EuJAP | Vol. 8 | No. 2 | 2012
1 Bayes, T.
1 Hume, D.
231.74
165.43

TALL TALES AND TESTIMONY TO THE MIRACULOUS

LYDIA McGREW
Kalamazoo, Michigan

ABSTRACT

In the debate over testimony to miracles, a common Humean move is to emphasize the prior improbability of miracles as the most important epistemic factor. Robert Fogelin uses the example of Henry, who tells multiple tall tales about meeting celebrities, to argue that low prior probabilities alone can render testimony unbelievable, with obvious implications for testimony to miracles. A detailed Bayesian analysis of Henry's stories shows instead that the fact that Henry tells multiple stories about events that occurred independently if they occurred at all is crucial to his loss of credibility. The epistemic structure is similar to that of a case of multiple lottery wins by the same person. Each of Henry's stories can confirm only one event, but all the stories confirm the hypothesis that Henry is a liar. This structure does not apply to testimony to just one event, however antecedently improbable. Such examples therefore do nothing to undermine a standard Bayesian analysis involving both priors and likelihoods in evaluating testimony to an improbable event.

Keywords: Bayesian probability, miracles, lotteries, Hume

1. Prior Probabilities and Testimony to Miracles

Ever since David Hume delivered in *Of Miracles* what he called an "everlasting check to all kinds of superstitious delusion" (Hume 1748/2000, 83), a major focus in the dismissal of testimony to the miraculous has been the established prior improbability of miracles. Hume set the stage by example:

> A miracle is a violation of the laws of nature; and as a firm and unalterable experience has established these laws, the proof against a miracle, from the very nature of the fact, is as entire as any argument from experience can possibly be imagined. (Hume 1748/2000, 86-7)

And he was not shy about drawing conclusions about the implications for testimony to the miraculous:

> When anyone tells me, that he saw a dead man restored to life, I immediately consider with myself whether it be more probable, that this person should either deceive or be deceived, or that the fact, which he relates, should really have happened. I weigh the one miracle against the other; and according to the superiority, which I discover, I pronounce my decision, and always reject the greater miracle. If the falsehood of his testimony would be more miraculous, than the event which he relates; then, and not till then, can he pretend to demand my belief or opinion. (Hume 1748/2000, 87-8)

Nor is there any question as to whether Hume thought that such a burden of "greater miraculousness" would ever be discharged. The Marquis de Laplace put the Humean point more pithily, though with less nuance: "There are things so extraordinary," he said, "that nothing can balance their improbability". (Laplace 1840/1951, 119)[1]

In contemporary times, J. L. Mackie has defended the general Humean position, while considering that Hume's case needs to be improved upon at points. Mackie emphasizes the impact of one's prior theistic or atheistic beliefs on the credibility of miracle claims.

> [W]e should distinguish two different contexts in which an alleged miracle might be discussed. One possible context would be where the parties in debate already both accept some general theistic doctrines, and the point at issue is whether a miracle has occurred which would enhance the authority of a specific sect or teacher. In this context supernatural intervention, though *prima facie* unlikely on any particular occasion, is, generally speaking, on the cards: ...But it is a very different matter if the context is that of fundamental debate about the truth of theism itself. Here one party to the debate is initially at least agnostic, and does not yet concede that there is a supernatural power at all. From this point of view the intrinsic improbability of a genuine miracle ... is very great, and 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.
>
> This entails that it is pretty well impossible that reported miracles should provide a worthwhile argument for theism addressed to those who are initially inclined to atheism or even to agnosticism.

1 I do not mean to enter into the purely historical controversy over whether a full Bayesian reconstruction of Hume's position is true to Hume's intentions, which would lie beyond the scope of this paper. My intent is to connect the discussion here to its context in the philosophy of religion and to show that Humeans such as Mackie and Fogelin who emphasize the antecedent improbability of miracles – which Bayesians would refer to as a low prior probability – are following the general trend of Hume's thought. For strenuous objections to any Bayesian modeling of Hume's argument, see Levine (2010). For a different view, see Earman. (2000, 43-8)

> … Not only are such reports unable to carry any rational conviction on their own, but also they are unable even to contribute independently to the kind of accumulation or battery of arguments referred to in the Introduction. To this extent Hume is right, despite the inaccuracies we have found in his statement of the case. (Mackie 1982, 27)

More recently, authors such as Rodney Holder (1998, 52ff) and John Earman (2000, 53-64), using the tools of Bayesian analysis, have shown that the Humean approach is facile, since any finite prior probability, however low, can in principle be overcome by testimonial evidence. Earman's conclusion is withering:

> In "Of Miracles," Hume pretends to stand on philosophical high ground, hurling down thunderbolts against miracle stories. The thunderbolts are supposed to issue from general principles about inductive inference and the credibility of eyewitness testimony. But when these principles are made explicit and examined under the lens of Bayesianism, they are found to be either vapid, specious, or at variance with actual scientific practice. (Earman 2000, 70)

