DAVID CORFIELD

BAYESIANISM IN MATHEMATICS

INTRODUCTION

I shall begin by giving an overview of the research programme named in the title of this paper. The term 'research programme' suggests perhaps a concerted effort by a group of researchers, so I should admit straight away that since I have started looking investigating the idea that plausible mathematical reasoning is illuminated by Bayesian ideas, I have not encountered in the literature anyone else who has thought to develop the views of the programme's founder, the Hungarian mathematician, George Pólya. I should further admit that Pólya never termed himself a Bayesian as such. Motivation for the programme may, therefore, be felt sorely necessary. Let us begin, then, with three reasons as to why one might want to explore the possibility of a Bayesian reconstruction of plausible mathematical reasoning:

(a) To acquire insight into a discipline one needs to understand how its practitioners reason plausibly. Understanding how mathematicians choose which problems to work on, how they formulate conjectures and the strategies they adopt to tackle them requires considerations of plausibility. Since Bayesianism is widely considered to offer a model of plausible reasoning, it provides a natural starting point. Furthermore, Pólya has already done much of the spadework with his informal, qualitative type of Bayesianism.

(b) The computer has only recently begun to make a serious impact on the way some branches of mathematics are conducted. A precise modelling of plausibility considerations might be expected to help in automated theorem proving and automated conjecture formation, by providing heuristics to guide the search and so prevent combinatorial explosion. Elsewhere, computers are used to provide enormous quantities of data. This raises the question of what sort of confirmation is provided by a vast number of verifications of a universal statement in an infinite domain. It also suggests that statistical treatments of data will become more important, and since the Bayesian approach to statistics is becoming increasingly popular, we might expect a Bayesian treatment of mathematical data, especially in view of its construal of probability in terms of states of knowledge, rather than random variables.

(c) The plausibility of scientific theories often depends on the plausibility of mathematical results. This has always been the case, but now we live in an era where for some physical theories the *only* testable predictions are mathematical ones. If we are to understand how physicists decide on the plausibility of their theories, this must involve paying due consideration to the effect of verifying mathematical predictions.

D. Corfield and J. Williamson (eds.), Foundations of Bayesianism, 175–201.

Now, if one decides to treat plausible and inductive reasoning in the sciences in Bayesian terms, it seems clear that one would want to do the same for mathematics. After all, it would appear a little extravagant to devise a second calculus. In any case, Bayesianism is usually presented by its proponents as capable of treating all forms of uncertain reasoning. This leads us to conclude that Bayesianism in science requires Bayesianism in mathematics. Once this is accepted, one must respond in two ways according to the discoveries one makes while examining Bayesianism in mathematics:

I Bayesianism cannot be made to work for mathematics, therefore Bayesianism cannot give a complete picture of scientific inference.

II Some forms of Bayesianism can be made to work for mathematics, therefore one of these must be adopted by Bayesian philosophers to give a more complete picture of scientific inference.

The arguments presented in this paper indicate that the antecedent of I is false and the antecedent of II true, opening the prospect of an expanded, but modified, Bayesianism.

In this paper there is only space to treat a part of the motivation given above. The first two sections question which varieties of the many forms of Bayesianism are able to accommodate mathematical reasoning. Many Bayesians hold it as a tenet that logically equivalent sentences should be believed with equal confidence and any evidence should have an equal impact on their degrees of belief. However, such an assumption plays havoc with any attempt to throw light on mathematical reasoning. In section 1 I argue that if a Bayesian modelling of plausible reasoning in mathematics is to work, then the assumption of logical omniscience must be dropped.

In Pólya's version, we have only the right to specify the direction of change in the credence we give to a statement on acquiring new information, not the magnitude. However, Edwin Jaynes demonstrated that one of the central grounds for this decision on the part of Pólya to avoid quantitative considerations was wrong. In section 2 I consider whether there is anything amiss with a quantitative form of Bayesianism in mathematics.

One criticism often made of Bayesian philosophy of science is that it does not help very much in anything beyond toy problems. While it can resolve simple issues, such as accounting for how observing a white tennis shoe provides no confirmation for the law 'all ravens are black', it provides no insight into real cases of theory appraisal and confirmation. Everything rests on the assignment of priors, but how an expert could be considered to go about this is enormously complicated. Recognising what is correct in this criticism, I think there is still useful work to be done. In section 3 I shall be looking in particular at: reasoning by analogy; choice of proof strategy (for automated theorem proving); and, large scale induction (particularly enumerative induction).

1 PROBABILITY THEORY AS LOGIC

In his *Mathematics and Plausible Reasoning* (Pólya [1954a; 1954b]), Pólya considers mathematics to be the perfect domain in which to devise a theory of plausible reasoning. After all, where else could you find such unequivocal instances of facts satisfying general laws? As a noted mathematician actively engaged in research, he delightfully conveys inferential patterns by means of examples of his own use of plausible reasoning to generate likely conjectures and workable strategies for their proof. Now, such plausible reasoning in mathematics is, of course, necessary only because mathematics does not emerge as it appears on the pages of a journal article or textbook, that is, in its semi-rigorous deductive plumage. Indeed, it is due to the failure of what we might call "logical omniscience", the capacity to know immediately the logical consequences of a set of hypotheses, that mathematicians are forced to resort to what might be called a guided process of trial and error, not so dissimilar to that employed in the natural sciences.

In the second of the two volumes mentioned above, Pólya works his account of plausible reasoning into a probabilistic mould. While he did not name himself as such, we can thus reasonably view Pólya as a member of the Bayesian camp and, indeed, as a pioneer who influenced some later prominent Bayesians. Certainly, Edwin Jaynes learned from his work, and it is clear that Judea Pearl has read him closely. So here we have something of a paradox: plausible mathematical reasoning, the subject of Pólya's analysis, was an important source of ideas for some of the leading figures of Bayesianism, and yet it is necessitated by the fact that people involved in this most rigorous branch of knowledge are not able to uphold one of the widely held tenets of Bayesianism, namely, that logically equivalent statements should receive identical degrees of belief, or alternatively, that tautologies should be believed with degree of belief set at 1.

Logical omniscience comes as part of a package which views Bayesianism as an extension of deductive logic, as for example in Howson (this volume). In another of its manifestations, we hear from Jaynes the motto 'probability theory as logic'. For him: "Aristotelian deductive logic is the limiting form of our rules for plausible reasoning, as the robot becomes more and more certain of its conclusions" [Jaynes, forthcoming, Ch 2, p. 11].[1] Here we are to imagine a robot who reasons perfectly in Bayesian terms, handicapped only by the imperfections of its data and the incompleteness of the set of hypotheses it is considering.

We have then a tension when it comes to mathematical reasoning: if Bayesianism is to be seen as an extension of deductive logic, in the sense that the premises are now not required to be known with certainty, then one should consider the two inferential calculi to be similar in as many respects as possible. Since deductive logic is held as a regulating ideal, as, for example, when we say:

(1) If A is true and A entails B, then B is true,

[1] References to Jaynes are from his unfinished book — *Probability Theory: The Logic of Science* - available at http://bayes.wustl.edu . This is soon to appear in print.

should we not have

(2) If $Pr(A) = p$ and A entails B, then $Pr(B) \geq p$?

However, making this assumption raises a few problems. For one thing it implies that any consequence of a given axiomatised mathematical theory should be believed at least as strongly as that theory. Then, assuming Wiles is correct, to mimic the ideal rational agent you must set Pr(Fermat's Last Theorem) no lower than Pr(ZFC set theory), indeed no lower than your degree of belief in whichever system you feel confident can cope with arithmetic. There is of course the question as to how one might want to interpret Pr(ZFC set theory), but for statements whose logical complexity is the same as that of Fermat's Last Theorem, all one needs is the consistency of ZFC for truth to entail proof. And if you were to pitch this at 0.5, say, then this would provide a minimum for all provable truths of arithmetic, along with those of just about any other branch of mathematics, of this logical complexity.

