A BAYESIAN BASELINE FOR BELIEF IN UNCOMMON EVENTS

VESA PALONEN
UNIVERSITY OF HELSINKI

Abstract. The plausibility of uncommon events and miracles based on testimony of such an event has been much discussed. When analyzing the probabilities involved, it has mostly been assumed that the common events can be taken as data in the calculations. However, we usually have only testimonies for the common events. While this difference does not have a significant effect on the inductive part of the inference, it has a large influence on how one should view the reliability of testimonies. In this work, a full Bayesian solution is given for the more realistic case, where one has a large number of testimonies for a common event and one testimony for an uncommon event. A free-running parameter is given for the unreliability of testimony, to be determined from data. It is seen that, in order for there to be a large amount of testimonies for a common event, the testimonies will probably be quite reliable. For this reason, because the testimonies are quite reliable based on the testimonies for the common events, the probability for the uncommon event, given a testimony for it, is also higher. Perhaps surprisingly, in the simple case, the increase in plausibility from testimony for the uncommon events is of the same magnitude as the decrease in plausibility from induction. In summary, one should be more open-minded when considering the plausibility of uncommon events.

INTRODUCTION

Is it reasonable to believe in a testimony of an uncommon event in the face of otherwise uniform contrary evidence from prior events? This question has been much discussed historically, with notable contributions from David Hume (Hume 1748), John Earman (Earman 2000), Peter Millican (Millican 2013), and many others.

David Hume's argument was not clearly formulated, but basically he argued that the evidence for common events (e.g., events compatible with

PP. 159–176
DOI: 10.24204/EJPR.V9I3.1909
AUTHOR: VESA.PALONEN@HELSINKI.FI
EUROPEAN JOURNAL FOR PHILOSOPHY OF RELIGION
VOL 9, NO 3 (2017)

the perceived laws of nature) is so strong that the testimony for uncommon events (e.g., miracles or exceptions to the perceived laws of nature) is usually not strong enough evidence for the uncommon event to be believable. "Extraordinary claims require extraordinary evidence" is the oft-used and dangerously poorly-defined phrase connected to Hume's position.

J. Earman (Earman 2000) does a systematic job of both trying to find a precise form for Hume's argument and then analyzing the argument in detail, leading to the view that Hume's argument is largely incorrect. Earman makes two important points:

(1) With a Bayesian calculation of inductive inference, the probability of an uncommon event does indeed go down with the amount of common events (as $1/(n+2)$), but never to zero. Hence, based on induction, one can hence never be certain that the uncommon will not happen.

(2) While discussing the role and reliability of testimonies for uncommon events, Earman shows that the testimony can often provide enough credibility for the uncommon event. Notably, in considering the evidential force of a testimony, one needs to consider, not just how often witnesses are wrong in general, but what is the probability that the witness would make just such a particular claim and be in error with that claim. For example, when a witness is testifying that John Doe won the lottery, it is not enough to suggest that a testimony is in general wrong with e.g. 10% probability, but one needs to take into account the probability that the claim would be made about John Doe in particular (why just him?) and also the probability that the claim indeed would be erroneous.

Coming to the present work, the calculations published on the topic up to now assume a large amount of common events. In reality we usually have a large amount of *testimonies* for the common events. That is, we do not have uniform *evidence* against the uncommon events. What we may have is uniform *testimony* against the uncommon events or miracles. In this sense, it can be said, that up to now, the problem of uncommon events and their believability based on testimony has not been fully analyzed, even on the basic level.

So, this paper will offer a full Bayesian solution for the more realistic case, namely, for the question: How believable is an uncommon event when we have a uniform mass of *testimony* to the contrary? The calculation can be seen as a baseline for further discussions on the topic, with nuances to be added as later as different additions and changes to the model are explored. Further consideration will involve considerations for several testimonies of the same event, the independence of the witnesses, the effects of prior beliefs against the uncommon event, and whether or not those testifying to a rare event are less trustworthy than those testifying to a common event.

In addition, the calculations will relax one common assumption in the discussions about uncommon events/miracles. We will not presume to know the probability of a testimony being wrong. Rather, we will assign a probability for the probability of a testimony being wrong, and let the data determine the most probable probability for a testimony being wrong.

The word miracle can have several definitions. For the purposes of this paper, a miracle is taken to be just one kind of uncommon event. Hence, the argument presented aims to be more general, the result providing a baseline for inference concerning uncommon events, without concentrating on the particulars of different cases. The results will then apply to miracles, testimonies of rare natural events like winning a lottery, and rare-event measurements in physics (e.g. the possibility of proton decay).