Yet it may seem that the Bayesian analysis must be missing something. It seems reasonable for Mackie to suggest that our independent evidence that some type of event is not "on the cards" plays a huge role in our evaluation of specific evidence that seems to favor the event. In charity, we should assume that Mackie means us to take it that the atheist or agnostic is an atheist or agnostic for reasons and not merely arbitrarily. And if this is so, then such a person has *ipso facto* reason to believe that miracles do not happen and reason to conclude that any testimony to the miraculous is "faulty in some way." We do sometimes, with apparently ample justification, dismiss evidence to a highly improbable event. Hume expressed this intuition when he wrote impatiently to Hugh Blair, "Does a man of sense run after every silly tale of witches or hobgoblins or fairies, and canvass particularly the evidence?" (Earman 2000, 59)

This intuition is further strengthened by the way that prior improbability reasonably affects our on-going estimate of the reliability of a witness. Would we not be inclined in some cases to think less of the messenger's veracity rather than accepting his story? Are there not cases in which the very fact that a witness says that a highly improbable event took place causes us to think that he is, at least in this instance, a liar or a practical joker? And if so, does this not mean that the testimony, rather than supporting the occurrence of the event, merely reflects negatively on the person who gives it? There might, in some cases, be reference class reasons that would limit the impact of a given tall tale on our evaluation of the person's general truthfulness; for example, if we suspect that he is playing a practical joke in connection with a birthday party, we will not conclude from this that he is likely to tell falsehoods in more sober contexts. But we can also imagine many plausible circumstances in which even a single instance of apparent lying, leg-pulling, or even credulity from a given witness will be negatively relevant to our evaluation of his later testimony. One cannot help wondering whether this impact of testimony to the improbable – leading us to wonder whether the witness is less level-headed or straightforward than we had previously thought – means that the Humeans are onto something.

2. Fogelin, Henry, and Tall Tales

In *A Defense of Hume on Miracles*, Robert Fogelin attempts to press the intuition in support of Hume's dismissal of testimony to the miraculous. Fogelin argues that there is a "reverse method" for evaluating testimony in which the effect of the testimony is to cause us reasonably to re-evaluate downwards our estimate of the witness's reliability rather than being moved to accept his story.

> If a person we take to be reliable tells us that a common sort of event has occurred, trusting to his reliability, we usually accept his report without hesitation. If, however, the very same person tells us that a perfectly fantastic event has occurred, we may then move in the other direction and reconsider our belief in his reliability. (Fogelin 2003, 10)

According to Fogelin, the difference between an ordinary case in which we accept testimony and a case in which we use the reverse method to dismiss testimony and reevaluate witness reliability lies entirely in the improbability of the event attested to.

> To see how these methods [the direct and reverse method] function in an ordinary, nonphilosophical, context, consider a report from a "normally reliable source" that President George W. Bush has been observed walking a tightrope over his swimming pool. Most people's initial reaction would be justified disbelief. The sheer bizarreness and improbability of such an event's taking place cast immediate doubt on the force of the testimony offered in its behalf. It seems more reasonable to treat the report as a hoax or perhaps as a misunderstanding of a political metaphor. This is a simple and I believe uncontroversial example of the reverse method at work. (Fogelin 2003, 10)

Fogelin does admit that the story of President Bush's tightrope walking might in the end be reasonably accepted if it were performed not only before witnesses but also before a video camera. But if the report were of a miracle, things would be different, and the reverse method would have to be applied to the testimony because of the improbability of the event:

> [I]nstead of being told that [Bush] tightrope-walked across his pool, we are told that he walked across the surface of the water. Surely

> we now have a story even more improbable than the original. ...[I]t would still be wrong to assign a probability of 0 to its occurrence. It is, however, a moral certainty – something amounting to a proof in Hume's use of this term – that no such event occurred. Given this, it would be perfectly reasonable to dismiss out of hand the testimony brought forward in its behalf. (Fogelin 2003, 12-3)

Fogelin argues in favor of his "reverse method" for evaluating testimony to the miraculous by way of the interesting example of a hypothetical raconteur named Henry.

> Henry...is full of stories about famous people he has met under unusual circumstances. Browsing in a bookstore in Greenwich Village, whom does he bump into but Woody Allen, thumbing through a copy of Baudelaire's poetry? ... Flying to London, whom does he find sitting next to him but Desmond Tutu? ... He was stuck in an elevator with Cindy Crawford. He fell in with Osama bin Laden at (of all places) a disco in Beirut. The list of such remarkable meetings runs on and on, including encounters with Oprah Winfrey, Mother Teresa, Donald Trump, Stephen Hawking, Mikhail Baryshnikov, and each of the last three popes. For years, Henry has enthralled (or bemused) dinner companions with stories of this kind.
>
> What are we to think of Henry and his stories? If we look at them individually, and if we have no prior reason to distrust Henry, his stories may strike us as remarkable, but still believable. ... What seems highly implausible is that Henry could have been involved in all (or even most) of these encounters. 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. ... As a result, we will not credit any single one of [the stories] if we have only Henry's word to go on. (Fogelin 2003, 11)

The intended application to the case of miracles is obvious: Fogelin uses the multiplication of Henry's stories merely as a way to generate a low prior probability for the truth of the conjunction of all of Henry's stories. He intends, therefore, to argue *a la* Hume that if the probability of some proposition is "too low," a witness's testimony to it should cause us to conclude that the witness is lying (or was duped, or has some other problem that makes him unreliable) rather than that the event took place.