If a mathematician suddenly became endowed with such omniscience, it would not be the end of mathematics, there is far more to mathematics than truth and provability, but one may safely predict that she would be much in demand. The logicistic conceptions of mathematics are accurate enough that the discipline would become unrecognisable. Without the sixty years leading up to Wiles' work, we would have known that (ZFC is consistent & Fermat's Last Theorem) is logically equivalent to (ZFC is consistent). And when finding a proof of a result we knew by omniscience to be correct, we could check up on the validity of lemmas rather than risk wasting time on false ones. How different a picture we gain from Pólya's representation of mathematics as a fallibly practised discipline and as the perfect place to investigate inductive and plausible reasoning.

So logical omniscience is an assumption that we cannot hold on to if we wish to investigate plausible reasoning in mathematics, which if Pólya was correct is perhaps what the Bayesian should be doing. But what prevents us from dropping this assumption? Two of the most common justifications for Bayesianism are Cox's theorem and the Dutch Book argument. Cox's theorem merely assumes that logical equivalence implies equality of probabilities. On the other hand, Dutch Book style arguments or those based on the preference for some linearly valued commodity attempt to justify it by claiming that if an agent offers different betting quotients on what are in fact logically equivalent sentences, then stakes can be set so that they will necessarily lose. But then isn't it surprising that there are many instances in the past where mathematicians have bet? Indeed, in view of the definitive way mathematical statements, even universal ones, may be settled, they would seem to make at least as good propositions to wager on as statements from the natural sciences.

Surely it is reasonable to prefer a bet on the trillionth decimal digit of π being between 0 and 8, than one at the same odds on its being 9. If, however, 9 is the correct digit, then it follows as a "mere" calculation from one of the series expansions for π. That is, "$\pi = 4(1 - 1/3 + 1/5 - 1/7 + \ldots)$" and "$\pi = 4(1 -$

$1/3 + 1/5 - 1/7 + \ldots$) & the trillionth decimal place of π is 9" would be logically equivalent and so to be believed with the same confidence, and so the second bet should be preferred. But, mathematicians spend their working lives making decisions on the basis of the level of their confidence in the truth of mathematical propositions. We would not want to brand them as irrational for devoting time to an attempted proof of their hunch that a certain statement follows from a set of assumptions merely because the hunch turns out to be wrong.

There is a suggestion in the writings of several Bayesians that (2) only holds when we *come to know* about the logical relationship between two propositions.

> Given two propositions A, B it may happen that one is true if and only if the other is true; we then say that they have the same *truth value*. This may be only a simple tautology (i.e., A and B are verbal statements which obviously say the same thing), or it may be that only after immense mathematical labors is it proved that A is the necessary and sufficient condition for B. From the standpoint of logic it does not matter; *once it is established*, by any means, that A and B have the same truth value, then they are logically equivalent propositions, in the sense that any evidence concerning the truth of one pertains equally well to the truth of the other, and they have the same implications for any further reasoning.
>
> Evidently, then, it must be the most primitive axiom of plausible reasoning that two propositions with the same truth-value are equally plausible. [Jaynes, forthcoming, Ch. 1, p. 6] (second emphasis mine)

In this less rigid framework we might say that if A is known by the agent to entail B, then she should ensure that she has $Pr(B) \geq Pr(A)$. In other words, we are generalising from an interpretation of deductive logic no stronger than:

(3) 'If I judge A to be true and I judge A to entail B, then I should judge B to be true.'

Opposed to the 'probability as logic' position are the subjectivists, whose number include followers of de Finetti. Here the accent is on uncertainty:

> The only relevant thing is uncertainty — the extent of our knowledge and ignorance. The actual fact of whether or not the events considered are in some sense *determined*, or known by other people, and so on, is of no consequence. [de Finetti, 1974, p. xi]

Since probability is seen as a measure of an individual's uncertainty, it is no wonder that de Finetti permits non-extreme degrees of belief about mathematical facts, even those which are decidable. Indeed, this probabilistic treatment seems to extend to even very accessible truths:

> Even in the field of tautology (i.e. of what is true or false by mere definition, independently of any contingent circumstances) we always

> find ourselves in a state of uncertainty. In fact, even a single verification of a tautological truth (for instance, of what is the seventh, or billionth, decimal place of π, or of what are the necessary or sufficient conditions for a given assertion) can turn out to be, at a given moment, to a greater or lesser extent accessible or affected with error, or to be just a doubtful memory. [de Finetti, 1974, p. 24]

Presumably then for de Finetti one may be rational and yet have a degree of belief in '91 is prime' less than 1. Perhaps you are unsure so you set it to 0.6. If so, when I ask you for $Pr(7 \times 13 = 91)$ you had better give me an answer no greater than 0.4. But then can't I force you to realise that you have an inconsistent betting quotient by making you see that 7×13 really is the same as 91, or is it just a case where I should allow you to alter your betting quotient after this lesson? More radically still, should one be expected to know that the correctness of this product contradicts the claim that 91 is prime?

In his article, *Slightly More Realistic Personal Probability*, Ian Hacking [1967] sets out a hierarchy of strengths of Bayesianism. These strengths he correlates with ways of saying whether a statement can be possibly true. At the weaker end we find a position he terms 'realistic personalism', where non-zero probabilities will be attributed by a subject to any statement not known by them to be false, knowledge being taken in a very strict sense: "a man can know how to use *modus ponens*, can know the rule is valid, can know p, and can know $p \supset q$, and yet not know q, simply because he has not thought of putting them together" [Hacking, 1967, p. 319]. At the stronger end we find logical omniscience and divine knowledge. Now clearly the coherence provided by realistic personalism is not enough to equip you for a life as a gambler. For instance, it is advisable not to advertise the odds at which you would accept either side of a wager on a decidable mathematical proposition on a mathematics electronic bulletin board. But Dutch Book arguments do not work to prove your irrationality on the grounds that someone may know more than you. If they do know more than you, you will lose whether the subject of your bet is mathematics, physics or the date of the next general election.

However, there is a point here: surely you can be criticised for betting on a proposition whose truth value you know you could discover with a modicum of effort, perhaps by the tap of a single button on the computer in front of you. As Hacking points out [Hacking, 1967, pp. 323–4], besides coherence one needs a principle which calls on you to maximize expected subjective utility. Information acquired for free can only help increase this, and so inexpensive reasoning or information gathering is generally a good thing. But this principle is not needed solely by a personalism weaker than that based on logical omniscience. Where the presupposition of logical omniscience forces you to reason, and indeed reason unreasonably much, it does not require you even to look down to note the colour of your socks before you bet on it. Only some principle of expected utility does this. But then surely you should allow this principle to be the rationale for your logical reasoning as well, rather than relying on the very unreasonable idealisation of logical omniscience which offers little more by way of advice than to be as perfect a mathematician as you can be.

Even admitting that we should not assume logical omniscience when we consider mathematics, it might be thought that this assumption is not too unrealistic for other walks of life. After all, doesn't the uncertainty which necessitates plausible reasoning in ordinary life and the natural sciences arise for other reasons – uncertainty in data, inaccessibility of object of study, incompleteness of background knowledge? You might think that it would count as the least of your worries that your logical and mathematical powers are not quite perfect. Hence, an assumption in standard Bayesian treatments of scientific inference that logically equivalent sentences should be accorded the same degree of belief. However, in many situations in science the uncertainty of mathematical knowledge plays an important part, as I have explained in a companion paper, not least in the area of mathematical predictions, a phenomenon as yet largely ignored by philosophers, where physicists gain confidence that they are on the right track when purely mathematical conjectures arising from their work turn out to be correct. Plausibility of scientific statements depends on uncertain mathematical knowledge.

To give briefly an indication of this, we hear of the mathematical physicist, Edward Witten, that he

> ...derived a formula for Donaldson invariants on Kähler manifolds using a twisted version of supersymmetric Yang-Mills theory in four dimensions. His argument depends on the existence of a mass gap, cluster decomposition, spontaneous symmetry breaking, asymptotic freedom, and gluino condensation. While none of this is rigorous by mathematical standards, the final formula is correct in all cases which can be checked against rigorous mathematical computations. [Freed and Uhlenbeck, 1995, p. 2]

Such confirmation increases your confidence in the power of your physical modelling. The more surprising the verified mathematical conjecture the greater the boost to your confidence.

