

A Justification of Empirical Inference

By Arnold Zuboff

(This is a fuller version of an article published in *Philosophy Now*)

The Challenge of Scepticism

How can you know that your present experience doesn't owe its existence to an artificial stimulation of your brain, disembodied in a vat, or to a merely chance and causeless occurrence of its pattern in the absence of any world outside of it? Either of these possibilities, like numberless others we could imagine, would involve exactly the consciousness that is yours at this moment. They, together with what you think to be your actual situation, would be completely indistinguishable from within this consciousness that they all would have within them. So what could legitimately count for you in favour of the sort of thing that you do think about the world? How, with any justification, can your thinking reach beyond the appearances that would be common to all these sceptical hypotheses?

The Problem of Induction

Before we directly confront such scepticism regarding the world external to appearances, I think it will be instructive for us first to take up David Hume's famous challenge to provide an intellectual justification for induction, for forming beliefs concerning repeatedly observed associations of qualities or things that they will continue into the future. Let me try to evoke for you Hume's classic 'problem of induction'.

On a newly discovered island we have so far observed 100 birds of the new species 'humbird'. And every one of these humbirds has been blue. After that many such observations we come to expect very confidently that the next new humbird we observe will also be blue. But is there any intellectual justification for this expectation?

Here is what Hume would have said about this case: It must be impossible for us to demonstrate *a priori*, through mere examination of the concepts involved, that it is logically necessary that the next bird be blue—in the way that it is logically necessary, for instance, that 2 and 3 be 5. We are indeed intellectually justified in thinking 2 and 3 will always be 5, because 2 and 3 are not distinct from but rather identical with 5. Therefore we can know that denying this claim and trying to think instead of 2 and 3 as not 5 brings us into the impossible mess of a contradiction. We can thus be intellectually justified in our confidence that there will be no cases of 2 added to 3 that are anything other than 5. In contrast to this, however, the next observation of a humbird must be distinct from all preceding observations. And blue is distinct from the other humbird qualities. So it must be impossible to discover in the concepts involved that there is any contradiction in the next humbird being instead some colour other than blue.

Hume gave up any hope for justifying intellectually the enormously important employment of induction. He was thrown back on seeing induction as an instinctive

development of habits of expectation arising out of repeated experiences of such logically unnecessary combinations of properties as the other properties of the humebird with blue. We see more and more humebirds being blue, and it is simply in our nature to come to expect that the future will resemble the past—that the next humebird will also be blue.

A Solution to the Problem of Induction

I think Hume's scepticism regarding induction is wrong. There is indeed, as Hume insisted, no logical necessity that the next humebird be blue; but there *is* a logical necessity that it is *probable* that the next humebird will be blue given this evidence. For it is necessarily probable that this collection of random samples has a similar proportion of blueness to that of the general population from which it has been taken. Let me explain.

I think that while we are observing the 100 humebirds, rather than forming a Humean habit of expectation, we are calculating implicitly the most probable hypothesis concerning the general population of humebirds from which these observed birds are being randomly sampled. That hypothesis regarding this population that we are justifiably coming to favour as most probable is the one that would make the occurrence of the evidence, our observations, as highly probable an event as it can be. For that high probability of the evidence within the hypothesis necessarily lends its weight to the probability of the hypothesis itself, as I shall explain.

It may help us in this to consider an example of a hypothesis that we would justifiably reject as improbable. The worst of these would be the hypothesis that the only birds that are blue among the population of humebirds we were sampling happened to be the 100 we have already seen. If that were true, then it would have been highly probable that non-blue birds would have got mixed into the first hundred observations. And our actual observations of only blue humebirds would have had to be an extremely improbable event. But, as I often observed in our earlier discussion, something improbable necessarily had a low probability of occurring. Hence the improbability of the evidence given this hypothesis makes the hypothesis combined with the evidence necessarily improbable.

That the observations were of nothing but blue humebirds, however, gets less and less improbable in those hypotheses that increase the proportion of humebirds that are blue. The truth of these hypotheses therefore, along with that of the evidence, involves less improbability. The least improbable hypothesis must be that the humebirds across the whole population being sampled were generally blue. That is the hypothesis we implicitly settle on as most probable. And it goes along with this hypothesis, of course, that the next humebird sampled from the same population (and under the same general conditions) should be expected to be blue. And this, I contend, is the implicit thinking that rightly makes us expect that the next humebird will be blue.

This reasoning constitutes an a priori justification of induction. But I have run into people with suspicions that I am cheating in claiming that all my reasoning about the observed evidence was done a priori, based on nothing but necessary truth whose denial would be a contradiction. Why? Well, I had used probabilities and they

thought that probabilities must have been empirically derived. But anyone who looks carefully will see that none of the probabilities appealed to are established empirically. For example, I claimed along the way that the first 100 random samples being purely blue would be improbable within a hypothesis that only 100 humebirds are blue and the rest not blue. For it would be a necessary, mathematical truth that there would be many more ways in which the first 100 observations could have gone differently, with other colours showing earlier, within the hypothesis of only 100 out of the whole humebird population being blue. Nothing there is established empirically.