> This account of Henry and his stories is important for the following reason: It shows that the application of the reverse method depends upon the *improbability* that an event, or set of events could occur. In this case, the extreme improbability that troubles us arises through the joint assertion of events that, taken individually, are *not* sufficiently

> improbable to trouble us. Such an extreme improbability can arise in other ways as well. In the application of the reverse test of testimony, it is the extreme improbability, not its source, that matters. ...[I]f the reverse method is sufficient for dismissing Henry's testimony, it is hard to see how it cannot be *at least* as applicable to reports of miracles. There is certainly no reason for being *more* tolerant of testimony concerning miracles than we are of Henry's reports. (Fogelin 2003, 12-3)

But Fogelin is wrong. The specifics of the Henry case are important to understanding the example correctly, and there is more to such a case than merely a low prior probability for the conjunction of all the stories. A Bayesian analysis explains both why it is reasonable in the Henry case to discount his testimony after he has told a number of such stories and also why the correct intuition in that case cannot be ported over to a single testimony to an event with a low prior probability.

3. Prior and likelihood – P-inductive and C-inductive

The first distinction we must make in analyzing any case of this kind is between the prior probability of a theory and the likelihood of a theory *vis a vis* a body of evidence. The prior probability of an hypothesis is, of course, the probability of that theory conditional on all evidence available to the subject other than the specific bit of evidence E whose impact we are considering. The *likelihood* of the hypothesis *vis a vis* E is the probability of E given the hypothesis. The Bayes factor is the ratio of the likelihoods of H and ~H for some E,

$$\frac{P(E|H)}{P(E|\sim H)},$$

and is a measure of the evidential impact of E upon H and ~H.

A Bayes factor that favors H will at least partially offset a low prior probability for H. For example, suppose that the prior probability of some scientific hypothesis H is .1 but that we make an observation, evidence E, which favors H by a ratio of 10 to 1. In other words, evidence E is ten times more probable if H is the case than if it is not. Here, the prior improbability of H is exactly balanced by the Bayes factor (that is, by the impact of E), which means that the posterior probability of H after we take E into account has risen to .5.

This distinction between priors and likelihoods is directly relevant to the further important distinction between what Richard Swinburne (2004, 6) calls a P-inductive argument and a C-inductive argument. If a theory H gives higher conditional probability to some piece of evidence E than ~H gives to E, then conditioning on E confirms H.

A C-inductive argument for H is just an argument that confirms H to some extent or other, giving H a higher probability than it would have had without that argument. But evidence E may not, of course, confirm H sufficiently to overwhelm a low prior probability for H, so the argument from E may not be P-inductive – that is, it may not raise the probability of H above .5. And in that case, *a fortiori* it will not lead one actually to believe H, since usually one would assume that evidence sufficient for belief must raise the probability of an event to something even better than merely "above .5." In the case of a story regarding some event with a low prior probability, a witness's testimony can easily be C-inductive but not P-inductive. The testimony may raise the probability that the event happened even though, because of the existence of other, contrary evidence, one still does not believe nor even consider believing that the event happened.

Fogelin clearly wants to say that an extremely low prior probability for some event renders the testimony of a witness to the event unbelievable, but he does nothing to tell us why his "reverse method" is supposed to work in terms of probability. In particular, he does not tell us whether, when we conclude that a witness is not telling the truth (and hence question the witness's reliability more than we did before), this means that the witness's testimony has *no* force in favor of the event or merely that it does not have *enough* force to overcome the low prior probability of the event. Fogelin's use of such phrases as "move in the other direction" and "cast immediate doubt on the force of the testimony" as well as his implication that we should never be able to be convinced that a man walked on the water seem to imply that the argument in such a case is not C-inductive merely because of the low prior probability of the event. But a low prior probability for the event to which a witness testifies cannot by itself guarantee that the witness's testimony is not C-inductive.[2]

Fogelin's inexplicitness on this point may arise from an understandable confusion: He may assume, though he does not say so, that if some testimony tends to confirm the proposition that there is something wrong with the witness (that he is lying or has been duped, for example), it cannot also confirm the occurrence of the event. That assumption would also explain his use of the phrase "reverse method" for the case in which we revise downward our estimate of a witness's reliability as a result of his testimony. But such an assumption regarding confirmation is false. It is, in fact, possible for both of two mutually exclusive propositions to be confirmed by a given piece of evidence, so long as they are not jointly exhaustive.

Consider, for example, the following disjunctive proposition:

2 On p. 7, Fogelin appears to allow that testimony to at least some improbable events has at least some force when he says that "the improbability of the event's occurring gives us some (though perhaps not decisive) grounds for challenging the force of the testimony." Even here, the phrase "challenging the force" is odd and confusing, and Fogelin's contrast elsewhere between the miraculous and non-miraculous scenarios seems to indicate that in all such cases the "challenge to the force of the testimony" should be taken to be "decisive" merely because of the low prior probability of the event.

D: Either event A happened, or event A did not happen and the witness is a liar.

