

Doomsday and the Extinction of Baseball

Leslie (1996: Chs.5 & 6)¹ argues that his Doomsday argument gives us reason to think that Doom is closer than we previously expected. There are a string of decent refutations of this argument in the literature, some of them directed at the papers from which Leslie developed his book. I take it that the arguments in Dieks (1992), Eckhardt (1993, 1997) and Korb and Oliver (1998) combine to show that Leslie's argument is unsound. I will briefly summarise the objections in section two, but there is a very simple way of stating the key point. Leslie's argument only goes through if we adopt a frequentist analysis of probability, and in particular if we adopt the 'finite frequentist' analysis of Russell (1948), but as Hájek (1997) shows the frequentist theory is, by philosophical standards, a howler, so Leslie's argument lacks cogency.

More on this later, because for now I want to focus on a surprising result. Despite these objections, something like a Doomsday argument is cogent in some (*some*) cases. As a consequence, the purely formal refutations of the Doomsday argument in the above papers are not successful. If we view the Doomsday argument as a schema, at least some of its instantiations are cogent, and not just by accident, so any argument which concludes that they are all lacking in cogency cannot work. These shows that some of the bullets aimed at Leslie in the above papers miss their targets. This should be about as reassuring to Doomsayers as the knowledge that one person in the firing squad will shoot to miss, but it does show that matters are more complicated than may have been expected.

The main purpose of this paper is to provide some more (philosophically) interesting examples of Doomsday and quasi-Doomsday arguments than have appeared in the literature thus far. There is a philosophical moral to the story, but it's not a particularly straightforward one. Many writers (including earlier time-slices of this author) seem to assume that to determine the impact of new evidence on our beliefs about the future of humanity, we just need to apply Bayes's theorem to our prior credences in each possible hypothesis about humanity's longevity. Not true. In fact we need to apply to Bayes's theorem to our prior credences in all the various hypotheses about the objective chance of each possible hypothesis about the fate of humanity turning out to be true. That won't win me any admirers for clarity of expression, but hopefully the relevant point will become more apparent after the examples. The paper is in two parts, the first providing the interesting examples, and defending this somewhat unwieldy moral, and the second defending my claim that frequentism is to blame for all of Leslie's errors.

Part I

Before I start, some definitions. For any class, let its *Total Population*, or TP, be the cardinality of its extension, that is, the number of such things there ever will be. Let the *Future Population*, or FP, be the number of those things which first become a member of that class at some time in the future. The difference between the TP and the FP, is the *Past Population*, or PP. Since I will be dealing with massively simplified examples, it will do the story no harm if we identify the past population with the Birth Rank of anything currently instantiating the class. The *Birth Rank*, or BR, of a thing in a particular class is normally defined to be one plus the number of creatures of that class born before it. So in practice my

¹ All page references, unless otherwise stated, to this book.

BR, *qua* human, is a billion or two less than the current PP of humanity. But such small differences make little difference in practice, and make even less difference in my examples.

For any variable, say TP for convenience, a rational agent may or may not have a precise belief about its value. Even if they do not, they may have a precise degree of belief in each of the propositions *The TP is x* for all possible values of *x*. If that is the case, we can work out the *expected value* of the TP for that agent, by multiplying each value of *x* by the probability they assign to it being the TP. We will call this the agent's *expected TP*, or ETP, remembering that saying an agent's ETP is *y* is not equivalent to saying she believes the TP is *y*. If the agent does not even have precise degrees of belief in some proposition of the form *The TP is x*, they will not even have a precise ETP. Nevertheless, we can still talk about the ETP rising or falling with the acquisition of new evidence.

One of the difficulties facing anyone trying to clarify the issues at stake in the Doomsday argument has been Leslie's insistence on using the future of humanity as the subject of his paradigm example. We really should aim for something a little more tractable, something about which clear and reliable intuitions can be discovered, but still of as much life and death importance as the extinction of our species. The obvious candidate is baseball.

Again in keeping with the spirit of the Doomsday argument, my examples are all of cases where people come to a subject matter with a large amount of evidence (I'll usually say fairly precisely how much) and then learn a fact about BR's. There should be, I hope, fairly wide agreement about how they should react, and I hope the lessons will carry over. So to begin.

Example One

Sean has incorrigible evidence that the FP of humanity is *x*. I don't particularly care how he got this evidence, maybe the Delphic Oracle told him, maybe he's simply that good a scientist, but he's got it. Question: What should the evidence that his BR is *y* to do his beliefs about humanity's TP? Answer: Make it precisely $x + y$.

