DOI: 10.1017/psa.2025.10091

This is a manuscript accepted for publication in *Philosophy of Science*.

This version may be subject to change during the production process.

Low Prior Probability and Witness Reliability

Lydia McGrew[1]*

[1]Kalamazoo, MI, USA

**Corresponding author*: lydiamcgrew@gmail.com

Abstract

It is plausible that a witness's credibility is lowered if he testifies to an outcome that has low probability, but this is not always true. If the outcome has a low prior due only to its specificity (as in a fair lottery), witness credibility is not lowered by attesting to it. Even when witness credibility does drop as a result of the testimony, it can substantially confirm its contents, and agreement between independent witnesses can "give back" credibility that one witness lost. Formal modeling and analysis shed light on when each of these situations obtains.

1. Introduction

There is something plausible about the claim that, if a witness attests to a low-probability event, his doing so lowers his own credibility. (For purposes of this article "credibility" and "reliability" will be used interchangeably, both referring to a rational view of the witness's truth-telling that is responsive to the relevant evidence we have. The formal definition of "minimal reliability" is given in section 2.) This intuition is relevant in contexts such as contradictory testimony (Schubert 2012b) and the debate over testimony to miracles.

Robert Fogelin (2003, 6–13) argues that there is a "direct" and a "reverse" method for evaluating testimony. In the direct method, one checks such things as the witness's appearance of honesty and whether his interests are on the side of what he attests. If the results of these tests are favorable, one grants the event more probability in view of the testimony. In the reverse method, one downgrades one's evaluation of the witness's reliability to a greater or lesser extent depending on the prior improbability of what he has said. In cases of extreme prior improbability, the reverse method operates instead of the direct method. Fogelin illustrates this reverse method by a thought experiment about Henry, who tells a series of mildly surprising stories about meeting celebrities. As Henry's stories multiply, "[T]he sheer improbability that all (or most) of these things happened is sufficient for discounting his testimony" (ibid., 12; see also McGrew 2012).

> As the list gets longer and longer, we move into the area of the utterly implausible and our opinion of the credibility of Henry's testimony correspondingly sinks. Invoking the reverse method, we will eventually conclude that the probability that all…of these stories are true is so low that we would not believe them even if they were told to us by Cato. [W]e will not credit any single one of them if we have only Henry's word to go upon…. This…shows that the application of the reverse method depends upon the *improbability* that an event, or set of events could occur…. In the application of the reverse test of testimony, it is the extreme improbability, not the source, that matters. (Fogelin 2003, 11–12; emphasis in original)

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

Fogelin argues that this reverse method applies to testimony to any miracle. Similarly, J. L. Mackie (1982, 27) has argued that the nontheist should think that the "intrinsic improbability" of any miracle is so great that, if he receives a report of one, "One or other of the alternative explanations...will always be much more likely—that is, either that the alleged event is not miraculous, or that it did not occur, that the testimony is faulty in some way."

Countering this Humean argument, some defenders of miracles have historically turned to the analogy of a fair lottery, arguing that, no matter how large the lottery, the testimony "of a newspaper, or of any common man" (Price 1767, 410–11) is sufficient to establish that a particular number has won (see Earman 2000, 49–53). If this is a good analogy for testimony to a miracle, it seems that it should not count *at all* against the witness's reliability, for if it did, the mere *attempt* to report the result of a fair lottery would downgrade the witness's reliability.[1]

Low absolute prior probability can arise in multiple ways. In a large, fair lottery *any* given number has a low absolute probability of being drawn, even though we know that someone must win. Previous testimony against a particular outcome is another source of low prior probability. Metaphysical and theological considerations constitute another possible source; for example, they can dictate that miracles are rare, as admitted by some advocates of the miraculous (Robert Boyle in Colie, 1963, 214). How should these different sources affect whether, and how much, testimony to a low-prior event lowers our estimate of a witness's reliability? And do varying likelihoods among possible outcomes affect posterior reliability? My model provides a way to tackle these questions.

Fogelin's and Mackie's treatment of this issue is informal. In section 2, I propose a formal model for defining and updating reliability, and in sections 3–6 I use this model to provide insight into the relationship between the contents of testimony and negative and positive effects upon witness reliability. Section 7 surveys issues that call for further research.

---

[1] I owe this point about the mere attempt to Timothy McGrew.

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

2. Updating reliability

We must define a witness's reliability before we can talk about updating it. Following John Earman (2000, 55), I consider minimal reliability to lie in the fact that a witness's testimony to some outcome has a top-heavy Bayes factor for what he attests. Where H is a hypothesis and tH is testimony that H has occurred, the Bayes factor representing the force of the witness's testimony is P(tH|H)/P(tH|~H). (In this article, the hypotheses and testimonies modeled formally will concern the drawing of numbers in lotteries.) This concept of reliability can be naturally extended to mean that reliability is higher as this Bayes factor is more top-heavy, which will depend on both the numerator and the denominator.

Though this is a significant simplification, for the examples in this article I will assume that the witness must say something about H; silence is not an option (see McGrew and McGrew 2012, 47, 59). And I will treat P(tH|H) as the probability that the witness speaks truthfully about the event in question. On pain of incoherence, we cannot give two different values to the previously mentioned Bayes factor within the same probability distribution. Strictly speaking, then, when we speak of a change in the witness's reliability in this instance, this refers to updating the probability that the witness *has* told the truth or *has* spoken falsely in that instance, which can then influence both the numerator and the denominator of the Bayes factor in some new case. If we have a *prior* probability that the witness tells the truth in this instance, how can we update it based on the fact that he has attested to this content?

I take it as a given that a witness's track record of telling truth and falsehood should be relevant to our estimate of his reliability. All else being equal, evidence that a witness has spoken truly in some instance should be positively relevant to our future view of his reliability, and evidence that he has spoken falsely should be negatively relevant. That all else is rarely equal, due to issues of reference class (McGrew and McGrew 2012, 52–53; Venn 1888, 399, 403, 415–16; Zabell 1988, 332–34), does not change the fact that track records should have *something* to do with ongoing reliability estimates. I take it as given that a rational prior probability that a witness speaks truly is informed by relevant evidence that I do not model, including independent evidence about previous instances where he and/or others like him have testified. This could include records from video, forensic evidence, personal observation of earlier outcomes, and more. It should include any relevant evidence about how specific issues (biases, physical limitations, etc.) are likely to affect this witness. Wherever a rational prior probability of truth-telling comes from, it is important to know if and when testimony, in itself, changes the probability that the witness has spoken truly in that case. In this study I construct a model for updating that probability, so that it can be counted as part of the witness's track record. The hope is that that updated probability can be used to represent the probability that the witness speaks truly and falsely in a new case, if the cases are relevantly similar, always a difficult judgment (see section 7).

