

Apocalypse Not Just Now

Mark Greenberg

The End of the World: The Science and Ethics of Human Extinction by John Leslie.
Routledge, 336 pp., £10.99, 26 March 1998, 0 415 18447 9

John Leslie comes to tell us that the end of the world is closer than we think. His book is no ordinary millennial manifesto, however. Leslie is a sophisticated philosopher of science, and the source of his message is not divine revelation, apocalyptic fantasy or anxiety about the year-2000 computer problem, but 'the Doomsday Argument' - an a priori argument that seeks support in probability theory. In fact, the most interesting questions The End of the World raises are not, despite its subtitle, about our eventual demise. Rather, they concern our susceptibility, when thinking about risk, uncertainty and probability, to a kind of cognitive illusion. The Doomsday Argument is a case-study in 'probabilistic illusion', for it rests on a web of insidious intuitions, hidden assumptions and seductive but imprecise analogies.

The Argument claims that the observation that we are alive now increases the probability that Homo sapiens will become extinct in the relatively near future. It does not predict Doom at a specific time or with a specific probability. Its conclusion is more abstract and puzzling: whatever our best estimate would be (based on all available evidence, including the latest scientific, historical or other research) of the probability that our species is relatively close to extinction, it must be revised upwards. In reaching this conclusion, the Argument does not rely on evidence in the ordinary sense or, indeed, on anything peculiar to our present situation; it would yield the same conclusion at any point in human history.

It may seem preposterous that such a conclusion could be reached by armchair reasoning from the mere fact of our being alive now. Yet it would be wrong to rush to judgment. The counter-intuitive nature of probability is itself a reason for caution; moreover, the Doomsday Argument involves issues about time and existence, which are themselves notoriously resistant to intuition. Many who encounter it immediately conclude that it is unsound (though, according to Leslie, cosmologists and philosophers of science find it convincing more often than other scientists and philosophers, and 'women tend to think [it is] obviously strong'). Their objections tend not to hold up under scrutiny, however, and some of the initially sceptical join the ranks of the converted, which now include several eminent scientists and philosophers. The debate has endured for nearly two decades, resurfacing regularly in philosophical, mathematical and scientific journals.

Leslie, the Argument's foremost proponent, has responded indefatigably and ingeniously to a host of objections. In *The End of the World*, he examines the debate at length and concludes that the Argument has not been successfully challenged.

The widespread sense that something is wrong with it is, however, correct, and its failings illuminate a main source of our trouble with understanding probability: its close connection to randomness or unpredictability. Randomness seems to confound us. For example, we have a tendency to infer non-randomness from apparent patterns in random events (witness the incorrigible optimists who spot trends in the spins of a roulette wheel or the ups and downs of the FT Share Index); at the same time, the history of statistics suggests that, when random samples are required, we often mistake the merely haphazard - or whatever happens to be near at hand - for the truly random. As I will show, the Doomsday Argument's fundamental mistake is to rely on the intuitive but misguided notion that we can in general take ourselves to be typical humans, and thus, in effect, random samples of the species.

Genuine randomness, as opposed to haphazardness, disorder, complexity or a degree of uncertainty, is not a salient feature of the natural world, at least at the macroscopic level. We are generally able to predict to a degree better than chance the behaviour of human beings and other animals, middle-sized physical objects, the heavenly bodies and even the weather. Thus, to assume that the world around us behaved randomly would be a bad strategy. A currently much-discussed idea is that our susceptibility to probabilistic illusion is the consequence of an evolutionary adaptation of our minds which enables them to deal with the kinds of non-random situations we normally encounter. Although such an adaptation might be advantageous overall (at least in our ancestral environment), it could lead to systematic mistakes in certain contexts.

Our paradigms of randomness - dice, roulette wheels, lotteries and the like - are not natural processes but devices designed to create random conditions. (The randomness that matters here is unpredictability relative to available knowledge, as opposed to inherent unpredictability, i.e. that stemming from indeterminacy in the world.) Tellingly, the earliest writings on probability concerned games of chance. According to tradition, in 1654, a gambler, the Chevalier de Méré, asked his friend Pascal for advice on the problem of how to divide a stake fairly when a gambling game is interrupted; Pascal corresponded on the question with Fermat and probability theory was born. It was not until the 19th century, however, that theorists took seriously the possibility of fully random phenomena.