For simplicity, the paper will concentrate on a binary case of mutually exclusive outcomes. One outcome is taken to be common, i.e. there are more testimonies for that outcome than for the uncommon outcome. In cases where there are more than two types of outcomes, they can be grouped into two mutually exclusive sets, the inference proceeding in the same vein, except for the possible modifications to the prior probabilities used.

A further complication in the field has been that the usage of probabilities in the discussion has been partial, with several authors dissecting the full formulas for partial arguments based on the full formulation, see e.g. (Ahmed 2015), with the full solution nowhere to be seen. The aim here will be to show the full solution for the simple default situation with few assumptions. From there, different assumptions can be added whenever the assumptions are well grounded.

The calculation will first be made for the case where we have several testimonies of common events and one testimony for an uncommon event. After obtaining the solution, its basic behavior as the number of testimonies for the common events increases will be discussed. We will then add known false testimonies for uncommon events to the calculation and show how the solution behaves for that case.

NOMENCLATURE AND TOOLS

Below is a table of notations used in the paper. For simplicity, the logical AND symbol, ∩, is usually dropped in the probability notation between propositions.

$\neg A$	Not A
$A \cap B$	A and B
$A \cup B$	A or B
$p(A \mid B)$	Conditional probability of A being true given that B is true
$p(A\ B)$	Probability of A and B.
B	A common event, (black ball drawn for an urn)
W	An uncommon event (white ball drawn from an urn), equivalent to $\neg B$ in our examples
B_i	The result of the i'th event is common (i'th ball was black)
B^l	l common events B
$t(\ldots)$	Testimony of an event
n	Number of testimonies of a common event
$t(B)^n$	n testimonies of a common event
C^n	A vector of n real events (W or B) behind the n testimonies
v	The (unknown) probability for the uncommon event to happen
d	The (unknown) probability for a testimony to be wrong, $d = p(W \mid t(B)) = p(B \mid t(W))$
J_W, J_B	Intermediate terms used in the calculation to simplify representation

We will be using general Bayesian methodology, which is basically finding out the joint probability distribution for all the parameters relevant to the case and calculating the wanted probability distribution from the joint distribution by using marginalization and the Bayes rule. (This approach is generally applicable and much used in the machine-learning community because from the joint distribution one can systematically calculate whatever probability one happens to need.)

The Bayes rule in one sense is a way to move the propositions to or from the condition-side inside the parenthesis of probabilities, that is, if a proposition is in the conditions-side, we can move it to the side for which probabilities are calculated for (Sivia 1996; Gelman et al. 2003):

$$p(A|B) = p(B|A)\frac{p(A)}{p(B)}.$$

With marginalisation, we can remove unwanted propositions from the parenthesis. That is, we sum over all possible values of the unwanted proposition or variable (Sivia 1996; Gelman et al. 2003):

$$p(A) = \sum_{Over\ all\ values\ of\ B} p(A\ B).$$

So, if we don't know *B*, we don't need to, and often should not, assume any particular value for it to estimate *A*. Instead, we take into account all relevant values for *B* by summing the joint probability of *A* and *B* over all values of *B*. With these two tools, in principle any probability can be calculated in a systematic way, as long as the joint probability for all the relevant propositions can be formulated.

THE BASELINE MODEL

It can be beneficial to think of the following model and discussion in terms of the following simple thought-experiment: We have an urn, which may contain only black (*B*) and white (*W*) balls. We have several testimonies of black balls (common events), drawn out of the urn, and one testimony of a white ball (an uncommon event). We will write down a probabilistic model for the case and be able to infer both the best estimates for the frequency of black and white balls in the urn and best estimates for the reliability of the witnesses. Based on these, we should be able answer the following: What is the probability that the one ball, testified to be white, was in fact white?

The general case is this: There are *n* testimonies *t* of a common event *B*, $t(B)^n$, and one testimony for an uncommon event *W*, $t(W)$. What is the probability of *W* being in fact true given the testimonies, $p(W \mid t(W)\ t(B)^n)$?

We will assume as little as we can about the probability of the witnesses being wrong, *d*, and about the real probability of the uncommon event happening, *v*. In effect, we will assign only reasonable prior probabilities for these probabilities and in the end let the data decide the most probable values for these probabilities. (These kinds of priors are often called hyperpriors in the data-analysis literature.)

For simplicity, we will use uniform priors $v \sim U(0,1)$, $d \sim U(0,0.2)$ where the notation $x \sim U(\mathrm{a},\mathrm{b})$ denotes that x is distributed uniformly between a and b, i.e. that the probability density for x is constant between a and b and zero elsewhere. With the latter prior we have assumed that in general the testimonies are over 80% reliable, an assumption which will be seen to matter less and less as n increases.

