



CHICAGO JOURNALS



---

No One Knows the Date or the Hour: An Unorthodox Application of Rev. Bayes's Theorem  
Author(s): Paul Bartha and Christopher Hitchcock  
Source: *Philosophy of Science*, Vol. 66, Supplement. Proceedings of the 1998 Biennial Meetings of the Philosophy of Science Association. Part I: Contributed Papers (Sep., 1999), pp. S339-S353  
Published by: [The University of Chicago Press](http://www.uchicago.edu) on behalf of the [Philosophy of Science Association](http://www.philosophyofscience.org)  
Stable URL: <http://www.jstor.org/stable/188782>  
Accessed: 08/04/2014 17:20

---

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



*The University of Chicago Press and Philosophy of Science Association are collaborating with JSTOR to digitize, preserve and extend access to Philosophy of Science.*

<http://www.jstor.org>

# No One Knows the Date or the Hour: An Unorthodox Application of Rev. Bayes's Theorem

Paul Bartha<sup>†</sup>

University of British Columbia

Christopher Hitchcock

Rice University

---

Carter and Leslie (1996) have argued, using Bayes's theorem, that our being alive now supports the hypothesis of an early 'Doomsday'. Unlike some critics (Eckhardt 1997), we accept their argument in part: given that we exist, our existence now indeed favors 'Doom sooner' over 'Doom later'. The very fact of our existence, however, favors 'Doom later'. In simple cases, a hypothetical approach to the problem of 'old evidence' shows that these two effects cancel out: our existence now yields no information about the coming of Doom. More complex cases suggest a move from countably additive to non-standard probability measures.

---

**1. Introduction.** As the millennium approaches, we are led to reflect on the ultimate fate of the universe, or perhaps just humankind. How will it all end? A meteor akin to the one that did in the dinosaurs? An epidemic that makes 1349 look like a mild flu season? For simplicity, let us collapse all such scenarios into just two hypotheses, which we call 'doom sooner' and 'doom later'. According to the 'doom sooner' hypothesis, Doomsday will occur in the next hundred years or so; in particular, it will occur at a time when approximately fifty billion people have lived. According to the 'doom later' hypothesis, Doomsday will not occur until the distant future, let's say after forty quadrillion people have lived on earth or in various parts of the galaxy.

<sup>†</sup>Bartha: Department of Philosophy, University of British Columbia, Vancouver, BC, Canada, V6T 1Z1; Hitchcock: Division of Humanities and Social Sciences 101-40, California Institute of Technology, Pasadena, CA 91125.

Philosophy of Science, 66 (Proceedings) pp. S339-S353. 0031-8248/99/66supp-0026\$0.00  
Copyright 1999 by the Philosophy of Science Association. All rights reserved.

Brandon Carter and John Leslie (1996) have argued, using Bayes's theorem, that the observation that we are alive now strongly supports the 'doom sooner' hypothesis. Leslie acknowledges that this result is counterintuitive. Indeed, when we reflect on probable doom scenarios, such as the consumption of the earth by the sun expanding under the force of fusion pressure, it is hard to see how our happening to live on the eve of the twenty-first century could have any evidential bearing whatsoever on the timing of Doom. We maintain that it does not, but that identifying the flaw in the Domsday Argument brings out interesting points about the selection of prior probabilities and the 'orthodox' use of Bayes's theorem.

The Domsday Argument runs roughly as follows:

1. I find myself among the first forty billion humans.
2. Bayes's theorem tells us that an observation provides most support for whichever hypothesis renders the observation least improbable.<sup>1</sup>
3. 'Doom sooner' makes it 80% probable that I will find myself among the first forty billion humans, since there will only ever be fifty billion. By contrast, 'doom later' renders my appearance among the first forty billion vastly improbable (one in a million, to be precise).
4. So the observation that I am alive now provides a strong argument for 'doom sooner'.