Example Two

Susie has incorrigible evidence that the FP of humanity will be determined once and for all by a chance event scheduled to happen tomorrow. There is an 0.5 chance that it will be *x*, and an 0.5 chance that it will be *z*. Question: What should the evidence that her BR is *y* to do her ETP? Answer: Make it equal $y + (x + z)/2$.

These examples are so uncontroversial that I haven't even needed to invoke baseball. And I presume, though this is controversial, that Leslie's comments about indeterminism on pp. 233-4 mean that he agrees. But they do bring out a crucial point. The Doomsday argument says that finding out a particular fact about the past, namely our BR, should change our beliefs about the future. Now that isn't particularly controversial, unless perhaps dear reader you think all inductive reasoning is irrational. But how, in practice, could this change be made? As Lewis (1980, 1986) makes clear, only by affecting our beliefs

about chances. If we have incorrigible evidence about chances, then no new evidence will affect them at all. So far I have assumed this, so the evidence has had little effect. It's time to broaden our horizons.

Example Three

Bob knows a little about baseball. He knows, for example, that major league baseball has been played for over 120 years now. He knows that there are such things as home runs, and he knows that one of the many statistical honours bestowed each year goes to the player who hits the most home runs. And, finally, he knows that for 1998 this honour went to Mark McGwire, of the St Louis Cardinals, who hit 70. Bob's been living in a cave somewhere, possibly Australia, so knows nothing else about McGwire's feat.

Now by hitting so many home-runs, McGwire brought into creation a rather abstract being, a 70 home-run season. Bob, being a philosopher and hence quite comfortable with abstract objects², immediately starts to wonder what the TP of 70 home-run seasons will be. It's rather hard to know what will go into Bob's considerations, but let's say he comes up with an ETP of x . Now he gets told a rather important piece of information: the BR of McGwire's 70 home-run season is 1. For those not following my jargon, and sharing Bob's ignorance, McGwire was the first person to ever hit seventy home-runs in a season.

This information will presumably decrease Bob's ETP of 70 home-run seasons, because it makes him realise that hitting 70 home-runs in a season is harder than he had thought. So Doomsday-type reasoning seems to work here. That alone doesn't show the Doomsday argument is not a fallacy. Even fallacious argument forms can have cogent instantiations, if they instantiate some other cogent form. But here the best explanation for why Bob's ETP should decrease is more or less the one provided by the Doomsayers. Given Bob's prior beliefs about 70 home-run seasons, which presumably included that they are a fairly regular occurrence, it would be very surprising if Bob happened to stumble across the very first one. Even more surprising given that major league baseball has been played for so long, though since this is a disanalogy with the case of humans perhaps the Doomsayers wouldn't stress it as much.

Bob's change of beliefs is entirely consistent with the general theory I am propounding. Given what he knew about home-run hitting and the longevity of baseball, he could form certain reasonable beliefs about the objective chance of 70 home-run seasons appearing. The new evidence makes those beliefs about the chances appear mistaken, and different beliefs assigning lower chances more plausible. And it is this change in beliefs about chances which drives down Bob's ETP.

There might be an objection to this example based on a little knowledge of how sporting records tend to work. When Roger Bannister ran the first four minute mile, after it had eluded athletes for many years of organised and timed competitions, we didn't immediately think *Oh well, it will be many years again until we see another four minute mile*. Rather we thought that it, like most other records, would soon be broken. Or at least that belief would have been reasonable and, as it turns out, correct. So finding out that McGwire just set the all-time record mightn't decrease Bob's ETP at all, because it might provide little evidence that 70 home-run seasons will be rare in the future, as is required.

² Remember that cave was in Australia!

However, for reasons set forth in Gould's excellent (1996), McGwire's achievement is much more remarkable than most new all-time records. We are all used to athletes setting all-time records in competitions against nature, such as running 100m. That is, we keep seeing the quality of athletes improve. But in running competitors are performing essentially the same task as their predecessors did. In fact they might be performing an easier task, given the massive improvements in tracks and equipment. But baseballers are not, in general, in the same boat. The spread of excellence we see most noticeably in running and swimming also extends to pitching. So hitting x home runs is much harder for McGwire today than it was for Babe Ruth between the wars, whereas running 100m is in principle easier for Donovan Bailey today than it was for Jesse Owens between the wars. So if anything our assumption should be batters will win the home run title with fewer and fewer homers. As Gould shows, that assumption holds remarkably well for batting averages, and before 1998 it held fairly well for single-season home-runs. The 'in general' above is because one statistical measure in baseball does measure a player's absolute performance, rather than his performance relative to his opponent, and that is fielding percentage. And Gould shows that just as we might expect, fielding percentages have improved in just the same way that times for the 100m sprint have. While Gould backs up his claims with masses of statistical data, the reasoning is essentially *a prioristic*. So we can expect Bob, with his knowledge that baseball is a competitive sport and thus that records are not, in general, made to be broken, to not infer from McGwire's effort that balls will keep flying out of baseball stadia at ever-increasing rates, and hence the objection fails.