In this case, the joint distribution factors as (see Appendix A for details)

$$p(\,W\,C^n\,t(W)\,t(B)^n\,v\,d\,) = p(v)\,p(d)\,p(W|v)\,p(C^n|v)\,p(t(W)\mid W\,d)\,p(t(B)^n|C^n d)$$

And the wanted probability is

$$p(W \mid t(W)\,t(B)^n) = \frac{p(W\,t(W)\,t(B)^n)}{p(t(W)\,t(B)^n)} = \frac{p(W\,t(W)\,t(\mathrm{B})^n)}{p(W\,t(W)\,t(B)^n) + p(B\,t(W)\,t(B)^n)} = \frac{J_W}{J_W + J_B}$$

where we have terms of the form (by marginalization)

$$J_W = p(W\,t(W)\,t(B)^n) = \sum_{\forall C^n} \int_0^1 \mathrm{d}v \int_0^{0.2} \mathrm{d}d\; p(W\,\mathrm{C}^n\,t(W)\,t(B)^n\,v\,d)$$

where the sum is over all possible combinations of the elements of C^n, that is, we marginalize over all the possibilities in $(W_1 \cup B_1) \cap (W_2 \cup B_2) \cap \ldots \cap (W_n \cup B_n)$. After calculations, the terms amount to (see the Appendix A for details)

$$J_W = \int_0^1 \mathrm{d}v \int_0^{0.2} \mathrm{d}d\; v(1-d)(2dv - d - v + 1)^n$$

and

$$J_B = \int_0^1 \mathrm{d}v \int_0^{0.2} \mathrm{d}d\;(1-v)\,d\,(2dv - d - v + 1)^n$$

which are numerical functions of n. With these terms in hand, we are now in the position to show some results.

Results of the baseline model

To reiterate, in previous works (see e.g. (Earman 2000)), it has been shown that when n common events are taken as data, simple Bayesian inference with reasonable priors assigns a $1/(n+2)$ probability for the uncommon event happening. This simple case of inductive inference does not take into account the testimonies for the events (common or uncommon), as is done in the current model.

Figure 1 shows the probability for the uncommon event, with one testimony for the uncommon event, as a function of n, the number of testimonies for the common event, $p(W \mid t(\mathrm{W})\, t(\mathrm{B})^n)$. Perhaps surprisingly, as the number of testimonies for the common event (n) grows large, the probability for the uncommon event given the testimonies approaches the value 0.5 asymptotically. This represents a significant difference to the results of the simple inductive inference mentioned above, where the probability approaches zero asymptotically.

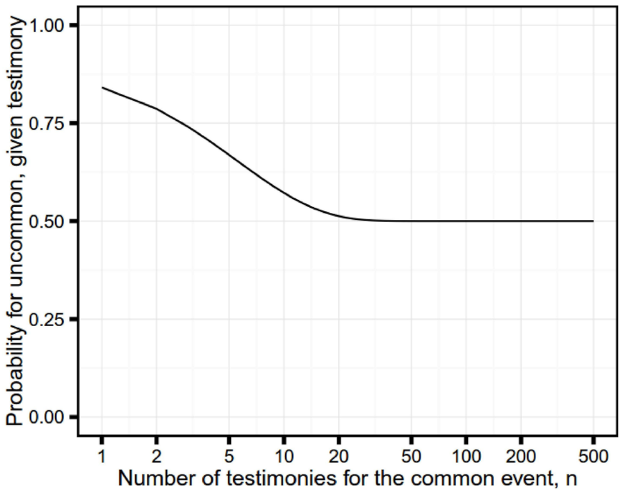

Figure 1. Probability for the uncommon event with one testimony for it, in the face of n testimonies for past common events. Note the logarithmic horizontal axis.

What is the reason for the difference of the results for the present more realistic model? Why does even one testimony for an uncommon event overcome the inductive part of the inference from the large amount of common events?

The basic reason is that, for there to be a large consistent amount of testimonies for the common events, the testimonies themselves have to be reliable. That is, if the testimonies were unreliable, it would be unlikely to have a uniform set of testimonies for the common case. Rather, there would likely be some testimonies for the uncommon event.