These alternatives are mutually exclusive but are not jointly exhaustive before we condition upon the evidence of the witness's testimony. Obviously, before we receive that testimony, a fairly large slice of probability goes to the negation of D – that is, to the proposition that the low-probability event did not take place and also that the witness is not a liar. In fact, that is the default assumption. One does not usually get up on a particular day thinking *either* that a friend is a liar *or* that some highly extraordinary event has taken place which he has witnessed. One assumes that neither of these is true, and for that very reason one does not expect to receive testimony to an extraordinary type of event. If, for simplicity's sake, we give zero probability, given ~D, to our receiving the witness's testimony, then ~D will be incompatible with the testimony and will disappear once we condition on the testimony. But the probability of both the hypothesis that the witness is a liar (and the event did not occur) and the hypothesis that the event did occur can rise as a result of that conditioning step – they can both be, as it were, the beneficiaries when we conclude that ~D must be false. Therefore, even if we must reevaluate the witness's reliability as a result of the testimony, it does not follow that the testimony does nothing to confirm the occurrence of the event.

4. The correct probabilistic analysis of the Henry case

Nonetheless, it does seem that after Henry has told story after story, we should in the end come to discount his testimony in the more thoroughgoing sense of not allowing it to confirm (or at least scarcely allowing it to confirm) the actual occurrence of further celebrity encounters. Is this intuition correct? And if it is, is the reason simply that adding many encounters dramatically lowers the probability of the conjunction of all of the encounters?

The intuition about eventually discounting Henry's testimony is correct. But a focus on the low prior probability of the conjunction alone will be misleading as to the real explanation for Henry's loss of credibility.

Let us model the Henry case by plugging in some concrete numbers. Suppose that the prior probability for each of the encounters Henry mentions is .01. This may be high for some of the stories and low for others, but for simplicity we will treat all the stories as if they each had the same relatively low but not extraordinarily low prior probability. Let L stand for "Henry is an extravagant liar." Let A stand for "Henry met celebrity A," B stand for "Henry met celebrity B," and so forth, where "met" includes the details that Henry gives. In other words, A, B, and C stand for events that would, if they occurred, make Henry's stories true. Let E_A be the evidence of Henry's telling his first story – that is, of Henry's saying that A, E_B be the evidence of Henry's saying that B, and so forth.

We will suppose, moreover, that the prior probability that Henry is an extravagant liar is independent of the probability that these things have really happened to him. That is to say, if such things do occasionally happen to someone, there is no prior reason to regard them as more or less likely to happen to someone who is inclined to tell extravagant lies than to a generally truthful person. To model this independence, we will take it that the prior probability of Henry's being a liar is .01 and that all priors that involve conjunctions with "Henry is a liar" or "Henry is not a liar" are calculated by multiplying individual priors. Hence, before Henry tells a story, the probability that he is a liar by nature but that the event in question did actually take place is .0001 (that is, .01 x .01). This small amount of probability is included in the .01 prior probability that the event (for example, A) took place. The prior probability that the event took place (which has an individual prior probability of .01) and that Henry is not a liar (which has an individual prior probability of .99) is .0099 (i.e., .01 x .99). And so forth. The prior probability that he is not a liar *and* that such an encounter has not taken place is, as we might expect, the highest of all, being .9801.

For simplicity, we will take it that the probability that Henry testifies to some event given both that the event did not take place and that Henry is not a liar is 0. This assumption is rendered more plausible if we take it that "liar" is shorthand for any of a number of possible scenarios – Henry's being a liar motivated by vanity, a practical joker who intends to enlighten people much later, an otherwise apparently sane maniac, a philosopher testing people's credulity, or anything of the kind. And let us treat the probability that Henry tells a story given either that he is a liar or that the event did take place as .8.

A word is in order about these numbers. I am assuming that the worst thing that happens to Henry for lying is exactly the kind of thing Fogelin describes: Namely, his friends will conclude that there is something wrong with him. Hence we have the high probability (.8) that if he is the kind of person inclined to tell fantastic lies (for whatever reason) he will tell these sorts of fantastic lies to his friends at parties.[3] One can imagine situations in which even a liar is sobered and moved to be truthful, but in our society we do not stone practical jokers or crucify them. The total prior probability that Henry is a liar assigned here (.01) is fairly low, but so, too, is the probability of meeting Osama bin Laden at a disco in Beirut. It may, in fact, be somewhat favorable to the bin Laden story to make its probability equal to that of Henry's being such a liar. The numbers assigned are intended chiefly for the sake of concreteness, but they do not seem wildly faulty, especially when viewed comparatively.