The only empirical part of the whole thing would be observing the colour of the actually sampled humebirds (which in our case, to help us raise Hume's problem, is simply given to have been blue). But the probabilities in our competing hypotheses are fixed a priori, within them. We then correctly judge to be most probable among them that hypothesis in which the observations would have been produced with the highest probability. For it is a necessary truth that that which is more probable is more probable to have occurred.

(And here, I think, is a clear demonstration that Hume's habit is not what would be responsible for forming such beliefs: Imagine two urns are presented to us. We are informed that one urn contains a million blue beads while the other contains only one blue bead in a million, but we so far have no way of knowing which description applies to which urn. We are then allowed to draw just one bead from one of the urns, without being able to look into either urn; and we find that the one bead we have drawn is blue. Surely we would that way have gained the same sort of confidence that we had about the birds, that the next bead drawn from that particular urn would also be blue. For we would think it extremely probable that the urn from which we randomly drew this one blue bead was not that with only one blue bead in a million but rather that in which all beads were blue. And, contrary to Hume's account of such beliefs, there would have been no chance here for any habitual expectation to form. Rather, again, we must be thinking that the probable population of beads in that urn from which we selected a bead is the one that would have made probable that single observation.

If instead we were asked to keep selecting beads from an urn about which we had been told nothing, and, if, having reached in and stirred them up to make sure our sampling was random, we then found that we were drawing out one blue bead after another, we would be acquiring with each additional observation the same sort of gradually strengthening conviction that we acquired with our observations of humebirds, that the next one observed would also be blue. For we would know for certain that it was *probable* that the population of beads in that urn was such as made probable what we were observing. With each additional one it would become less probable that the result of pure blue would have occurred if the randomly sampled population were not generally blue. Our confidence would be rightly based fully on our implicit inference to the probable general character of the population. No habit need enter into this process.

What helps further to see the non-involvement of habit is the fact that if the same number of beads had been gathered instead during one dip of a hand—with the hand reaching round randomly to various locations in the urn and accumulating beads from each in a fist—if, when the withdrawn fist was opened, it was seen all at

once that all of the beads in it were blue, the belief that it was probable that the beads in the urn were generally blue (and that the next bead taken from the urn would therefore also be blue) would be exactly as strong as at the end of the series of one-by-one observations we considered above. But in this case, of course, there has been no repetition of any sort that might have been thought to have formed a habit.)

Similarly, if a coin has landed nothing but heads up in many consecutive tosses, we are rationally required to think that the most probable hypothesis is not that the coin was fair but that it was loaded or double headed or otherwise fixed to land that way. The fair coin hypothesis could be true with that evidence only if something inherently improbable in that hypothesis had occurred, and the occurrence of something improbable is itself improbable. Let's apply this to Hume's famous example: If the sun has repeatedly risen in the morning, we are required to think that it is highly probable that it did so because somehow, like a coin repeatedly landing heads, it was in some stable fashion fixed to do so; and therefore we are also required to expect that that sun will keep rising in the morning for some time to come. This stability of coin or sun would, of course, be essential to the hypothesis that made such observations probable.

Based on the conceptual distinctness from each other of successive observations and of contingently cohering properties, like that of blueness from the other humebird characteristics, Hume had argued correctly that it would be impossible to demonstrate an *a priori* necessity for such combinations and therefore impossible to justify our inductive expectations in that way. But I have argued that Hume overlooked a proper *a priori* justification of induction, the one on which our expectations actually do depend, which is the rationally required assignment of more or less probability to the occurrence of competing hypotheses based on whether they make the occurrence of the evidence more or less probable, as this is discovered in the concepts involved.

An Answer to the Challenge of Scepticism

Let me now return to our earlier question: How can you know that your present experience doesn't owe its existence to an artificial stimulation of your brain, disembodied in a vat, or to a merely chance and causeless occurrence of its pattern in the absence of any world or even any time outside of it? The classic scepticism regarding the possibility of intellectual justification for judgments about the character of the world beyond the present appearances in a mind, including the rest of time outside this moment's impressions of memory and anticipations, shows the same inspiration as Hume's scepticism about induction. Based on the conceptual distinctness of a current impression of the world from the world and times external to that impression (which may include causes of the impression), the sceptic argues correctly for the impossibility of discovering an *a priori* necessity for any combination of the impression with any particular character of that external world or even with its existence.

I maintain that the sceptic, however, in Humean fashion, is overlooking the *a priori* justification of empirical inference on which our judgments about the external world actually depend. This is the rational requirement of an assignment of more or less

probability to the occurrence of competing hypotheses based on whether the hypotheses make the occurrence of the evidence, the overall pattern of the impression (including apparent memories and apparently previously formed beliefs and anticipations), more or less probable, as is discoverable in our concepts of the hypotheses and the evidence.