The Doomsday Argument tells us to adjust the probability of imminent Doom in the light of the observation that we are alive at this point in human history. The effect of a new observation

on pre-existing probabilities is specified by a mathematical formula known as Bayes's theorem, after the 18th-century English clergyman who discovered it. Roughly speaking, the idea is that the probability of a hypothesis is increased to the extent that the observation would be more likely if the hypothesis were true than if it were false. We infer from raised umbrellas in the street that it is raining because that hypothesis makes the umbrellas more likely than the alternative hypothesis, that the sky is clear. The degree of support that an observation gives to a hypothesis depends on a comparison between how likely the observation would be under different possibilities, at least one of which has to be counterfactual: to evaluate the evidence that Jarvis's fingerprints are on the knife, it is not enough to know that were Jarvis the murderer, he might well have left his prints there; we also need to know the likelihood of his prints being on the knife if he were innocent - perhaps he polishes the silver every evening.

Leslie is cavalier about Bayesian reasoning, encouraging us to think that there is no need to 'dress up' non-mathematical 'common sense'. Human beings tend to be poor Bayesians, however. For example, we often neglect to take into account how likely an observation would be if the hypothesis were false and how likely the hypothesis is to be false. Consequently, the results of correct Bayesian reasoning can be extremely counter-intuitive. To take just one cautionary example: suppose that a test for a particular disease is 99 per cent reliable (meaning that only 1 per cent of healthy people test positive) and that the disease affects 1 in 10,000 people in the relevant age group. If Jane, randomly selected, tests positive, roughly how probable is it that she actually has the disease? Most people - including a disturbingly high proportion of doctors - conclude that she is likely to have the disease; in fact, her chance of having it is less than 1 per cent. (Out of every 10,000 people, it would be expected that 1 would have the disease and that about 100 would falsely test positive; the chance that a positive test result indicates the presence of the disease is therefore no more than 1 out of 101.)

The Doomsday Argument begins with two hypotheses: Doom Soon, according to which human history will end relatively soon (by 2150, say); and Doom Delayed, according to which the species will survive long into the future, with the ultimate total population of all humans who will ever live numbering perhaps thousands of billions. (Leslie assumes, dubiously, that the birth rate will remain near or above its present level for most of the rest of human history, but this assumption is not critical to the overall argument and may be accepted for present purposes.) Using all the evidence available, we, or our best experts, estimate the probabilities of Doom Soon and Doom Delayed.

Leslie's key premise is that the likelihood of your living at your present position in human history - approximately sixty thousand years and fifty billion humans after the origin of our

species - would be greater if Doom Soon were true than if Doom Delayed were true. From this it follows, by Bayesian reasoning, that we should revise upwards our estimate of the probability of Doom Soon. But why should we believe that Doom Soon would make our present position more likely than would Doom Delayed?

The Argument's chief support derives from the application of Bayesian reasoning to that staple of probability theory, the drawing of lottery balls from urns. Suppose you are to select a ball at random from an urn. You know that the urn contains either 10 or 100 balls, numbered sequentially starting from 1. Your selection turns out to be ball number 3. This increases the probability that the urn contains ten balls, because the likelihood of your selecting ball 3 is (10 times) greater if the urn contains 10 balls than if it contains 100. Whatever the prior probabilities, the selection of ball 3 shifts them in favour of there being fewer balls. (Call the underlying principle the 'ball-3 principle'.) According to Leslie, 'calculations on similar lines can suggest that the risk that the human race will end soon has been regularly underestimated, perhaps very severely.' Roughly, his idea is that your observation that you are approximately human number 50 billion is analogous to the observation that the selected ball is ball number 3. Just as the latter observation makes it more likely that the total number of balls is only 10, Leslie thinks that the former makes it more likely that the total number of humans who will ever live is relatively small: 100 billion, say, as opposed to thousands of billions.

The ball-3 principle depends on the argument that the fewer balls there are in the urn, the more likely each one is to be selected. However, this is valid only if two conditions are satisfied. The first of these is the random selection condition. A procedure for selecting an item is random only if each member of the pool has an equal chance of being selected on each occasion. Without such a procedure, it cannot be assumed that the fewer the items, the more likely each one is to be selected. Consider a variant of the urn example, in which the random selection condition is not satisfied because balls 11 to 100 have no chance of being selected (they might be heavier than the others, and always settle to the unreachable bottom of the urn). Since, in this case, the likelihood that ball 3 will be selected is the same regardless of whether the urn contains 10 or 100 balls, its selection would have no effect on the probabilities.