When evaluating the probability that a witness has spoken truly in a given instance, I propose that the Bayes factor should be P(tH|truth)/P(tH|falsehood), where "truth" is defined as "the witness speaks truly in this instance" and "falsehood" as "the witness speaks falsely in this instance." Note that this Bayes factor is different from the Bayes factor modeling the impact of the testimony on H.

There are some overlaps between my model of updating witness reliability and that of Bovens and Hartmann (2003, 56–75). In their model, as in mine, what we are updating is the reliability of the witness in the current instance (ibid., 17). And in cases of complete epistemic symmetry such as a fair lottery, their system and mine agree that the witness's reliability does not change (ibid., 60; see my section 3). But there are also differences between the models, where my model provides an advantage in efficiently and naturally getting at what is of most interest in actual cases. I will note these differences in the course of the article.

3. Equal priors, equal likelihoods

Suppose that W1 is announcing the winning number in a fair 100-ticket lottery. Suppose that the prior probability that W1 speaks truthfully concerning this lottery is .75 and that he speaks falsely is .25. W1 must make some specific report about the winning number. Suppose moreover that all possible outcomes are symmetrical in all epistemically relevant respects. As far as our evidence goes, W1 is no more likely to tell the truth in favor of or against any number. W1 is not infallible, but if all we know is that he speaks falsely, this tells us nothing about what number he is likely to report. He is, as far as our knowledge extends, no more likely to say one number falsely (or truly) than any other.

Let "1" stand for "The winning number is 1." Let "t1" stand for "W1 testifies that 1 is the winning number." Let "truth" stand for "The witness speaks truly in this case," and "falsehood" stand for "The witness speaks falsely in this case." Then,

$$P(t1|truth) = P(t1|falsehood) = P(1) = .01,$$

and the same for testimony to any other number. It follows that t1 is probabilistically irrelevant to the proposition that W1 has spoken the truth in this instance, leaving that probability exactly the same as its prior.

$$P(truth|t1) = P(truth) = .75 \text{ and}$$

$$P(falsehood|t1) = P(falsehood) = .25$$

In other words, despite the fact that the probability that 1 is the winning number is .01, which is rather low in absolute terms, this testimony does not confirm that W1 has spoken falsely because his testimony is no more probable given falsehood than given truth. And this would be the case for any finite lottery with the same symmetry conditions, no matter how large, even though a still larger lottery would mean that the outcome attested to would have a still lower absolute prior.[2]

When Bovens and Hartmann (2003, 57) model the impact (or nonimpact) of specific testimony upon witness reliability (which they call REL), they define "the witness is reliable" as meaning that the witness is perfectly reliable—P(tH|H & REL) = 1 and P(tH|~H & REL) = 0. A "reliable" witness on this definition is guaranteed to tell the truth. P(REL) would (correctly) not *change* in the present example in their system, but a

[2] Because there are 100 possible stories, this result is the same if the prior probability of 1's winning is 1/100, even if the prior probabilities among the other 99 options differ among themselves. If all other symmetries obtain (the probability that W1 speaks the truth or falsehood is the same for all options, and each option is 1/99 of the possible stories that could be told if he speaks falsely) testimony to a number with a .01 prior (where .01 is 1/n possible options) will neither confirm nor disconfirm the truthfulness of the witness. Proof omitted.

high prior probability for REL is incorrect in virtually all real-world examples, even with a good track record in hand, making REL practically irrelevant. In contrast, there are many cases in which, based on background evidence, we have a probability above .5 that an imperfect witness will testify truly in a given instance, and my formalism models this fact easily.

Nor does their model help to model and update the picture of the witness under the supposition that the witness is imperfect. They define ~REL as meaning that the witness is a randomizer over a number of possible stories, including the truth. Their randomization parameter (called *a*) must either ignore the nature of the witness or must combine the nature of the situation (e.g., how many entries in a lottery) with the nature of an imperfect witness (the witness's competence and desire to tell the truth). Stefan Schubert and Erik Olsson (2007) and Schubert (2012a, 2012b) construct similar systems in which reliability is defined as perfection and unreliability as evidential irrelevance. But in most real cases we are nearly certain that an imperfect witness's testimony bears some nonrandom relation to the truth. For example, he may be trying to speak truthfully but be limited by the imperfection of his senses or memory. Even if he is trying to avoid telling the truth, this is nonrandom. If *a* is just the total number of tickets, without any consideration of the imperfect witness's relation to truth, both P(truth) and P(t1|1) will usually be *wrong* because the product of P(~REL), which should realistically be high, and *a* will contribute heavily to both of these. If one tries to pack other considerations into *a*, no mechanism is provided for disentangling very different aspects of *a*. In my model, P(truth) is the true probabilistic complement of P(falsehood). Both are fully compatible with the assumption that the witness is imperfect, and both are *prima facie* relevant to the nature of the witness.

In the case under consideration here, t1 does confirm the proposition that 1 is the winning number.

$$P(1|t1) = P(1 \& t1)/P(t1) = .0075/(.0075 + .0025) = .75.$$

This confirmation arises from the fact that t1 has a top-heavy Bayes factor for its contents. Because W1 is bound to report, the numerator of the Bayes factor, P(t1|1), is simply the prior probability that he speaks truly, .75.