Consider, for example, the sceptical hypothesis that there simply is no external world. This would make it terrifically improbable that my therefore uncontrolled experience, merely by chance, as this would have to have been, had taken on the seemingly disciplined patterns I find it now has. Combined with that evidence—such a pattern—such a hypothesis makes up a package inherently improbable to have occurred (like the combination of the hypothesis that a coin was fair with the evidence of its landing consistently heads), as we can discover in the very concepts involved. And that which is improbable to have occurred is, indeed, improbable to have occurred.

We might call the single *a priori* principle that thus governs one's overall empirical thinking the 'highest probability principle'. It requires us always to favour in our beliefs, as most probable, that overall context of our current experience that would, as discovered purely in our concepts of it, have had inherent in it the highest probability of having produced the pattern of our current experience. We must do so because it is a necessary truth that the pattern's being produced in the most probable way is an event that was in itself more probable than the pattern's being produced in any less probable way. Let me just add that sometimes, of course, we must believe that an event which was locally improbable is the most probable to have occurred; but we can only properly do so when this local improbability has been needed in strictly the most probable overall hypothesis.

Ad hoc sceptical hypotheses, like that of a tricky powerful demon as the sole source of all my experience, must be rejected as extremely improbable because they contain causes that in their *general* character would have made the evidence improbable and can only *seem* to have made the evidence probable because of arbitrary and therefore inherently improbable specification in the detail of the hypothesis. Such would be the specification of a powerful spirit's specific interest in producing in me an impression of a world that would far more naturally have flowed from the general characters of the sorts of innumerable varied causes that I rightly think to be vastly more probable as sources of the impression. Such *ad hoc* elaboration in the demon hypothesis is no better at increasing its probability than would be such an *ad hoc* specification regarding a fair coin—that it happens to be one, in its detailed description, that is landing all heads many times—at increasing the probability of that incredible hypothesis. In both cases, although the specification is guaranteed conceptually to get us the evidence, the same specification can also be conceptually discovered to be utterly arbitrary and therefore extremely improbable given the general character of the causes within these hypotheses.

Among the things that you experience in the physical world that you believe in are human bodies that make motions and sounds that you interpret as behaviour and speech, as caused by minds other than your own. You interpret the marks on this paper as writing, as a product of a mind. The sceptic about other minds questions in particular the step beyond these bodily motions and sounds or marks to their

conceptually distinct causes, the beliefs and desires that are, you believe, their inspirations.

The answer to this scepticism regarding other minds is, of course, once again the reasoning that gave us the probability of the loaded coin and of the rest of the external world—the inference to the highest probability. Consider the hypothesis that there was no mind responsible for such movements and sounds or marks, that, for example, they resulted from random electro-chemical discharges in brains. But this would have been extremely improbable to have produced such patterns, ones that would have been made probable by only a mind intending behaviour and speech. (Whether it is probable that the minds are sincere or deceiving in their communication and other behaviour is a further question of the pattern of the evidence.)

A Bad Response to Scepticism

Recall our initial humebird problem. Imagine now a rather desperate attempt to establish a priori that the next humebird we observed would be blue by merely *defining* ‘humebirds’ as blue. That would indeed assure us that the next ‘humebird’ would have to be blue, but it would pretty obviously leave us with the substantive question unanswered, the question, as it now might be put, of whether the next bird otherwise like a humebird will be blue.

We have found, of course, that we can deal appropriately with this substantive question by establishing a priori that, given our evidence, it was necessarily probable that the next such bird would be blue. (And we can leave the definition of ‘humebird’ open with regard to what colour one might have. After all, it is far more useful to allow nature to teach us about the natural kinds of the world than to fix them as having lots of their characteristics trivially by definition.)

Well, philosophers trying to meet the challenge of scepticism have often resembled desperate ‘humebird’ redefiners in that they have arrived at theories of the external world and other minds that amounted to something like defining the external world as real or other people as behaving and speaking in the way that we would normally think them to be doing. Their answers to scepticism were thus curiously trivial, restricted to mere presuppositions, of experience or of interpretation of the speech of others, or established in the rules, the criteria, for the use of our language. And these philosophers missed the proper answer to scepticism, which is the establishment of the tremendous probability of the objective reality of the external world and other minds given the evidence of the pattern of experience.

For example, Kant argued for the existence of an external world ‘transcendentally’ by claiming that our momentary present pattern of mental content deserved to be called ‘experience’ only if it was viewed as caused by an external ‘phenomenal world’ that gave a seemingly objective significance to that otherwise merely subjective pattern.

But then the sceptic’s challenging question must simply be changed from why we should think that ‘experience’ is caused by an external world of roughly the sort in which we naturally believe to the question of why we should think this momentary pattern of mental content *does* deserve to be called ‘experience’.

