy and Phenomenological Research*, 106(2), 394–401.

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press

Vogel, J. (2012). The enduring trouble with tracking. In K. Becker & T. Black (Eds.), *The sensitivity principle in epistemology* (pp. 122–151). Cambridge University Press.
Warfield, T. (2005). Knowledge from falsehood. *Philosophical Perspectives*, 19(1), 405–416.
Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.
Williamson, T. (2009). Replies to critics. In D. Pritchard & P. Greenough (Eds.), *Williamson on knowledge* (pp. 279–384). Oxford University Press.
Zalabardo, J. (2012). *Scepticism and reliable belief*. Oxford University Press.
Zhao, H. (2023). How to play the lottery safely? *Episteme*, 20(1), 23–38.
Zhao, H. (2024). Why better safe than sensitive. *Philosophy and Phenomenological Research*, 109(3), 838–855.
Zhao, H. (2026). Revisiting epistemic sensitivity and evidence. *Inquiry*, 1–13.
Zhao, H., & Baumann, P. (2021). Inductive knowledge and lotteries: Could one explain both 'safely'? *Ratio*, 34(2), 118–126.

Cite this article: Zhao, H. 2025. Solving the Lottery Problem: From Modal Accounts to Explanationism. *Canadian Journal of Philosophy* 55: 387–404, doi:10.1017/can.2026.10059

https://doi.org/10.1017/can.2026.10059 Published online by Cambridge University Press