Originally published in *Journal of Philosophy*, Vol. 94, No. 94
May 1997

A Shooting Room View of Doomsday

William Eckhardt

1) The Doomsday Argument

Recently H.B. Nielsen (1989, pp. 454-59), John Leslie (1990, 1992, 1993, 1996), and J. Richard Gott III (1993) have published similar arguments purporting to show that the time left until the end of the human race may be shorter than we generally suppose. In its simplest form this reasoning can be summarized as follows: We may think of ourselves as randomly selected from among all who ever live. Our randomness makes it unlikely that we are among the earliest humans. A calculation demonstrates that if we are not among the earliest humans, then the human race has less time left than is usually allotted to it.

The treatments differ mainly in emphasis. Gott's article has more mathematical trappings and brims with speculation; he also considers questions other than that of human survival. Leslie presents a dazzling array of illustrations and analogies, evidently as an aid to negotiating this treacherous terrain. He is far and away the most prolific on this topic having written a dozen articles (for references see Leslie (1996)) as well as a book length treatment of issues relating to the end of the world. Following Leslie, we call this kind of reasoning the *Doomsday argument*.

The ease with which the Doomsday argument extracts from a quite modest investment of current fact substantial information about the remote future has provoked suspicion in most quarters; however, critiques that have not been snidely dismissive have tended to be as mystifying as the Doomsday argument itself. The purpose of this report is to show that the fallacy inherent in the argument can easily be discerned by focusing on correct application of

Bayes' Theorem. It facilitates the analysis that Doomsday reasoning closely parallels a certain fallacious train of thought associated with a probability theory puzzle called the Shooting Room.

Leslie reviews a variety of objections to the Doomsday argument, but I believe these distract from the core issue.[1] Many concern some intrinsic unknowability of the future, e.g., Doomsday scenarios inhabit a limbo of unrealized possibilities, thereby evading all current attempts at divination. The difficulty in mounting an attack on the Doomsday argument solely on the basis of some alleged intractability of future events — that they are indeterminate, unactualized, multi-potential, etc. — is that success would necessitate discarding the whole project of making future projections based on current information. When a demographer makes a projection of the world population in the year A.D. 2100, it is quite reasonable to challenge the data, the model or the computation, but does it make sense to maintain the estimate is invalid because physics may be indeterministic or because these projections are merely of possible people whereas in the year A.D. 2100 the world will be populated with real people? It would be calamitous if we had to refrain from employing probabilistic or statistical reasoning regarding future populations until elementary particle physicists come to a decision as to whether the world is indeterministic or until philosophers settle the numerosity of possible futures or the nature of unborn identity. If the Doomsday argument makes questionable projections, it is on account of

---

[1] As an example, Leslie expends not a little effort countering the somewhat fatuous objection that future humans are not alive to observe anything [(1996) pp. 19-21, 214-218, 246-247 and (1993) pp. 489-490 in which, unaccountably, this objection is ascribed to me]. Surely it would suffice to point out that the issue is not what unborn humans can or cannot do, but what we can infer about their numerosity from our birth rank.

flaws in the procedure, not because of some generic defect in the idea of estimating unborn populations. Instead of wrangling over the metaphysics of measuring a fate that has not yet been sealed, we choose to examine the validity of the probabilistic reasoning underlying the Doomsday argument.

Leslie's views on the regulating role of determinism in the Doomsday and kindred arguments are, I believe, wholly unjustified. According to him, this kind of reasoning works excellently under determinism, but as we slide along the scale of increasing indeterminism Doomsday arguments become progressively undermined until for the case of radical indeterminism they may fizzle altogether [Leslie (1992) p. 537; (1996) pp. 188, 233-34]. However, if there existed a mode of statistical inference that were valid according to the extent that determinism were true, then by repeatedly testing the accuracy of this type of statistical inference, one could gauge the correctness of determinism. Since this conclusion is highly implausible it is a safe bet that statistical inferences, including those that underlie the Doomsday argument do not hinge on the truth of determinism. (Except for certain quantum experiments, uncertainties from other sources swamp any uncertainty originating in indeterminism; this is why the determinism question is not a burning issue among say, insurance companies.) Similar counterarguments can be fashioned around some of the other philosophical issues upon which certain critics have claimed the Doomsday argument's validity rests — the multiplicity of possible futures, presentist vs. eternalist theories of time, and the nature of human identity. It would be a tremendous boon to philosophical investigation if solutions to such problems were connected to the power or accuracy of certain statistical techniques or of applications of Bayes' Theorem. Lacking this philosopher's stone, the decision maker must "play" against uncertainty due to ignorance and uncertainty due to indeterminism in the same manner; this applies also to

Doomsday reasoning. (One might protest that the Doomsday situation is unique and cannot be modeled experimentally; however, Leslie is explicit that his novel reasoning about determinism applies to certain games, some of which can be played or simulated. See sec. 4 for an example.)