If we begin iterating conditioning steps on stories (see Table 1), labeling the stories E_A, E_B, and $E_{C,}$ where E_A is Henry's testimony to celebrity encounter A, E_B to encounter B,

3 One objection to this high likelihood is that Henry is not likely to say, e.g., that he has met Bin Laden at some given time, since there are many possible lies he could tell about meeting celebrities. But Fogelin implies that we have known Henry for quite some time and have had opportunities to hear many stories. Therefore, we are not simply considering the probability of his telling a given story at one point in time. The .8 likelihood also accords well with the intuition that, if Henry really is an attention-seeking liar, this constitutes a good explanation for his stories.

and so forth, we find that, at the first conditioning step, the probability of encounter A is significantly confirmed and ends up as just above .5; that posterior probability includes the (improbable but possible) sub-hypothesis that Henry is indeed a liar by nature, though this particular thing did really happen to him.[4] This initial probability for A after receiving only E_A corresponds roughly to Fogelin's claim that "his stories may strike us as remarkable, but still believable." L itself is also just over .5 probable. That both A and L are a little more probable than .5 is possible because of that small area of possible overlap between them. The hypothesis that Henry is a liar and that A did not happen now comes in at about .49. As discussed above, Henry's credibility has been called into question, because his being a liar would in this case be a good explanation of the evidence, though the probability that the event occurred has also risen.

Table 1 – Some effects of repeated conditioning on Henry's stories

P: Prior probabilities	P*: After conditioning on E_A – H. says that A.	P**: After conditioning on E_B – H. says that B.	P***: After conditioning on E_C – H. says that C.
P (A) = .01	P* (A)=just over .5	P** (A) ≈ .02	P***(A) = just over .01
	P*(B) (No evidence yet from H. about B.) = .01	P** (B) ≈ .02	P***(B) = just over .01
		P** (C) (No evidence yet from H. about C.) = .01	P***(C) = just over .01
			P***(A & B & C) ≈.0001
P (A & L) =.0001	P*(A &L)≈.005		
P (L) = .01	P* (L)= just over .5	P** (L) ≈ .99	P***(L) ≈ .9999
P (~A & L) = .0099	P* (~A & L)≈.4975		
P (~A & ~L)=.9801	P* (~A & ~L) = 0		
	P* (A & ~L & ~B)≈.49	P** (A & ~L & ~B) = 0	
	P* (~A & ~B & L)≈.49	P** (~A & ~B & L) ≈ .97	
			P***(~A & ~B & ~C &L) ≈.97

All likelihoods are either 0 or .8. P (E_A|L) = .8; P(E_B|L) = .8; P (E_C|L) = .8; P (E_A|A) = .8; P (E_A|~A & ~L) = 0. The same is true *mutatis mutandis* for E_B and B and for E_C and C. Some probabilities are approximate, as noted. Since these are only partial distributions involving approximate probabilities and some propositions (such as L and A) that are not mutually exclusive, probabilities given do not sum to one.

But as we continue iterating steps of conditioning on similar evidence, we find that the confirmation for the events in question is considerably altered by the addition of more stories. After we condition on E_B, we find that both event A and event B are now only at approximately .02, which is still better than their original .01 but considerably

4 The confirmation of A to just above .5 can be obtained as follows: The prior probability of A is stipulated as .01. $P(E_A|A)$ is stipulated as .8. The regular form of Bayes's Theorem is $P(H|E) = [P(E|H) \times P(H)]/ P(E)$. We can calculate $P(E_A) = P(A) \times P(E_A|A) + P(\sim A \,\&\, \sim L) \times P(E_A|\sim A \,\&\, \sim L) + P(\sim A \,\&\, L) \times P(E_A|\sim A \,\&\, L) = .008 + 0 + .00792 = .01592$. Plugging this prior and likelihood for A and this prior probability for E_A into Bayes's Theorem yields P*(A) equals approximately .5025. A similar calculation, *mutatis mutandis*, yields the same posterior probability for L in this first step. It is possible for both L and A to have posterior probabilities over .5 because they are not mutually exclusive.

less probable than event A was when Henry had told only one story. The reason for the disconfirmation of A by E_B and its regression towards its prior probability is fairly obvious. The probability that A is true, that Henry is not a liar, and that B is false – which was a completely live option (at about .49) before conditioning on E_B – is now 0, because it is incompatible with E_B. This fact makes for significant disconfirmation of A by E_B, while L – that Henry is a liar, is becoming more and more probable all the time. Indeed, the probability that neither A nor B occurred and that Henry is a liar jumps from about .49 to about .97.[5]

By the time we condition on a third story, E_C, to a third encounter, C, the situation Fogelin sketches is emerging clearly. The probability of each of A, B, and C is just barely higher than its prior of .01. C has received almost no confirmation from Henry's saying that C occurred. The probability that all three stories are false is about .97 and is just a tiny bit lower than .99 cubed – the probability that none of the events happened before conditioning on any evidence at all. And the probability that Henry is a liar is overwhelming, being about .9999. The probability that he has lied about at least two of the stories (that is, that he is a liar and that at least two stories are false) is over .999.[6]

What is happening is clear: As Henry tells more stories and we become more and more convinced that he is a liar, we reasonably discount what he says. His stories are losing their power to confer confirmation upon the events to which they attest, since they can be just as easily explained by the hypothesis that he is once more lying, which is itself now highly probable.