To make the argument more vivid, Leslie asks us to imagine a lottery in which names printed on slips of paper are successively drawn from an urn without replacement. I am told that exactly one slip has my name on it, and that every name in the urn appears just once; furthermore, each name has an equal chance of being selected on a given draw. This reflects a key assumption of the argument: that my own position or 'birth rank' in the sequence of all humans is more or less random. I further learn that one of two hypotheses is true: the urn contains either fifty names or forty million names. (The ratio is the same as in the example above, but the numbers are reduced by a factor of one billion.) Now suppose that my name appears among the first forty drawn. Then an application of Bayes's theorem shows that, whatever my prior probabilities for the two hypotheses, I should now revise them dramatically in favor of the hypothesis that there are fifty names.

Call this the 'lottery' version of the Domsday Argument. We shall

1. This is an informal characterization of the role that likelihoods play in Bayes's Theorem, not an endorsement of any particular measure of degree of confirmation; see Fitelson 1999.

formulate and evaluate it more precisely in Section 3. There is also a second version of the Doomsday Argument which we call the 'shooting-room' version, after a different analogy also developed by Leslie. This argument depends upon the supposition that the population increases at something like a geometric rate, so that a high proportion of those who have ever lived are alive at any given time. A current estimate is approximately 15%. Similarly, 15% of all those who ever live will be in the final, doomed generation which we shall call 'generation Nix'. Since we have no reason to regard our own position in the sequence of human lives as special, we should put our own likelihood of being in generation Nix at 15%.

In this paper, we focus primarily on the 'lottery' version of the Doomsday Argument, but we wish to make a few remarks about the much-discussed shooting-room version by way of preparation.

**2. Shooting-Room Version of the Doomsday Argument.** Imagine that a judge summons people to a place known as the shooting-room. He rolls two ordinary six-sided dice. If the result is double six, all those in the room are shot. On any other roll, the occupants of the shooting-room are free to leave, and the next group is summoned.

This game has two unusual features. First, the groups are summoned to the room in increasingly large sizes:

1, 9, 90, 900, 9000, 90000, . . .

Second, the game ends as soon as double six is rolled and the participants in the current round of the game are shot. From these two facts it is apparent that (with probability one) at least 90% of those who ever participate in this game will be shot. (Apparently a 15% death rate was not high enough to attract philosophical attention!) Hence, concludes the argument, for any participant in the game, the probability of being in the final, fatal round is 0.9.

There is an obvious objection to this argument. Assuming that participants in the game know that (1) they will be shot if and only if the dice come up double six, and (2) the chance of that is one in thirty-six, their subjective probability of being shot is one in thirty-six, rather than 0.9.<sup>2</sup>

It is worth noting that Leslie first presents the shooting-room analogy as an *objection* to the Doomsday Argument (Leslie 1996, 235–236). Indeed, he acknowledges that the probability of death is one in thirty-six rather than .9 so long as the dice tosses are genuinely indetermin-

2. Cognoscenti will note that this argument presupposes that subjective probabilities are in accord with Lewis's Principal Principle (see Lewis 1980).

istic. Curiously, he maintains that in a *deterministic* version of the shooting-room set-up, in which the outcomes of the dice throws are determined from the beginning, the argument for a 0.9 probability of being shot is valid. We shall not pursue this issue further here (see Leslie 1996, 254ff.).

It is instructive to see what is wrong with the inference from ‘90% of those who enter the room will die’ to ‘my probability of dying (given that I have entered the room) is .9.’ This inference rests on the crucial assumption that we are not special with respect to the order in which we are summoned to the room. The problem here is that this assumption requires a uniform prior distribution over a countable infinity of possibilities, since we might be assigned any natural number as our selection position. If we assume that our probability measure is countably additive, then there is no such uniform prior distribution.

Analogously, hypotheses postulating an early doom must be more probable than hypotheses postulating later doom. Let  $D_j$  represent the hypothesis that doom occurs in generation  $j$ . Let  $P(D_j)$  represent the probability of doom in generation  $j$ , where  $P$  is a standard countably additive probability measure. By assumption,  $\sum_{j=1}^{\infty} P(D_j) = 1$ . It follows that for any small probability  $\varepsilon > 0$ , there will be some  $n$  such that the probability that doom occurs in a generation later than  $n$  will be less than  $\varepsilon$ . This is clearly not the sort of argument in favor of early doom that Leslie has in mind. For one thing, it in no way depends upon our observing that we are alive now.