2) Bayes' Theorem

Proponents of the Doomsday argument base their reasoning on Bayes' Theorem, an uncontroversial, almost trivial proposition of the probability calculus. (See the formula on p. 20; virtually any textbook of probability theory treats this topic, e.g., Feller (1968) p. 124.) Use of Bayes' theorem should not be confused with Bayesianism, a controversial viewpoint on the nature of probability associated with subjectivism and the claim that correct statistical inference depends crucially on assessment of prior distributions. The most avid anti-Bayesian has no quarrel with Bayes' Theorem proper. Leslie blurs this distinction in a manner that allows him to portray the Doomsday argument as an application of a controversial doctrine (Bayesianism) to a straightforward fact of our existence (birth rank) with the suggestion that the argument, although a little dubious, may, along with Bayesianism, turn out to be fundamentally sound. However, Bayes' Theorem is a rigorously demonstrable statement of the probability calculus; if the premises of the theorem are fulfilled, the conclusion follows with the force of logic. As with other rigorous results, the sticking point in applications is not whether the theorem is correct, but the extent to which the premises are fulfilled. The Doomsday argument should be seen as a straightforward application of an uncontroversial rule (Bayes') to data produced from a highly questionable assumption (the Human Randomness assumption, described below).

Leslie seems to believe these matters can be settled through sheer weight of accumulated analogies. Although of possible pedagogic or heuristic value, such analogical reasoning can at best support only the preliminary stages of investigation, whereafter it becomes incumbent to find nonanalogical evidence or to investigate the validity of the analogies themselves. Our approach is to concentrate on whether in its Doomsday use, the preconditions for valid application of Bayes' Theorem have been fulfilled.

We begin with a set of alternatives A_1, A_2, ..., A_n...and a piece of information R. For the Doomsday argument we interpret *R* as knowledge of one's own birth rank[2] and the A_n's to be various Doomsday scenarios. To apply Bayes' Theorem we need two sets of probabilities:

i. *Prior probabilities* for the alternatives A_n. These probabilities are *logically* prior to R in the sense that they do not take R into account; in applications the A_n's often refer to events that are also temporally prior to R, but this is not essential.

ii. *Conditional probabilities* connecting R to the A_n's; these tell how likely R is under each of the relevant alternatives.

Bayes' Theorem then permits the derivation of a new set of probabilities revealing how likely each of the A_n's are after the information R has been factored in; these new probabilities represent a revision of the prior probabilities in light of the additional information R. If no

---

2 *Birth rank* refers to one's place in the birth order, numbered consecutively, of all humans who ever live. Of course we know our own birth rank only very roughly. By assuming knowledge of an exact rank, we simplify the analysis while making a stipulation that is favorable to the argument we are critiquing. If an exact rank is worthless for projecting Doomsday, then an approximate one is also.

revision occurs we speak of a *trivial* application of the theorem; intuitively this means R has nothing to do with the A_n's.

3) The Fallacy

In the literature on Bayes' Theorem most attention has been paid to securing the often problematic prior probabilities. By contrast the role of the conditional probabilities is abundantly clear: if the information R is not probabilistically connected to the alternatives, then learning R could not possibly tell us anything new about these alternatives, and applying Bayes' Theorem would be fruitless.

In the Doomsday argument, the conditional probabilities address the question of how Doomsday is probabilistically connected to one's own birth rank. The natural, and I believe correct answer is that there exists no connection whatever. In that case Bayes' Theorem applies trivially and there should be no shift in our assessments of Doomsday upon learning our own birth rank. If the Doomsdayers draw the opposite conclusion, it is because they possess an additional assumption, namely, *the Human Randomness (HR) Assumption*: *We can validly consider our birth rank as generated by random or equiprobable sampling from the collection of all persons who ever live.*

Bayes' Theorem permits the derivation of a conditional dependence from the converse dependence, in the presence of the right kind of background information. If our birth rank can tell us via Bayes' Theorem something about the likelihood of Doomsday, then it has to be because Doomsday can tell us something about the likelihood of our birth rank; there is no way around this. The HR assumption provides the needed connection since in that case the probability one has a given birth rank is inversely proportional to how many ranks exist: the sooner Doomsday comes, the fewer the available birth ranks, the more likely a random human has a given rank. It is this that powers the Doomsday argument, since under the HR assumption, chancing to have a birth rank as low as we do makes it likely there are relatively few ranks available overall.