Example Four

Ted also knows a little about professional sport, though he is very rusty on its history. But Ted is one of that evil breed whose only interest in sport is its commercial value. So he has done a detailed study of the viability of baseball's World Series and football's Superbowl into the next century. He has not uncovered in his studies the BR of either the World Series or the Superbowl³, but he does know a lot about ticket sales, television ratings, international licensing agreements, merchandise, event costs and so forth. Based on this data he forms a fairly precise EFP for both the World Series and the Superbowl. Question: What should the information that the PP of the World Series is more than three times that of the Superbowl do to this EFP? Answer: Intuitively, not a lot.

Example Five

Dwight is visiting Australia, where they play a rather odd sport called cricket. He is attending a match in the premier domestic competition, the Sheffield Shield. He knows that the teams are spread across the country, so travel costs in putting on games are high. He knows that games go for four days at a time, so player payments are not cheap. And he knows, despite the cheap price of tickets, that attendances for games are in the hundreds. Sometimes. Dwight forms the quite reasonable belief that the competition is simply not viable, and his EFP for Sheffield Shield seasons is thus rather low.

³ Since in reality each Superbowl normally is named after its BR I am obviously using a little dramatic licence here, but I hope no more than is allowed.

Then Dwight is told the remarkable news that the BR of this season is over 100. His most natural reaction to this news would be, it might seem, to raise his EFP. This might seem to be good news for the Domsday argument. It might be thought that it would predict that his EFP will rise because it would be remarkable were he seeing the very last season of such a long-running competition, though see Objection Two below for a careful discussion of this point. In any case, the Domsday argument doesn't quite nail the reason for Dwight's change of view. Dwight's initial evidence makes the Sheffield Shield look a lot like a start-up competition doomed to failure, and his initial estimates should have been based on that assumption. The fact that it is over 100 years old destroys that reasoning, and gives him cause to think there must be some reason, other than commercial viability, that keeps the competition in place. (In practice it's the role the competition plays as a 'feeder' to the national team.) Hence he has no reason to think it will fold immediately, so his EFP rises.

If my diagnosis of the last three examples is correct, then there is no systematic pattern to what should happen to EFP's, or ETP's, on learning BR's. This shows there should be no formal disproof of the Domsday argument, but similarly that it cannot be formally proven correct. After all, Example Four (if not examples One and Two) show that evidence about BR's can have no effect whatsoever on EFP's, and in rare cases may have no effect on ETP's. Example Six confirms these suspicions.

Example Six

Following Leslie, call the proposition that humanity's FP is 10^{11} Quick Doom, or QD, and the proposition that humanity's FP is 10^{12} Doom Deferred, or DD. An agent who knows a fair bit about the world, including the fact that its population is growing but not her BR, has credence 0.5 in QD and 0.5 in DD. Hence her ETP is $5.5 \cdot 10^{11}$. Question: What should the discovery that her BR is $6 \cdot 10^{10}$ do to her ETP? Answer: There isn't enough evidence to say.

Proof. Alice has evidence, which she reasonably takes to be incorrigible, that when humanity's PP reaches 10^{11} , as it certainly will, a chance event may take place which will lead to the immediate destruction of humanity, but if it doesn't humanity's TP will end up at 10^{12} . She believes the chance of this event occurring is 0.5, so she satisfies the description in the question. But in her case it should be clear enough that the evidence causes no change at all. Her case is quite like Susie's; she has evidence about chances she takes to be incorrigible, so it will not be dislodged by anything consistent with it.

Belinda also believes that the TP of humanity will be determined by a chance event which happens when humanity's PP reaches 10^{11} . But she believes that the chance humanity will then end is 50/51. However, she is unsure what humanity's current PP is. Her credence that it is less than 10^{11} is 0.51, and that it is greater than 10^{11} , and hence that we already survived the chance event, is 0.49. Either way, she is certain that if humanity does not die out when its PP reaches 10^{11} , it will when it reaches 10^{12} . Given these beliefs, Belinda satisfies the description of an agent mentioned in the question. But when she finds out that we haven't yet crossed the threshold, her credence that the TP is 10^{11} moves *up* to 50/51. It seems here we have a Domsday argument that works.