We should be careful not to make the mistake of assuming that the probability that all the stories are false will continue to rise indefinitely with the addition of more stories. Even if we were simply considering in the abstract the possibility of Henry's encountering celebrities, without any testimony from him to the effect that he had really met any celebrities, each such encounter has a non-zero prior probability. The more such possible encounters we consider, the more likely it is that at least one of them has actually taken place by sheer coincidence. On the principle that a stopped clock is right twice a day, we should take it that Henry, like anyone else, might occasionally actually run into a celebrity by accident, and that in that case he would be likely to tell the story. This means that the most that can happen as we have a higher and higher probability that Henry is a liar is that our posteriors after receiving Henry's evidence will come, in essence, to equal our priors. As more stories are added, the probability that *all* the stories are false will very slowly drop, just as the probability would drop if we multiplied .99 (the prior probability that some given story is false) by itself as many times as the

5 Because of the need to calculate probabilities for such complex hypotheses as (~B & L & A), it is not possible here to give many calculations in complete detail, but here is another example at a somewhat greater level of detail than given in the text and chart: We can calculate P*(~A & L) using Bayes's Theorem and the stipulated numbers as approximately .4975. (The process is similar to that given for calculating P*(A) in footnote 4.) Since B and A are independent, we multiply by the .99 probability of ~B in P* to get P*(~B & ~A & L) equals approximately .492525. Giving all sub-hypotheses their due, we can calculate $P^*(E_B)$ equals approximately .40598. $P^*(E_B|{\sim}B \,\&\, {\sim}A \,\&\, L)$ equals .8, as always. Plugging these numbers into Bayes's Theorem yields P**(~A & ~B & L) equals approximately .97054.

6 This last statement, about the probability that Henry has lied about at least two stories, is not represented in Table 1 for reasons of space.

number of encounters under consideration. To say this is simply to recognize that, just as Fogelin implies, we are moving rapidly toward a situation in which Henry's celebrity stories make no difference whatsoever to our beliefs about Henry's meeting celebrities. They become background noise rather than effective evidence.

Here we need to recognize something absolutely crucial to the understanding of the entire case: L was confirmed by E_A, going from .01 to just over .5. That confirmation *carries over* when we prepare to condition on E_B, for it is simply the probability that L has after we condition on E_A. Hence, L has just over .5 probability even before Harry tells his second story, and, with each new story that Henry tells, L is confirmed again. After E_B, L is approximately .99, and of course higher still after E_C. But there is no similar carry-over effect to B of the confirmation that accrues to A from Henry's testimony to A. B, prior to Henry's story E_B, just has the same .01 prior probability that any of the celebrity encounters has before we receive testimony about it. The fact that, after receiving E_A, we have more reason to believe A than we had before does not mean that we have any more reason to believe B. The occurrence of B is independent of the occurrence of A. And the same is true of C. Unlike the hypothesis that Henry is a liar, which is confirmed by each of the stories, each celebrity encounter "starts all over again" with its own ordinary prior probability, before Henry attests to it. In short, Henry's being a liar explains all of the evidence (E_A-E_C), but each event's occurrence explains only one part of the evidence. Relatedly, in order to account for the evidence on the assumption that Henry is not a liar, we must hypothesize for each story a new, independent encounter, so that the only "hypothesis" consistent both with the evidence and with Henry's not being a liar is a conjunction of independent events which grows longer each time Henry tells another story.[7] Nothing similar is true of L, which is able to account for all the stories Henry tells quite well without being elaborated in any way.

There is a parallel here to cases involving multiple lotteries. If the hypothesis that the lottery is fair does not have a prior probability of 1, a lottery win by N confirms to some degree the theory that effective cheating on N's behalf has taken place (since effective cheating favoring other participants is ruled out by N's win). Under plausible background assumptions, the same hypothesis of cheating-to-favor-N is confirmed by all of N's wins in a series of lotteries, but our expectation that N will win the next lottery fairly is not rationally increased by his winning the first lottery, since the lotteries are independent if they are fair. This is why it is more reasonable to believe that cheating is occurring if N wins three different 100-ticket lotteries than if N wins a single 1,000,000-ticket lottery.

The analysis thus far might raise the question as to whether conditioning on E_A, E_B, and E_C *seriatim* is being treated as probabilistically different from conditioning on

7 Fogelin makes this independence quite clear when he fleshes out Henry's stories as involving chance encounters. If Henry is telling the truth, he is not a celebrity chaser, nor does he have an advantageous job or a well-connected friend who helps him to meet celebrities. Rather, Henry presents himself as someone who has experienced an amazing series of coincidences. And it is upon this independence that Fogelin relies for the low prior so important to his analysis.

their conjunction – (E_A & E_B & E_C). From a Bayesian perspective, these should not differ. Nor do they, in the abstract and for a logically omniscient being. However, it is extremely difficult epistemically to access in the abstract the necessary information for making the shift in a single conditioning step. The probability distributions P*, P**, and P*** contain probabilities for complex conjunctions not represented, for reasons of space and clarity, in Table 1, as well as for those conjunctions represented in the table. These include, for example, the conjunction (L & A & ~B & ~C) – that is, the proposition that Henry is indeed a liar but that some particular one of the three celebrity encounters did take place, the proposition (~L & A & ~B & ~C) – the proposition that Henry is *not* a liar and that some particular one of the celebrity encounters just happened to take place, and so forth. Each of these different complex hypotheses constitutes a portion of the probability space in the initial distribution (though some can be bundled together initially), and those that have not been ruled out by the evidence receive a slice of the probability space in the final distribution. But the final probability for, e.g., (L & A & ~B & ~C) differs in specific ways from its initial probability as a result of the force of *each part* of the evidence – E_A, $E_{B,}$ and E_C – and its final probability must therefore be calculated carefully by taking into account *each piece* of evidence. Indeed, the final probability of the simple (and interesting) proposition L, conditional on all three stories, is evident only by way of calculations that take into account the effects of the various parts of the evidence on the complex sub-hypotheses contained in L, involving the truth and falsehood of various stories. It is not accessible by some simple calculation from the initial distribution P.