De Finetti (1975) has famously argued against the condition of countable additivity on grounds similar to these. He argued that it ought to be rational to believe that every ticket in a countably infinite lottery has an equal chance of winning. We have shown in (Bartha and Hitchcock 1999) that if we abandon countable additivity and adopt a non-standard measure, then we can indeed find such a uniform prior distribution. Even then, however, it can be demonstrated that the subjective probability for death should remain 1 in 36.

In short, we claim that the shooting-room version of the argument does not succeed; the detailed arguments are presented elsewhere. However, this does not undermine Leslie’s main argument. In particular, the ‘lottery’ version makes no illicit appeal to a uniform prior probability distribution. Indeed, our contention is that that argument involves a valid probabilistic inference, but is mistaken in its eschatological application.

### 3. Lottery Version of the Doomsday Argument.

*3.1. Formulation of the Argument.* We begin by making the argument sketched in the introduction more precise. Let  $D_1$  stand for ‘doom

sooner' (fifty billion humans) and  $D_2$  for 'doom later' (forty quadrillion humans), and let  $F$  stand for the observation that I am among the first forty billion humans. Let us assume that the prior probabilities  $P(D_1)$  and  $P(D_2)$  are 1% and 99%, respectively, so that we begin with the assumption that doom is much more likely to be remote. Then Bayes's theorem tells us:

$$\begin{aligned} P(D_1/F) &= \frac{P(F/D_1) \cdot P(D_1)}{P(F/D_1) \cdot P(D_1) + P(F/D_2) \cdot P(D_2)} \\ &= \frac{(0.8)(0.01)}{(0.8)(0.01) + (0.000001)(0.99)} \\ &= 0.9999 \end{aligned}$$

and by a similar calculation,  $P(D_2/F) = 0.0001$ . Conditioning on  $F$  dramatically increases the probability of  $D_1$  at the expense of  $D_2$ , no matter what prior probabilities we started with. So we have what Leslie calls a Bayesian shift in favor of early doom.

To see more clearly how the argument works, let us generalize to allow for more than two hypotheses about Doomsday. We adopt the following notation:

$D_j$ : Doom occurs after precisely  $j$  people have been born.

$B_i$ : My birth rank is  $i$ ; i.e., I am the  $i$ 'th person born.

We make just two assumptions:

1.  $\sum_{j=1}^{\infty} P(D_j) = 1$  [Doom occurs]
2. For each  $j$  and each  $i \leq j$ ,  $P(B_i/D_j) = 1/j$ ; if  $i > j$ , then  $P(B_i/D_j) = 0$ . [Lottery Assumption]

The first assumption, that with probability 1 Doomsday (the end of the human race) will occur at some point, is unproblematic. The second assumption is the crucial one. It states the following:

*Conditional upon the assumption* that exactly  $j$  people will be born before Doomsday arrives, I ought to assign a subjective probability  $1/j$  to my having any birth rank between 1 and  $j$ , and 0 for any higher rank.

We call this the 'lottery assumption' because it implies that our situation vis-à-vis Doomsday is precisely analogous to a lottery of unknown size in which each ticket has an equal chance of winning. Ticket number is analogous to birth rank;  $D_j$  is analogous to "j tickets are issued" and  $B_i$  to "my ticket number is  $i$ ." The lottery assumption is quite plausible. It is just the idea that none of us is special with respect

to which ticket we get, or in the original problem, that none of us is special with respect to birth rank.