Carol also believes that a random event, this time with chance $2/3$, determines the TP of humanity. But on her beliefs the time humans have in the universe is fixed in advance, and if the chance event is to occur, it occurs at the start of time. If it does occur, then after the 50 billionth person is born, there is an immediate and drastic thinning out of the population of humanity so only 50 billion people are born in the time we might otherwise have expected 950 billion to have been born. Carol's current credence that the PP of humanity is less than 50 billion is $3/4$, and that it is greater than 50 billion is $1/4$. Since Carol can see that the population has not thinned out drastically, she believes that if the PP is above 50 billion the TP is 10^{12} . Again, these beliefs suffice for Carol to satisfy the description in the question. Now when Carol finds out that her BR is 60 billion, she will come to believe, more or less for certain, that the TP of humanity will be 10^{12} .

So three agents, admittedly with highly stylised beliefs, can have exactly the same attitudes towards propositions about humanity's TP, be told the same fact about their BR, and react in different ways. And there is a good theoretical explanation for this. Beliefs about humanity's TP are determined by, *but do not determine*, beliefs about the chances of each value of x being the TP. Alice is certain about what the chances are; Carol is certain about what the chances were, and Belinda is completely unsure. By different paths they end up with the same doxastic attitudes towards hypothesis about the TP. But new evidence does not impact directly on an agent's beliefs about the TP. Rather, as Lewis showed, if it is to have any impact it must do so by going *via* her beliefs about the chance structure which underlie those beliefs about the TP. I once thought that it was possible to use Alice as a paradigm reasoner⁴ and formally show Doomsday reasoning to be mistaken. I saw this was mistaken because of Example Three. The reason it was mistaken was because I mistakenly assumed that everyone who shared Alice's credences about the TP must share her beliefs about the chance structure which generated those credences. My only consolation is that I appear to have not been alone in making this mistake.

In practice, I think the best conclusions to draw are these. First, Doomsday style reasoning can be important in cases where we may have misjudged the chances. And in those cases it can have pretty much the effects that its proponents claim for it. But secondly, those cases are confined to ones where the BR forms a large part of our evidence. The more we approximate Sean and Susie, with their incorrigible knowledge about the chances of each value being our species' FP, the less and less importance Doomsday reasoning plays. In practice, this means that we cannot say that the risk of annihilation due to nuclear war is greater than we thought it was because of the Doomsday argument, though perhaps we *maybe* able to say that the risk of extinction from a cause we haven't yet thought of is *a little* greater than we may have expected because of the Doomsday argument. More particularly, it's hard to believe that someone like Leslie, who has conducted a book-length study into the objective chances of extinction, has any right at all to wield the argument. At best the argument provides some way for the totally ignorant to focus what little evidence they have; a better approach would be to go and get more evidence and remove one's reliance on it. If the evidence comes in that we are likely to be in the first 0.01% of humans that ever lived, well, that's the evidence.

⁴ With perhaps the ever so slightly more plausible prior used in Korb and Oliver (1998) instead of Leslie's highly stylised one; but even Korb and Oliver's prior is fantastic, and it's a good thing to use simple examples to make a philosophical point. Indeed, if I'm right in this paper, Leslie's fault is not oversimplification, as Korb and Oliver allege, but the opposite flaw of failing to start with the very simplest of examples.

If that seems too much to concede to the Doomsayers, remember our situation is very different to those people for whom a Doomsday argument works. Although finding out our (approximate) BR can come as a surprise to many people, there is no clear analogy between our situation and that of Belinda or Carol. Whatever evidence we have about the objective chance of the world ending is about the chance of it ending at a time relative to our current position. We have no justification for adopting a belief state anything like Belinda's or Carol's. Our position is much more like that of Susie; we know quite a lot about the objective chances of various FP's, new information about the PP simply shifts our ETP while leaving our EFP unchanged. If one wants this could be called Doomsday reasoning, but only a conceptually confused person would be more scared of holocausts because of it.

Part II

Whenever I try and write something about Leslie's book, I am always reminded of a cautionary tale about reviewers. After Maynard Keynes's groundbreaking *General Theory of Employment, Interest and Money* appeared, there were more than the usual flurry of reviews in leading journals. The most hostile was by Frank Knight in the *Canadian Journal of Economics*. As was customary, the journal's editor offered Keynes the right of reply. Keynes declined, noting that, "with Professor Knight's two main conclusions, namely, that my book had caused him intense irritation, and that he had great difficulty in understanding it, I am in agreement."⁵ I have little desire to see this part of history repeated, and even less desire to reprise Knight's role, so I will try to afford Leslie that measure of 'good will' which an author is entitled to expect from a reader.