And indeed Kant agrees that what he establishes to be true transcendentally is merely true among the presuppositions of experience; it is not true of the world as it is in itself.

What Kant needed was an explicit recognition of the a priori judgment that we all implicitly make, that, given the pattern of our experience, it is overwhelmingly probable that there exists, as a full reality, an external world beyond our present mental content, a world more or less of the sort in which we normally believe.

Sometimes what has been redefined in an effort to avoid scepticism is 'truth'. The trick is to lower the standard enough to allow us to reach 'truth' despite the sceptical nagging that this is impossible. The high standard requires for its truth that a belief or claim, if it is to be regarded as true, correspond with reality. The sceptic challenges our justification for thinking that there ever is this correspondence.

Theories concerning truth that are meant to disarm such scepticism by lowering what counts as truth are self-refuting because it turns out that they can make no sense except as claims that these theories themselves are true according to the rejected higher standard. But then they make no sense at all since it is essential to them that nothing *can* be true in the way that they would have to be.

One important example of this is pragmatism, the doctrine that beliefs are true if, and only if, they are useful. How is this thought to defeat scepticism? Well, according to pragmatism if our usual beliefs in an external world are useful then we can know them to be true despite any sceptical worries about whether such beliefs correspond to a world that is actually there.

But the pragmatist is failing to notice that he is refuting himself. Pragmatism itself can only be true in a way that it cannot allow. What brings this out is to notice that a pragmatist must hold that if a rejection of pragmatism became useful then that rejection would therein be true—but only *because* it was useful. Hence the pragmatist would still be regarding pragmatism as the underlying truth about truth quite *apart from* whether pragmatism was useful. In other words, pragmatism cannot see *itself* as true merely pragmatically.

Another important example would be a theory that insisted that the truth of what we say depends on nothing but rules for when it is right to say such a thing in our 'language game'. This insistence can seem to defeat skepticism because we usually allow ourselves to say that we know all manner of things to be true that a sceptic is maintaining we don't really know.

But it could make no sense for a holder of this theory to regard the truth of this theory itself as decided by the rules of a language game. For the view must hold that if a rejection of this language game theory itself were to be incorporated into the rules of the language game then that rejection would therein be true—but only *because* it was in the rules of the language game to say that it was. Hence this theory would still be regarding itself as the underlying truth about truth quite apart from whether it accorded with our linguistic practices to say that it was. And thus it cannot be regarding *its* truth as fixed by the rules of a language game.

Much attention has been given to how the definition of 'knowledge' might allow us to say that we can know things when the sceptic says we can't. The classic definition of knowledge is 'justified true belief'. What the sceptic questions is whether we are properly *justified* in our fundamental beliefs. Without that justification, he then might go on to say, we cannot truly be said to have knowledge. But this judgment that we don't have knowledge is secondary for the sceptic to the problem that we don't have justification. Some philosophers try to avoid that secondary sceptical conclusion about knowledge by pointing to our actual attributions of knowledge in which they see the bar of justification as lower than that set by the sceptic. Well, they say, if justification according to that lower standard still gives us 'knowledge' according to our usage of the term, who are the skeptics to be asking us to change the meaning of the word?

I have argued that we simply *are* justified in the beliefs that we actually do form, on the basis of mathematical probability. We have the very justification that the sceptics were requiring, and therefore according to these high standards we have justified true beliefs about the world. They are fundamentally all beliefs about probabilities. But the belief that something is extremely probable we can speak of also as a belief that it simply is the case. My knowledge that there is an external world that is at least roughly of the sort that I believe it to be depends for its justification on the knowledge that it is most probable that my pattern of experience is caused in a way that would have made this evidence most probable to occur. This justifying knowledge is not in itself knowledge *that there is an external world*, however. And for me to be said to have knowledge of *that* it must also be *true that there is an external world* (as distinguished from its being true that its existence is highly probable, which is what justifies the belief).

I think the involvement of probability sheds light on a famous problem that has emerged from the concentration on the definition of 'knowledge', that of the Gettier cases. In a Gettier case a belief is true and justified yet there is no knowledge. An example would be something true that I come to believe because I read it in a reputable newspaper but, unknown to me, the reporting of it in the paper is somehow due only to an error that purely by coincidence caused the report to agree with the truth. Here most people would be reluctant to say my belief was proper knowledge even though it would be both true and justified. What has proved a puzzle is the question of just what makes such a case fall short of being knowledge.

Well, I *would* have proper knowledge in this case of the high probability of this truth based on the report's being true as the most probable cause of my seeing the report in this reputable paper (which is my evidence). My belief is thus justified. But the evidence and the truth I believe on the basis of it would in reality here be linked only through an improbable coincidence. And I believe that this fact, of the improbability in the coupling of my justification to the truth, would be what kept the belief from having the status of knowledge.