It is interesting to wonder why nobody (at least to my knowledge) has taken Pólya up on his Bayesianism in mathematics. What is the underlying intuition behind the avoidance of a Bayesian treatment of plausible and inductive reasoning in mathematics? We can begin to understand what is at stake when we read Mary Hesse's claim that "...since mathematical theorems, unlike scientific laws, are matters of *proof*, it is not likely that our degree of belief in Goldbach's conjecture is happily explicated by probability functions." [Hesse, 1974, p. 191]. There are two responses to this. First, while it is true that the nature of mathematics is characterised like no other discipline by its possession of deductive proof as a means of attaining the highest confidence in the trustworthiness of its results, proofs are never perfectly secure. Second, and more importantly, what gets overlooked here is the prevalence in mathematics of factors other than proof for changing degrees of belief.

The lack of attention plausible mathematical reasoning has received reflects the refusal of most anglophone philosophers of mathematics to consider the way mathematical research is conducted and assessed. On the basis of this refusal, it is very easy then to persist in thinking of mathematics merely as a body of established truths. As classical deductive logic may be captured from a probability calculus which allows propositions to have probabilities either 0 or 1, the belief that mathematics is some kind of elaboration of logic and that the mathematical statements to be considered philosophically are those known to be right or wrong go hand in hand. We could say in fact that mathematics has suffered philosophically from its success at accumulating knowledge since this has deflected philosophers' attention from mathematics as it is being developed. But one has only to glance at one of the many survey articles in which mathematicians discuss the state of play in their field, to realise the vastness of what they know to be unknown but are very eager to know, and about which they may be thought to have degrees of belief equal neither to 0 nor to 1.[2]

We shall see in section 3 how mathematical evidence comes in very different shapes and sizes. But even remaining with 'proved' or well-established statements, although there would appear to be little scope for plausible reasoning, there are a number of ways that less than certain degrees of belief can be attributed to these results. David Hume described this lack of certainty well:

> There is no Algebraist nor mathematician so expert in his science, as to place entire confidence in his proof immediately on his discovery of it, or regard it as any thing, but a mere probability. Every time he runs over his proofs, his confidence encreases; but still more by the approbation of his friends; and is rais'd to its utmost perfection by the universal assent and applauses of the learned world. Now 'tis evident, that this gradual encrease of assurance is nothing but the addition of new probabilities, and is deriv'd from the constant union of causes and effects, according to past experience and observation. [Hume, 1739, pp. 180–1]

Perfect credibility may be difficult to achieve for proofs taking one of a number of non-standard forms, from humanly-generated unsurveyable proofs to computer-assisted proofs to probabilistic proofs. These latter include tests for the primality of a natural number, n. Due to the fact that around half of the numbers less than n are easily computed "witnesses" to its being composite, if such is the case, a small sample can quickly show beyond any set level of doubt whether n is prime.

While a certain amount of suspicion surrounds the latter type of 'proof', from the Bayesian perspective, one can claim that all evidence shares the property that it

[2] I mean to exclude here the immense tracts of totally uninteresting statements expressible in the language of ZFC in which one will never care to have a degree of belief. An idea of the plans in place for the expansion of (interesting) mathematics can be gleaned from the following claim: 'It is clear... that the set-based mathematics we know and love is just the tip of an immense iceberg of n-categorical, and ultimately ω-categorical, mathematics. The prospect of exploring this huge body of new mathematics is both exhilarating and daunting.' [Baez and Dolan, 1999, p. 32].

produces changes in some degrees of belief. The absence of any qualitative difference in the epistemic import of different types of proof has recently been noted by Don Fallis [1997], who considers many possible ways of distinguishing epistemically between deductive proofs and probabilistic proofs and finds none of them adequate. He draws the conclusion, therefore, that there is no such difference. Fallis centres his discussion around 'proofs' which involve clever ways of getting strands of DNA to react to model searches for paths through graphs, putting beyond reasonable doubt the existence or non-existence of such paths. Despite there being here a reliance on biochemical knowledge, Fallis still sees no qualitative difference as regards the justificatory power of this type of proof. Confidence in mathematical statements is being determined by natural scientific theory. This appears less surprising when you consider how complicated, yet well-modelled, configurations of silicon can be used to generate evidence for mathematical propositions.

Fallis's point may be expressed in Bayesian terms as follows.[3] The acceptability of a mathematical statement is dependent solely on your rational degree of belief in that statement conditionalised on all the relevant evidence. Whatever level you set yourself (0.99 or 0.99999) the type of evidence which has led you there is irrelevant. A ten thousand page proof may provide as much support as a probabilistic proof or the non-appearance of a counter-example. To contemplate the reliability of a result in a particular field we should think of someone from outside the field asking a specialist for their advice. If the trustworthy expert says she is very certain that the result may be relied upon, does it matter to the enquirer how the specialist's confidence arises? This depiction could be taken as part of a larger Bayesian picture. The very strong evidence we glorify with the name 'proof' is just as much a piece of evidence as is a verification of a single consequence. Bayesianism treats in a uniform manner not just the very strong evidence that Fallis considers, but all varieties of partial evidence. Let us now see what we are to make of this partial evidence.

2 QUANTITATIVE BAYESIANISM

Pólya understood plausible inference to be quite different from deductive logic. In his eyes [Pólya, 1954b, pp. 112–116], deductive logic is:

(a) Impersonal — independent of the reasoner;

(b) Universal – independent of the subject matter;

(c) Self-sufficient – nothing beyond the premises is needed;

(d) Definitive – the premises may be discarded at the end of the argument.

[3] He points out (private communication), however, that he is not necessarily committed to a Bayesian analysis of his position which assumes that one's rational degree of belief is all that really matters in mathematical justification.

On the other hand, plausible inference is characterised by the following properties:

(a) The direction of change in credibility is impersonal, but the strength may be personal;

(b) It can be applied universally, but domain knowledge becomes important for the strength of change, so there are practical limitations;

(c) New information may have a bearing on a plausible inference, causing one to revise it;

(d) The work of plausible inference is never finished as one cannot predict what new relevant information may arise.

One of the principal differences seems to be that in the deductive case nobody requires of you that you maximise the set of deductive consequences of what you hold to be certain. If you are asked whether you know the truth status of a statement, you search about for a proof or disproof of it from what you already know. If you find nothing, you just admit your ignorance, and no-one can accuse you of anything worse than stupidity if you have overlooked such a proof or disproof. We do not go around blaming ourselves for not having known before Wiles that Fermat's Last Theorem is provable, even though the resources were in some sense available. Deductive logic is there to safeguard you from taking a false step, not from omitting to take a correct step. On the other hand, we may use plausible inference to argue about the plausibility of *any* statement based on what we know at present.[4] The question is how to think about the way we go about arriving at degrees of belief on the basis of what we already know.

It is clear that the strength of a mathematician's belief in the correctness of a result has an impact on their practice: Andrew Wiles would hardly have devoted seven years to Fermat's Last Theorem had he not had great faith in its veracity. No doubt we could give a complicated Bayesian reconstruction of his decision to do so in terms of the utility of success, the expected utility of lemmas derived in a failed attempt, and so on. For a more simple example, we may give a Bayesian reconstruction of the following decision of the French Academy:

> The impossibility of squaring the circle was shown in 1885, but before that date all geometers considered this impossibility as so "probable" that the Académie des Sciences rejected without examination the, alas!, too numerous memoirs on this subject that a few unhappy madmen sent in every year. Was the Académie wrong? Evidently not,

[4]Jaynes [forthcoming, Ch. 10, p. 21] has a similar view on the difference between deductive logic and probability theory as logic: "Nothing in our past experience could have prepared us for this; it is a situation without parallel in any other field. In other applications of mathematics, if we fail to use all of the relevant data of a problem, the result will be that we are unable to get any answer at all. But probability theory cannot have any such built-in safety device, because in principle, the theory must be able to operate no matter what our incomplete information might be".

> and it knew perfectly well that by acting in this manner it did not run the least risk of stifling a discovery of moment. The Académie could not have proved that it was right, but it knew well that its instincts did not deceive it. If you had asked the Academicians, they would have answered: "We have compared the probability that an unknown scientist should have found out what has been vainly sought for so long, with the probability that there is one madman the more on earth, and the latter has appeared to us the greater. [Poincaré, 1905, pp. 191–2]

These alternatives, being mad and being right, were hardly exhaustive. Leaving aside the person's sanity we can contrast the probability that their proof is correct with the probability that it is incorrect.