It is easy to verify that, on either the lottery or Doomsday interpretation, we have

$$P(B_i) = \sum_{j=i}^{\infty} (1/j)P(D_j).^3 \quad (1)$$

From assumptions 1 and 2, it follows via Bayes's Theorem that if  $i \leq j$ ,

$$\begin{aligned} \frac{P(D_j/B_i)}{P(D_j)} &= \frac{P(B_i/D_j)}{P(B_i)} \\ &= \frac{(1/j)}{P(B_i)}; \end{aligned}$$

that is, given the information that my birth rank is  $i$ , each doom hypothesis  $D_j$  has its probability shifted by a factor  $\frac{(1/j)}{P(B_i)}$ . Since  $P(B_i)$  is independent of  $j$ , this value is inversely proportional to  $j$ . Note that the probability of the hypotheses  $D_j$  will increase for all values of  $j$  up to  $\frac{1}{P(B_i)}$ , and decrease beyond that point. On finding that you have a low birth rank, you should revise your estimates in favor of early doom—just as, on finding a low number on your winning ticket, you should revise your estimates in favor of a small lottery.

Some critics of the Doomsday Argument, most notably Eckhardt, have argued that there is no way to make sense of the 'lottery' assumption, which Eckhardt calls the *human randomness (HR) assumption*. He puts the assumption in the following form:

We can validly consider our birth rank as generated by random or equiprobable sampling from the collection of all persons who ever live. (1997, 248)

More precisely, his interpretation is this:

$$P(B_i) = 1/M \text{ for } i \leq M, \text{ and } 0 \text{ for } i > M,$$

where exactly  $M$  people will ever be born before Doomsday arrives. Such an assignment is obviously objectionable, since we do not know in advance what  $M$  is. But it should be clear from the earlier discussion

3. Eckhardt (1997, fn. 8) mentions this distribution as a possible "doomsday" prior.

that the lottery assumption 2 is quite different. It requires only that *conditional on the assumption* that exactly  $M$  people will be born before the coming of Doom, I ought to assign a subjective probability of  $1/M$  for my having birth rank between 1 and  $M$ , and probability zero for any higher birth rank. In this case, my subjective probability is a *mixture* of probabilities conditional upon different numbers of people being born before doom. There is nothing incoherent about this assumption, and this is all that Leslie needs to make the argument work.

Some might argue that even if Eckhardt is wrong about the impossibility of a lottery assumption, nevertheless the flaw lies in the unjustifiability of the lottery prior distribution, i.e., assumption 2. We disagree. Granted, one can always criticize any choice of priors. We wish to emphasize, however, that the Doomsday Argument is a sufficiently interesting result if it can be shown to succeed for a class of well-motivated prior distributions, namely 'lottery' priors consistent with assumptions 1 and 2. Unlike some critics of the Doomsday Argument, we accept that the lottery priors are well-motivated. In particular, the assumption 2 is permissible because it does not require a uniform measure over hypotheses about birth order, which would be ruled out by countable additivity. Furthermore, the lottery priors are based on the intuitively plausible idea that none of us is special with respect to birth rank. Thus, the argument as presented is much more interesting than the mere assertion that a shift in favor of early doom occurs for some (possibly implausible) choices of priors.

Leslie does maintain (1996, 203) that the Doomsday Argument requires nothing more than an application of Bayes's Theorem to our initial Doomsday probabilities. In fact, the formalized version of the argument just presented shows that the argument's success depends crucially upon the selection of a 'lottery' prior, i.e., a prior distribution consistent with assumptions 1 and 2. So a commitment to Bayes's Theorem does not by itself ensure the 'Doomsday shift'. Insofar as he allows great latitude in assigning the prior probabilities  $P(D_j)$  but appears to regard the 'lottery' prior as the only natural choice, Leslie's views are less stringent than those of logical probabilists such as Keynes or Carnap, but more stringent than those of pure subjective Bayesians.

What, then, is the problem with the Doomsday Argument? The first step is to realize that a small, but ultimately significant, modification should be made for the lottery assumption to be legitimate. Observe that assumptions 1 and 2 entail that I exist (with probability 1). The assumption  $E$  that I exist is just  $i \leq j$ —i.e., my birth rank is less than or equal to the index of doom. More formally,  $E$  is equivalent to  $\bigvee_{i \leq j} B_i \cdot D_j$ . Assuming the possibilities  $B_i$  are exclusive, assumption 2 entails that



$$P(E/D_j) = \sum_{i=1}^j P(B_i/D_j) = 1$$

for each  $j$ . Combining this with assumption 1 yields  $P(E) = 1$ .