Objection One - Frequentism

After my introduction the usual, some would say decent, practice would be to carefully catalogue the evidence that Leslie is indeed a frequentist, and perhaps glance in the direction of arguments showing this to be a mistake. Leslie, unfortunately, never says, "I am a frequentist", so this could be quite time-consuming. It might be thought the discussion of reference classes on pages 212 and 256-263 give the game away, but perhaps Leslie has a more sophisticated theory of probability in mind⁶. And the most crucial step in the Doomsday argument, the claim that the conditional probability that we are the n th person given that there are precisely m people in all of history is $1/m$, for any value of n , seems to be indefensible unless you (a) adopt an inconsistent principle of indifference or (b) happen to be a frequentist. I will say more about option (a) under Objection Four below, but I hope this points make it clear that there is a *prima facie* case of frequentism against Leslie, but we would need a detailed exegesis to settle the matter beyond all doubt. So I have decided to let the part tell the whole, and will focus on the one passage which displays most clearly how frequentist Leslie is, and how absurd frequentism really is.

The now-classic shooting room example describes the Devil herding a group of people into a room, rolling a pair of dice, and if the dice land double-six shooting all those in the room. If the dice land any

⁵ The quotation is from pages 217-18 of volume 29 of Keynes (1973-89).

⁶ The most obvious candidate is the theory developed by Kyburg (1974). He adopts a logical analysis of probability, following Keynes and Carnap, whereby probability is a relationship between some evidence and a hypothesis, and represents roughly the reasonable degree of belief in the hypothesis given the evidence. The twist in Kyburg is that the only evidence that counts is evidence of frequencies, so reference classes start to matter.

other way, he lets those people go and draws in ten times as many people as before, and repeats the procedure. He keeps on doing this until he gets to shoot someone, thus guaranteeing that at least ninety percent of those who enter the room will be shot⁷. Now it is assumed that (a) the dice roll is an indeterministic event, and in fact the dice are fair, and (b) there are enough people around to ensure the Devil can keep on finding victims for as long as it takes. Presumably (b) means there are at least aleph-nought persons around, which will complicate some calculations⁸, but hopefully won't disrupt our reasoning. Now Leslie accepts (pp. 252-3) that the probability a person walking into the room should assign to "I'll be shot" is $1/36$.

However, Leslie then goes on to discuss a deterministic variant of the Shooting Room. The Devil has a computer which is calculating the decimal expansion of pi. This will be used in place of the dice to determine who lives and who dies. He picks a number d at random⁹, Leslie says 7 million, ignores the first d digits, reveals the next two, and if they are 66, shoots everyone. Otherwise he lets everyone go, brings in 10 times as many victims, reveals the next two digits, and if they are 66 shoots everyone. And again the process continues until we get some shooting. What probability should a victim in this little game assign to "I'll be shot" on entering the room, on the assumption that you can't quickly compute that many digits of pi? There is a slight complication here for those of us who follow Savage's (1967) principle that the only rational probability to assign to mathematical truths is 1, but we'll try and set that to one side. Leslie's answer is that it should be at least $9/10$, because at least $9/10$ ths of the people who enter the room will be shot. His argument against the quite sensible answer that it should be $1/100$ is this paragraph.

In this deterministic case it would be silly to argue that doom - shooting 'oughtn't to be expected since *at most points in time* the people who expected it would be wrong'. For, given determinism, what's crucial is that if everybody expected it, then *most people* who expected it would be right (255, emphasis in original).

There is one correct point in this paragraph: Leslie's imaginary interlocutor would be silly, because they would be assuming frequentism. But it's all downhill from there. Eckhardt (1997) gives three good reasons for rejecting Leslie's conclusions. First, the determinism / indeterminism distinction is not capable of supporting the different conclusions Leslie derives. Secondly, if we say that 'victims' are not shot, but in fact receive a small reward, we can apply standard Bayesian analysis to the repeated shooting room problem, and the right answer turns out to be roughly $1/100$. Thirdly, Leslie's argument seems to

⁷ In Leslie (1992) the example is credited to personal communication with David Lewis, but it appears uncredited in the book.

⁸ Specifically, it throws doubt on Leslie's discussion of the insurance agent on pg 251. He approvingly quotes van Fraassen's claim that an insurance agent offering insurance on the assumption that $35/36$ ths of those entering the room would leave would be making 'a costly mistake'. Well this isn't obviously correct. The agent would stand at best an infinitesimal chance of making a profit, but if they did show a profit it would be infinitely large, whereas any potential loss would be finite. Given the assumptions about the population, this agent presumably works for a company with infinitely many shareholders, so each shareholder has a potential real-valued gain against at worst an infinitesimal loss. Is this that bad a deal?

⁹ Presumably with an equal chance of choosing each number.

stand and fall with well-known 'multiplicative betting schemes', and it's equally well-known that those schemes fall.