Pr(proof correct| author unknown, I) =

Pr(proof correct| author unknown, true, I)·Pr(true|author unknown, I) +
Pr(proof correct| author unknown, false, I)·Pr(false|author unknown, I)

= Pr(proof correct| author unknown, true, I)·Pr(true| I)

where I is the background knowledge.

Substituting reasonable estimates of the Académie's degrees of belief will lead to a very small value for this last expression because its two factors are small. On the other hand, a submitted proof of the possibility of squaring the circle by a known mathematician, or a submitted proof of its impossibility by an unknown author would presumably have been dealt with more tolerantly.

Notice that this reconstruction would not seem to require one to go beyond vague talk of very high or very low probabilities. By contrast, when it comes to offering a betting ratio for the trillionth decimal digit of π being 9, it would seem to be eminently reasonable to propose precisely 1/10, and yet neither the coherence of realistic personalism nor any requirement to maximize expected subjective utility imposes this value upon you. What appears to determine this value is some form of the principle of indifference based on our background knowledge. With a simple grasp of the idea of a decimal expansion we simply have no reason to believe any single digit more likely than any other. Those who know a little more may have heard that there is no statistical evidence to date for any lack of uniformity in the known portion of the expansion, probably rendering them much less likely to be swayed in their betting ratio by a spate of 9s occurring shortly before the trillionth place. So, unless some dramatic piece of theoretical evidence is found, it seems that most mathematicians would stick with the same betting ratio until the point when they hear that computers have calculated the trillionth place.[5]

The issue to be treated in this section is whether we require a quantitative, or even algorithmic, form of Bayesianism to allow us to explicate plausible mathematical reasoning, or whether, like Pólya, we can make do with a qualitative form

[5] As of 1999 they had reached the 206 billionth.

of it. First, it will be helpful for us to contrast Pólya's position with that of Jaynes. For Jaynes, Pólya was an inspiration. Indeed, he

> ... was the original source of many of the ideas underlying the present work. We show how Pólya's principles may be made quantitative, with resulting useful applications. [Jaynes, forthcoming, Ch. 1, p. 3]

As is well known, Jaynes was at the objectivist end of the Bayesian spectrum. In other words, his aim was to establish principles (maximum entropy, transformation groups, etc.) applicable in as many situations as possible, in which a reasonable being could rationally decide on their prior probabilities. Pólya, on the other hand, reckoned that one would have to stay with a qualitative treatment (e.g., if A is analogous to B and B becomes more likely, then A becomes somewhat more likely), in that the direction of changes to confidence might be determined but not their strength. But Jaynes claimed that this decision was based on a faulty calculation made by Pólya when he was considering the support provided to Newton's theory of gravitation by its prediction of the existence and location of a new planet, now called Neptune. The incorrect calculation occurred when Pólya was discussing the boost to confidence in Newtonian gravitation brought about by the observation of a previously unknown planet precisely where calculations predicted it to be, based on observed deviations in Uranus's orbit.

Pólya takes Bayes theorem in the form,

$$\text{Pr(Newt. Grav.}|\text{Neptune)} = \frac{\text{Pr(Newt. Grav.)}\cdot\text{Pr(Neptune}|\text{Newt. Grav.)}}{\text{Pr(Neptune)},}$$

where Pr(Neptune) corresponds to a scientist's degree of belief that the proposed planet lie in the predicted direction. For the purposes of the calculation, he estimates Pr(Neptune) in two ways. First, he calculates the probability of a point lying within one degree of solid angle of the predicted direction, and arrives at a figure of $0.00007615 \approx 1/13100$. Second, on the grounds that the new planet might have been expected to lie on the ecliptic, he uses the probability of a point on a circle lying within one degree of the specified position, yielding a value for Pr(Neptune) of 1/180. He then argues that Pr(Newtonian Gravitation) must be less than Pr(Neptune), otherwise Bayes' theorem will lead to a posterior probability greater then 1, but that it is unreasonable to imagine a scientist's degree of belief being less than even the larger figure of 1/180, since Newtonian Gravitation was already well-confirmed by that point. He concludes, "We may be tempted to regard this as a refutation of the proposed inequality." [Pólya, 1954b, p. 132] and suggests we return to a safer qualitative treatment.

However, as Jaynes points out, Pólya's calculations were in fact of the prior to posterior odds ratio of two theories: on the one hand, Newtonian gravitation, and on the other, a theory which predicts merely that there be another planet, firstly anywhere and secondly on the ecliptic. Indeed, from the confirmation, Newtonian gravitation is receiving a boost of 13100 or 180 *relative* to the theory that there is one more planet somewhere. Pólya had forgotten that if Pr(Newtonian Gravitation) is already high then so too would Pr(Neptune) be.

We are told by Jaynes that Pólya realised his mistake and went on to participate vigorously in the former's lectures at Stanford University in the 1950s. However, Pólya had given several further arguments against quantitative plausible reasoning, so even if Jaynes could block this particular argument, one would need to confront the others. Reading through them, however, one notes that Pólya is making fairly standard points: the incomparability of evidence and conjectures, problems with the principle of indifference, etc.

Could it be that your background predisposes you to adopt a certain type of Bayesianism? The physicist relies on symmetry considerations pertaining to the mechanisms producing the data, the philosopher of science on vaguer considerations of theory evaluation, while the economist must integrate a mass of data with her qualitative, quasi-causal understanding of the economy. Are disputes among Bayesians like the blind men feeling different parts of an elephant?

Bayesianism applied to reasoning in the natural sciences appears to fall into two rather distinct categories:

(i) analysis of data from, say, nuclear magnetic resonance experiments or astrophysical observations;

(ii) plausible reasoning of scientists by philosophers of science (e.g., [Franklin, 1986]).

We may wonder how strong the relation is between them. Rosenkrantz [1977] attempted a unified treatment, and he indicates by his subtitle 'Towards a Bayesian Philosophy of Science' that a treatment of history and philosophy of science issues alongside statistical issues should be 'mutually enriching' [Rosenkrantz, 1977, p. xi].

Jaynes himself was less sure about how far one could take the historical reconstructions of scientific inference down a Bayesian route. After his discussion of Pólya's attempt to quantify Neptune discovery he claims:

> But the example also shows clearly that in practice the situation faced by the scientist is so complicated that there is little hope of applying Bayes' theorem to give quantitative results about the relative status of theories. Also there is no need to do this, because the real difficulty of the scientist is not in the reasoning process itself; his common sense is quite adequate for that. The real difficulty is in learning how to formulate new alternatives which fit better the facts. Usually, when one succeeds in doing this, the evidence for the new theory soon becomes so overwhelming that nobody needs probability theory to tell him what conclusions to draw. [Jaynes, forthcoming, Ch. 5, p. 17]

This note occurs in a chapter entitled 'Queer uses of probability', by which he intends that at present we have no rational means for ascribing priors. So, despite his professed debt to *Mathematics and Plausible Reasoning*, we find two poles of Bayesianism represented by Jaynes and Pólya. For Jaynes, any rational agent possessing the same information will assign identical probability functions. For Pólya, two experts with the same training may accord different changes to their degrees of belief on discovery of the same fact. One imagines a machine making plausible inferences, the other emphasises the human aspect.

Jaynes:

> ... instead of asking, "How can we build a mathematical model of human common sense?" let us ask, "How could we build a machine which would carry out useful plausible reasoning, following clearly defined principles expressing an idealized common sense? [Jaynes, forthcoming, Ch. 1, p. 5]

Pólya:

> A person has a background, a machine has not. Indeed, you can build a machine to draw demonstrative conclusions for you, but I think you can never build a machine that will draw plausible inferences. [Pólya, 1954b, p. 116]

Perhaps it is the lack of exactitude which steers Jaynes away from modelling scientific reasoning. After a lifetime investigating how symmetry considerations allow the derivation of the principles of statistical mechanics, it must be difficult to adapt to thinking about plausibility in complex situations.