Mere Description of Empirical Inference

There have been descriptions of empirical reasoning that in some fashion resemble the one I have given. A positing of causes like that I have described features in the

theories that call such reasoning ‘inference to the best explanation’, ‘abduction’ or ‘postulation’; and something like this is also represented in the ‘hypothetico-deductive’ model of scientific reasoning. But, to my knowledge, in none of these descriptions of, and prescriptions for, empirical reasoning has that reasoning been *justified*, as it has been here, because none have been explicitly shaped and animated by the highest probability principle, which is the essence of empirical thought.

Rather those with this approach have seemed to think one could do no better than merely to note, without any explanation or justification, the general features of the theories of the world that we tend to favour. Good empirical theory, for example, may well be described by them as possessing the unexplained virtue of simplicity. But I believe, by contrast, that simplicity and the other virtues of empirical thought are wholly explainable as aspects of an inference to the highest probability.

Imagine we are trying to infer the nature of a curve by plotting on a graph points on that curve that are being randomly reported to us. A simple curve will, of course, always yield points that can be interpreted as lying on a simple curve. Though any one of the numberless complex curves that could be drawn through this same set of points could indeed also have yielded them, it would be an increasingly improbable coincidence if points randomly selected from a complex curve were continuing to lie as well on some simple curve. Thus probability requires us to favour a simple interpretation if a simple one is possible.

But, let me stress, the simplicity of the curve must be understood as resulting from some regularity in its cause. A hypothesis that would genuinely make our evidence probable must display a principled connection between the inherent character of the hypothesis and the production of that evidence. The hypothesis of a simple curve produced according to a regular principle would do this. The hypothesis of a curve that was merely simple through chance, however, would be useless for making any evidence of its simplicity probable.

In their descriptions of and prescriptions for empirical reasoning, philosophers have sometimes invoked ‘Ockham’s razor’ and the ‘principle of sufficient reason’ as unexplained axiomatic strictures. But we can explain them as aspects of the inference to the highest probability. For the most probable account of our experience is that which gives us both nothing less and nothing more than that which would make our experience most probable to occur.

The hypothesis that there is no external world gives us obviously less than we need to make our evidence probable; and it thus fails to conform to the principle of sufficient reason. But that hypothesis also gives us something more and other than what we need, since it claims that the world has a general character that is opposed by our evidence. Such posits against the evidence must be shaved away with Ockham’s razor. Anything in our theory of the world that is not needed to make the evidence probable is a probability risk—it would be something whose presence in the world could be nothing but an improbable coincidence.

The hypothesis of the tricky powerful demon is perhaps more obvious as a candidate for Ockham’s razor than was the hypothesis that there is no external world, because

the deceiver hypothesis more obviously asserts the positive existence of something against the evidence—a controller of experience who does not distinctively reveal himself within its pattern. But therein the hidden controller also violates the principle of sufficient reason; for despite the *ad hoc* specifications that would have him with certainty producing deceptive evidence that was like our experience, an agent conceived of more generally is extremely improbable to be doing so.

Goodman's 'New Riddle of Induction'¹

Nelson Goodman thought he had already dealt with Hume through making the sort of response to him that I labeled as bad. He had simply defined induction as rational despite his admission that he could give no justification for it in terms of necessary truth (like the justification I provide). It is in relation to his own solution that Goodman raises his famous 'new riddle of induction'. Induction, as he has defended it, merely generalises the predicates that we have observed. We have always observed emeralds to be green. It seems then, according to his understanding of induction, that we would be therefore entitled to conclude rationally, by merely generalising this predicate, that *all* emeralds are green.

But, he then asks, how can we know just which predicates to generalise? The emeralds we have observed so far could instead be thought of as possessing the strange predicate 'grue', which means that those already observed before a certain date are green but those that will be observed after that date will be blue. Our observations would be just as consistent with this. (Note that this is logical consistency—there's not a thought about probability.) But if the predicate we are to generalise could just as well be 'grue', we could not use induction to predict that newly observed emeralds will be green after the crucial date.

I don't believe that this can even get started as a riddle for induction once we have introduced our probability considerations. But we'll do our best.

Let's go back to the case of the two urns, one containing a million blue beads and the other containing a million beads but only one of them blue. We could add that all the non-blue beads are green. But in deference to Nelson Goodman's example of inferring that all emeralds are green (not blue) we shall reverse these colours and stipulate rather that one urn has a million *green* beads and the other only one green bead, with the rest of its million being blue.

A fair coin is tossed to determine which urn is pushed forward for sampling. If the single randomly sampled bead is green, we conclude that it is a million times more probable that the urn from which it came was that with all green beads instead of the urn with only one green and the rest blue. And, of course, based on this we would predict that the second bead removed from the urn would also be green and not blue.