I want to make a much more radical objection. Leslie's preferred position is not just mistaken, it is inconsistent. To see this, imagine a victim walking into the room who agreed with his reasoning. Let p be the proposition that person will be shot, q be the proposition that the next two digits of pi to be revealed will be 66, and r the material biconditional p iff q . Leslie says that $Pr(p) \approx 9/10$. Now presumably $Pr(q) \approx 1/100$. If we apply Leslie's frequentism we get that result, we get it by applying the principle of indifference, and indeed we will get it on almost any theory of probability that doesn't make it come out as 0 or 1, *a la* Savage. Finally, on anyone's theory of probability, $Pr(r) = 1$. If we construe r as a subject-predicate sentence, it is true of all actual subjects, it is true of all subjects in the long run, we know that it is true, and if we give the Devil certain nomic powers, as seems appropriate for an anti-deity, then surely the objective chance of r is 1. But these three probabilities are inconsistent, since it is a theorem of the probability calculus that $Pr(q) \geq Pr(p) + Pr(p \leftrightarrow q) - 1$. Finally, the quoted passage seems strong evidence of frequentism, since it assumes that what is probable is what happens to most people.

Objection Two - The Use of the 'Typical'

Secondly, there is an interesting difference between the formal Doomsday argument Leslie gives and his informal gloss on it. Since the argument seems to derive some of its punch by the interplay between its informal and formal parts, discovering they do not go hand-in-hand does weaken each of them somewhat. Speaking informally, Leslie says the idea motivating the Doomsday argument is that we should think of ourselves as typical¹⁰. If humanity is to continue for thousands of millennia, and continue growing, then we are amongst the first 0.01% of all humans who will ever live. On the other hand if it ends in the next three or four centuries, we will be around the middle of all humans who ever lived. Since the latter hypothesis makes us more 'typical', we should prefer it.

If the informal gloss really were the Doomsday argument, it would be a howler. For one thing, this apparent concept of 'typicalness' is rather hard to analyse formally, and it may turn out to be a chimera. Even if we can make sense of the concept, the argument appears to rest on confusing the platitude that we can expect results to be typical with the false claim that any result which would count as a typical result should be more probable. The platitude is explained not by the particular typical results being more probable, but generally by there just being more of them.

As Eckhardt (1997: 248-9) observes, however, the formal version of the Doomsday argument does not provide most support to theories according to which we are typical. On the contrary, it provides most support to theories which make us atypically late. According to the Doomsday argument, the best explanation for why our BR is so low is that there are no higher BR's to be had. This is because the argument says that the probability of us having the BR we actually have on the assumption of a TP of m is $1/m$. As Eckhardt points out a consequence of this, the formal Doomsday argument always makes us decrease ETP's, not raise them as might sometimes be expected. So while Leslie's informal reasoning might support the conclusions of Example Five, his formal argument does not.

¹⁰ See, *inter alia*, pages 2, 15 and 206.

Leslie seems dimly aware of this odd consequence of the formal version of the argument, and has a response to it (pg. 195). He claims that even were humanity to die out today, we would still be a long way from the last 0.01% of humans because of the number of humans born after us. But this point is far from decisive. The objection is to the general form of Doomsday arguments, so we cannot take solace from facts about a particular case. The situation is much as if Leslie had responded to a charge of affirming the consequent by providing separate evidence for his conclusion, or evidence that the conditional being used should really be a biconditional. Even if this move is effective, it does nothing to rehabilitate the original argument.

Objection Three - Old Evidence

The Doomsday Argument purports to be an abductive argument, an inference to the best explanation. The explanandum is never precisely stated, but at least in some passages Leslie takes it to be the fact that we exist in the late 20th century, for example at pp 214-9 where the issues about new and old evidence are discussed. Now of course we knew well before we read Leslie's book that we existed in the late 20th century, so it might be wondered what impact Leslie could make on our beliefs by merely pointing this fact out to us. More generally, why should we ever change our beliefs because of old evidence?

Only a straw man would seriously endorse the argument behind this rhetorical question, and Leslie has a reply to the straw man. Abductive reasoning sometimes works, and when it does work it doesn't do so because we discovered a new fact, but because we saw a new way of explaining old facts. Now we can concede this to Leslie without conceding any of the points he takes to follow from it. In particular, the fact that old evidence must be considered for abductive inferences in no way justifies the claim that we should conditionalise on this old evidence, as is done in the Doomsday Argument. More generally, it does look like a conceptual confusion to talk about conditionalising on old evidence, as if we are running together different and incompatible concepts of probability. More on this in the next objection.