But if a physicist might be excused, what of a philosopher? John Earman, while discussing how a physicist's degrees of belief in cosmological propositions were affected by the appearance of General Relativity on the scene, tells us

> But the problem we are now facing is quite unlike those allegedly solved by classical principles of indifference or modern variants thereof, such as E. T. Jaynes's maximum entropy principle, where it assumed that we know nothing or very little about the possibilities in question. In typical cases the scientific community will possess a vast store of relevant experimental and theoretical information. Using that information to inform the redistribution of probabilities over the competing theories on the occasion of the introduction of the new theory or theories is a process that is, in the strict sense of the term, *a*rational: it cannot be accomplished by some neat formal rules, or, to use Kuhn's term, by an algorithm. On the other hand, the process is far from *ir*rational, since it is 'informed by reasons. But the reasons, as Kuhn has emphasized, come in the form of persuasions rather than proof. In Bayesian terms, the reasons are marshalled in the guise of plausibility arguments. The deployment of plausibility arguments is an art form for which there currently exists no taxonomy. And in view of the limitless variety of such arguments, it is unlikely that anything

> more than a superficial taxonomy can be developed. [Earman, 1992, p. 197]

This seems a rather pessimistic analysis for a professed Bayesian. Does the 'limitless variety' of these arguments mean that we should not expect to find patterns among them? Despite the talk of their deployment being an 'art form', Earman does allow himself to talk about the objective quality of these plausibility arguments. Indeed, he claims that:

> Part of what it means to be an "expert" in a field is to possess the ability to recognize when such persuasions are good and when they are not. [Earman, 1992, p. 140]

Interestingly, it is Pólya the "expert" in mathematics who believes that it is possible to extract the patterns of good plausibility arguments from his field.

So, out of the three, Jaynes, Pólya and Earman, representatives of three different types of Bayesianism, it is Pólya who believes one can say something quite concrete about plausible reasoning. All realise that plausible reasoning is a very complex process. Neither Jaynes nor Earman cannot see a way forward with plausible scientific reasoning. This leaves Pólya who gets involved with real cases of (his own) mathematical reasoning, which he goes on to relate to juridical reasoning and reasoning about one's neighbour's behaviour. Is he right to claim that mathematics provides a better launch pad to tackle everyday reasoning than does science?

If we want a fourth Bayesian to complete the square, we might look to the computer scientist Judea Pearl. Like Pólya, Pearl believes we can formulate the principles of everyday common sense reasoning, and like Jaynes he thinks Bayesian inference can be conducted algorithmically. To be able to do the latter requires a way of encoding prior information efficiently to allow Bayesian inference to occur. For Pearl [Pearl, 2000], (*this volume*) humans store their background information efficiently in the form of causal knowledge. The representation of this causal knowledge in a reasonably sparse Bayesian network is the means by which a machine can be made to carry out plausible reasoning and so extend our powers of uncertain reasoning.

In his earlier work Pearl [1988] expressed his appreciation of Pólya's ideas, and yet found fault with his restriction to the elucidation of *patterns* of plausible reasoning rather than a *logic*. He considers Pólya's loose characterisation of these patterns not to have distinguished between evidence factoring through consequences and evidence factoring through causes. For instance, Pólya asserts that when B is known to be a consequence of A, the discovery that B holds makes it more likely that A holds. This, however, is a well known fallacy of causal reasoning. I see that the sprinkler on my lawn is running and that the grass is wet, but this does make it more probable to me that it has rained recently even though wet grass is a consequence of it having done so. But one need not remain with causal stories to reveal this fallacy. A consequence of a natural number being divisible by four is that it is even. I find that a number I seek is either 2 or 6. Although I have learnt that it is even, this discovery reduces the probability of its being divisible by 4 to zero. Essentially, what Pólya overlooked was the web-like nature of our beliefs, only departing from patterns involving two propositions when he considered the possibility of two facts having a common ground. In Bayesian networks, converging arrows are equally important but must be treated differently.

It remains to be seen whether the techniques of Bayesian networks may illuminate scientific inference. Now we shall turn our attention to examine what Bayesianism has to say about certain aspects of mathematical reasoning.

3 WHAT MIGHT BE ACHIEVED BY BAYESIANISM IN MATHEMATICS

Varieties of mathematical evidence may be very subtle, lending support to Earman and Jaynes' scepticism. Pólya [1954b, p. 111] himself had the intuition that two mathematicians with apparently similar expertise in a field might have different degrees of belief in the truth of a result and treat evidence for that result differently. Even though each found a consequence of the result equally plausible, the establishment of this consequence could have an unequal effect on their ratings of the likelihood of the first result being correct. The complex blending of the various kinds of evidence experienced through a mathematician's career would explain the differences in these reactions, some of which might be attributable to aspirations on the part of each of them either to prove or disprove the theorem. But Pólya goes further to suggest that such differences of judgement are based on "still more obscure, scarcely formulated, inarticulate grounds" [Pólya, 1954b, p. 111].

Such appeals to the inexpressible, or at least to the imprecisely expressed, are not at all uncommon. For example, the mathematician Sir Michael Atiyah asserts that

> ... it is hard to communicate understanding because that is something you get by living with a problem for a long time. You study it, perhaps for years. You get the feel of it and it is in your bones. [Atiyah, 1984, p. 305]

Such comments may have been devised by mathematicians to give an air of mystery to their practice. A sceptic could point out that doctors have done likewise in the past by alleging that diagnosis requires some profound intuitive faculty of divination, an attractive image shattered by the successful construction of expert systems which have shown physicians to be replaceable in some situations, by machines using propositionally encoded evidence. However, the success of artificial intelligence in some areas of medical diagnosis may be contrasted with the extreme difficulty in getting computers to do anything that might be termed creative in mathematics.[6] The essential point does not concern whether or not

[6] A possible exception is the recent successful automated solution of the Robbins problem (see http://www.mcs.anl.gov/~mccune/), drawn to my attention by Don Fallis.

mathematicians in fact rely on non-propositional knowledge, so much as whether there might be something about this type of knowledge which is indispensable to doing mathematics.

Certainly, evidence for the correctness of a statement may be very subtle. It may even arise through an experience of failure. In [Corfield, 1997] I pointed out the inaccuracy on Lakatos's part of his notion of lemma-incorporation, the idea that faulty proofs are generally repaired by some simple insertion of a lemma. As I explained there, while proving the so-called 'duality theorem' Poincaré had come to realise that an assumption he was making about the way differential manifolds intersect was invalid in general. However, he still believed that the general strategy could be made to work of constructing for a given set of manifolds of equal dimension a manifold of complementary dimension which intersected each of the members of the set exactly once. He just needed to have the intersections occur in a more controlled fashion. One can only guess how this experience impacted on his degree of belief in the duality theorem. It is quite probable that even though the initial proof was found to be wrong, the experience of near success with a variant of a strategy gave him hope that another variant would work. It must also happen, however, that mathematicians are discouraged by such setbacks.

Evidence can also involve the non-discovery of something as Sherlock Holmes well knew when he built his case on the observation of a dog that did not bark. The classic example of the unsurveyable human-generated kind of proof at the present time is the proof of the classification of finite simple groups into 5 infinite families and 26 sporadic outsiders. How does one's degree of belief in this result depend on such potentially flawed lengthy evidence? Fallis [1997] has Gorenstein, the driving force behind the collective proof, confessing that confidence is boosted less by the proof itself than by the fact that no other such groups have been found. Similarly, remarks are often to be heard concerning the consistency of ZFC that one would have expected to have encountered a contradiction by now.

We should also remember that evidence for mathematical propositions comes from sources which have only recently become available. The use of computers to fill in the gaps of human proofs has become acceptable, but computers are used in many other ways in mathematics. For example, they provide evidence for conjectures via calculations made on samples, and they produce visual evidence in dynamical systems theory, as in the drawing of attractors or power spectra. Reliance on computer evidence raises some novel issues. Oscar Lanford is attributed with pointing out that

> ... in order to justify a computer calcūlation as part of a proof..., you must not only prove that the program is correct (and how often is that done?) but you must understand how the computer rounds numbers, and how the operating system functions, including how the time-sharing system works. [Hirsch, 1994, p. 188]

Moreover, if more than one piece of computer evidence is being considered, how do we judge how similar they are for conditionalising purposes? This would require one to know the mathematics behind any similarities between the algorithms utilised.