Let's try to apply Goodman's puzzle to this. Well, our first observation would have been equally (logically) consistent with the beads in the urn being not green but grue,

¹ Nelson Goodman, 'The New Riddle of Induction'. In *Fact, Fiction and Forecast*, 59-83. Cambridge: Harvard UP, 1955.

where in this case 'grue' will mean green if drawn from the urn at the time of the first selection but blue if drawn out at a later time. And if we couldn't decide between the beads being green and the beads being grue based on this observation we could have no justification for predicting that the next bead drawn will be green instead of blue.

But it is necessarily *improbable* that the beads in the urn are grue.

For the beads to be grue, the urn sampled would have to have been that containing only the single green bead—and that one bead would have to have been drawn out first. But it had to be immensely improbable that the first bead drawn from *that* urn *would* be the green bead and therefore immensely improbable that the beads in that urn would qualify as being grue. Of course a specification that this happened to be a case where that *was* the bead drawn first could *seem* to make the drawing out of the green bead probable even if the beads were grue, but that specification—of the green bead being the one observed at the time of the first drawing—would be merely ad hoc and inherently improbable given the stipulated general character of that urn.

When we considered our earlier one-in-a-million case, I followed that with a case more like our usual induction. I'll next quote what I said there but substitute 'green' for 'blue' and 'emeralds' for 'humebirds'

'If...we were asked to keep selecting beads from an urn about which we had been told nothing, and, if, having reached in and stirred them up to make sure our sampling was random, we then found that we were drawing out one green bead after another, we would be acquiring with each additional observation the same sort of gradually strengthening conviction that we have acquired with our actual observations of emeralds, that the next one observed would also be green. For we would know for certain that it *was probable* that the population of beads in that urn was such as made probable what we were observing. With each additional one it would become less probable that the result of pure green would have occurred if the randomly sampled population were not generally green.'

Hypotheses that construed these beads or emeralds to be grue—and therefore the green to be abruptly succeeded by blue after this time or that time or another time—could only be improbable, since they would require that our observations be conforming to an improbable order—first missing all the blues and hitting all the greens and then hitting all the blues.

Sometimes, in later versions of the riddle, 'grue' means green before some arbitrary future time (close enough to upset predictions) at which all, including those already seen, will *change* to blue. But, as with our earlier coin landing heads and the sun rising, stability is required in a hypothesis that would make evidence probable. That a loaded coin was somehow threatening to become fair would have made less probable that it had so far landed heads a thousand times in a row. And the threat being future would just be an ad hoc specification.

The Perspectival Nature of Probability

I have been arguing that probability is the justified basis of empirical reasoning. The remaining two sections of this paper are focused on the nature of probability itself.

It is a mistake to think simply that if a deuce of spades is drawn randomly from a deck something has just happened with the low probability of one in fifty-two. For this event at once has many other descriptions that are associated with differing probabilities:

1. 'selecting the deuce of spades', with a probability of one in fifty-two,
2. 'selecting a deuce', whose probability was one in thirteen,
3. 'selecting a spade', with the probability of one in four,
4. 'not selecting the queen of hearts', with the probability of fifty-one in fifty-two,
5. simply 'selecting a card in the deck', whose probability, within what was given in the case, was one.

The salience of a description and therein the probability of the event will depend on a person's relationship to it. If each person in a large group had made a prediction of which card would be drawn, with the predictions pretty much covering the deck, this event would have involved a low one in fifty-two probability for someone who had predicted the deuce of spades (as it would have done for the deuce of spades itself if the card were conscious) but a high probability of fifty-one in fifty-two for someone who had predicted the queen of hearts. For someone just watching there would be no improbability at all in the result.

Improbability depends always on coincidence, the coincidence of two descriptions that have been independently designated—such as that in the only prediction that was mine and that of the card that was actually drawn. When we routinely state that the deuce of spades has a one in fifty-two probability of being drawn we are implicitly but crucially imagining the deuce of spades being designated independently of its being drawn. If it is merely read off as the card that was drawn its probability of being drawn was one in one.

If only one person in the group mentioned earlier had made a prediction and had predicted the deuce of spades, the event would have been improbable to the tune of one in fifty-two for the whole group. But for an observer of many such groups there would have been no improbability in one of the groups getting it right. That event would remain an improbable coincidence for everyone in the group while at the same time involving no improbability for that observer. Similarly, for the winner of a lottery an improbable coincidence has occurred in the winner being him, whereas in this same event there was nothing improbable for an uninvolved observer. (Notice, by the way, that all these probability judgments hold objectively within each of these perspectives. They are perspectival but not subjective.)

Now, there is an a priori principle of inference that has already featured in my justification of empirical thought: It is most probable that the most probable thing will occur, or, in the past, *has* occurred. And thus, when we have evidence and competing hypotheses, that hypothesis must be judged most probable that would have made most probable the occurrence of the evidence.