As an aside, it might be wondered what Leslie is doing conditionalising at all. I know several good arguments to the conclusion that updating should occur by conditionalisation, and one impeccable one, but none of them stand up if we adopt a frequentist analysis of probability. End of digression.

The important objection that is being made here is best brought out by a concrete example. This one is taken from pg. 201. After looking at all the risks associated with nuclear war, environmental catastrophe, philosopher kings and the like, we decide that the probability of Doom Soon is 1%. Call this state the prior. By applying the Doomsday Argument, the probability of Doom Soon shifts to 50%. Call this state the posterior. I have seen no argument in Leslie that the kinds of states we were in before we read his book, or indeed heard of his argument, was analogous to the prior state, and not to the posterior.

But perhaps this shouldn't matter. Maybe in whatever state at all we are in, a Doomsday argument is applicable. Perhaps, that is, we could defend an iterated Doomsday argument. By one application of the Doomsday argument the probability of Doom Soon shifts to 50%, by another it shifts to 99.99%, by another it goes higher again, and so on. This just must be ridiculous, so there must be some property of

mental states which makes Doomsday arguments inapplicable to them. Because of Leslie's response to the problem of old evidence, the property cannot be that state includes the knowledge that it is instantiated in the late 20th century. Leslie argues that we could apply the Doomsday argument even if we were certain of that all along. So if the Doomsday argument is to have any importance, Leslie must show there is a theoretically important property which distinguishes those mental states for which the argument is applicable from those for which it is not, and show that the majority of minds we happen to bump into on the street fall into the former category. I'm sceptical about his chances of completing either task.

But even that isn't the real problem raised by the discussion of old evidence. The real problem, as was first noticed by Dieks (1992), is that there are several explananda waiting for an explanation. Let's accept that one of them is that we exist in the late 20th century, and even accept (as is needed for the Doomsday argument) that the best explanation of this is that the world will end soon. We also need to explain the fact that we are human. Since smaller TP's make this more surprising, hypotheses about larger TP's will be better explanations. And given some simplifying assumptions, we can show that the two abductive arguments in play, from our temporal position to Doom Soon and from our humanity to Doom Delayed, pretty much cancel each other out.

Leslie attempts to respond to this point in a few ways on pp. 224-30, but none of the responses seem to work. In particular, none of the responses provide a reason for distinguishing the 'evidence' that we are human from the 'evidence' that we live in the 20th century. For example, he points out on page 227 that the alternative abductive argument leads to absurd consequences when taken on its own. True enough, but we only ever said it was as good an argument as the Doomsday argument, so it's no objection to say the argument is ridiculous. And the passages where he argues that the relevant probabilities are those that are worked out given we are humans just misses the point. After all, we could turn the tables and work out probabilities given that we are existing in the late 20th century which would produce opposite results.

Finally, it might be noted there is a revealing modal fallacy on the bottom of page 226. Leslie says that the only 'opportunities to be a human' are at places where humans actually existed. The assumption seems to be that the only opportunity to ϕ is somewhere that someone actually ϕ 'd. But this is clearly false; when we say Hitler had an opportunity to invade England in September 1940 we don't thereby commit ourselves to the existence of an invasion of England in September 1940. There seems to be a hostility to modal talk which leads Leslie to this bizarre theory about opportunities, and which I suspect motivates his frequentism.

Objection Four - Objective and Subjective

The final two objections are not as devastating to the Doomsday argument as the earlier three, and are more profitably viewed as symptoms of an underlying malaise than as being the malaise itself. In each case the central issue is Leslie's rather odd understanding of the import of indeterminism.

“The distinction between objective and subjective [probabilities] can be hard to draw, anyhow, because (a) there are often strong excuses for declaring the probabilities are simply ‘subjective’ expressions of ignorance ...” (243)

There are two important errors in this half-sentence. First, there is a very clear distinction between what Leslie is calling the subjective and objective probabilities. It is just the distinction between beliefs and truths. When philosophers say that the ‘subjective probability’ of p is x , they just mean that someone’s degree of belief in p is x . There is less agreement about what is meant by when philosophers say the objective probability of p is x , but save in the special case where p is about mental states, objective probabilities do not depend on anyone’s beliefs. This is all set out very clearly in Lewis (1980) where subjective probabilities are called credences, and objective probabilities chances.

Some confusion might be engendered by the fact that neither of these philosophers’ concepts correspond particularly well to the folk concept of probability. That is best understood in terms of *reasonable* degrees of belief. And of course calling this a folk concept does not mean all philosophers have been blind to it. For example I take it Keynes (1921) and Carnap (1950) had the analysis of probability at least roughly right. But this just shows that the concepts picked out by ‘subjective probability’ and ‘objective probability’ are misleadingly labelled, not that they fail to pick out important and distinct concepts.