It is clear then that any account of mathematical inference will require a very expressive language to represent all the various forms of evidence which impact on belief in mathematical propositions. The Bayesian wishing to treat only propositions couched in the language of the object level might hope to be able to resort to Jeffrey conditionalisation, but this comes at the price of glossing over interesting features of learning. Concerning scientific inference, Earman [1992, pp. 196–8] asserts that many experiences will cause the scientist to undergo *non-Bayesian* shifts in their degrees of belief, i.e., ones unaccountable for by any form of algorithmic conditionalisation. These shifts, the resetting of initial probabilities, are very common, he claims, arising from the expansion of the theoretic framework or from the experience of events such as "[n]ew observations, even of familiar scenes; conversations with friends; idle speculations; dreams..." [Earman, 1992, p. 198]. One might despair of making any headway, but taking Pólya as a guide we may be able to achieve something. While recognising that making sense of plausible reasoning in mathematics will not be easy, I believe that three key areas of promise for this kind of Bayesianism in mathematics are analogy, strategy and enumerative induction.

3.1 Analogy

Before turning to a probabilistic analysis of plausible reasoning in the second volume of *Plausible Reasoning*, Pólya had devoted the first volume [1954a], as its subtitle suggests, to the themes of analogy and induction. Analogies vary as to their precision. When vague they contribute to what he called the *general atmosphere* surrounding a mathematical conjecture, which he contrasts to pertinent *clear facts*. While verifications of particular consequences are straightforwardly relevant facts, the pertinence of analogical constructions may be hard to discern precisely. Nevertheless, mathematicians, succh as Gian-Carlo Rota, take them to be vitally important:

> The enrapturing discoveries of our field systematically conceal, like footprints erased in the sand, the analogical train of thought that is the authentic life of mathematics. [Kac *et al.*, 1986, p. ix]

Let us illustrate this with an example. At the present time the vast majority of mathematicians have a high degree of belief in the Riemann Hypothesis. Recall that the Riemann zeta function is defined as the analytic continuation of $\zeta(s) = \Sigma n^{-s}$ summed over the natural numbers, and that the hypothesis claims that if s is a zero of $\zeta(s)$, then either $s = -2, -4, -6, \ldots$, or the real part of s equals 1/2. Many roots have been calculated (including the first 1.5 billion zeros in the upper complex plane along with other blocks of zeros), all confirming the theory, but despite this "overwhelming numerical evidence, no mathematical proof is in sight." [Cartier, 1992, p. 15]. As Bayesians have explained, there are limits to the value of showing that your theory passes tests which are conceived to be very similar. If, for example, a further 100 million zeros of the zeta function are found to have their real part equal to 1/2, then little change will occur in mathematicians' degrees of belief, although a little more credibility would be gained if this were true of 100 million zeros around the 10^{20} th, which is precisely what has happened.

In this example the clear facts making up the numerical evidence can lend only limited credence by themselves. After all, there are 'natural' properties of the natural numbers which are known to hold for exceedingly long initial sequences. What counts in addition beyond evidential facts, however numerous, is the credibility of stronger results, general consequences and analogies. Indeed, if an analogy is deemed strong enough, results holding for one side of it are thought to provide considerable support for their parallels. Concerning the Riemann conjecture (RC), we are told that:

> There is impressive numerical evidence in its favour but certainly the best reason to believe that it is true comes from the analogy of number fields with function fields of curves over finite fields where the analogue of RC has first been proved by A. Weil. [Deninger, 1994, p. 493]

This analogy[7] was postulated early in this century as a useful way of providing a halfway house across an older analogy, developed by Dedekind and Weber, from algebraic number fields to function fields over the complex numbers. However, the translation techniques between the three domains have still not been perfected. The more geometric side of the analogy Deninger mentions was able to absorb cohomological techniques, allowing Weil to prove the Riemann hypothesis analogue in 1940. An extraordinary amount of effort has since been expended trying to apply cohomology to number theory (Weil, Grothendieck, Deligne, etc.) with the establishment of the standard Riemann hypothesis as one of its aims.

How should we judge how analogous two propositions, A and B, are to each other? For Pólya [1954b, p. 27] it correlates to the strength of your "hope" for a common ground from which they both would naturally follow. Increase in confidence in A will then feed up to the common ground, H, and back down to B.[8] If analogy is to be treated anywhere, I believe mathematics will provide a good location, since there are plenty of excellent examples to be found there. In Pólya's principal example, Euler noticed that the function sin x/x resembles a polynomial in several respects: it has no poles; it has the right number of zeros, which do not accumulate; it behaves symmetrically at $\pm\infty$. On the other hand, unlike a polynomial, sin x/x remains bounded. Even with this disanalogy, it seemed plausible that polynomials and sin x/x would share other properties. One notable feature of complex polynomials is that any one of them may be expressed as a product

[7] See also [Katz and Sarnak, 1999], in particular the table on page 12.

[8] Notice here the flavour of a Bayesian network: H pointing to both A and B.

of factors of the form (1 − x/root), taken over all of its roots. Might this property also apply to sin x/x? Well, the roots of this function are $\pm\pi, \pm 2\pi, \pm 3\pi, \ldots$, suggesting that we should have

$$\frac{\sin x}{x} = \left(1 - \frac{x^2}{\pi^2}\right)\left(1 - \frac{x^2}{4\pi^2}\right)\left(1 - \frac{x^2}{9\pi^2}\right)\ldots,$$

On the other hand, expanding the function as a Taylor series, we have

$$\sin x/x = 1 - x^2/6 + x^4/120 - \ldots$$

Equating coefficients of x^2 suggests then that

$$1 + 1/4 + 1/9 + 1/16 + \ldots = \pi^2/6.$$

Even after checking this sum to several decimal places Euler was not absolutely confident in the result, but he had in fact solved a famous problem by analogical reasoning.

It might be that what is happening here is something similar to what Pearl [2000] has termed the "transfer of robust mechanisms to new situations". We have a mechanism that links factorisation of a function to its zeros. We find it applies for complex polynomials and wonder whether it may be extended. Features of polynomials that may be required in the new setting are that they have the right number of zeros, they remain bounded on compact sets, they behave similarly at $\pm\infty$. Might the mechanism be expected to work for a non-polynomial function possessing these features, such as sin x/x? What if you force the variable measuring the number of roots to be infinite? We may find it hard to estimate quantitatively the similarity between a function like sin x/x and a complex polynomial, but it is clear that tanx/x or exp x are less similar, the former having poles, the latter having no zeros and asymmetric behaviour at $\pm\infty$, and indeed the mechanism does fail for them.

In this case, once the realisation that an analogy was possible, it didn't cost much to work through the particular example. Euler hardly needed to weigh up the degree of similarity since calculations of the sum quickly convinced him. However, to develop a general theory of the expansion of complex functions did require greater faith in the analogy. This paid off when further exploration into this mechanism allowed mathematicians to form a very general result concerning entire complex functions, providing the "common ground" for the analogues.

3.2 Strategy

Moving on now to strategy, the title of Deninger's paper—*Evidence for a Cohomological Approach to Analytic Number Theory*—is also relevant to us. His aim in this paper is to increase our degree of belief that a particular means of thinking about a field will lead to new results in that field. This is a question of strategy.

At a finer level one talks of tactics. Researchers from the AI community working on automated theorem proving, have borrowed these terms. One tactic devised by Larry Wos [Wos and Pieper, 1999] involves thinking in terms of how probable it is that the computer can reach the target theorem from a particular formula generated from the hypotheses during the running of the programme. This tactic takes the form of a weighting in the search algorithm in favour of formulas which have a syntactical form matching the target.