The second condition is easier to miss. The reasoning that the random selection of any one ball would be more likely with fewer balls is valid only if we can be sure that that particular ball would still be in the urn if it contained fewer balls. This 'in-the-pool' condition thus requires that it be certain that the item selected would still be in the pool if the hypothesis of fewer items were true. In my original example, the condition is satisfied for balls 1 to 10, but not for balls 11 to 100 (since the latter would not be in the urn on the 10-ball hypothesis). More

generally, the condition cannot be satisfied for every ball because it cannot be true that every ball would still be in the urn if there were fewer of them. Thus, it cannot be assumed that a randomly selected item would be in the pool if it contained fewer items. Take the following example: as before, the urn contains either 10 or 100 balls; if it contains 10, however, they can be any 10 of those between 1 and 100. There is no ball that is certain to be in the urn if it turns out to contain only 10, so it is not true of any ball that its selection would be more likely if the urn contained only 10. A selection therefore cannot in this case increase the probability of the urn's containing 10 balls.

If we overlooked the in-the-pool condition, we might well come away with the impression that whenever an item is selected, there is reason to decrease our estimate of the size of the pool, which is absurd. In order for an item's selection to be able to increase the probability that the selection pool is smaller than previously believed, it must be possible for a different item's selection to decrease that probability.

Leslie's explanation of why the selection of ball 3 increases the probability that there are 10, rather than 100, balls is that, if there were 100, ball 3 would be 'exceptionally early', as opposed to more 'typical' or 'ordinary'. Thereafter, he takes it to be established that a hypothesis that would make a selected item 'exceptionally early' should be regarded as implausible, and he relies on this to support the applicability of the ball-3 principle in the case of Doom: 'If there were many trillion humans scattered through space-time, then a human who had been born before roughly the end of the 20th century would have been born very, very exceptionally early.' This explanation manifests Leslie's basic assumption that, in general, any human being we happen to consider is likely to be typical or ordinary, in a sense he does not explain.

The reasoning behind the ball-3 principle has nothing to do, however, with the notion that 3 is 'exceptionally early' in a pool of 100 but more 'typical' in a pool of 10. Random selection implies that, given 100 balls, ball 3 is exactly as likely to be selected as ball 46; or, to take another example, that in a lottery with 100 names, your name is exactly as likely to be chosen first, 46th, or at any other rank, regardless of whether that rank is 'exceptionally early' or a 'typical' middle rank. Consequently, the ball-3 principle applies even when the selected item is precisely in the middle of the larger pool, which is the reverse of what Leslie's typicality notion would dictate: if your name is in an urn that contains either 20 or 40 names, its being the 20th selected will increase the odds that the urn contains only 20 names, despite 20's being squarely in the middle of 40.

Leslie's reliance on vague intuitions of typicality leads him into confusion. He makes no attempt to specify precisely the analogy between the actual human case and the selection of balls from urns, nor to determine whether the necessary

conditions of the ball-3 principle are satisfied in the human case. Instead, he urges repeatedly that Doom Delayed is implausible because the existence in the future of thousands of billions of humans would make us 'exceptionally early'. The probability of Doom Soon would not be increased, however, unless our position would be more likely on Doom Soon than on Doom Delayed. The only relevant notion of typicality concerns how likely the observation would be on different hypotheses, some of which cannot be actual. The vagueness of terms such as 'exceptional', 'typical' and 'ordinary' obscures this basic point, tempting us to think that, whether or not my observation could have been different, I should expect the future to unfold so as to make that observation seem 'typical'. Thus, we may slip into thinking that we can limit consideration to how things actually are.

Setting aside intuitive expectations of typicality, does the ball-3 principle support the Doomsday Argument? The answer is no, because the random selection and in-the-pool conditions are not satisfied. The Argument does not rely on the actual random selection of a human being from all those who will ever live. We have not selected someone by a method that gives each past, present and future human an equal chance, and could not do so. Nor can we compensate for the extent to which our sample is unrepresentative, for we can have no idea what proportion of the total pool we have excluded.

Leslie seems unconcerned by the impossibility of random selection, however. His underlying assumption is that we can, in general, take ourselves to be typical. He takes for granted that, as a rule, the probability that an individual in a population will meet with a particular outcome corresponds to the overall proportion of that outcome in the population. For example, he assumes that, if a very high proportion of all humans will live near Doom, we are very likely to live near Doom. Similarly, he repeatedly tries to show that some claim is probably true of you by pointing out that if everyone bet the claim was true of himself, there would be more winners than losers. Since the probability that a randomly selected member of the population will have any particular characteristic necessarily corresponds to the proportion of the population possessing that characteristic, to assume that we are typical in this sense is tantamount to assuming that we can be treated as if we were randomly selected.