But whether an event is probable or improbable has been shown to be perspectival. Therefore such inferences will also have to be perspectival.

Consider again how the event of the winning of a lottery is improbable for the winner but not for an uninvolved observer. The circumstances of most lotteries will force the winner simply to swallow happily the indigestible improbability of his winning (after he's pinched himself to make sure he isn't dreaming) on account of other parts of his overall evidence having to be even more improbable if he was not the winner. (By the way, the lottery having been fixed would not have made his winning more probable unless somehow it would independently have been probable that he rather than another entrant would have been the beneficiary.)

But in a lottery he's won where each entrant has been isolated from the others and where there are two rival hypotheses equally available—only his winning or every entrant winning—the improbability of the coincidence from his perspective if only he has won will make it by that much more probable that every entrant winning is the case. Informing an uninvolved observer of this winning, however, would be giving him no evidence at all for preferring as more probable that all the entrants have won rather than only this one he's been informed of. Since there would be no improbable coincidence for that uninvolved observer in the hypothesis of only one winner, he can have no reason to infer its improbability.

The ascription to something of improbability is like the ascription to something of dangerousness. There is a way in which everyone in every situation could properly say that a tiger is dangerous. But it is also importantly true that for you, unless you are close to him (perhaps having fallen into his enclosure), he actually isn't dangerous at all. In fact, the first, the unconditional attribution of dangerousness to a tiger can be nothing but shorthand for the tiger's being a danger to those he is close enough to hurt. And this has great practical implications. No action is required to save those not threatened by the tiger. And just so, we do not look for inferences to relieve an alleged 'improbability' of merely drawing an ace of spades from a deck. No, the improbability that could drive inferences would exist only for those who were in some sense close to the card—for whom, that is, there was a coincidence in that card's selection (as there would be, again, for the card itself always but also for someone predicting it). So I would not *simply* describe the result of picking a card from a deck as 'improbable'. And such 'improbability', should one still wish to ascribe it, would have to be inferentially inert for all those not especially connected to the event. It could have no force to drive an inference.

A failure to understand these points will, of course, badly distort one's view of what probability inferences consist in. Some probability theorists who do not understand this built-in perspectival character of objective probability—that is, of the mathematical probability of chances, based purely on the number of ways that events can develop—feel forced to abandon an objective account and instead to

shape their understanding of probability inference around the shifting credences (degrees of belief) of those who are making the inferences.

For in the thinking of such theorists every selection of a card from a deck that is innocent of coincidence nevertheless objectively possesses only the low probability of one in 52. And yet, they see, we have not the slightest problem in believing that such a selection has occurred. So it must be, they think, that *following* an objectively improbable event like that one, what we do is to adjust our credences and then ascribe to the event a merely subjective probability of one, of certainty, in any further reckoning we may do. (What *I* think, of course, is that a selection of a card free of coincidence has an *objective* probability of one before, during and after it.)

It is as though an objective reckoning of the dangerousness of tigers had been abandoned because it was thought that the description of a tiger as dangerous simply always applied and instead there had been nothing but a charting of the pattern of degrees of fear that tigers aroused. If a tiger is equally dangerous from every perspective, then we can only describe and never explain or justify the differing reactions of someone who has just fallen into the tiger's enclosure and someone at home reading in bed about tigers. Why should it be that one of these and not the other feels tremendous fear along with a need to brace himself for a tiger's attack? Well, all we can say is that these differing sorts of reactions to tigers are typical of people in situations like those. They are subjective reactions that cannot be justified by objective differences in how much danger a tiger is posing, since a tiger is always dangerous.

And such probability theorists must also be inclined to say that if we were to attempt to understand our probability inferences as simply based on the principle that what is objectively improbable is objectively improbable we should have to find ourselves repeatedly engaged in wild and futile inferences to try to push away an ever surrounding flood of objective improbability. So, as I've described, instead of understanding our beliefs as objectively justified responses to facts, they merely chart the developing pattern of the beliefs themselves without any proper explanation of them.

There's yet another bad reason for adopting a subjective approach to probability; and we'll be discussing that in the next and final section, on prior probabilities.

Prior Probabilities

Let us go back to imagining two enormous urns, this time each containing a *trillion* beads. In one urn all trillion beads are blue, whereas in the other urn only one bead of its trillion is blue. This second urn has been well stirred so that the single blue bead has nestled into a random location among the other beads.

First, let us say, a toss of a fair coin decides which of the two urns is pushed forward for sampling. Then a single bead that is randomly drawn from that urn is shown to an observer who has no other basis for judging what it contains and who understands all the circumstances I have described.