Now the Doomsday Argument is supposed to rest on the intuition that “I am nobody special”; but surely the fact that I exist is extraordinary—just as extraordinary, in the lottery analogy, as finding that I happen to be one of the lucky people who have a lottery ticket. We should modify our assumptions to allow for the possibility that I might never have existed. We can do this by making the two assumptions conditional on  $E$ :

- 1'  $\sum_{j=1}^{\infty} P(D_j/E) = 1$   
 2' For each  $j$  and each  $i \leq j$ ,  $P(B_i/D_j \cdot E) = 1/j$ ; if  $i > j$ , then  $P(B_i/D_j \cdot E) = 0$ .

This step is innocuous in one sense, since if the assumptions 1 and 2 are true and we take our existence for granted, then the modified versions 1' and 2' must also be true. The main advantage of the modification is that the assumptions remain plausible even if we allow for the possibility of our own non-existence (i.e.,  $P(E) \neq 1$ ). This lets us carry out the Doomsday Argument even in this more general case. Replacing (1), we have

$$P(B_i/E) = \sum_{j=i}^{\infty} (1/j)P(D_j/E). \quad (2)$$

The Bayesian shift must now be seen not as a relation between the absolute prior  $P(D_j)$  and  $P(D_j/B_i)$ , but rather as relating  $P(D_j/E)$  and  $P(D_j/B_i \cdot E)$ . For  $i \leq j$ , we have:

$$\begin{aligned} \frac{P(D_j/B_i \cdot E)}{P(D_j/E)} &= \frac{P(B_i/D_j \cdot E)}{P(B_i/E)} \\ &= \frac{(1/j)}{P(B_i/E)}. \end{aligned} \quad (3)$$

**3.2 Critique.** If the lottery distribution is plausible given the information that I exist, and if the ensuing model entails a massive shift in favor of ‘doom sooner’, then where is the flaw in the Doomsday Argument? We suggest that the flaw lies in ignoring the fact that there is an earlier, equally massive shift in favor of ‘doom later’ that occurs when we first conditionalize upon the information that we exist. This shift from  $P(D_j)$  to  $P(D_j/E)$  virtually cancels out the second ‘Doomsday shift’ from  $P(D_j/E)$  to  $P(D_j/B_i \cdot E)$ ; the reason for the qualifier “vir-

tually" will shortly be made precise. We will refer to these successive probabilities as the 'initial', 'intermediate', and 'final' probabilities of  $D_j$ . Note that the intermediate probability  $P(D_j/E)$  is the posterior probability of the first shift and the prior probability of the second.

To begin with, note that the derivation of (3) goes through equally well if we impose an upper bound  $M$  on the number of humans that might ever be born. We can let  $M$  be as large as we please—for instance, a generous estimate of the number of elementary particles in the universe times a generous estimate of the number of seconds between the big bang and the heat death of the universe. In this case, we need only replace assumption 1' with  $\sum_{j=1}^M P(D_j/E) = 1$  and restrict assumption 2' to cases where the conditional probability is well defined, i.e., where  $j \leq M$ . Imposing such a cap simplifies much of the mathematics; we will remove this restriction in Section 3.3.

Let us return to the case of the lottery of unknown size. It might help to think instead of the "Publishers Clearing House Awards" that we all receive occasionally, though imagine that the award notices are numbered. Here, the analogue of the proposition  $E$  that I exist is the proposition that I receive a notice. If I discover that I have received such a notice, then the probability for finding any particular number on the notice is given by the distribution (2). That conclusion is supported by the thought that I should treat myself as no different from anybody else as regards my chance of receiving any particular notice.