Elsewhere, researchers in Edinburgh are interested in the idea of the choice of tactics [Bundy, 1999]. There is an idea of likening mathematics to a game of bridge where the mathematician, like the declarer, has some information and a range of strategies to achieve their goal (finesse, draw trumps, squeeze). Of course, there is a difference. In bridge, you are in the dynamic situation where you cannot try out every strategy, as the cards get played. This forces you to pay very close attention to which tactics have the best chance of working. In mathematics, on the other hand, with a computer it does not cost you much to try things out, although one does risk combinatorial explosion. At present, probabilities are being used for their computer bridge player, they are not yet being used for their automated theorem prover. While the computer has a small repertoire of syntactical tactics (rippling, resonance, heat, etc.) there is less need for an assessment of the chance of each working, but presumably the number of proof techniques will grow.

These automated proof strategies are at present syntactically based. Naturally, semantic considerations play the dominant role for human mathematician. Pólya was active in this area too. To give a brief flavour of his ideas, if when planning to solve a problem, any of the following should increase your confidence in your plan [Pólya, 1954b, pp. 152–153]:

- Your plan takes all relevant information into account.
- Your plan provides for a connection between what is known and what is unknown.
- Your plan resembles some which have been successful in problems of this kind.
- Your plan is similar to one that succeeded in solving an analogous problem.
- Your plan succeeded in solving a particular case of the problem.
- Your plan succeeded in solving a part of the problem.

3.3 Enumerative induction

Besides the incorrect Bayesian calculation of the confirmation provided by the observation of Neptune, Pólya does resort to a quantitative sketch in another place [Pólya, 1954b, pp. 96–7]. Here he outlines how one might think through the boost to the credibility of Euler's formula for a polyhedron (vertices − edges +

faces = 2) known to hold for some simple cases, when it is found to be true of the icosahedron. ($12 - 30 + 20 = 2$). His approach is to reduce the problem to the chances of finding three numbers in the range 1 to 30 with the property that the second is equal to the sum of the other two, i.e., $(V - 1) + (F - 1) = E$. The proportion of these triples is around 1 in 60, providing, he argues, a boost of approximately 60 to the prior probability of Euler's conjecture. Here again we see the same problem that Jaynes located in the Neptune calculation. The ratio of the likelihood of the Euler conjecture compared to that of its negation is 60.

In any case Pólya's construction can only be viewed as sketchy. It is not hard to see that the number of edges will always be at least as great as one and a half times the number of faces or the number of vertices. (For the latter, for example, note that each edge has two ends, but at least three of these ends coincide at a vertex). Thus one should have realised that there are further constraints on the possible triples and hence that the likelihood ratio due to the evidence for the Euler formula should have been in comparison to better informed rival conjecture, and so not so large. But the interesting point is that Pólya goes on to say that:

> If the verifications continue without interruption, there comes a moment, sooner or later when we feel obliged to reject the explanation by chance. [Pólya, 1954b, p. 97]

The question then arises as to whether one is justified in saying such a thing on the basis of a finite number of verifications of a law covering an infinite number of cases. This will hinge on the issue of the prior probability of such a law.

Now, consider Laplace's rule of succession. If you imagine yourself drawing with replacement from a bag of some unknown mixture of white and black balls, and you have seen m white balls, but no black balls, the standard use of the principle of indifference suggests that the probability that the next n will be positive is

$$(m + 1)/(m + n + 1).$$

As $n \to \infty$, this probability tends to zero. In other words, if verifying a mathematical conjecture could be modelled in this fashion, no amount of verification could help you raise your degree of belief above zero.

This accords with the way Rosenkrantz [1977] views the situation. He considers the particular case of the twin prime conjecture: that there are an infinite number of pairs of primes with difference 2. He mentions that beyond the verification of many cases, there are arguments in analytic number theory which suggest that you can form an estimate for the number of twin primes less than n and show that it diverges. He then continues:

> Now if Popper's point is that no examination of 'positive cases' could ever raise the probability of such a conjecture to a finite positive value, I cannot but agree. Instances alone cannot sway us! But if his claim is that *evidence of any kind* (short of proof) can raise the probability of a

> general law to a finite positive value, I emphatically disagree. On the cited evidence for the twin prime conjecture, for example, it would seem to me quite rational to accept a bet on the truth of the conjecture at odds of, say 100:1, that is to stake say $ 100 against a return of $ 10 000 should the conjecture prove true. [Rosenkrantz, 1977, p. 132]

So for Rosenkrantz, with no background knowledge, the principle of indifference forces a universal to have zero, or perhaps an infinitesimal (something also considered by Pólya), prior probability. However, other considerations may determine a positive probability.

> Subject-specific arguments usually underlie probability assessments in mathematics. [Rosenkrantz, 1977, p. 90]

In support of this view, returning to the Euler conjecture, we should note that there was background knowledge. For polygons, it is a trivial fact that there is a linear relation between the number of vertices and the number of edges, namely, V = E. Hence, a simple linear relation might be expected one dimension higher.

Is it always this kind of background knowledge which gives the prior probability of a conjecture a 'leg-up'? Do we ever have a situation with no background knowledge, i.e., where a *general atmosphere* is lacking? Consider the case of John Conway's 'Monstrous Moonshine', the conjectured link between the j-function and the monster simple group. The j-function arose in the nineteenth century from the study of the parameterisation of elliptic curves. It has a Fourier expansion in $q = \exp(2\pi i\tau)$:

$$j(\tau) = 1/q + 744 + 196884q + 21493760q^2 + 864299970q^3 + \ldots$$

One day while leafing through a book containing this expansion, a mathematician named John Mackay observed that there was something familiar about the third coefficient of this series. He recalled that 196,883 was the dimension of the smallest non-trivial irreducible representation of what was to become known as the monster group, later confirmed to be the largest of the 26 sporadic finite simple groups. Better still, adding on the 1 dimension of the trivial representation of the monster group results in equality.

In view of the very different origins of these entities, the j-function from nineteenth century work on elliptic curves and the monster group from contemporary work in finite group theory, if one had asked a mathematician how likely she thought it that there be some substantial conceptual connection between them or common ground explaining them both, the answer would presumably have been "vanishingly small". In Bayesian terms, Pr(connection/ numerical observation) is considerably greater than Pr(connection), but the latter is so low that even this unlikely coincidence does not bolster it sufficiently to make it credible. Naturally, McKay was told that he was 'talking nonsense'. He then went on, however, to observe that the second nontrivial representation has dimension 21296876. A quick calculation revealed that the fourth coefficient of the j-function could be expressed as: 21493760 = 21296876 + 196883 + 1. In fact every further coefficient of the j-function turns out to be a simple sum of the dimensions of the monster's representations. At this point the question of whether there is some connection has been all but answered—it has become a near certainty. Conway challenged the mathematics community to resolve this puzzle.

Fourier expansion in $q = \exp(2\pi i\tau)$:

$$j(\tau) = 1/q + 744 + 196884q + 21493760q^2 + 864299970q^3 + ...$$

196884	$196883 + 1$
21493760	$21296876 + 196883 + 1$
864299970	$842609326 + 21296876 + 196883 + 196883 + 1 + 1$
$\ldots$	

The answer eventually arrived through a construction by Richard Borcherds, a student of Conway, which earned him a Fields' Medal. Borcherds managed to spin a thread from the j-function to the 24-dimensional Leech lattice, and from there to a 26-dimensional space-time inhabited by a string theory whose vertex algebra has the monster as its symmetry group.

So why does the monster group–j-function connection become so likely by the time you have seen three or four of the sums, even with a minuscule prior, when other inductions are less certain after billions of verifications? Would we find consensus on how the number of instances affects one's confidence? Surely most people would agree that it was a little reckless on Fermat's part to conjecture publicly that $2^{2^n} + 1$ is prime after verifying only 5 cases (and perhaps a check on divisibility by low primes for the sixth).

n	0	1	2	3	4	5
$2^{2^n} + 1$	3	5	17	257	65537	$4294967297 = 641 \times 6700417$

Is it possible to use Bayes' theorem, even merely suggestively?

Let us return to the case of the Riemann hypothesis (RH). If we have a prior degree of belief for RH, how can 1.5 billion verifications affect it? Surely they must, but then is there some asymptotic limit? One might want to factor the RH as follows

$$\Pr(\text{RH} \mid \text{Data}) = \Pr(\text{RH}|p = 1, \text{Data}) \cdot \Pr(p = 1|\text{Data}),$$

where p denotes the limiting proportion, if this exists, of the zeros that lie on the line, taking the zeros in the order of increasing size of modulus.