Leslie claims that Doomsday reasoning "doesn't generate risk estimates *all by itself*. It argues for Bayesian shifts which magnify any risk-estimates that have been reached by other means." [Leslie (1996) p. 204 (his italics); see also p. 213]. If lacking risk assessments means, among other things, openmindedness about the possibilities for very long term survival of the human race, then acceptance of the HR assumption radically curtails such liberality. Furthermore, in the form outlined at the beginning of this article, the argument employs no prior risk assessments. To settle this matter, we describe a procedure in the appendix for deriving precise doom probabilities for cases with absolutely no prior risk estimates. It is not true that the Doomsday argument never pulls projections "out of the air."

There are other indications that Leslie has not fully faced up to the perversities of Doomsday reasoning when applied to human destiny rather than to lotteries. For explicitness we begin with a lottery that has either 400 or 600 tickets. A valid random drawing yields ticket number 399. What should be concluded? A Doomsdayer might reason as follows: "If the lottery has 400 tickets the drawing is so unusual as to produce a ticket in the top 1/2% of all tickets, whereas if the lottery has 600 tickets, number 399 is much more reasonably placed in the 66th percentile. So randomly drawing ticket 399 strongly favors the larger pool." This reasoning is demonstrably incorrect. There is one chance in four hundred of randomly drawing ticket 399 from a pool of 400 and only one chance in six hundred of doing so from a pool of 600. By Bayes' Theorem you should adjust your estimates in favor of the smaller pool. This example not only exposes a peril inherent in this kind of reasoning, but highlights its most abiding peculiarity: when correct allowance is made for the randomness assumption, Doomsday reasoning has an unalterable bias toward *earlier* doom. (See the appendix for verification of this.) The revision brought on by the HR assumption is not a redress of some optimistic bias in earlier estimates; it is an invariant feature of the argument that any and all revisions *shorten* time left until doom. It matters not whether prior assessments predict early or late doom, whether current populations are large or small, growing or shrinking, or whether many or few have been born to date; there exists no constellation of circumstances under which a user of the argument concludes that prior assessments of doom ought to recede. If you are convinced doom is a week away, the argument calls for earlier arrival; if you think doom is likely in ten minutes, the argument presses for sooner. There is no room in this complex for the idea that original projections might err through earliness as well as lateness and that corrective action might require upward as well as downward revision. Can a procedure be valid that holds all other estimates of survival time to be too generous, irrespective of what they are or on what basis they were obtained?

Without the HR assumption the proposed application of Bayes' Theorem is trivial and fruitless, with the assumption the reasoning runs smoothly to its alarming conclusion. For a true lottery with an unknown quantity of tickets consecutively numbered starting with 1, the random drawing of a numbered ticket does indeed give us information about the probable size of the ticket pool. However, the issue in the Doomsday argument is whether such random lotteries constitute appropriate models — whether the HR assumption holds. In light of this Leslie's copious urn and lottery examples can all be seen to be question begging; each one assumes equiprobability of sampling, precisely what must be established to validate Doomsday reasoning. (In response to my objection to the assumption of random human sampling in the Doomsday argument, Leslie produced additional urn and lottery analogies presupposing random sampling [Leslie (1993)]).

If one does not assume equiprobable selection, then natural examples cut the other way: suppose on each trial the *Consecutive Token Dispenser* expels either 50 (early doom) or 100 (late doom) consecutively numbered tokens at the rate of one per minute. The sampler's task is to select a randomly ranked token *as it is expelled*. A moment's reflection reveals that this cannot be done without prior knowledge of the pool size for each trial; the sampler must contrive to choose randomly from a domain of 1 through 50 when the machine is going to expel only 50, and randomly from 1 through 100 in other cases, without knowing which cases are which. It makes not the least difference whether the pool size for each trial is decided by a deterministic or indeterministic process; the obstacle is ignorance, not indeterminateness. Without knowledge of the trial's termination it is impossible to select random ranks.