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

The calculation for the denominator of the Bayes factor, P(t1|~1), is a little more complex; it is not merely the probability of W1's reporting falsely, .25. If some number other than 1 is drawn, and W1 reports falsely, what is the probability that he will report specifically that 1 was drawn? Given that the situation is completely epistemically symmetrical, 1 is just one of 99 possible false stories, and there is no more reason to think that he will report *that* false story than any other. Randomization in my model thus arises entirely from the nature of the situation and from our having no reason to think that a false-speaking witness would choose one option over another. When the likelihoods are all equal (as in the examples in this section and section 4), the probability of the report given that the witness speaks falsely is P(~H)/n, where n is the number of epistemically equiprobable false stories the witness could tell. (On epistemic randomness, see McGrew and McGrew 2007, 148–53 and Kyburg, 1977.) Hence, P(t1|~1) = (.25)(1/99), or 1/396. This yields 297/1 odds in favor of the contents of W1's testimony. Putting that together with the prior odds of 1/99 yields posterior odds of 3/1, or a probability of .75. (See Holder 1998, 52, on the "many possible false stories" calculation. This is an improvement on de Condorcet, 1783, where the probability of the testimony given the falsehood of its contents is merely the complement of the probability given the truth of its contents, with no consideration of many possible false stories.)

This also means that in the posterior distribution, the other 99 possible numbers have each been significantly disconfirmed, for the posterior probability of their disjunction is .25, and each of them is just 1/99 of that .25, meaning that for any *specific* number other than 1, the posterior probability, conditional on testimony to 1, is a meager 1/396. Back-solving shows us that each of these specific numbers, having moved from a prior of 1/100 to a posterior of 1/396, has been disconfirmed by odds against it of $\frac{3.\overline{98}}{1}$, a point that will be important in section 4.

The confirmation of what W1 attests can thus be significantly greater than the confirmation of the proposition that W1 has told the truth. In this case, the posterior probability that W1 has told the truth is exactly the same as the prior, due to the symmetries involved. But the confirmation that the winning number is 1 is significant. The posterior probability of "W speaks the truth in this instance" is the same as the posterior probability of "1 is the winning number," but the former remains what it was while the latter rises to meet it. It is necessarily true that these posterior probabilities are identical, since the only way for W1's testimony to be true, given the contents of that testimony, is if 1 is the winning number.[3] This calculation shows that, contra Fogelin and Mackie, testimony to a low-prior proposition *in and of itself* does not disconfirm a witness's reliability. Moreover, contra Fogelin, the identity of the posterior probabilities (of the contents and of the proposition that W speaks truly) need not result from a reduction in the probability that W speaks truly. Even when the prior of the contents is low, this identity of the posteriors can result from a significant increase in the probability of the outcome.

4. Unequal priors, equal likelihoods

Despite the result in the previous section, it is difficult to dispel the intuition that there is *something* about testimony to a low-prior event that lowers credibility, at least a little. We can recognize this intuition formally by considering a case in which the prior probability of some possible outcomes is asymmetrically lower than others.

Suppose that we obtain testimony from another witness (W2) concerning the same 100-ticket lottery discussed in the previous section. For heuristic reasons, it is useful to imagine that W2's testimony comes in subsequent to that of W1, though order does not matter because these are Bayesian conditionalizations on the testimony, not Jeffrey conditionalizations. Suppose that the testimonies are independent in the sense that their contents, and the negation of their contents, screen off the reports from one another (Bovens and Hartmann 2003, 16). We can picture this concretely by stipulating that neither witness knows of the other's testimony.

Because W1 has already testified, in the new distribution, P*, the probability of "The winning number is 1," which is the prior probability vis á vis the testimony of W2, is .75. The following analysis also applies to a case where a lottery is not fair but is so

[3] A reviewer points out that this point also applies to cases in which the prior probability that the witness speaks truthfully is low, e.g., .1, and where the prior probability of the outcome is higher, e.g., .5, as when there are only two possible, symmetrical outcomes. The posterior probability of the contents of the testimony would then decrease to .1. *Note*: The fact that in this example there is only one true story and one false story is crucial.

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

weighted in favor of 1 that it has a prior probability of .75, while the disjunction of the remaining 99 numbers has probability .25.

Suppose that all other epistemic symmetries remain in place: W2 has no personal bias that we know of in favor of any number. There is a .75 probability that W2 speaks truly, and this probability applies to all 100 possible outcomes. Given that the specific content of the testimony is false, there is no reason to think that W2 speaks falsely in favor of one number rather than any other. This means that the Bayes factor calculated in the last section for the testimony of W1 is the same for W2—namely, .75/(.25)(1/99) or odds of 297/1.

But there is a crucial difference. Because W1 has already spoken, the proposition that 1 is the winner has a probability (for those who know about W1's testimony) of .75, and all other numbers have a probability of 1/396 (namely, 1/99 of .25). What difference does this make?

Suppose that W2 testifies in a way that directly contradicts W1. W2 says that 2 is the winner. This testimony confirms its own content but simultaneously confirms falsehood on the part of W2. It also lowers the probability that W1 spoke truly by lowering the probability that the winning number is 1. (The probability that W1 spoke truly will be identical to the probability that W2 spoke truly when both testimonies are taken into account.)

How is it possible for testimony simultaneously to confirm its own content and also false speaking on the part of the witness? This may seem a bit paradoxical, though there is some intuitive force to it as well. If someone whom you have good reason to trust testifies to something that seems highly implausible, this will make you wonder about your friend. For example, it might make you wonder if he is given to practical jokes. But it will also have some tendency to make you think that there just might be "something to" the content of his testimony.

The testimony of W2 confirms its own content by "taking away" some probability from all other numbers, especially the number 1. In the P* distribution, there is some probability that W2 will attest that 1 is the winner, but when he testifies to 2 instead, all of that probability is redistributed. Some of this change occurs by increasing the probability that W2 speaks falsely in this instance. Our estimate of W2's reliability drops

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

below .75 while the probability of “The winning number is 2” rises to meet it, and the probability of “The winning number is 1” is also disconfirmed to the point that it is equal to “The winning number is 2.”

The change in the credibility of W2 does not occur only because the proposition that 2 is the winning number has a low prior probability. In the P* distribution, “2 is the winning number” *does* have a low probability—1/396, as already stated. But that is not the crucial issue. As we saw in section 3, in a perfectly symmetrical lottery with 396 tickets, testimony to one outcome would not lower the estimate of the witness’s reliability. The problem for the number 2 in P* is that it is one of a class of not-1 numbers that have an *asymmetrically* low prior probability as compared with the number 1. All these not-1 numbers have a prior probability that is lower than 1/n, where n = 100.