Note that the available testimonies affect the likely unreliability of testimonies in the model. So, if the testimonies are homogenous, it will be very likely that testimonies are reliable (low d), and the testimonies have a uniform source (low v). So, if there are e.g. 10^6 testimonies for the common event, the probability for an uncommon event will be of the order of 10^{-6}, this being the probability for an uncommon event without a testimony for it. However, if there is one testimony for an uncommon event, the likelihood for that testimony being wrong will also be of the order of 10^{-6}.

Figure 2 shows the mean values of the probability of the uncommon event happening (v) and of the probability of a testimony being false (d) for n testimonies for common events and one uncommon event. It is seen that as the number of testimonies (n) for the common event increases, the probability of the uncommon event decreases as expected, but at the same time the probability of a false testimony also decreases, and roughly at the same rate. Hence, even one testimony for an uncommon event is able to balance out the inductive part of the inference and make the uncommon event almost believable.

On the other hand, if there are more past testimonies for the uncommon event, the inductive part of the inference will not be so strong against the uncommon events, resulting in a not-so-low v.

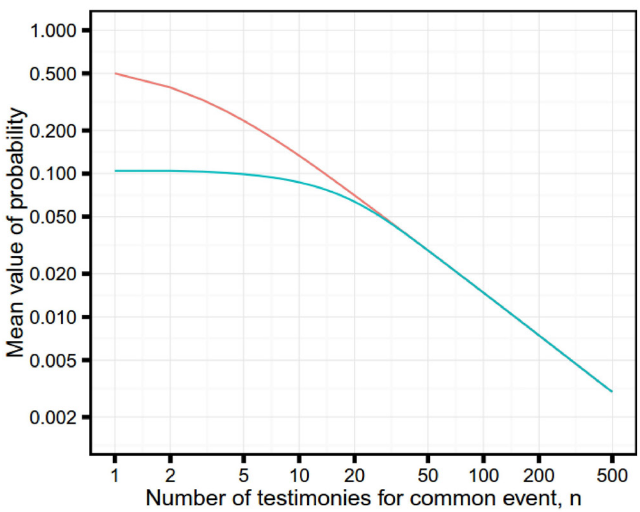

Figure 2. Mean values for the probabilities for the uncommon event (red) and false testimony (blue). Note the logarithmic axes.

APPENDED CASE WITH KNOWN ERRONEOUS TESTIMONIES

Let us now append the previous case by including an l amount of false testimonies for the uncommon event. Our additional data is then $B^l \cap t(W)^l$. The probability we will be interested in is $p(W \mid t(W)\ \ t(B)^n\ B^l\ t(W)^l)$.

The joint distribution will now factor as (see Appendix B for more details)

$$p\big(W\ C^n\ t(W)\ t(B)^n\ B^l\ t(W)^l\ v\ d\big) = p(v)\ p(d)\ p(W|v)\ p(C^n|v)p\big(B^l\big|v\big)\ p(t(W)\mid W\ d)\ p\big(t(W)^l\big|B^l\ d\big)\ p(\ t(B)^n|C^n\ d)$$

The calculations will proceed as before, with some additional terms. The wanted probability is again of the form

$$p(W \mid t(W)\ \ t(B)^n\ \ B^l\ t(W)^l) = \frac{J_W{}'}{J_W{}' + J_B{}'}$$

where

$$J_W{}' = \int_0^1 \mathrm{d}v \int_0^{0.2} \mathrm{d}d\ v(1-v)^l(1-d)d^l(2dv-d-v+1)^n$$

$$J_B{}' = \int_0^1 \mathrm{d}v \int_0^{0.2} \mathrm{d}d\ (1-v)^{l+1}\ d^{l+1}\ (2dv-d-v+1)^n$$

Results for the case with erroneous testimonies

Figure 3 shows the probability for the uncommon event given the testimonies for the appended case. Shown are cases with the number of known false testimonies $l = 0$, 1, 3, 10, and 50.

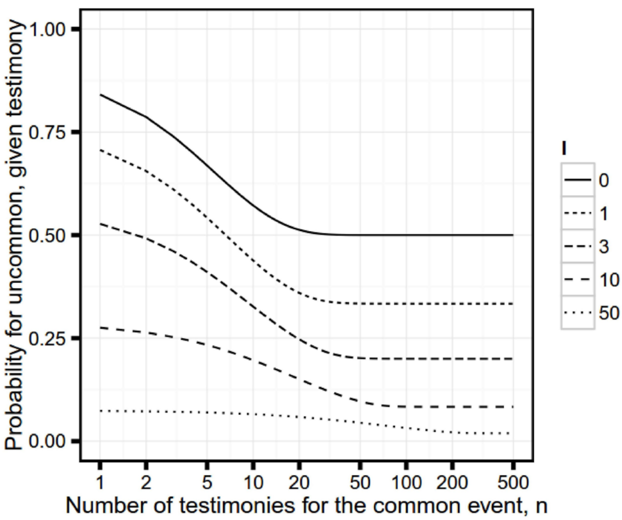

Figure 3. Probability for the uncommon event with one testimony for it, in the face of n testimonies for a common event, and different amounts of known false testimonies. Note the logarithmic horizontal axis.