Bayes' Theorem cannot give us something from nothing. If the theorem is to tell us what consequence our birth rank has for Doomsday, it can only be because we have fed in information as to what consequence Doomsday has for our birth rank. Lacking knowledge of the latter, we cannot hope to apply Bayes' Theorem profitably.

4) The Shooting Room

Successive groups of individuals are brought into a room and given the same highly favorable wager, say betting $100 that the "House," with fair dice, rolls anything but double sixes. (In the original formulation, losing players are shot, but this added gruesomeness, if nothing else, complicates the question of how one should bet.) Whenever the room occupants win their bets, ten times as many people are recruited for the next round. Once the House wins, the game series is over. So the House can truthfully announce before any games are played that, in spite of the highly favorable odds, at least 90% of all players will lose. The puzzle is that these bets appear to be both favorable and unfavorable, favorable because double sixes are rare, unfavorable because the overwhelming majority of players lose.

Leslie analyzes the Shooting Room in detail [Leslie (1996) pp. 235-36, 251-56]. His treatment closely parallels his approach to the Doomsday argument. What are suspicious maneuvers in the case of the Doomsday argument translate into demonstrable errors in the more explicit Shooting Room story. Leslie considers two Shooting Room variants, a deterministic and an indeterministic one, which we call the *D-series* and the *I-series* respectively. In the D-series, for each round the one chance in thirty-six of losing is simulated by consulting the decimal expansion of π. The outcome of these games are completely pre-determined, although players do not know these outcomes. In the I-series, identical odds are delivered by a device that is presumed to be truly indeterministic. Leslie concludes that in the I-series, players ought always to bet, since they have about a 97% chance of winning, but in the D-series players ought *not* to bet since it is predetermined that at least 90% shall lose.

Leslie states that "there is nothing too paradoxical" in his approach (p. 256). In fact it requires only minor changes to the narrative to nudge this novel reasoning about the role of determinism into self-refutation. Two examples:

i) Add a spectator who can bet in any of the D-series games. It is easy to show that the spectator should opt to play all games in the series, either for the customary reason that in each game there is a better than 97% chance of winning, or for the peculiar outcome based reason that the spectator is thereby fated to a string of wins capped with a single loss. The most the spectator can lose is $100 but her expected winnings are measured in the thousands. By any accounting, the spectator ought to judge these games as highly favorable, and it is surely no argument against this that the precise moment at which her winning streak will end is preordained by the fortunes of π's expansion. Then for any game in the D-series, Leslie's reasoning dictates that the spectator's bet is favorable and the room player's bet is unfavorable, even though they both make the same bet on the same outcome of the same game.

Should it be maintained, for reasons I cannot fathom, that the spectators' statistics cannot be legitimately segregated from those of the room players; we also offer an example that requires no additional personnel.

ii) Play the I-series games first and record the results. These results, which are now a completely pre-determined sequence just like π's decimal expansion, are then used for the D-series games. In other words, in corresponding rounds, I-players and D-players bet on exactly the same "roll" except D-players bet after the fact, but without knowledge of the earlier results. Once again, this reasoning collides with itself. The I-bets are favorable and should be taken, the D-bets are unfavorable and should not be taken, even though in this case these bets are on exactly the same outcomes of the same games.

Our stance is that the Shooting Room is not a paradox at all; rather it is a cogent line of reasoning alongside an utterly spurious one, masquerading as horns of a dilemma. The bets in all Shooting Room games are highly favorable; no other conclusion is possible. The first thing to note is the solid dependence among outcomes for the room occupants in a single round — they all win or they all lose together. If five players lose independently at a *given kind* of wager, that may constitute, in a loose sense, five reasons not to play, but if a million players all lose the *same* wager at once, that constitutes, in the above sense, *one* reason not to play. (It has been remarked that an insurance company would go broke insuring all the players in the Shooting Room as though they had a 97% chance of winning. And so might any insurance company that treats highly dependent events as though independent — a diversified company that finds it rarely receives simultaneous flood claims and decides to insure everyone living on the banks of the Mississippi.) It has long been known that by successively increasing bet size in a sequence of unfavorable bets one can theoretically obtain winning results (e.g., Epstein (1977) pp 52-56) this is the basis of various infamous doubling systems in Roulette and other games; in the Shooting Room game it is the House that carries out the "multiplicative" betting scheme. Irrespective of whether it is a D-game or an I-game, each player ought to reason thus: 90% of all players will lose, but I have less than a 3% chance of belonging to that losing majority. This is no paradox; each player is prospectively likely to be in the minority, since he or she is prospectively likely to win and winning itself causes there to be enough subsequent players to guarantee the winner minority status.