The probability that W2 says “2 is the winning number,” given that W2 tells the truth, is 1/396. This is just the prior probability that 2 really is the winning number.

The probability of testimony that 2 is the winner, given that W2 speaks falsely, requires a somewhat more complex calculation. Given all the other symmetries,

$$P^*(t2 \ \& \ \text{falsehood}) = 395/396 \times (.25/99) \approx .002519$$

395/396 is the probability of ~2, i.e., that t2 is false. Because the probability of false testimony across the board is .25, and because 2 is just one of 99 possible false stories that might be told, the probability of t2 given ~2 is (.25/99). So, the probability of t2 and false testimony is the product of these two numbers. Thus,

$$P^*(t2|\text{falsehood}) = P^*(t2 \ \& \ \text{falsehood})/P^*(\text{falsehood}) = P^*(\sim 2) \times 1/99 \approx .002519/.25 \approx .01008$$

Because the probability of t2 given truth (1/396) is considerably lower than the probability given falsehood, falsehood is confirmed and truth-speaking disconfirmed.

A simple way to think of this is that, without knowing what W2 attests, we should have a considerably higher prior expectation that he will say that 1 won than that any other number won if he speaks truly because “The number 1 won” is more likely to *be* true than the statement that any other number won. The statement that any other number

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

won, given truth-speaking, is especially low because 1 is "hogging" so much of the truth-speaking probability.

In fact, W2's testimony disconfirms both his own truth-speaking and the specific number 1 by the very same odds of $\frac{3.\overline{98}}{1}$ that we saw in the previous section. Because both of these start out in P* with a probability of .75, their probability after both testimonies is approximately .429.

At the same time, in P* we have no more reason to think that, given ~2, W2 will testify falsely to the number 2 than to any other number. The probability of t2, given that 1 is the true number drawn, is

$$1/99 \times .25 = .00\overline{25}$$

and the same if any of 3-100 is the winning number. This means that the probability of W2's testimony to 2 given the falsehood of its own contents (i.e., a false positive) is $.00\overline{25}$, just as it was for the testimony of W1 to the number 1. In other words, $P(t2|{\sim}2) < P(t2|\text{falsehood})$ because $.00\overline{25} < .01008$. So even though truth-speaking in this instance is somewhat disconfirmed, W2's stated number (2) is confirmed by the very same 297/1 odds by which the testimony of W1 confirmed the number 1. In this case, of course, the number 2 starts at a great disadvantage, due to the previous testimony of W1.

When the testimony of W2 is taken into account and the number 2 is confirmed, it ends up at the same point as the number 1: $P^*(2|t2) \approx .429$. The estimate of W2's reliability falls from .75 while the probability of what W2 attests rises to meet it. After both testimonies, the probability that 1 is the winning number is the same as the probability that 2 is the winning number, and the probability that each witness spoke truly is also approximately .429. This is both the calculated and the intuitively correct result.

This result highlights another advantage of my model: Both the probability that W1 speaks truly about this lottery and that W2 does so have changed, as they should. We should be able to change our minds, on the basis of new evidence, about the probability that a witness is both honest and competent in a given case. In Bovens and Hartmann's model (2003, 58–75), the only place where such considerations could possibly be taken

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

into account, for an imperfect witness, is the randomization parameter *a*, but they provide no mechanism for updating *a*.

Both of the preceding results yield epistemic insight. The result in section 3 illustrates the effect of a story that is improbable *only* in virtue of its specificity. Any other equally specific story would be equally improbable. Such a story should not reduce our estimate of a witness's reliability, for any state of affairs described in sufficient detail will be improbable.

In contrast, testimony to a result that is improbable "by type"—that is, *asymmetrically improbable in comparison to other possible outcomes*—has a negative effect on the witness's credibility. (As we shall see in later sections, the strength of this effect will vary widely, depending on specifics.) Where there are n possible outcomes, any outcome with a prior probability less than 1/n is asymmetrically improbable. Such an outcome will be not only relatively improbable (compared to others) but also absolutely improbable in the sense of having a probability less than .5, even if there are only two possible outcomes; the absolute improbability will fall increasingly below .5 to the extent that n > 2.

Suppose that I state that I recently met a neighbor named June walking a collie named Daisy. Given ordinary background information about my location in space and time, this is about as probable as that I met a neighbor named June walking a Pekingese named Ralph. These are both examples of the kind of thing that one might easily see on a walk, though they differ in particulars. But if I say that I met a neighbor named June walking an elephant named Daisy, this is asymmetrically improbable as compared with the kinds of things that one would normally see on a walk.[4]

[4] A reviewer raises the following objection: Suppose that lottery tickets 1–100 are all equiprobable but that 1–98 are blue while 99 and 100 are red. W reports that ticket 99 won. Because he is reporting a specific number, the fact that this is one of only two red tickets seems irrelevant. Does this mean that speaking of improbability by type is misleading? I answer that the content of W's testimony is not improbable by type even though it entails that the ticket is one of only two that are red. The priors of all specific numbers are equal, and if we assume either that W is unaware of the setup or that his testimony to the number is uninfluenced by the color, the situation is epistemically symmetrical. The numbers 99 and 100 would be improbable by type if *each* red ticket had a prior probability less than 1/100. A sighting of an exotic animal in an American neighborhood has an asymmetrically low prior compared to a single instance of an ordinary type of neighborhood sight. The phrase "improbable by type" is meant to capture these concepts.

In this way we come back to the issue of testimony to a miracle. It is intuitively plausible that miracles are *relatively* uncommon. Even some Christians would say that a miracle requires a stable backdrop of natural order to stand out as a sign (Robert Boyle in Colie, 1963, 214). Fogelin, in explaining his "reverse method" for evaluating testimony, does contrast a "common sort of event" with a "perfectly fantastic event" (2003, 10), but the only probabilistic explanation he gives of the latter is "sheer improbability" (ibid., 12), which as we have seen is wrong; he also fails to realize that, even when the prior probability is *very* low, the witness's testimony can confirm its content even when his own credibility is reduced by a "reverse" effect. The idea that miracles are improbable by type, as defined here, provides a better understanding than bare references to low prior probability.