Figure 4 shows the mean values of the probabilities of the uncommon event happening (v) and for a testimony being false (d) for cases with different number of known false testimonies for the uncommon event.

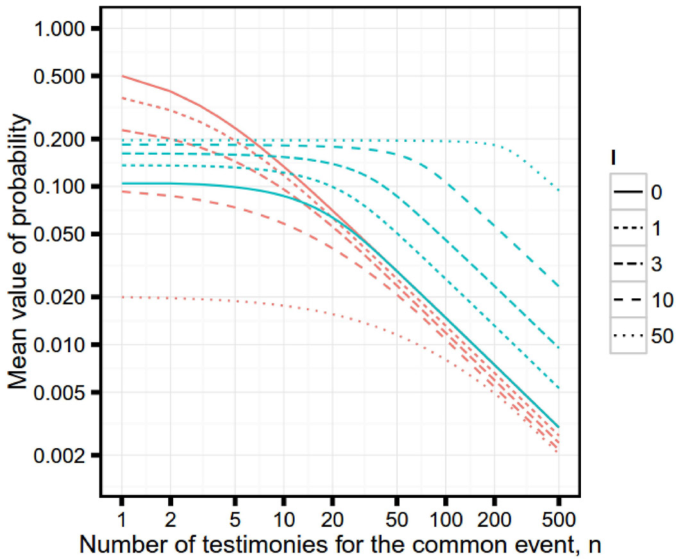

Figure 4. Mean values for the probabilities for the uncommon event (red) and false testimony (blue), given a number l of known false testimonies. Note the logarithmic axes.

One can see from the results that a small amount of known false testimonies for uncommon events does not significantly alter the believability of an uncommon event for which one testimony is not known to be false. For example, with three known false testimonies for an uncommon event and a large number of testimonies for common events, the probability for an uncommon event given one testimony for it is still roughly 0.2.

DISCUSSION

In the present model only a few assumptions were made and e.g. the probabilities for an uncommon event and the testimonies were left open and decided based on the available data. Yet, and importantly, it was assumed that the probability of a false testimony is symmetric, that is, that it is as likely for a person to make a mistake in the testimony for an uncommon as in the testimony for a common event. Hence, the number of testimonies for a common event had a bearing on the reliability on testimonies in general and hence also for the testimony for the uncommon event. It might be tempting to disconnect the two probabilities or to assume that a testimony for an uncommon event is more likely to be false than a testimony for a common event. While the former is possible, it would be hard to maintain that there is no connection between the reliability for the testimonies of uncommon and common events, the disconnection possibly leading to absurd results for low values of n. The latter option of assuming that the testimonies for uncommon events are less reliable seems biased. Because such an assumption would equate bringing more information to bear on the case, there should be a clear and agreed-on grounding for making this assumption. The author suspects that such an assumption is not sustainable, but leaves that for further, more nuanced, discussions.

For simplicity, the present model has two outcomes, a common and an uncommon event. In several applications, there will be many different kinds of common and uncommon outcomes. In this case, as noted earlier, different kinds of common or uncommon events can be combined into a single set of outcomes and the analysis will then proceed the same as for the binary case. One may want to use a different prior distribution for v, e.g. a distribution concentrated more on values above 0.5 would give more probability for common outcomes, which might be reasonable if they are more numerous in type. However, as long as the prior is reasonable, giving a non-zero possibility for both outcomes, the effect of this choice will dwindle as the number of testimonies increases. In this case one might also relax the assumption about the symmetricity of probabilities for testimonies being wrong. Whether or not one should relax the assumption here, depends on the model of testimony one might entertain. If the source of the testimony is taken to be in some sense intrinsically random, with errors being made uniformly from the truth to random outcomes, distributed uniformly to all possible outcomes, the option with more numerous outcomes will be more falsely testified-to. In cases where the types of common outcomes are more numerous, this would make uncommon events more plausible given a testimony for it, and vice versa. Note that the present model does not use probabilities for testimonies, but instead probabilities for a testimony when the true event is known. That is, the model does not use terms of the form $p(t(B))$, which are highly dependent on frequencies of occurrence, but terms of the form $p(t(B)|W)$, which are not dependent on frequencies of occurrence (v). For cases where testimony is less random in this sense, i.e. not heavily dependent on the number of different possible outcomes, the present simple assumption of symmetric probability for false testimonies is quite plausible.

