

The Perspectival Nature of Probability and Inference

Arnold Zuboff

University College London

It is argued that two observers with the same information may *rightly* disagree about the probability of an event that they are both observing. This is a correct way of describing the view of a lottery outcome from the perspective of a winner and from the perspective of an observer not connected with the winner – the outcome is improbable for the winner and not improbable for the unconnected observer. This claim is both argued for and extended by developing a case in which a probabilistic inference is supported for one observer and not for another, though they relevantly differ only in perspective, not in any information that they have. It is pointed out, finally, that all probabilities are in this way dependent on perspective.

It has recently been argued in this journal that a case may be constructed in which probabilities are observer-relative.¹ In this article I attempt to demonstrate that probability is always perspectival.²

Let me start with a question: When is it likely that an unlikely event has occurred? This is a question with two importantly different answers. Firstly, there may be what we could call the cases of ‘strained’ improbability. Imagine that a coin, in only one sequence of tosses observed randomly by us, has fallen heads many times, yet we are convinced by other extremely powerful evidence that it would be even more improbable that the coin was not fair than that a fair coin landed heads so many times in a row. In such strained cases of improbability we have evidence that, if it is to be made likely, forces us to accept that other evidence must have come about in a way that was inherently unlikely to have happened.

Then there will be the cases of ‘relaxed’ improbability, in which the locally improbable has actually been made probable for us from a certain perspective because of a combination of the overall context and the conditions of our observation. Imagine there are countless tossings of a fair coin. It must be expected that in a sequence of countless tossings of a fair coin there will be runs of 1000 consecutive heads. It will still be as unlikely as ever that any randomly observed sequence of 1000 tosses within this process of countless tossings will be all heads, and how many tosses there are overall will be absolutely irrelevant to what happens in any such sequence of 1000 tosses. So if these countless tosses had all been filmed and we viewed a randomly selected piece of film of just 1000 of these tosses, it would be incredibly unlikely that this piece of film would show us all heads. But if we were

presented with a film clip of 1000 consecutive heads that had been deliberately snipped from the rest for the very reason that it *was* all heads, the inherently unlikely event of our observing all heads with a fair coin would have been rendered likely by the conditions of our observation. So a relaxed case of improbability is one in which enough events have occurred so that the locally unlikely has actually been made likely to have occurred, while our freedom or direction in observing these events is adequate to have made it likely that we observed the event in question.

That improbabilities can be either strained or relaxed makes probability and probabilistic inference perspectival. For the same event can be improbable for one observer and probable for another within the same hypothesis because the condition under which one observer is viewing the event would give the occurrence of such an event an unavoidably strained improbability for him while the condition under which the other is viewing the event would relax that improbability, would make the overall event probable for him. Thus for the winner of a lottery the unlikely coincidence between *his* entry and the *winning* entry makes the event of his winning improbable, while for an uninvolved observer this is merely the event of *someone* winning, an event which will actually be certain if there is bound to be a winner. This perspectival effect can be obscured for us by the fact that the circumstances of the usual lottery will force the winner to accept that something unlikely for him has indeed occurred. Yet he may still rightly describe the result as 'unbelievable' (as it would not at all be for an uninvolved observer) and may at first have flirted with the idea that his winning was a practical joke or a dream – he may have pinched himself – though he finally does believe it. And we may confusedly think that the winner ought himself, like us, to regard his winning as not at all improbable. We may think that reflection on the mere fact that it was likely, or even certain, that *someone* was going to win ought to satisfy him, as well as it does us, that nothing unexpected has happened. But we are mistaken if we think this. There is for him in this an unavoidably strained improbability, while for us any improbability in an entrant winning has been relaxed by the number of entrants, or the guarantee of a win, and our stipulated lack of any independent involvement with the winner.

But we can clearly see the force of improbability from such a winner's perspective (improbability completely undiminished by the fact that 'someone had to win') when we see that improbability supporting an inference for him to an alternative hypothesis in a case that is free from the obscuring factors of forced acceptance in the usual lottery.

Imagine a hotel with a million rooms. In the hotel lobby is a large urn containing a million beads. Each of a million players takes a turn to draw one bead from the urn without looking at it and keeps his bead in his clenched fist till he gets to his own private room and shuts the door. Then, alone, he looks at

the bead he has drawn. If the bead is blue, he is a winner. But he knows there are two kinds of game. In the 'hard' game there would have been only one blue bead among the million beads in the urn, and the prize would be large. In the 'easy' game, every bead of the million would have been blue, and each of the million winners would receive a much smaller prize.

Any of these players must regard himself as having a random sample among the beads, since it was not somehow drawn out by him *because* it was blue. He observes a bead that has the completely independent, random designation that it was simply the one that *he* drew out. Now, in the hard game a random sample of one of the beads was unlikely to have been blue, and the sampler was therefore unlikely to be a winner (unlikely to the tune of one in a million); but in the easy game the bead was bound to be blue and every player a winner. Thus any player who finds that his bead is blue, that he has won, should figure that the easy game was a million times more likely to be the one he had played.³

If the hard game had been played, *someone* would still have drawn the blue bead and been a winner. Yet that rare someone would then still have been right to infer the greater likelihood that the easy game had been played; but in doing so he would have been misled into thinking that the easy game actually was the one that was played. The point is that it would have been enormously unlikely that any randomly designated player (as any player must be for himself) would have been that rare someone who was a winner in the hard game, and thus it would have been enormously unlikely that any player, stuck with just the bead that *he* drew, win or lose, would have had the blue bead in his hand and therefore been in the position to make the rare misleading inference to the easy game in the event of the hard game being played. This inference has only a one in a million probability of misleading even in the hard game, since to be so misled a player would have had to have drawn the hard game's one blue bead in a million. So no player should be bothered about being the someone with the blue bead in the hard game. A player should think that he *could* just be the someone with the blue bead in a hard game, but it's incredibly improbable that he *is*.