This analysis serves as a caution to advocates of miracles: A specific result of a large, fair lottery is improbable, but the analogy to a miracle is imperfect, for it does not account for the intuition that testimony to a miracle plausibly causes the witness's credibility to take a hit to at least some degree. I am not saying that *any* miracle is more improbable a priori than *any* nonmiraculous event. In a specific context, there could be other evidence that would raise the prior probability of a miracle performed by a specific person. But I am suggesting that advocates of miracles should not assume that they can capture all of the negative probabilistic considerations by the analogy of a large enough fair lottery.

Yet even events that are improbable "by type" lie within the reach of testimony (Earman 2000), as later examples in this study show.

5. Equal priors, unequal likelihoods

It might seem that the discussion thus far highlights a mere academic curiosity. It is a bit counterintuitive to say that it matters *why* the contents of testimony have a low prior probability. But when we say that a person's claim to have been abducted by aliens is improbable, we tacitly mean that (based on our previous experience) alien abductions as such are unlikely; the claim does seem to have an asymmetrically low prior, and this is often intuitively the case, including with reports of miracles. One may wonder how often

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

the concept of improbability merely in virtue of specificity will help the advocate of an improbable event.

But the insight that testimony can simultaneously decrease our estimate of a witness's reliability and significantly increase the probability of its contents is *widely* applicable, including in cases of reports of miracles. When we turn to a more complicated situation in which possible stories have different likelihoods, this insight and the methods used in the previous sections yield even more interesting results.

Suppose that W3 will announce the winning ticket in a 100-ticket lottery. All numbers have an equal prior probability of being drawn, but there is a significant asymmetry in the probability that W3 will testify to a given winner. W3 is the sister of the person who owns ticket 1, and W3 knows that her brother owns that ticket. Making matters more difficult for her, her brother has no commitment to truth but has a great interest in winning the lottery, and he has uttered threatened to harm her if she doesn't announce his number. However, W3 may be caught and suffer negative legal consequences if she lies about the ticket number. She also has personal scruples about lying. So, her motives are mixed, but she has some self-interest in announcing that ticket 1 is the winner, even if it is not. She also has an interest in avoiding a mere error that would cause her to announce some other number, if ticket 1 really is the winner. If she can truthfully announce that ticket 1 won, this will satisfy her brother, the law, and her conscience simultaneously. We know of nothing else that makes the false negative rate or false positive rate differ among the numbers; they are epistemically symmetrical in other respects.

Let t1 represent testimony that 1 is the winning number and let 1 represent the proposition "Ticket 1 is the winning number." Let tN represent testimony that some other *specific* number (not 1) won, without specifying what that number is. Let N represent the fact that some other specific number won, while not specifying what that number is. (This notation helps to avoid confusion between N and ~1. The latter is the disjunction of *all* possibilities other than 1.) As before, I make the substantive assumption that W3 cannot remain silent. Let the true positive rate for 1 differ from the true positive rate for any other number of the 99 other possibilities, such that

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

$P(t1|1) = .85$

$P(tN|N) = .75$

I have kept this difference fairly small because W3 is not infallible and I am assuming that some false reports (in favor of any option) result from errors that are hard for W3 to guard against. But there is more pressure for her to lie if 1 is not the winning number and more pressure for her to get it right if possible to announce truly that 1 did win. Therefore,

$P(\text{falsehood}|1) = .15$

$P(\text{falsehood}|{\sim}1) = .25$

The overall probability that W3 tells the truth about this lottery, prior to our receiving her testimony, is

$(.01 \times .85) + (.99 \times .75) = .751$

And because W3 must speak, the probability of false speaking is the complement,

$1 - .751 = .249$

If 1 did not win and W3 speaks falsely, she is more likely to do so in favor of the number 1 than any other number because that is where her interest lies. Thus, to the extent that false speaking is deliberate rather than a mere error, it favors false speaking in favor of the number 1. I model this as follows:

$P(t1|{\sim}1 \ \& \ \text{falsehood}) = .4$

$P(tN|{\sim}1 \ \& \ \text{falsehood}) = (.6) \times 1/98$

The latter is the probability, given that ticket 1 did *not* win and that W3 speaks *falsely*, that she announces some specific number that is not 1. The rationale here is that if 1 did not win, some other specific number won. If W3 speaks falsely, she does not testify to the number that won; if she says that some nonwinning, non-1 number won, she will testify to one of 98 possible false stories. (I note in passing that Bovens and Hartmann provide no guidance for modeling cases of this kind.)

Because all other numbers are epistemologically symmetrical, we have

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

$$P(tN|1 \text{ \& falsehood}) = 1/99 = .\overline{01}$$

If 1 won and W3 falsely announces that some other number won, there is no reason to expect her to announce any one false number more than any other, and in that case there are 99 possible false stories.

With all this in place, if W3 does announce that her brother's number won, this disconfirms her own truthfulness.

$$P(t1|\text{truth}) = P(t1 \text{ \& truth})/P(\text{truth}) = .0085/.751 \approx .0113$$

$$P(t1 \text{ \& falsehood}) = .99 \times .25 \times .4 = .099, \text{ so}$$

$$P(t1|\text{falsehood}) = P(t1 \text{ \& falsehood})/P(\text{falsehood}) = .099/.249 \approx .39759$$

So t1 is more probable given falsehood than given truth.

The overall probability of t1 is

$$.099 + .0085 = .1075$$

The overall probability that W3 has spoken truth if she testifies to 1 as the winner is

$$P(\text{truth}|t1) = P(t1 \text{ \& truth})/P(t1) = .0085/.1075 \approx .0791,$$

which is a quite dramatic decrease from the prior probability of .751.

As before, the same calculation gives us the probability that 1 was really the winning number. The probability that 1 won rises, on the basis of the testimony of W3, from its former .01 to meet the new probability that W3 speaks truly in this case.

$$P(1|t1) = P(1 \text{ \& } t1)/P(t1) \approx .0791$$

Though higher than the prior, this is not a terribly exciting posterior.