But it is a consequence of elementary statistics - and a notorious source of confusion - that the overall proportions in a population may not reflect the relevant probabilities for any individual, even an average one. The overall distribution of outcomes can be skewed by factors irrelevant to each individual's chances. As a result, the overall rates of different outcomes can actually be inverse to the probabilities for every individual. It can be true that a higher proportion of Floridians die each year than Georgians, yet that any individual's chances of dying are lower in Florida than in Georgia. To see why, suppose that in every age group Floridians have a lower death

rate than Georgians, but that a much higher proportion are elderly and therefore in the age group with the highest death rate. The overall death rate in Florida could thus be higher, despite the fact that every individual would minimise his chances of dying by living in Florida rather than Georgia, because the overall proportions depend not only on the individual probabilities but also on the distribution of individuals. It is a serious though common error to treat those who merely happen to be at hand - the people in the street or at the next table, or ourselves - as if they were a random sample, and thus to assume that the probability of an outcome corresponds to its overall rate in the population.

We are necessarily alive at the time we consider our position in human history, so the Doomsday Argument excludes from the selection pool everyone who is not alive now. A randomly selected person, however, would be likely to live later than now, at least if Doom Delayed is true, which drives home the point that we cannot assume that we are randomly selected. It might be thought that the exclusion of everyone who is not alive now does not matter because there is nothing special about the present time, i.e. because now is effectively a random time. On the contrary, however, if now is a random time, we cannot be random in the way that the Argument requires. Because population levels are not randomly distributed over time, a selection procedure that is random with respect to time cannot be random with respect to 'birth rank' - a person's position in the ordering of all human births. Because the population of England is not randomly distributed across the country, someone selected by a method that is random with respect to spatial location - e.g. which selects the nearest person to a randomly chosen square foot - would not be random in other respects. Thus, although a high proportion of English people live in towns, someone random with respect to location would probably live in a rural area. Similarly, a selection procedure that is random with respect to time gives each interval an equal chance of selection, thus creating a disproportionately great chance of selection for individuals living in periods of low population. So someone selected randomly with respect to time is more likely than someone selected randomly with respect to birth rank to live in a time of relatively low population. The claim that we are random with respect to time would thus undermine, rather than support, the Argument.

Even if someone who merely happens to live at a particular time could legitimately be treated as random with respect to birth rank, the Doomsday Argument would still fail, since, regardless of when that someone's position in human history is observed, he will always be in the same position relative to Doom Soon and Doom Delayed. The Argument uses all the information about the species that is available at the time, including how long it has so far survived and how many humans have so far lived, to reformulate the hypotheses and recalculate their prior probabilities. Accordingly, both hypotheses take the current time as the starting point, before

which there is no possibility of the species becoming extinct. Since we take ourselves as the sample, the person selected is guaranteed to be alive at the starting point of both Doom Soon and Doom Delayed, regardless of which is true.

Consequently, even if every human's birth rank were randomly assigned, we could not learn anything about the likelihood of Doom Soon or Doom Delayed from the sample person's birth rank. It would be wrong to reason that, if Doom Delayed were true, the person selected would be likely to have been born much later, after the species had become extinct on Doom Soon, and thus that Doom Soon is more likely than we had previously believed. That the person selected is alive at the starting point is an inevitable artefact of the procedure, and thus cannot affect the probabilities of the hypotheses. Although Leslie urges that 'a human observer' would be unlikely to find himself alive now if there are to be thousands of billions of future humans, it is certain that the human observer the Domsday Argument takes as a sample will live at the starting point of the hypotheses.

That we can choose someone from among those who are alive now tells us only the obvious fact that someone is alive at the time the argument is considered. It might therefore seem promising to take into account the fact that a particular person - you - happens to be alive now. Although everyone has to have a birth rank, it could be more likely that you would have your actual birth rank if Doom Soon, as opposed to Doom Delayed, were the case. Just as the certainty that some ball will be selected does not prevent ball 3's selection from supporting the hypothesis that there are only 10 balls in the urn, perhaps a particular person's living now could support Doom Soon. Leslie takes this tack, arguing that you would not expect to find yourself alive in the late 20th century if the human race were going to survive for a vast length of time thereafter. He deploys a series of thought experiments, the intended import of which is that, if Doom Delayed were true, you would very likely have been born after the first, say, hundred billion humans; and, consequently, your observation that you are roughly the 50 billionth human provides strong confirmation for Doom Soon.