It is, of course, relatively simple to calculate *in hindsight* the Bayes factor for the partition of L and ~L and for the conjuncted evidence (E_A & E_B & E_C):

$$\frac{P(E_A \;\&\; E_B \;\&\; E_C | L)}{P(E_A \;\&\; E_B \;\&\; E_C | {\sim}L)}$$

This Bayes factor, given the probabilistic assumptions stated, must be approximately 10^6.[8] And it is clear from the reversion to approximate prior probabilities for the events in the final distribution that the Bayes factor for the conjunctive evidence *vis a vis* each event (A, B, or C) must be very little better than a one to one ratio – in other words, that the evidence of all three stories together scarcely favors the occurrence of any one of the tales at all, for reasons we have already discussed. But these Bayes factors may well not be those we would have estimated were we attempting to look at the evidence in a single lump at the outset. In particular, the very high Bayes factor for the conjunction in favor of L may come as a surprise. It is quite easy to underestimate the cumulative force of evidence if we do not look at the evidence in its component parts. The epistemic path from P to P*** is thus made clear and accurate by not attempting to jump from one

8 As implied in the text, this Bayes factor is backsolved. To do so, we use the odds form of Bayes's Theorem, the stipulated probabilities for L and ~L prior to all the pieces of evidence, and the final posteriors calculated for L and ~L in P***. In moving from a probability for L of .01 to one of .9999, we are moving from odds of 99 to 1 against L to odds of 9999 to 1 in favor of L. This is very nearly a shift from 100 to 1 odds against L to 10,000 to 1 odds in favor of L – six orders of magnitude.

distribution to the other in one conditioning step but rather by considering in detail the effect of the evidence one piece at a time. Calculation by iterated conditioning steps on the separate parts of the evidence makes perspicuous the epistemic relations among the many sub-hypotheses and the various parts of the evidence which must obtain and which a logically omniscient being would understand readily. The iterated conditioning enables us to see the probabilistic relevance of the considerations already discussed – the carryover of confirmation to L and the independence of A, B, and C and the consequent need for a more and more complex conjunction of independent encounters if Henry is to be entirely believed. In this way, too, the iterated conditioning is true to Fogelin's description of the informal situation in which our opinion of Henry changes as he tells more and more tall tales.

5. The difference between the Henry case and a single improbable story

The analysis thus far should make it clear that no testimony to some single event, however initially improbable, can have all the probabilistically crucial features of the Henry case. This is because, contra Fogelin, the multiplicity of stories in the case of Henry does more than simply drive down the prior probability of their conjunction. The fact that Henry tells many different, independent stories also fragments the evidence. Each story Henry tells confirms only one event, while all the stories confirm the hypothesis that Henry is a liar. The liar hypothesis can explain the stories just as well as does the occurrence of the events, and since it can explain *all* of the stories, while each event explains only *one* story, the liar hypothesis is overwhelmingly confirmed by the entire evidence set.

There is no parallel to this situation when a witness testifies to some single event, such as, for example, to having met the resurrected Jesus. In any single case, we have to deal with a much simpler probabilistic situation as we set up a conditioning step: In order to decide whether and to what extent the testimony confirms the event, we need to decide whether and to what extent the testimony is more to be expected if that single event occurred than if it didn't. Where M is some ostensibly miraculous event and T is testimony to that event, the question in the individual case for deciding the strength of the evidence is the size of the Bayes factor, the ratio

$$\frac{P(T|M)}{P(T|\sim M)}.$$

This Bayes factor may favor the event overwhelmingly, slightly, or not at all, but its weight is not determined by the prior probability of M.

To say that this situation is probabilistically "simple" is, of course, to speak only of the formal representation in the form of a single Bayes factor for a single event. The evaluation or even estimation of that Bayes factor is not automatically a simple epistemic matter. Indeed, once it is acknowledged that priors are not the whole story in the evaluation of testimony to extraordinary stories, the focus naturally shifts to the question of how we should rightly estimate the Bayes factor. In cases of testimony from long ago, skeptics will understandably raise the additional question of whether witnesses really testified to the event at all or whether the putative testimony is a later, invented story by someone who could not have been a witness. The example of the resurrection is used here merely as an example of a single miraculous (hence, antecedently improbable) event to which we might receive testimony, as opposed to a series of independent improbable events, and specific questions of New Testament textual scholarship obviously lie beyond the scope of this paper. (See McGrew and McGrew 2009, 597-604 for some discussion of these issues.)