I can also ask for my prior probability of receiving such a notice at all, if an unknown number of notices have been mailed out. Employing the same reasoning, this probability should depend solely upon the size of the "Publishers" lottery and the size of the population (i.e., the maximum number of people who could possibly receive a notice). If the total population is  $M$ , and as before  $D_j$  signifies a lottery of size  $j$ , then  $P(E/D_j) = j/M$ . So my prior probability  $P(E)$  of receiving a notice is given by

$$\begin{aligned} P(E) &= \sum_{j=1}^M P(E/D_j)P(D_j) \\ &= \sum_{j=1}^M (j/M)P(D_j). \end{aligned}$$

It follows from Bayes's Theorem that for  $i \leq j$ ,

$$\begin{aligned} \frac{P(D_j/E)}{P(D_j)} &= \frac{P(E/D_j)P(D_j)}{P(E)P(D_j)} \\ &= \frac{(j/M)}{P(E)}. \end{aligned}$$

This will obviously be much higher for large values of  $j$ . When we put this together with the calculation (3), we can see that the two shifts virtually cancel out if we learn both that we have a notice and that its number is  $i$ :

$$\begin{aligned} \frac{P(D_j/B_i \cdot E)}{P(D_j)} &= \frac{P(D_j/B_i \cdot E)}{P(D_j/E)} \frac{P(D_j/E)}{P(D_j)} \\ &= \frac{(1/j)}{P(B_i/E)} \frac{(j/M)}{P(E)} \\ &= \frac{1}{MP(B_i \cdot E)}, \end{aligned} \tag{4}$$

which is independent of  $j$ .

None of this is surprising. Imagine that we start with some notion of how many ‘‘Publishers Clearing House’’ award notices are sent out in a typical mailing. When we find an award notice in the afternoon mail, our natural reaction is to revise the initial estimate upwards: it seems unlikely that we would have received anything unless the mailing list was huge. But we then discover a relatively low number on our notice, and revise our estimate downwards. The two revisions virtually cancel out, leaving us with no more information than that at least as many people received notices as the number printed on ours.<sup>4</sup>

Back to the Doomsday Argument. Here, the intermediate probabilities  $P(D_j/E)$  (which serve as the priors in Leslie’s argument) involve conditionalizing on my being alive, but not on my being alive ‘now’ with birth rank  $i$ . In such a scenario, our contention is that we ought indeed to assign tiny intermediate probabilities to early doom. It seems queer to talk this way, since it requires us to imagine ourselves in a position where we first have initial subjective probabilities of the form  $P(D_j)$ , and then revise them upon learning that we actually exist (while still ignorant of when we exist). The prior probabilities are ones we literally could never have had—unless we take seriously the possibility of Cartesian doubts about our own existence (though not seriously enough to worry about whose priors are then at issue).

The difficulty here is the well-known problem of ‘old evidence’ for Bayesian confirmation theory: if I have known some proposition  $E$  all along, then conditionalizing on it cannot change the probability of another proposition  $H$ . Therefore,  $E$  cannot serve as evidence for or

4. Our thanks to Bernie Linsky for pointing out the analogy to this sort of phenomenon.

against  $H$ . One line of response to this problem, suggested by Howson and Urbach (1989, 270–271), is to understand the evidential bearing of  $E$  upon  $H$  in terms of conditionalization upon  $E$  using some hypothetical prior probability distribution that does not assign to  $E$  a probability of one. This is precisely what we are doing here with the initial probability distribution  $P(D_j)$ .

We do not pretend to have shown that the use of hypothetical priors solves the problem of old evidence to complete satisfaction, but we make two observations. First, *ad hominem*, Leslie himself seems to endorse a solution to the problem of old evidence along these lines (1996, 218ff.). Second, even if some such solution to the problem of old evidence does not work, that does not undermine our argument. If no solution to the problem of old evidence is in the offing, then Bayesian confirmation theory will be inadequate for capturing relations of evidential bearing. Our contention is that our being alive at all does have evidential bearing upon the number of people that will ever live. This may be true regardless of whether the relevant notion of evidential bearing is captured in Bayesian terms.

Here is a 'just-so story' that might ground our hypothetical initial probability distribution. There is a large number ( $M$ ) of souls in heaven, and they are all patiently waiting in line to be embodied down on earth. Until my soul is embodied, I have no idea where it is in line. After person  $j$  is born, there is a possibility that doom will enshroud the earth, and no more souls will be embodied. We do not know when this will happen, but we have certain subjective probabilities  $P(D_j)$ .