For the model with known false testimonies for the uncommon event, the false testimonies might be viewed as a reason to relax the symmetry of the reliability of the testimonies of common and uncommon events. This exercise and grounding thereof is also left for further study on the matter.

For most readers, the present result will be intuitively acceptable in well-defined and simple cases. Take the urn with black and possibly some white balls. However many testimonies of black balls drawn from the urn, already one testimony of a white ball is enough to make the claim of a white ball almost believable. Similarly, while most testimonies report white swans and black crows, a bird watcher's testimony of a black swan or a white crow should be met with an open mind (Taleb 2007). Indeed, in the words of John Stuart Mill: "No amount of observations of white swans can allow the inference that all swans are white, but the observation of a single black swan is sufficient to refute that conclusion". Also along these lines, the decades of crystallographic experiments and articles observing only periodic crystals should not have been taken as strong evidence against D. Shechtman's reports about quasicrystals (D. Shechtman et al. 1984; Daniel Shechtman 2013).

In these easy-to-agree cases we do not have additional underlying assumptions or information against the uncommon event happening, or against the reliability of the particular witness. The inference is based purely on induction from testimonies of many common events. As was demonstrated, the strength of proper inductive inference from testimonies is rather weak and can be (almost) overcome with one (average) witness to the contrary.

Because the baseline results is 0.5 probability for an uncommon event based on one (normal) testimony, it is often the particulars of each case that decide whether or not we should believe the claim. Hence, it is not so much induction but additional information or assumptions about reality or about the reliability or unreliability of a witness that make claims about uncommon events believable or unbelievable. In practice, our background knowledge leads to an intuition about the believability of a claim. Upon enlarging on the present baseline result, our aim will be to codify relevant parts from our background information to probabilistic form, enabling inference with transparent assumptions in each particular case.

Let us briefly consider a claim of an uncommon event which is intuitively almost certainly false: Let's say you walk down the street, and a man in a tin-foil hat tells you that the government is using orbital mind control rays to make everyone an unwitting slave to Beyoncé. This is the only testimony you have received to this effect and every other piece of testimony you have supports the negation of this claim. Should you now adopt a surprisingly high credence that the man in the tin foil hat is correct?

Here, we are apt to disbelieve the claim, not only because of induction, which would only take us to a tie, but because the particulars of the case make the claim unbelievable: the unfashionable hat, the lack of motive for the government to make Beyoncé a puppet-master of the world, most people in fact not being unwitting slaves to Beyoncé, the intuitive and scientific difficulties related to controlling particular thoughts with physical processes, let alone with "rays", and so on. Each of these particulars would lower the probability of the claim by one to several orders of magnitude. So, we should disbelieve the claim, partly based on induction, and importantly because of the several improbable particulars.

This is of course analogous to the criminal cases in the courts. Whether or not a defendant has a criminal history matters only somewhat. While the defense attorney could bring thousands of testimonies about events of non-criminal conduct on part of the defendant, this, the inductive part of inference, would matter little. The particulars of the case will matter much more.

Similarly, in the case of supernatural miracles, it is likely not be the inductive part of inference which makes a miracle claim intuitively unbelievable to many, but additional assumptions (or lack thereof) about underlying reality. Thus, in the case of supernatural miracles, it is likely that background assumptions about unbreakable natural laws and closed systems (*ceteris paribus* assumptions), are the main reason for miracle claims being intuitively unbelievable.

CONCLUSIONS

The main result of the paper is that, when we have a large amount of testimonies for a common event and even only one testimony for an uncommon event, the probability we should assign for the uncommon event is surprisingly large, namely 0.5. This is assuming that, without information to the contrary, we are treating all the testimonies the same way, and we are not assuming additional structure for reality behind the events.

The result underlines the fact that trying to put limits on what can happen or will happen in the future, based on what has been testified to have happened in the past, is always an uncertain form of inference, far from the comforting certainty often ascribed to it. Sweeping Humean claims against uncommon events based on numerous common events are, in fact, incorrect.

In addition, for the case with some known-false testimonies for the uncommon event, the probability for the uncommon event is lower but not significantly so. Hence, the additional Humean argument against uncommon events based on some false testimonies of uncommon events does not seem to have much force either.

The result is also relevant for science; We should be more open to testimonies for "weird" empirical results which may not be in line with previous measurements or the current theoretical understanding. Daniel Shechtman's discovery of quasicrystals represents one such case.