What if, contrary to her own interest, W3 testifies that some number other than her brother's—say, two—is the winner? First, consider whether this testimony confirms

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

that she has told the truth. I will use t2 to indicate testimony that the number two won and 2 to indicate that the number two won.

The calculation for the probability of t2 given truth is

$P(t2|truth) = P(t2 \& truth)/P(truth) = .0075/.751 \approx .00999$

The calculation for P(t2|falsehood) is somewhat more complex.

$P(t2 \& 1) = .01 \times (.15 \times 1/99) = .0000\overline{15}$

$P(t2 \& {\sim}2 \& {\sim}1) = 98/100 \times .25 \times (.6 \times 1/98) = .0015$

The sum of these is $.00\overline{15}$. Therefore,

$P(t2|falsehood) = .00\overline{15}/.249 \approx .006085$

In this case, the probability of W3's testimony given that she speaks truthfully is scarcely less than .01, while the probability of her testimony given falsehood is notably lower; her testimony against interest does raise (though not dramatically) the probability that she has spoken truly.

But the effect of her testimony on the probability that 2 was the winner is very noticeable. The overall probability of t2 is

$.0075 + .00\overline{15} \approx .009015$, so

$P(2|t2) = P(2 \& t2)/P(t2) \approx .0075/.009015 \approx .83195$,

which is a striking improvement on the prior of .01. Although in any real-life case everything will depend on specifics, this case illustrates the power of testimony against interest.

6. Unequal priors, unequal likelihoods

The results in section 5 implicitly introduce the notion of an offset. The witness was especially likely to be truthful if her brother's number was drawn and was tempted to lie

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

if some other number was drawn. Hence the probability of truth-speaking given that her brother's number was drawn was greater than if some other number was drawn. If all else were equal, this consideration would mean that testifying to her brother's number as opposed to any other would raise the probability that she spoke truly. But the witness was *much* more likely to testify to her brother's number than to any other number if she spoke falsely, which more than offset the slight advantage of her being more likely to speak truthfully if her brother's number was drawn. Her testimony to her brother's number thus would confirm falsehood (disconfirm truth) quite dramatically. In contrast, her testimony against interest would rather dramatically confirm its own contents and would modestly confirm her own truthfulness.

How might the concept of an offset affect a case in which the content of testimony has an asymmetrically low prior but where testifying to that content is against one's self-interest? As we saw in section 4, the former consideration (testimony to a number with an asymmetrically low prior) *taken by itself* would tend to disconfirm the witness's truthfulness. But in section 5 we saw that the fact that testimony is against interest tends to confirm the witness's truthfulness. And in both cases, the testimony can confirm its own contents.

Suppose that the number 1 has an asymmetrically high prior probability of .75, as in section 3. (This could arise either because we have other testimony that 1 was the winning number or due to an asymmetry in the lottery setup.) Suppose also that it is in the witness's interest to testify to the number 1. We will call the witness here W4. This special interest in testifying to 1 will be modeled by the likelihood considerations exactly as in section 4. The probability that W4 speaks truthfully is higher (.85) if 1 is drawn than if any other number is drawn (.75). If some other number is drawn and W4 speaks falsely, the probability that she testifies to 1 is .4, while all the other possible false testimonies are crowded into .6 of the falsehood space.

With all this in place, what happens if W4 testifies against interest to one of the low-prior outcomes, say, to the number two, which has an asymmetrically low prior of 1/396? The short answer is that the asymmetrically low prior in this particular setup, combined with the somewhat higher probability that the witness speaks truthfully if 1 is drawn, outweighs the fact that she is testifying against interest *just in the sense that* the

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

testimony disconfirms her truthfulness. The prior probability that W4 speaks truly in this case is .825, due in no small part to the high prior probability that 1 is the winner and the fact that the witness is especially likely to speak truthfully if the number 1 is drawn. As always, the posterior probability that W4 has spoken the truth is identical to the posterior probability of the contents of the testimony. Here this amounts to a disconfirmation of truthfulness from its prior of .825 to about .5556. (The details in this scenario are so complex that I have placed them in the appendix.)

But the fact that W4 is testifying against interest outweighs the asymmetrically low prior *in a different sense*: The posterior probability of the number 2, given W4's testimony that 2 won, is necessarily also about .5556, which is a fairly impressive confirmation of the contents. Even more important, when W4 testifies that the number 2 won, her brother's number is disconfirmed to the point that its posterior probability is *less than* the probability that 2 won. From a prior probability of .75, the drawing of the number 1 is disconfirmed by W4's testimony against interest to approximately .3332.

The permutations, of course, are endless once we are varying both priors and likelihoods among possible outcomes. It thus becomes especially important to stress that, even when W4's credibility takes a hit due to testimony to an especially low-prior outcome, this is compatible in principle with a *very* high posterior probability that W4 speaks truly in this instance. When the high-prior option (the number 1) is the brother's number, truthful speaking has an especially high prior probability. This is due to the fact that we expect W4 to be especially motivated to testify truly if that number is drawn, and that that number is also especially likely to be drawn. This makes it, in a sense, "easier" to disconfirm truth by testimony to some other number because truth starts with such a high probability due to the special circumstances of the case. These factors also contribute to the low false positive rate for other outcomes because if the number 1 is drawn, W4 is very unlikely to speak falsely, and therefore if she does speak falsely the probability of her testifying to any particular other outcome is very low.[5] Under different

[5] I do not mean to imply that it is impossible to construct a distribution in which testimony to a number with an asymmetrically low prior, against interest, confirms both witness truth-speaking *and* its own contents simultaneously. But such a distribution would be even more complex. A crucial feature of such a distribution, if "interest" is modeled as it is here, would be disassociating the "interest" number from the high-prior number. E.g., the number with an especially high prior could be 1 while the brother's number could be 3, one of the numbers with an especially low prior. Proof omitted.

probabilistic assumptions, the number 2 could have *such* a low false positive rate that the probability of truthfulness would be still very high, conditional on testimony to 2. This makes the disconfirmation of truth-speaking a somewhat less urgent matter than one might think.