The processes that determine when doom will occur are independent of those that determine where my soul is in line. Thus, the crucial assumptions about initial probabilities are the following, where as usual  $D_j$  means doom after the  $j$ 'th person is born, and  $B_i$  means that my birth rank (or position in line) is  $i$ .

1.  $\sum_{j=1}^M P(D_j) = 1$ , i.e., doom happens eventually.
2.  $P(B_i) = 1/M$ : I am equally likely to be assigned any position in line.
3.  $P(B_i \cdot D_j) = (1/M)P(D_j)$ : my position in line and the coming of doom are independent.

Given the representation for  $E$  cited earlier, these assumptions ground the 'lottery' distribution. We have, for  $i \leq j$ ,

$$\begin{aligned}
 P(B_i/E \cdot D_j) &= \frac{P(B_i \cdot E \cdot D_j)}{P(E \cdot D_j)} \\
 &= \frac{P(B_i \cdot D_j)}{P(E \cdot D_j)} \\
 &= \frac{(1/M)P(D_j)}{\sum_{i=1}^j P(B_i \cdot D_j)} \\
 &= \frac{(1/M)P(D_j)}{(j/M)P(D_j)} \\
 &= 1/j.
 \end{aligned} \tag{5}$$

So this imaginary set-up permits reasoning analogous to that which leads to (4) in the lottery case. The conclusion is the same: the two shifts virtually cancel out. All that happens when I conditionalize on *both* my existence and my early birth rank is that the ‘probability surplus’ that arises from finding out that Doomsday has not happened yet is divided proportionately among all the remaining Doomsday hypotheses. In other words, the probabilities of the unfalsified Doomsday hypotheses are renormalized so that they sum to one. It is this proportionate shift that we wish to signify when we say that the two shifts virtually cancel out.

The ‘just-so’ story is not intended as a serious piece of metaphysics or theology, of course, but rather as a way of making graphic a certain sort of hypothetical probability distribution  $P(D_j)$  that reflects the possibility that we might never have existed. Leslie’s argument begins with priors of the form  $P(D_j/E)$ . We have shown that there is a clear sense in which these ‘priors’ ought to already reflect a massive shift toward ‘doom far’ hypotheses.

Leslie does consider this sort of objection:

The bigger our race is in its temporal entirety, the more opportunities there are of being born into it. This counterbalances the greater unlikelihood of being born early. (1996, 225)

To counter the objection, he considers an extreme case. God tosses a coin; on a result of heads, he creates ninety million people, one of whom is named ‘Dr. Black’, and on a result of tails, he creates just one person named ‘Dr. Green’. On the reasoning above, upon learning that you exist, you ought to favor the hypothesis that the coin landed heads. But, Leslie argues, you surely ought to bet that you are Dr. Green rather than Dr. Black.

This response misses the point, as it fails to take account of the possibility that you might not have existed. It assumes, for example,

that if the coin lands tails, then not only would Dr. Green exist, but *you* would be Dr. Green. If instead we imagine that there are ninety million plus one souls waiting for the result of the coin toss, ninety million hoping for heads and one lonely soul hoping for tails, then the fact of your existence makes it almost a certainty that the result of the toss was heads, and a straightforward calculation shows that you have an equal chance ( $1/90,000,001$ ) of being Dr. Green or Dr. Black.

*3.3 Removing the Cap.* The argument of the preceding section presupposed an a priori cap on the possible size of the human population. Nothing in Leslie's argument prohibits the imposition of such a cap: the shift in favor of early doom hypotheses was supposed to follow from Bayes's theorem, supplemented only by assumptions 1' and 2' above. Nonetheless, it is worth exploring the consequences of removing the cap for Leslie's argument. It turns out that if we remove the cap, the argument of the previous section can no longer go through.<sup>5</sup> In particular if we make the following assumptions:

1.  $P(B_i \cdot D_j) = P(B_i)P(D_j)$ , and
2.  $P(B_i/E \cdot D_j) = 1/j$ , for  $i \leq j$ ,

then it follows that the  $P(B_i)$  are all equal, in violation of countable additivity. So we cannot find an initial probability measure according to which birth rank and timing of doom are independent, and which yields the lottery distribution as an intermediate distribution. Put another way, we can find initial probability distributions, those that yield the lottery distribution as an intermediate, such that my being born with rank  $i$  does have evidential bearing upon the timing of Doom.