BIBLIOGRAPHY

Ahmed, Arif. 2015. "Hume and the Independent Witnesses." *Mind* 124 (496): 1013–44. doi:10.1093/mind/fzv076.

Earman, John. 2000. *Hume's Abject Failure : The Argument Against Miracles*. Oxford University Press, USA.

Gelman, Andrew, John B Carlin, Hal S Stern, and Donald B Rubin. 2003. *Bayesian Data Analysis, Second Edition*. Boca Raton: Chapman & Hall/CRC.

Hume, David. 1748. *An Enquiry Concerning Human Understanding*. London: A. Millar. https://ebooks.adelaide.edu.au/h/hume/david/h92e/.

Millican, Peter. 2013. "Earman on Hume on Miracles." In *Debates in Modern Philosophy: Essential Readings and Contemporary Responses*, edited by Stewart Duncan and Antonia LoLordo. Routledge.

Neapolitan, Richard E. 2004. *Learning Bayesian Networks*. Pearson Prentice Hall.

Pearl, Judea. 1997. "Bayesian Networks." *UCLA Computer Science Department, Technical Report* R246: 1–5.

Shechtman, D., I. Blech, D. Gratias, and J. W. Cahn. 1984. "Metallic Phase with Long-Range Orientational Order and No Translational Symmetry." *Physical Review Letters* 53 (20): 1951–53. doi:10.1103/PhysRevLett.53.1951.

Shechtman, Daniel. 2013. "Quasi-Periodic Crystals-the Long Road from Discovery to Acceptance." *Rambam Maimonides Medical Journal* 4 (1). Rambam Health Care Campus: e0002. doi:10.5041/RMMJ.10102.

Sivia, D S. 1996. *Data Analysis - A Bayesian Tutorial*. Oxford: Clarendon Press.

Taleb, Nassim Nicholas. 2007. *The Black Swan : The Impact of the Highly Improbable*. *Penguin Books Ltd*. doi:10.5465/AMP.2011.61020810.

APPENDIX A DETAILED CALCULATIONS FOR THE BASELINE MODEL

Joint factorization

Figure A1 gives the dependencies between the parameters of the model as a directed acyclic graph (Pearl 1997; Neapolitan 2004)

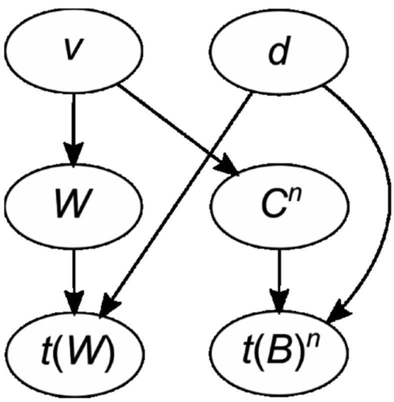

Figure A1. Directed acyclic graph of the case.

The arrows in the graph represent direct probabilistic dependencies between the parameters of the model. The natural factorization of the joint distribution can be read from the DAG (Neapolitan 2004) to be

$$p(\,W\,C^n\,t(W)\,t(B)^n\,v\,d\,) = p(v)\,p(d)\,p(W|v)\,p(C^n|v)\,p(t(W)\mid W\,d)\,p(t(B)^n|C^n d)$$

Summation over possibilities of C^n

Recall that in the simple model we have two terms of the form

$$J_W = p(W\,t(W)\,t(B)^n) = \sum_{\forall C^n}\int_0^1 \mathrm{d}v \int_0^{0.2} \mathrm{d}d\, p(W\,\mathrm{C}^n\,t(W)\,t(B)^n\,v\,d)$$

In this section we will calculate this term, notably the sum over all possibilities of C^n. Now

$$\begin{aligned} J_W &= \sum_{\forall C^n}\int_0^1 \mathrm{d}v \int_0^{0.2} \mathrm{d}d\, p(v)\,p(d)\,p(W|v)\,p(C^n|v)\,p(t(W)\mid W\,d)\,p(t(B)^n|C^n d) \\ &= \int_0^1 \mathrm{d}v\, p(v)\,p(W|v) \int_0^{0.2} \mathrm{d}d\, p(d)\,p(t(W)\mid W\,d)\, \sum_{\forall C^n} p(C^n|v)\,p(t(B)^n|C^n d) \\ &= c\int_0^1 \mathrm{d}v\;\, p(W|v) \int_0^{0.2} \mathrm{d}d\, p(t(W)\mid W\,d)\;\, S_n\;, \end{aligned}$$

where the constant c is a product of the constant priors of v and d, and

$$p(W|v) = v$$

$$p(t(W)\mid W\,d) = 1-d$$

$$S_n = \sum_{\forall C^n} p(C^n|v)\,p(t(B)^n|C^n d) = (1-v-d+2vd)^n$$

The following is an inductive proof for the last identity:

For $S2$, the sum is over the possibilities $(W_1 \cup B_1) \cap (W_2 \cup B_2)$

$$\begin{aligned} S_2 &= p(W_1\,W_2 \mid v)\,p(t(B_1)t(B_2) \mid W_1\,W_2\,d) + p(W_1\,B_2 \mid v)\,p(t(B_1)t(B_2) \mid W_1\,B_2\,d) \\ &\quad + p(B_1\,W_2 \mid v)\,p(t(B_1)t(B_2) \mid B_1\,W_2\,d) + p(B_1\,B_2 \mid v)\,p(t(B_1)t(B_2)\,|B_1\,B_2\,d) \\ &= v^2d^2 + 2v(1-v)d(1-d) + (1-v)^2(1-d)^2 = (1-v-d+2vd)^2. \end{aligned}$$

Next, with a lower case c_i we will denote the i'th element of C^n and similarly for $t(B)^n$. For S_n+1, we have

$$\begin{aligned} S_{n+1} &= \sum_{\forall C^{n+1}} p(C^{n+1}|v)\,p(t(B)^{n+1}|C^{n+1}d) = \sum_{\forall c_{n+1}}\sum_{\forall C^n} p(C^n c_{n+1}|v)\,p(t(B)^n t(B)_{n+1}|C^n c_{n+1}d) \\ &= \sum_{\forall c_{n+1}}\sum_{\forall C^n} p(C^n|v)\,p(c_{n+1}|v)\,p(t(B)^n|C^n d)\,p(t(B)_{n+1}|c_{n+1}d) \\ &= S_n \sum_{\forall c_{n+1}} p(c_{n+1}|v)p(t(B)_{n+1}|c_{n+1}d) = (1-v-d+2vd)^{n+1}\,\blacksquare \end{aligned}$$

APPENDIX B DETAILED CALCULATIONS FOR THE CASE WITH ERRONEOUS TESTIMONIES

Figure B1 gives the dependencies between the parameters of the model.

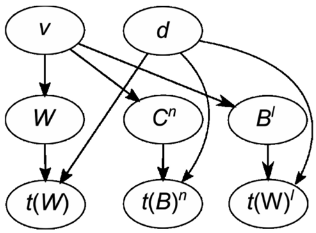

Figure B1. A directed acyclic graph of the model with erraneous testimonies.

Again, the joint distribution can be read from the graph to be

$$p\big(W\ C^n\ t(W)\ t(B)^n\ B^l\ t(W)^l\ v\ d\big)$$
$$= p(v)\ p(d)\ p(W|v)\ p(C^n|v)p\big(B^l\big|v\big)\ p(t(W)\mid W\ d)\ p\big(t(W)^l\big|B^l\ d\big)\ p(\ t(B)^n|C^n\ d)$$

And the wanted probability is

$$p(W\mid t(W)\ \ t(B)^n\ \ B^l\ t(W)^l) = \frac{p(W\ t(W)\ t(B)^n\ \ B^l\ t(W)^l)}{p(t(W)\ t(B)^n\ \ B^l\ t(W)^l)} = \frac{J_W{}'}{J_W{}' + J_B{}'}$$

Where

$$J'_W = \sum_{\forall C^n}\int_0^1 \mathrm{d}v \int_0^{0.2} \mathrm{d}d\ p\big(W\ C^n\ t(W)\ t(B)^n\ B^l\ t(W)^l\ v\ d\big)$$
$$= \int_0^1 \mathrm{d}v\ p(v)\ p(W|v)p\big(B^l\big|v\big) \int_0^{0.2} \mathrm{d}d\ p(d)\ p(t(W)\mid W\ d)\ p\big(t(W)^l\big|B^l\ d\big) \sum_{\forall C^n} p(C^n|v)\ p(\ t(B)^n|C^n\ d)$$
$$= c\int_0^1 \mathrm{d}v\ \ v(1-v)^l \int_0^{0.2} \mathrm{d}d\ \ (1-d)d^l\ S_n.$$

And similarly

$$J'_B = c\int_0^1 \mathrm{d}v\ \ (1-v)^{l+1} \int_0^{0.2} \mathrm{d}d\ \ d^{l+1}\ S_n.$$