Next imagine an external observer who will be told of the result in just one hotel room that has been selected in some way that is random relative to whether it contains a winner. If this observer hears that the occupant of that room has a blue bead, he ought to infer the likelihood of the easy game having been played. For he is in the same position regarding these probabilities as the room's occupant, since both are connected with a bead in that urn in a way that designated it randomly relative to its being blue. The hard game hypothesis would make highly improbable the evidence that is known to both this external observer and that one player. Hence both can infer that it is unlikely that what they know has resulted from the hard game.

But now imagine that an external observer is going to be told about what

happens in only the room of a single *winner* no matter which game has been played. Then this guaranteed observer of a winner, unlike either the external or the internal random observer, cannot use his news of a winning player, guaranteed as it is in either game, to infer any greater likelihood that one game was played rather than the other. He would be hearing about a blue bead and a winner either way. Yet, if he is told about the winner's inference that the easy game had been played, he can totally sympathize with that inference; he should agree that that inference should indeed be made from that position; but he cannot himself join that winner in making it. For he has learned about this evidence under the wrong conditions for it to constitute the random sample that he would need for making that inference.

Now imagine that there are, in the million rooms of a second hotel, a million guaranteed observers, each isolated in his own room and looking at (and listening to) a winner of the game on a television receiver. If the hard game is played, each observer will be viewing the same single winner. If the easy game is played, each will be viewing a different one of the million winners.

One of these observers sees a winner looking into the camera and hears him say:

I can't know for certain whether I am the single lucky winner of the hard game addressing a million guaranteed observers or just one of a million winners addressing the *one* guaranteed observer who happens to be watching *me* in the easy game. But I must infer, given what I *do* know, the far greater likelihood of the easy game having been played because that makes what I am observing, my winning, much more likely to have happened from my perspective. But I realize that this must be an 'egocentric' inference, to be made by me (and by any other winner) for myself (or for himself) alone. I realize that though you see this very same winning that I see, this same blue bead I am now holding to the camera, you cannot use this as evidence for the playing of the easy as opposed to the hard game. For you see me and this bead under a crucially different condition from that under which I see myself. I see myself simply as myself, win or lose. You see me only as a winner, in a way that is not random relative to that which makes different the two hypotheses about the game.

And the viewer must agree; this winner, whether he is the single lucky winner of the hard game or merely one of the million of the easy game, is absolutely right to infer that the easy game is the one that was likely to be played. It *is* the likely game from that perspective. But this guaranteed observer, though he has no information that could change the mind of that winner, still cannot adopt as his own the inference he endorses for the other.

(Why have I bothered to restrict communication between the winner and any observer of him to a one-way television broadcast? The answer is that I wished to eliminate extraneous grounds for inference to the hard or the easy game based on the winner's consideration of the number of those who are observing him. If he was the winner in the hard game, he could of course

know this if a million observers were communicating with him. If the easy game was played, he could know this too if he found himself observed by only one observer. But why have I made the observers number a million rather than one? If there were only one observer of the game – and he knew this – then he would be in possession of information that would disturb the inference of the winner he observes if that winner was to have that information. For the single observer would know that something improbable would have occurred – not for him but for the winner he observed – in case the easy game had been played: that this winner would have been selected as the only one observed by the single observer. In the case as I described it above, however, none of the million observers individually possesses any information that the winner observed by him could have used to change his inference. And yet, as we've seen, a player and an observer who do thus share each other's evidence are prevented by their difference in perspective from joining together in making that inference.)

Let me stress that though our hotel inference is perspectival it is not in any sense merely subjective. It is *objectively* true, in terms of the frequency of events, that the player in this game is unlikely, from his perspective, to be a winner in the hard game. And it is equally objectively true that the guaranteed observer is as likely to find a winner in one kind of game as in the other.

Let me add that any event has at least potential perspectives attached to it that would make that event improbable for possessors of those perspectives. If a single card is selected from a deck, for example, there is nothing improbable in that event for the uninvolved observer. But if the card itself were somehow conscious, that card itself would rightly find the selection of only *it* an unlikely coincidence. The card itself could infer, if the hypothesis were available to it, the likelihood that more than one card was selected. For this would make the event of its having been selected less improbable from its perspective. But since for the uninvolved observer the selection of any one card was not at all improbable, for him no inference to either a one-card or a many-card hypothesis could be supported by this evidence.

NOTES

- 1 John Leslie, 'Observer-relative Chances and the Doomsday Argument', *Inquiry* 40 (1997), pp. 427–36. I think Leslie's claim of observer relativity is based on a fallacy, but that is not my concern in this paper.
- 2 The account I give here echoes a discussion in an earlier paper focused mainly on Personal Identity ('One Self: The Logic of Experience', *Inquiry* 33 [1990], pp. 39–68).

3 We could make the prior probabilities for this inference equal by adding the knowledge that which kind of game is played is decided by a single toss of a fair coin.

Received 24 January 2000

Arnold Zuboff, Department of Philosophy, University College London, Gower Street, London WC1E 6BT, England, UK. E-mail: a.zuboff@uc.ac.uk