The second error is that we cannot infer probability of any kind at all just from ignorance. Or at the very least if we can do so it will not be in anything like the way Leslie seems to think that we can. As has been well known since at least von Kries (1886), distributing probabilities equally over the range of our ignorance leads to inconsistency. The examples showing this are well set out in Keynes (1921: Ch. 4) and van Fraassen (1989: Ch. ??). This point casts further doubt on Leslie’s fundamental assumption, that the probability of being the n th person is the inverse of humanity’s TP.

Objection Five - Determinism

I mentioned under Objection One Leslie’s rather odd view about the import of determinism. In short, he agrees with the analysis of Example Two I gave above, so that if we know the objective chance of Doom Soon is x , then our credence in Doom Soon should be x , whatever the Doomsday argument might have to say. Eckhardt (1997) gives a number of good reasons why Leslie cannot consistently hold this and hold that the Doomsday argument works when we think the world is determinate. Given the effectiveness of the objections mentioned above (many of which are due, it is worth recalling, to Eckhardt) I think the most charitable thing to say about Leslie is that his views about the determinate and indeterminate case really are incoherent. That is, the Doomsday argument should work just as effectively in either case. But in the indeterminate case, Leslie’s good common sense has managed to overcome his bad theory. This seems more charitable than saying, for example, he has consistently applied a bad theory, and by luck got the right result in one of the cases.

Good common sense, however, can only take one so far in the philosophy of probability. Leslie says on page 250 that the Doomsday argument, even if it is impotent against the theory that the world is indeterminate, can work against the theory that it is virtually determined that humanity will last a long

time. I'm not entirely sure what it would be for this to be virtually determined, and at least on some obvious interpretations this is clearly false. Imagine an agent who assigns equal credence 0.5 to the propositions that (a) the world is determinate, and it is determined that human life will continue for a long time and (b) the world is indeterminate, and the objective chance that the world will continue for a long time is 0.9999. It seemed to be common ground that if she was certain that (a), or that she was certain that (b), the Doomsday argument would have no effect. So how could it have an effect on her if she is uncertain which of these possibilities holds? Leslie's informal style of presentation may eliminate the initial impenetrability of many technical works, but in the end it seems to obscure more than it enlightens.

References

- Carnap, Rudolf (1950) *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Dieks, Dennis (1992) "Doomsday – or, The Dangers of Statistics" *Philosophical Quarterly* 42: 78-84.
- Eckhardt, William (1993) "Probability Theory and the Doomsday Argument" *Mind* 102: 483-488.
- Eckhardt, William (1997) "A Shooting-Room View of Doomsday" *Journal of Philosophy* 94: 244-259.
- van Fraassen, Bas (1989) *Laws and Symmetry*. Oxford: Clarendon Press.
- Gould, Stephen Jay (1996) *Life's Grandeur*. London: Vintage
- Hájek, Alan (1997) "Mises Redux – Redux: Fifteen Arguments Against Finite Frequentism" *Erkenntnis* 45: 209-227.
- Keynes, John Maynard (1921) *Treatise on Probability*. London: Macmillan. Reprinted as vol. VIII of Keynes (1971-1989).
- Keynes, John Maynard (1971-89) *The Collected Works of John Maynard Keynes*. 30 vols. edited by Donald Moggridge and Elizabeth Johnson. London: Macmillan.
- Korb, Kevin and Oliver, Jon (1998) "A Refutation of the Doomsday Argument" *Mind* 107: 403-410.
- von Kries, J. (1886) *Die Principien der Wahrscheinlichkeitsrechnung Eine logische Untersuchung* Freiburg.
- Kyburg, Henry (1974) *The Logical Foundations of Statistical Inference*. Dordrecht: Reidel.
- Leslie, John (1992) "Time and the Anthropic Principle" *Mind* 101: 521-540.
- Leslie, John (1996) *The End of the World*. London: Routledge.
- Lewis, David (1980) "A Subjectivist's Guide to Objective Chance" In R. C. Jeffery (ed) *Studies in Inductive Logic and Probability* Vol. 2, Berkeley: University of California Press, Reprinted with additional postscripts in his (1986), pp 83 - 132.
- Lewis, David (1986) *Philosophical Papers* Vol. 2, New York: Oxford University Press
- Russell, Bertrand (1948) *Human Knowledge: Its Scope and Limits*. London: Allen and Unwin.
- Savage, Leonard (1967) "Difficulties in the Theory of Personal Probability" *Philosophy of Science* 34: 305-310.