Indeed, it is precisely when we have recognized the importance of the force of specific evidence, as opposed to mere prior improbability, that purely historical investigation comes into its own. This recognition, too, should make clear the insufficiency of any “trial by proxy,” such as Hume indulges in in Part II of his essay, for evaluating testimony to some improbable tale. That is to say, it is not enough to show that *other* stories, vaguely similar to the story one really wishes to attack, are not credible (see McGrew and McGrew 2009, 653-8). The strength of a Bayes factor for some specific testimonial evidence will depend upon a host of particulars whose evaluation cannot be avoided (see McGrew and McGrew 2012, 57-62).[9] Thus the evaluation of the evidence in the case of the resurrection or any other specific miracle claim must be the subject of actual investigation of those claims, not of claims to entirely different miracles.

Returning to Henry’s stories, we can see that Henry’s first story confirmed both the event and the hypothesis that Henry is a liar, though that confirmation did not move either hypothesis to a very high posterior probability. In some other case, including cases involving testimony to the miraculous, the Bayes factor might favor the occurrence of the event even more strongly over its negation than it does here.[10] Everything depends on the specifics of the evidential situation, but since only one conditioning step is involved concerning only one putative event M, there is no question of multiple confirmations of the hypothesis that the witness is a liar by various pieces of evidence which can be accounted for (on the side of M) only by further elaborations of M.

9 For example, if witnesses claim to have seen and spoken with someone who was previously dead, we would like to evaluate how close they were to the person when they supposedly recognized him and under what other conditions (physical and psychological) these meetings supposedly occurred, how extensive and detailed these alleged encounters were, how many people were allegedly present at them, how well they knew the person before, what resources and motives someone else would have had for tricking them, and what motives the witnesses have for lying, telling the truth, or remaining silent. If an extraordinary healing is claimed, we understandably want to examine the evidence concerning whether the person was really ill in the way stated prior to the alleged healing, how closely he was examined (and by whom) afterwards, whether the alleged healing was permanent, and what natural explanations might account for the healing.

10 As noted above, Henry suffers no extreme penalty for his stories, even when people disbelieve him. If a witness is under severe duress to recant but sticks to his story, this usually makes his story at least somewhat more valuable than testimony from a witness under no duress, especially if he was in a position to know about the event to which he attests.

6. Conclusion

It is unfortunate that defenders of Hume try so hard to place their entire case on the priors. It seems that one possible reason for this is a distaste for the nitty-gritty of historical particularity. The Humean approach presents the opponent of miracles with the opportunity to avoid what Earman (2000, 61) calls "difficult and delicate empirical investigations." But while Bayesian probability theory does not all by itself tell us what probability – either prior or posterior – we should rationally assign to the occurrence of any putative miracle, it does tell us that prior probabilities are not the whole story.

The Henry case is, as Fogelin uses it, yet another attempt to sidestep the importance of anything other than the prior improbability of a miracle and to imply by vague reference to a "reverse method" that any sufficiently improbable story merely casts doubt on the veracity of the person telling it and may therefore be disregarded. The above analysis shows that the Henry example does not fulfill that purpose, for two reasons. First, testimonial evidence can lower our confidence in a witness's credibility while at the same time confirming the event testified to. Second, the proper analysis of the Henry example depends critically on the complex nature of the evidence as it relates both to the separate encounters and to the hypothesis that Henry is a liar, and that analysis will not apply to an attestation to a single event.

Acknowledgements: Thanks to Timothy McGrew for technical help during the course of this project as well as for noting the parallel to cases of multiple lotteries.

REFERENCES

Earman, J. 2000. *Hume's Abject Failure*. Oxford: Oxford University Press.

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

Holder, R. 1998. Hume on Miracles: Bayesian Interpretation, Multiple Testimony, and the Existence of God. *British Journal for the Philosophy of Science* 49,1:49-65.

Hume, D. 1748. *An Enquiry concerning Human Understanding: A Critical Edition*. Ed. T. Beauchamp. Oxford: Oxford University Press. (Cited from Oxford edition, 2000)

Laplace, P. S. 1840. *A Philosophical Essay on Probabilities*. Trans. F. W. Truscott and F. L. Emory. New York: Dover. (Cited from Dover edition, 1951)

Levine, M. 2010. Miracles. In *The Stanford Encyclopedia of Philosophy (Fall 2010 Edition)*, ed. E. N. Zalta. http://plato.stanford.edu/archives/fall2010/entries/miracles/.

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

McGrew, L. and T. McGrew. 2009. The Argument from Miracles: A Cumulative Case for the Resurrection of Jesus of Nazareth. In *The Blackwell Companion to Natural Theology*, eds. W. L. Craig and J. P. Moreland, 593-662. Oxford: Wiley-Blackwell.

McGrew, L. and T. McGrew. 2012. The Reliability of Witnesses and Testimony to the Miraculous. In *Probability in the Philosophy of Religion,* eds. J. Chandler and V. Harrison, 46-63. Oxford: Oxford University Press.

Swinburne, R. 2004 *The Existence of God*. Oxford: Oxford University Press.

Received: August 12, 2010
Accepted: May 15, 2011

Kalamazoo, Michigan
USA
lydiamcgrew@yahoo.com