Bayes on trial by fire

By DOUG SHAVER
March 2016

Late last year, the *Wall Street Journal* had an article (now available only to subscribers) in which the authors reported a discovery that some people in Medieval times seemed to have miraculously survived trials by fire. Specifically, in one Hungarian city during the 13th century, defendants in certain criminal cases were ordered by the civil courts to prove their innocence in an ecclesiastical court by grabbing a red-hot iron bar that had just been taken out of a fire. In more than half of roughly 200 cases, the authors said, the defendants were acquitted because they survived the ordeal without injury. The likeliest explanation, they suggest, is that the trials were rigged: the officiating priests faked the ordeals so that the defendants only seemed to have been at risk of serious injury.

Joe Heschmeyer, a contributor to a Christian apologetic website was, uh, very skeptical of this explanation. For the authors' scenario to be credible, he said, among other things:

- All the witnesses had to be "stupid enough to believe that a piece of iron is smoldering hot when it isn't";
- There had to be "a massive conspiracy of priests to fake miracles"; and
- There had to be "a global conspiracy that left no paper trail, and apparently raised no eyebrows."

What really bothered Heschmeyer, of course, was the authors' "[refusal] to accept the evidence, no matter how strong, because of a prior commitment to rejecting miracles." He concluded his essay thus: "Moderns think of trial by ordeal as proof positive of the irrational dogmatism of our religious ancestors' culture. In fact, the story seems to reveal a great deal more about the irrational dogmatism of our own irreligious culture."

Well, some atheists certainly are irrationally dogmatic. We're just as human as any Christian, and irrational dogmatism comes with the human territory. But I don't think that's the only way to account for our skepticism about these stories.

Let's begin by stipulating that we should not presuppose the impossibility of miracles, that we must admit that they could happen. That admission does not epistemically obligate us to believe every report about a miracle. Christians, who don't doubt that some miracles have actually happened, don't believe every miracle story they hear. Indeed, I suspect that most of them are inclined to disbelieve most such stories that they hear about outside of church. And they disbelieve them for the same reason atheists do: They are antecedently skeptical, and in their judgment, the stories are insufficient evidence to overcome their skepticism. What Heschmeyer is arguing in effect, then, is that in the case of these Medieval trials, the evidence ought to be enough to overcome any justified degree of skepticism.

A key premise of his argument is that the evidence is grossly inconsistent with the hypothesis of trial-rigging. Well, that sort of begs for a Bayesian analysis, and so let's see how his argument fares under that sort of scrutiny. Heschmeyer posted [his comments on the Strange Notions website](), and the following is a revised (substantially in some parts) version of comments I posted in response to him.

There are various formulas available for Bayesian calculations, but some not-very-advanced algebra shows them to be all mathematically equivalent. My preferred formula is bit more complicated than the others, but it has the advantage that all relevant assumptions must be explicitly accounted for. When used properly, it doesn't allow any hidden assumptions to slip in. Here it is:

$$P(H|E) = \frac{P(H) * P(E|H)}{P(H) * P(E|H) + P(\sim H) * P(E|\sim H)}$$

Even this is a slightly abbreviated version. Every term is assumed to terminate with a "$|B$" or "$.B$", meaning "on background knowledge." Thus the numerator, written in full, would be

$$P(H|B) * P(E|H.B)$$

In other words, we're not supposed to ignore, or to treat as irrelevant, everything else we know or think we know about how the world operates. All probabilities are to be estimated, and the evidence evaluated, in light of that background knowledge. We don't treat this as an epistemically isolated incident.

$P(H|E)$, which I will call the consequent probability, is the probability that the hypothesis is true *given the evidence*. This is what we're trying to calculate. The formula presents it as a function of four variables, but really only three because one of them is determined by another. $P(H)$, the prior probability, is the probability that we should have assigned to the hypothesis *before discovering the evidence*. $P(\sim H)$ is the prior probability that the hypothesis is false, and this has to be $1 - P(H)$. This leaves just two more variables to estimate. $P(E|H)$ is the probability of the evidence obtaining given a true hypothesis, and $P(E|\sim H)$ is the probability of the evidence obtaining given a false hypothesis.

We can schematize this a bit to get a better idea of the relationships among the three estimated variables. With appropriate substitutions to eliminate a few keystrokes, we get the schematic

$$P = \frac{AB}{AB + CD}$$

But then we recall that $C = 1 - A$, and now we have

$$P = \frac{AB}{AB + (1 - A)D}$$

A mathematical consequence of this is that if $B = D$, then $P = A$. In other words, whenever the evidence is just as likely whether the hypothesis is true or false, then the consequent probability equals the prior probability: $P(H|E) = P(H)$. It doesn't matter whether the assigned values are high or low. As long as they're about the same, then the evidence is epistemically irrelevant. Or, we can say that it's not really evidence after all. Whatever justification we had for believing or rejecting the hypothesis before seeing the evidence, we have neither more nor less justification afterward.

Whether the consequent probability is more or less than the prior depends strictly on the difference between B and D, i.e. the value of $P(E|H) - P(E|{\sim}H)$. A positive value produces $P(H|E) > P(H)$, a negative value yields $P(H|E) < P(H)$, and larger absolute values produce a larger difference between prior and consequent. (A negative value, if that's what we get, just means the evidence is against the hypothesis: It reduces whatever prior justification we might have had for believing it.)