For the second factor we might then have started with a prior distribution over p according to the weighted sum of the exhaustive set of hypotheses about p: non-convergent p; p in $[0, 1); p = 1$. Then if one can imagine some element of independence between the zeros, e.g., the fact that the nth zero lies on the line provides no information on the $(n + 1)$th, then the confirmation provided by the 1.5 billion zeros should push the posterior of $p = 1$ to take up nearly all the probability accorded to convergent p. This kind of assumption of independence has been used by mathematicians to make conjectures about the distribution of primes, so may be appropriate here. We might also consider that 1.5 billion positive instances provides an indication that p is convergent. Again, however, this consideration would depend on experience in similar situations.

For the first factor, out of all the functions you have met for which their zeros have held for a large initial section and the proportion of cases is 1, you are wondering what proportion are universally true. It is clear, then, that again much would depend on prior experience. For example, something that would be kept in mind is that the function $\pi(x)$, defined as the number of primes less than x, is known to be less than a certain function, denoted $\mathrm{li}(x)$, up to 10^{12}, and that there is good evidence that this is so up to 10^{30}. But it is known not to hold somewhere before 10^{400}. Indeed, there appears to be a change close to 1.4×10^{316}.

Returning finally to 'Monstrous Moonshine', perhaps we should look harder for a reliance on background knowledge. First, it is worth remembering that the dimensions of the monster group's representations and the coefficients of the j-function were not 'made up'. They come from 'natural' mathematical considerations. Imagine in the monstrous moonshine case if the two sides were not 'interesting' entities or that you knew for a fact that these numbers were randomly generated, wouldn't you take more convincing? Similar considerations are discussed by Paris [Paris *et al.*, 2000], who wish to justify some 'natural' prior distribution of probability functions over n variables.

> ...what in practice I might claim to know, or at least feel justified in believing, is that the data I shall receive will come from some real world 'experiment', some natural probability function; it will not simply have been made up. And in this case, according to my modeling, I do have a prior distribution for such functions. [Paris *et al.*, 2000, p. 313]

Evidence for the fact that background knowledge is coming into play in this case is provided by the fact that on presentation of the example to an audience of non-mathematicians they found the numerical coincidences not at all convincing. Despite the fact that a mathematician has no knowledge of a reason for a connection between these two mathematical entities, some slight considerations must play a role. Indeed, what seemed to disappoint the non-mathematicians was the need to include multiples of the dimensions of the irreducible representations. A mathematician, on the other hand, is well aware that in general a group representation is a sum of copies of irreducible ones. For example, the right regular representation, where the group acts on a vector space with basis its own elements, is such a sum where the number of copies of each irreducible representation is equal to its dimension. Behind the addition of dimensions are sums of vector spaces. Second, a mathematician would know that the j-function arises as a basic function, invariant under the action of the modular group. This offers the possibility that group theory might shed some light on the connection.

4 CONCLUSION

We have covered a considerable stretch of ground here. Clearly much work remains to be done on Pólya's research programme, but I think we can allow ourselves a little more optimism than Earman. I have isolated the following areas as potentially useful to study in a Bayesian light: (1) Analogy; (2) Strategy choice; and, (3) The use of large computations to increase plausibility of conjectures. Elsewhere I shall consider two additional areas: (4) Mathematical predictions in physics; and, (5) The use of stochastic ideas in mathematics (random graphs, random matrices, etc.). It is important to note that we need not necessarily arrive at some quantitative, algorithmic Bayesian procedure to have made progress. If Bayesianism in mathematics suggests interesting questions in the philosophy of mathematics, then I think we can say that it has served its purpose.

Department of Philosophy, King's College London.

BIBLIOGRAPHY

[Atiyah, 1984] M. Atiyah. An Interview with Michael Atiyah. *Mathematical Intelligencer*, **6**(1), 1984. Reprinted in *Collected Works, vol. 1: Early Papers, General Papers*, pp. 297–307, Oxford: Oxford University Press, 1988.

[Baez and Dolan, 1999] J. Baez and J. Dolan. Categorification. In *Higher Order Category Theory*, E. Tetzler and M. Kapranov, eds. pp. 1–36. American Mathematical Society, Providence, RI, 1999.

[Bundy, 1999] A. Bundy. Proof planning methods as schemas. *Journal of Symbolic Computation* **11**, 1–25, 1999.

[Cartier, 1992] P. Cartier. An introduction to compact Riemann surfaces. In *From Number Theory to Physics*,M. Waldschmidt, P. Moussa, J.-M. Luck and C. Itzykson, eds. Springer-Verlag, Berlin, 1992.

[Corfield, 1997] D. Corfield. Assaying Lakatos's philosophy of mathematics. *Studies in History and Philosophy of Science*, **28**, 99–121, 1997.

[Deninger, 1994] C. Deninger. Evidence for a cohomological approach to analytic number theory. In *First European Congress of Mathematics, Vol. 1*,, A. Joseph *et al.* eds. pp. 491–510. Birkhaüser, Basel, 1994.

[Earman, 1992] J. Earman. *Bayes or Bust?: A Critical Examination of Bayesian Confirmation Theory*, MIT Press, Cambridge, MA, 1992.

[Fallis, 1997] D. Fallis. The epistemic status of probabilistic proof. *Journal of Philosophy* **94**, 165–186, 1997.

[de Finetti, 1974] B. de Finetti. *Theory of Probability: A Critical Introductory Treatment*. Translated by A. Machi and A. Smith. Wiley, London, 1974.

[Franklin, 1986] A. Franklin. *The Neglect of Experiment*, Cambridge University Press, Cambridge, 1986.

[Freed and Uhlenbeck, 1995] D. Freed and K. Uhlenbeck, eds. *Geometry and Quantum Field Theory*, American Mathematical Society, 1995.

[Hacking, 1967] I. Hacking. Slightly more realistic personal probability, *Philosophy of Science* **34**, 311–325, 1967.

[Hesse, 1974] M. Hesse. *The Structure of Scientific Inference*. MacMillan, London, 1974.

[Hirsch, 1994] M. Hirsch. Responses to "Theoretical Mathematics", by A. Jaffe and F. Quinn', *Bulletin of the American Mathematical Society* **30**, 187–191, 1994.

[Hume, 1739] D. Hume. *A Treatise of Human Nature*. Clarendon Press, Oxford, 1739.

[Jaynes, forthcoming] E. Jaynes. *Probability Theory: The Logic of Science*, Cambridge University Press, forthcoming.

[Kac *et al.*, 1986] M. Kac, G.-C. Rota and J. Schwartz. *Discrete Thoughts: Essays in Mathematics, Science, and Philosophy.* Birkhaüser, Boston, 1986.

[Katz and Sarnak, 1999] N. Katz and P. Sarnak. Zeroes of zeta functions and symmetry. *Bulletin of the American Mathematical Society*, 36(1): 1–26, 1999.

[Paris *et al.*, 2000] J. Paris, P. Watton and G. Wilmers. On the structure of probability functions in the natural world. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 2000.

[Pearl, 1988] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan Kaufman, 1988.

[Pearl, 2000] J. Pearl. *Causality: Models, Reasoning and Inference*, Cambridge University Press, 2000.

[Poincaré, 1905] H. Poincaré. *Science and Hypothesis*. Dover Publications, New York, 1905.

[Pólya, 1954a] G. Pólya. *Mathematics and Plausible Reasoning: Induction and analogy in mathematics, vol. 1*, Princeton University Press, Princeton, 1954.

[Pólya, 1954b] G. Pólya. *Mathematics and Plausible Reasoning: Patterns of plausible inference, vol. 2*, Princeton University Press, Princeton, 1954.

[Rosenkrantz, 1977] R. Rosenkrantz. *Inference, Method and Decision. Towards a Bayesian Philosophy of Science*. Reidel, Boston, 1977.

[Wos and Pieper, 1999] L. Wos and G. W. Pieper. *A fascinating country in the world of computing: Your guide to automated reasoning*. Singapore: World Scientific, 1999.