If the bead that is shown is blue, the observer should infer that it is a trillion times more probable that the urn being sampled is the urn with beads that were all of them blue. If it were instead the urn with only one blue bead, then this random drawing of a bead that was blue would have had to be something overwhelmingly improbable. But it is overwhelmingly improbable that something overwhelmingly improbable is what has occurred. Hence that hypothesis, combined with this evidence, is in itself overwhelmingly improbable and we must infer that the other hypothesis, of the urn being that with all blue beads, is overwhelmingly more probable to be true. We should expect this inference to give us the wrong answer in something like once in every trillion times this is tried. But it is overwhelmingly improbable that this is such a time. And even then it would surely have been the rational inference to make.

In this case a fair coin was used to decide which urn would be pushed forward. This meant that what are called the 'prior probabilities' of the competing hypotheses were equal. That is, apart from consideration of which hypothesis made the evidence more probable, it was equally probable that either was true. (Some may notice that I am not defining prior probabilities in the usual way—in terms of the credences of those making the inference. As I've explained in the previous section and will be further explaining below, my approach to probability is strictly objective.)

But what should we think in a case where we don't know such objective prior probabilities? Let's try one.

In this new case there are once again the two urns, each with a trillion beads, all blue beads in one and only one blue bead in the other. But this time we don't know what decided which would be pushed forward for sampling. A bead is randomly selected, and it is a blue bead. Can we not still, as in our earlier uncontroversial case, confidently say that it is overwhelmingly more probable that the urn pushed forward was that with all blue beads because if it had been the other urn something overwhelmingly improbable must just have occurred and it is overwhelmingly improbable that something overwhelmingly improbable occurred?

But some people think there is a problem with judging the probabilities of hypotheses in light of the evidence when their objective prior probabilities are unknown. They even demote such probability judgments to being merely 'subjective' or 'inductive'. I am claiming that this is mistaken; but let's look at what worries them.

What if, unknown to us, the objective prior probabilities had made it overwhelmingly improbable that the urn with all blue beads was pushed forward. For example, 'for all we know' (as these worriers might say), behind the scenes it could have been that pushing forward the urn with all blue beads depended on pulling out the blue bead by chance in a single selection from the urn containing only one blue bead in a trillion. If a non-blue bead had been drawn, then that same urn with only one blue bead (and the drawn bead thrown back into it) would then have been pushed forward. So it would have been a trillion times *more* probable that the urn that was pushed forward was the one that, in turn, would have made it a trillion times *less* probable that a blue bead would be drawn in the selection that we were to witness. The two improbabilities, of the good urn being pushed forward and of the bad urn yielding a blue bead in a random selection, would in that case have precisely

equaled each other so that each hypothesis would have been equally probable (or, more to the point, equally improbable) *given* the evidence of the blue bead.

But I say that we can easily infer that it was overwhelmingly improbable that the prior probabilities themselves were anything like that. Our principle that we must consider it improbable that something improbable occurred reaches back unstoppably to what is happening behind the scenes as well as in them, to what is happening, objectively, in the determination of the prior probabilities.

If the objective prior probabilities were bad for the hypothesis that is good for producing the evidence, then our evidence of the blue bead being drawn would have been condemned to be improbable *whichever* hypothesis was true—a condemnation that is itself improbable. For within such an *overall* hypothesis (including within it a theory of what the prior probabilities were), the hypothesis that *favoured* the evidence would be supposed to involve an improbability while, of course, the occurrence of the evidence within the hypothesis that did *not* favour it would have to remain improbable (however much that hypothesis may be favoured by the prior probabilities). With this wretched overall hypothesis we have to swallow an improbability whichever hypothesis is true. Well, we must then simply regard a hypothesis that the unknown objective prior probabilities made the evidence improbable as itself improbable.

A powerful source of confusion is this: There is a perfectly fine equation for figuring such probabilities, which has an unquestioned objective a priori status; but prior probabilities must be plugged into it. This is the equation given in Bayes' theorem. But what if the objective prior probabilities are not known? Then the hungry equation seems to require some feeding. And according to Bayes' postulate, unwisely, I think, tacked on to Bayes' theorem, what one should do is put equal prior probabilities for each hypothesis into the equation to represent their being equally unknown. In other words, in such cases we are being advised to calculate probability based on our subjective credences. This is surely a category mistake born of frustration. This is not as obviously bad as would be the advice to treat all the variables in algebraic equations whose values one hasn't yet discovered as being equal—in order to represent one's equal state of ignorance regarding them—but it's the same kind of muddle.

The right view of the mathematics, I think, is that weighing the hypotheses simply in terms of their favourability to the evidence already gives you their objective probabilities when combined with that evidence. Then the objective prior probabilities would, if not directly known, be merely missing further information that one should anyway expect to favour the hypothesis that is favourable to the evidence—because otherwise something less probable would have to have been what occurred, which must be less probably the case.

Where the difference in how much the hypotheses would favour the evidence is small, the unknown prior probabilities might have been decisive, and not knowing them would be important. But not so in the cases like that of the trillion-bead urns, in which there is a massive difference in how much each hypothesis would have favoured the evidence.