The effect of W4's testimony against interest also sets the stage for us to gain further information if there are other independent witnesses to the number drawn. As we saw above, after W4's testimony that 2 is the winner, 2 is no longer a disfavored outcome. In fact, its probability, though a modest .5556, is higher than that of any other outcome, including the previously favored 1. This means that further, independent testimony that 2 was the winner from a credible witness can restore W4's credibility. In this scenario, if there is a W5, with no particular axe to grind and with (say) a .75 probability of speaking the truth, who testifies independently that 2 is the winning number, this would more than restore the credibility of W4. In fact, the probability that W4 spoke the truth after testimony from W5 in favor of 2 would be approximately .9973. (See the appendix.)

7. Remaining issues

There are other important issues in modeling witness testimony and reliability that lie beyond the scope of this article. One is the relationship between induction and Bayesian models in coming up with an initial probability that a witness tells the truth in a given situation, which can then be updated. Where does that initial probability come from? To the extent that it incorporates a track record based on sampling a witness's statements, it appears to be inductive, that is non-Bayesian (see Kyburg 1977 on direct inference). And taking the posterior probability that the witness has spoken truly in one instance to be the prior that he speaks truly in another instance, even a similar one, is not a Bayesian update. So, the Bayesian needs induction. In induction, the issue of the right reference class looms large: How should we use information about probable truth and falsehood in one instance (as modeled here) to project probable truth and falsehood in a scenario that differs in some relevant respects?

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

Because an inductive track record of truthfulness requires information about individual instances, and this study has shown that Bayesian updating can be intuitively applied to combine evidence about truth-telling in an individual case, it seems that the inductive reasoner also needs Bayesianism. We must continue to seek a way to combine inductive and inverse inference into a single, rational model.

Yet another issue is the modeling of silence. It is often incorrect to treat a person's silence as testimony to *any* outcome, and it is often simply false that a witness must say something (McGrew and McGrew 2012, 59). Silence may indicate epistemologically unproblematic ignorance, a taciturn personality, lack of interest in a topic, tact or prudence, true or false belief about an alleged event, and more. The calculations here have been simplified by treating the probability that a witness does not speak falsely as equivalent to P(tH|H). If silence were an option, this assumption would be wrong. The possibility of silence would also further complicate the (already complicated) calculation of P(tH|~H). That probability might be even lower if the witness can remain silent, or he might be especially likely to remain silent rather than lying about specific ~H options.

One more issue that deserves more thought is the legitimacy of the "many possible false stories" assumption that has been so prevalent in the examples here. At times we have only a vague idea that there are many things that a speaker could say if speaking falsely, but no clear idea of how many of them there are or whether they are true alternatives. Real-life scenarios are seldom as probabilistically well-behaved as lottery announcements.

Despite these remaining issues, this study represents progress on the question of when and how testimony confirms or disconfirms truth-speaking in a given report. I have nuanced and partially endorsed the intuition that testimony to a proposition with a high prior should increase our confidence in the speaker and that testimony to a low-prior outcome should decrease the speaker's credibility.

But there is a potential pitfall here: We might be tempted to use these results to bolster a mere confirmation bias, whereby we believe testimony that confirms our previous beliefs and disbelieve testimony that challenges them. This would be a misapplication of the formalism. Because testimony to a low-prior proposition can raise the probability of that proposition, while decreasing our confidence that the witness

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

spoke truly, the testimony can rightly cause us to reconsider our inclination to disbelieve the proposition. Even if the prior probability of a proposition is well below .5, the posterior can be above .5, even on the testimony of a single witness. And as others have noted (Earman 2000, 53–56; Babbage 1838, 192–203; Holder 1998, 50–53), and as the second example in the appendix shows, multiple, independent testimonies to the *same* low-prior proposition can raise its probability to something very high indeed.

In light of this study, we can see that agreeing testimony can "give back" all the credibility (and more) that a witness lost by being the first to testify to an event with a low prior. In this way, the formalism reveals the truth implicit in the common inclination to discredit witnesses who testify against our previous beliefs, the danger it represents, and a remedy for that danger.

Appendix

1. W4's testimony against interest to a low-prior outcome

Let t1 represent testimony that 1 is the winning number. Let 1 represent "1 is the winning number." Let t2 represent testimony that 2 is the winning number. Let 2 represent "2 is the winning number."

$P(1) = .75$

Probability of the disjunction of all other 99 numbers $(\sim 1) = .25$

All other 99 numbers, including 2, have the same prior probability of being the winning number.

$1/99$ of $.25 = 1/396$

It is in the interest of W4 to testify that 1 is the winning number. It is not in the interest of W4 to testify that any other number is the winning number, and all other numbers are, as far as we know, symmetrical in this respect. This situation is modeled as follows:

$P(\text{truth}|1) = .85$

$P(\text{truth} \,\&\, 1) = .85 \times .75 = .6375$

$P(\text{falsehood}|1) = .15$

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

P(t2|1 & falsehood) = 1/99 of .15 = $.00\overline{15}$ (and the same for testimony to any other non-1 number)

P(t2 & 1) = .75 × $.00\overline{15}$ = $.0011\overline{36}$

P(truth|2) = .75 (and the same for any other non-1 number)

P(falsehood|2) = .25 (and the same for any other non-1 number)

P(truth & ~1) = .75 × .25 = .1875

P(t2 & 2) = 1/396 × .75 ≈ .001894

P(truth) = .1875 + .6375 = .825

P(falsehood) = 1 - .825 = .175

P(t1|2) = .4 (and the same given any other non-1 number)

The probability that the witness testifies falsely to some number other than 1, given that some (other) number other than 1 is drawn = .6. Therefore,

P(t2|~2 & ~1) = 1/98 of .6 ≈ .006122448 (and the same for any other non-1 number)

Rationale: If the number drawn is not 1 and is also not 2, and if the witness testifies that 2, the witness is speaking falsely. If a number other than 1 is drawn and the witness testifies to some number other than the true number but does not say that 1 is the winning number, there are 98 possible false numbers to which the witness could attest. We have no reason to believe that the witness prefers one of these to any other. Therefore, the probability of testifying to 2 is 1/98 of the .6 probability that the witness speaks falsely by saying some number other than 1.