It might help to consider some extreme cases. If we assign a prior probability of zero, then by mathematical necessity, the consequent probability will be zero. If our minds are made up that the hypothesis is impossible, then no evidence will convince us, and according to Bayes, it should not convince us. With a prior of zero, evidence is irrelevant. This means that Bayes forces us, if we claim to be open to persuasion by sufficient evidence, to admit that the hypothesis is not an impossibility: We cannot, in good faith, assigned it a probability of zero. Likewise, if the prior probability is 1, then the consequent probability is also 1, regardless of the evidence, because then $P({\sim}H) = 0$, meaning the hypothesis cannot possibly be false. In either case—prior of impossibility or prior of certainty—evidence becomes irrelevant.

With that background, let's consider some numbers for the current problem. Our evidence, E, is certain documents reporting that some innocent people underwent trial by ordeal conducted by certain priests and, on subsequent examination, were found to be uninjured, resulting in their acquittal of the charges against them. The hypothesis, H, is that they were protected from injury by divine intervention, i.e. a miracle occurred. The contrary hypothesis, ${\sim}H$, is that some natural occurrence prevented the defendants from being detectably injured. A thorough analysis would have to separately evaluate all possible naturalistic alternatives to divine intervention, but for simplicity we'll assume only one is worth considering: priestly complicity in a sham ordeal, i.e. the trials were faked by those responsible for conducting them, presumably because those priests were antecedently convinced that those accused were innocent. We need estimates for: $P(H)$, the prior probability that a miracle occurred; $P(E|H)$, the probability that if the miracle had occurred, we would have this evidence; and $P(E|{\sim}H)$, the probability that we would have the same evidence if something other than a miracle had occurred.

An atheist is going to think that $P(H)$ is, if not zero, so close as to make no difference, but let's try to really hard to avoid begging the question. Fairness might suggest that we ask a theist what he thinks the prior probability is, but it might be hard to find a theist who will give us a straight answer. For starters, though, let's suppose we don't have any better reason to doubt the report than to believe it. Almost by definition, then, this gives us a prior probability of 0.5. This is not realistic to any atheist, but bear with me.

What is the probability that we would have this evidence if the miracle had really occurred? We need first the probability that a record would have been produced. We can assume for the present discussion that they almost always were, and so let's go with 0.99. But those records didn't always survive, and the probability we're looking for is the probability that we would have those records now, 700 years later. What percentage of church records from 13th century Hungary are still extant? I have no idea, but let's guess half of them. As it turns out in this case, any guess will do, but when you're doing Bayes, you have to ask these questions. The probabilities multiply, so we get $P(E|H) = 0.5$, rounding to one significant figure.

And what if there was no real miracle? Would that have made a difference in the probability that the records (a) would have been produced and (b) would have survived into modern times so that we'd know about them? I don't see why. The trial was supposed to be conducted the same as any authentic trial, and so the record-producing process would have been identical for the sake of appearances. Afterward, the authenticity of the ordeal, i.e. whether or not a miracle really happened, should have made no difference to the survival of the documents. Once they were produced, subsequent custodians, even if they'd been in the know, would have had no reason to treat them differently from records of authentic ordeals. The best way to avoid suspicion is to act as if you have nothing to hide. Thus we get $P(E|{\sim}H) = P(E|H)$, and so we are just as likely to have this evidence whether the ordeal was authentic or faked.

This means, as noted above, that regardless of our estimate for $P(H)$, we get the same value for $P(H|E)$. Whatever antecedent reason we had for believing that a miracle happened on this occasion, we have no better reason, just because we have these documents saying that it happened, for believing that it did happen. Thus, to whatever extent I am justified in thinking that a miracle is more likely to be faked than to have actually occurred, to that extent I am justified in believing that this particular miracle was faked.

Maybe I'm not justified at all. Maybe I should believe that most of the time, when someone says a miracle happened, it really did happen. But that is a separate issue and must be supported by its own argument. In this particular case, the evidence is just as likely to have existed whether or not the miracle really happened, and so this evidence does nothing to enhance the credibility of this particular miracle, and I see no reason at all to suppose that the odds of its happening were anywhere near 50-50.

Heschmeyer claimed that the probability of fakery was extremely low, but we were considering only two hypotheses, assuming that they exhaust the possibilities, and the sum of their assigned probabilities must be 1.0. In that case, an argument for $P(\sim H) = 0.1$, let us say, would entail $P(H) = 0.9$. But the assumption of a single alternative hypothesis was made for the sake of simplicity. An argument for a $P(H)$ significantly above 0.5 would have to demonstrate that all possible naturalistic explanations had a combined probability significantly below 0.5.

Now, a believer might argue that every naturalistic alternative to trial-rigging is so improbable as to be negligible. In that case, the believer might claim a low antecedent probability for the rigged-trial hypothesis and thus infer a high antecedent probability of a miracle. Very well. As I currently understand Bayesian theory, it doesn't tell us whether, when there are only two alternative hypotheses, we must estimate $P(H)$ first and then derive $P(\sim H)$ or may do it vice versa. In that case, the reasonableness of doubt vs. belief depends on the reasonableness of supposing that there was a greater than 50-50 chance that some priests during the Middle Ages were sufficiently compassionate to fake a miracle in order to avoid torturing some people who, in their judgment, had done nothing to deserve torture. To me, that looks like a pretty reasonable supposition, but any Christian who wants to argument against is certainly welcome to do so. So far, though, I have seen no such argument.

(Last update: March 12, 2016.)

[Back to site home.]