It is not at all obvious, in the first place, that there is an appropriate sense in which it is true that you could have been born at a different point in human history; and if you couldn't, the likelihood of your living now is the same on Doom Soon as on Doom Delayed, and your living now has no effect on the probabilities. But even if we grant Leslie the randomness assumption he needs - that you had an equal chance of finding that you had been born at any birth rank - the Domsday Argument fails because this assumption turns out to be inconsistent with the in-the-pool condition (that you would certainly have existed even if Doom Soon were true).

If Doom Delayed is true as things actually are, which the Argument assumes to be possible, you are one of perhaps thousands of billions of humans. As the in-the-pool condition

specifies, you can infer that you would more probably have been alive now if Doom Soon were true only if you know that you would definitely exist if Doom Soon were true. A human randomly selected from all humans who would exist on Doom Delayed obviously would not be guaranteed to exist on Doom Soon since it could not be true that every one of the thousands of billions of Doom Delayed would be among the roughly 100 billion who would exist if Doom Soon were true. Thus, the in-the-pool condition is inconsistent with your being randomly selected; for it to be satisfied, you would have to be special among the humans of Doom Delayed.

Given the assumption that you could have been born at any birth rank, the only thing that sets you apart from the other thousands of billions of humans of Doom Delayed is that, unlike the great majority of them, you were born among the first 100 billion. Thus, to ensure that you would exist if there were to be only 100 billion humans, we would need a guarantee that those who happen to be among the first 100 billion would still be among the first 100 billion even if there were fewer humans in total, even if, in particular, the first 100 billion were the only 100 billion. In other words, Doom Delayed would then simply add a large number of different humans after the time when humanity would have become extinct on Doom Soon; the earlier humans would necessarily be the same ones who would have lived had Doom Soon been true. But if the first 100 billion humans would be the same regardless of whether Doom Soon or Doom Delayed is true, it is false that your living now would be more likely on Doom Soon than on Doom Delayed, and the probabilities of the two hypotheses are unaffected.

Intuition and probability collide openly when Leslie glimpses the problems and attempts to repair the damage. He protests that his intuitions are offended by the idea that our chances of being born increase as the number of humans increases. Given the counter-intuitive nature of the Doomsday Argument, it is ironic that he should now make intuition his last resort. More importantly, however, what offends Leslie's intuitions is the consequence of the very assumption he needs. That your chance of being born would have been smaller if there were not to be so many humans - and correspondingly greater if there were to be more - follows from his assumption that your position in human history is random. The intuition that your chance of being born would not have been affected by the total number of humans may well be correct. If it is, your position in human history cannot be random. On the other hand, if it is incorrect, the in-the-pool condition is not satisfied. Either way, the Argument collapses.

Leslie's attempt to reconcile intuition and probability fails. He urges that 'the only thing relevant to a 20th-century human considering the Doomsday Argument would be the probability of finding oneself in the 20th century granted that one had in fact been born into the human species.' When I pressed him on this point in a recent public exchange, Leslie insisted that

we need consider only the actual course of human history, in which you definitely exist, and fell back on the assertion that a vast number of humans in the future would still make you 'exceptionally early', even if you could not have been born later. As we have seen, however, whether the probabilistic reasoning behind the urn examples supports the Doomsday Argument depends on a comparison between how likely it would be for you to be alive at this point on different possible courses of the history of our species. The idea that you are unlikely to be exceptionally early is valid only if you had an equal chance of being born at any birth rank. But if you had no chance of being born at any other position in human history - it was certain that you would be in your actual position - then the claim that you are unlikely to be exceptionally early amounts merely to the assertion that it is unlikely that a high proportion of all humans will live in the future - which begs the question. Probability theory provides no support for the notion that the future is likely to turn out in such a way as to make us unexceptional.

Our being alive now does not increase the probability of our imminent extinction, and this is no surprise. Because the Argument does not make it possible for an observation to decrease that probability - whenever we happened to be alive, the Argument would be the same - our being alive now cannot affect it. In thinking about probability, we must be wary of intuition and, in particular, of the tendency to assume that we, or our observations, are typical or representative. This core assumption of the Argument is of a piece with more common and less respectable millenarian myths, and no more legitimate. If it were true that the majority of humans would live late in the history of the species, a human who was genuinely selected at random from all humans who will ever live would be likely to live relatively close to Doomsday. But, though such a truly randomly selected human might be doomed, for the rest of us, it's not the end of the world.

Mark Greenberg is Deputy Assistant Attorney General in the US Department of Justice. In the academic year that has just ended he was a member of the philosophy and law faculties at Oxford

[[current issue](#) | [archive](#) | [subscriptions](#) | [about the LRB](#)
| [advertising](#) | [site index](#) | [search](#) | [write to us](#)] all material
copyright © London Review of Books 1997-99