P(t2 & ~2 & ~1) = (98/99 × .25) × (.25 × (.6 × 1/98)) ≈ .000379

Rationale: The prior probability of (~2 & ~1) is 98/99 of .25. (The prior probability of ~1—that is, the disjunction of all possible outcomes other than 1—is .25. A total of 99 possible outcomes, 2–100, occupy this probability space. If 2 is *also* not the number drawn, the prior probability is 98/99 of that .25 prior space.) The probability of falsehood

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

given that neither 1 nor 2 is the winning number is .25 (because it is the probability that the witness speaks falsely if any number other than 1 is drawn). The probability of t2 given that neither 2 nor 1 is the winning number is 1/98 of .6, for the reasons explained above. So,

$$P(t2) \approx .001894 + .001136 + .000379 = .003409$$

$$P(t2|truth) = P(t2 \& truth)/P(truth) \approx .001894/.825 \approx .0022958$$

$$P(t2|falsehood) = P(t2 \& falsehood)/P(falsehood) \approx .001515/.175 \approx .008657$$

The probability of t2 given falsehood is higher than the probability of t2 given truth. Hence, t2 confirms that W4 has spoken falsely in this instance. But

$P(2|t2) = P(t2 \& 2)/P(t2) \approx .001894/.003409 \approx .5556$ (This is also the proper calculation for $P(truth|t2)$.)

Hence, t2 confirms 2 from a prior of 1/396 to a posterior of approximately .5556 and disconfirms truth-telling (in this instance) from a prior of .825 to a posterior of approximately .5556.

Testimony to 2 also disconfirms 1.

$$P(1|t2) = P(1 \& t2)/P(t2) \approx .001136/.003409 \approx .3332$$

2. Restoring W4's credibility (and then some)

Let t2* be testimony from an independent W5 that 2 is the winning number after W4 has testified. W5 has no particular interest in testifying to any number and does have some credibility for what he attests. The probability that W5 speaks truthfully about the outcome is .75 for all numbers. If W5 speaks falsely, he is no more likely to name any of the 99 possible false stories than any other. Call the distribution induced by W4's testimony P_{W4}. So that the relevant numbers sum to 1, we will dispense with the approximations at the outset.

$$P_{W4}(1) = .3332$$

$$P_{W4}(2) = .5556$$

$$P_{W4}(\sim 1 \& \sim 2) = .1112$$

https://doi.org/10.1017/psa.2025.10091 Published online by Cambridge University Press

$P_{W4}(2 \text{ \& } t2^*) = .5556 \times .75 = .4167$

$P_{W4}(1 \text{ \& } t2^*) = .3332 \times .25 \times 1/99 = .000841413$

$P_{W4}(t2^* \text{ \& } {\sim}1 \text{ \& } {\sim}2) = .1112 \times .25 \times 1/99 \approx .00028$

Therefore, $P_{W4}(t2^*) = .417821413$

Therefore, $P_{W4}(2|t2^*) = P_{W4}(2 \text{ \& } t2^*)/P_{W4}(t2^*) = P_{W4}(\text{truth}|t2^*) = .4167/.417821413 \approx .9973$

This is the probability both that W5 spoke the truth and that W4 spoke the truth. Therefore W5's independent testimony more than restores the credibility of W4 that was lost by attesting to a low-prior outcome.

Competing interests. None to declare.

Funding. None to declare.

References

Babbage, Charles. 1838. *The Ninth Bridgewater Treatise: A Fragment*. London: John Murray.

Bovens, Luc, and Stephan Hartmann. 2003. *Bayesian Epistemology*. Oxford: Oxford University Press. https://doi.org/10.1093/mind/fzi394.

Colie, Rosalie. 1963. "Spinoza in England: 1665–1730," *Proceedings of the American Philosophical Society* 107 (3)183–219. https://doi.org/10.1038/200407d0.

De Condorcet, Nicolas. 1783. "Sur le probabilité des faits extraordinaires." In *Histoire de l'Académie Royale des Sciences*. Paris: DuPont.

Earman, John. 2000. *Hume's Abject Failure.* Oxford: Oxford University Press. https://doi.org/10.1093/0195127382.001.0001.

Fogelin, Robert J. 2003. *A Defense of Hume on Miracles*. Princeton: Princeton University Press.

Holder, Rodney D. 1998. “Hume on Miracles: Bayesian Interpretation, Multiple Testimony, and the Existence of God.” *The British Journal for the Philosophy of Science* 49 (1):49 –65. https://doi.org/10.1093/bjps/49.1.49.

Kyburg, Henry E. 1977. “Randomness and the Right Reference Class.” *Journal of Philosophy* 74:501–21. https://doi.org/10.2307/2025794.

Mackie, J. L. 1982. *The Miracle of Theism: Arguments for and against the Existence of God*. Oxford: Clarendon Press.

McGrew, Lydia. 2012. “Tall Tales and Testimonies to the Miraculous.” *European Journal of Analytic Philosophy* 8:39–55.

McGrew, Lydia, and Timothy McGrew. 2007. *Internalism and Epistemology: The Architecture of Reason*. Abingdon: Routledge.

McGrew, Lydia, and Timothy McGrew. 2012. “The Reliability of Witnesses and Testimony to the Miraculous.” In *Probability in the Philosophy of Religion*, edited by Jake Chandler and Victoria Harrison, 46–63. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199604760.001.0001.

Price, Richard. 1767. *Four Dissertations*. London: A. Miller and T. Cadell.

Schubert, Stefan. 2012a. “Coherence Reasoning and Reliability: A Defense of the Shogenji Measure.” *Synthese* 187:305–19. https://doi.org/10.1007/s11229-010-9856-6.

Schubert, Stefan. 2012b. “Is Coherence Conducive to Reliability?” *Synthese* 187:607–21. https://doi.org/10.1007/s11229-010-9865-5.

Schubert, Stefan, and Olsson, Erik J. 2007. “Reliability Conducive Measures of Coherence.” *Synthese* 157:297–308. https://doi.org/10.1007/s11229-006-9056-6.

Venn, John. 1888. *The Logic of Chance*. New York: Chelsea Publishing Co.

Zabell, S. L. 1988. “The Probabilistic Analysis of Testimony.” *Journal of Statistical Planning and Inference* 20:327–54. https://doi.org/10.1016/0378-3758(88)90096-1.