Of course, one can always choose an initial probability distribution such that anything can count as evidence for anything else we choose (except elevating an initial probability of zero or lowering an initial probability of one). The mere existence of such distributions does not itself prove very much. The question is: which initial distributions are sufficiently well-motivated to merit serious attention? We know something of the candidates: we can have initial probabilities satisfying assumption 1 or assumption 2, but not both. How seriously should we take candidates that satisfy 2 but not 1? As noted earlier, the lottery distribution has appeal on the grounds of symmetry. On the other

5. Note also that the argument does not go through for all capped distributions, just those satisfying the further constraints described above. But again: Leslie's argument was supposed to follow from Bayes's theorem together with assumptions 1' and 2' alone, and not to require further assumptions about the structure of the relevant distribution. Thanks to Peter Vranas for discussion on this point.

hand, any initial probability that satisfies assumption 2 will yield some version of the Doomsday Argument, a consequence most people find extraordinarily counterintuitive. Starting with an initial probability that satisfies condition 1 will allow us to sidestep the Doomsday Argument, although we will have to forego the lottery distribution. That may not be so bad—no worse, perhaps, than giving up an infinite lottery where each ticket is equally likely to win. One person's *modus ponens* is another's *modus tollens*.

But before we find ourselves trapped by a false dichotomy, note that we can accept both assumptions 1 and 2 if we are willing to give up countable additivity. Perhaps the simplest way to do this would be to use a non-standard measure. If we let  $M$  stand for some infinite—or more precisely, hyperfinite—integer, and let  $P$  be a hyperfinitely additive measure yielding non-standard values on an outcome space generated by events of the form  $B_i \cdot D_j$ , where  $1 \leq i, j \leq M$ , then the derivation of the previous section remains valid.<sup>6</sup> This would allow us to remove the finite cap on the size of the population, while retaining the virtual cancellation of the shifts resulting from the 'discoveries' that I exist, and that my birth rank is  $i$ .

**4. Conclusion.** Given that we exist, our existence now does indeed favor 'Doom sooner' over 'Doom later'. On the other hand, we might not have existed at all, and our existence favors 'Doom later'. In the simplest cases, these two effects cancel one another out, yielding the happy result that our existence now tells us nothing whatsoever about the coming of Doom, except that it has not yet happened. Two factors have helped to obscure this rather straightforward remedy to our anxieties about Doom. First, we never actually have subjective probabilities in which our existence is not taken as a given. This is just a version of the familiar problem of old evidence. Second, an interesting range of cases are not 'the simplest'. Yet even in these more complex cases, where there is no upper bound on the possible size of the human population, Bayes's theorem in no way compels us to regard our present existence as evidence for the imminence of doom.

#### REFERENCES

- Bartha, P. and C. Hitchcock (1999), "The Shooting-Room Paradox and Conditionalizing on 'Measurably Challenged' Sets," Forthcoming in *Synthese*.  
 de Finetti, B. (1975), *Theory of Probability*, vols. 1 and 2. Translated by Antonio Machì and Adrian Smith. New York: Wiley.

6. We carry out a construction of this sort for the "shooting-room" problem in Bartha and Hitchcock 1999.

- Eckhardt, W. (1997), "A Shooting-Room View of Doomsday," *Journal of Philosophy* XCIII (5), 244-259.
- Fitelson, B. (1999). The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity. *Philosophy of Science* 66:3 (Supplement), (this issue).
- Howson, C. and P. Urbach (1989), *Scientific Reasoning: the Bayesian Approach*. La Salle: Open Court Press.
- Leslie, J. (1996), *The End of the World*. New York: Routledge.
- Lewis, D. (1980), "A Subjectivist's Guide to Objective Chance," in R. C. Jeffrey (ed.), *Studies in Inductive Logic and Probability*. Berkeley: University of California Press, 263-293.