The multiplicative betting system that the House exploits to secure a preponderantly losing final pool is not such as should make a difference to a player's decision. A player knows that the game is structured so that the majority ends up losing. This is to be achieved in one of two ways: a) about a 3% chance the player will lose, along with enough others to assure the right preponderance of losers. b) about a 97% chance the player will win, and enough other players will lose in some future round to secure the right proportion of losers. Surely there is no reason to be found in (b) not to play. Nor should (a), a 3% chance that the player will lose along with many others, be an inhibiting factor. The bet is clearly favorable. (It may come as a relief to theoretical physicists that Leslie's alleged "destruction" of the Many Worlds interpretation of quantum theory [Leslie (1966) pp. 264-266] relies on the same kind of reasoning that prompts him to judge betting even money against double sixes to be unfavorable.)

To explain how the bets can be so favorable for the players yet the House can be ultimately guaranteed an enormous profit, it may help to consider the consequences of playing in more than one Shooting Room Series. Call a sequence of ordinary Shooting Room series a *mega-series*. Players for these consecutive game series are drawn from an unlimited store of numbered players. A player can only play once in any series; some series employ more players than others, but each new series begins with player 1 again, etc. Let us examine the likely career of a particular player, *player n*. Lower numbered players in general play more often than higher numbered players, but it can be demonstrated that if the mega-series continues long enough player n, no matter how high her rank, eventually plays any preassigned number of games. Suppose she has played 100 times. How many times should we expect that player n had to play against double-sixes? Recall she is any player and these are 100 distinct rolls of the dice or simulator. To suppose player n has faced double sixes about 90% of the time is to suppose the absurdity that 90% of all rolls of the dice in the games have been double sixes. One can only conclude that player n has had to contend with double sixes on average about three times out of the hundred, and that the more often she bets the greater her expected net gain becomes. This confirms that the Shooting Room games are highly favorable to players. It remains only to square the rising expected gain of each player with the fact that the House collects in a net sense enormous amounts from the store of players by the end of every game series. Examining the composition of net winners and net losers at virtually any stage of the mega series, we would find the overwhelming majority of players who had bet more than once to be net winners. The House's advantage would be financed almost exclusively by vast numbers who had bet *once* and lost. As the mega-series continued these net losers would have additional opportunities to play and, with freakish exceptions, would become net winners, while the ever expanding group of mostly one time losers would drift ever higher in average rank. Note that the House's advantage depends crucially on the unlimitedness of the player store. With a finite store of players, no matter how large, the House would repeatedly find itself having to pay off all the players without there being any larger group from whom to reclaim the losses. (It is a feature of all multiplicative betting systems that they eventually fail disastrously.) Although theoretically assured of astronomical gains, any real House that undertook to run a Shooting Room would go broke absorbing bad bets while waiting for a huge payoff. It is most assuredly the House and not the player that makes "crazy" bets in this game.

5) Room and Doom

The Doomsday argument relies on a subtle appeal to temporally reversed causation. We begin by exhibiting similar reliance in Leslie's treatment of the more explicit (deterministic) Shooting Room story. Consider a variant with a limited number of rounds say, five, even if no player loss occurs. We do not know the final frequencies for this variant with as much precision as in the standard Shooting Room, but we do know that roughly seven times out of eight the final pool consists exclusively of winners. These bets are without doubt highly favorable for the players. In Leslie's approach, although the version limited to five rounds is favorable for the player, the first five rounds of the standard version are unfavorable (all the games in that version are allegedly unfavorable). However, the sole difference between the two variants concerns what happens after the fifth round and to different players, i.e., it is the possibility of later rounds with other players that allegedly renders the first five rounds of the standard series unfavorable.

The Shooting Room story is structured so that the causal factors that determine whether a player wins or loses all factor through the roll of the dice; the composition of the final pool cannot make itself felt without reverse causation. Suppose you are a candidate player scheduled for a particular round. The dice frequencies influence your membership in the final player pool in at least two ways: *i*) you are caused to become a player and hence to be a member of the final player pool by a player win on the *previous* round; *ii*) you are caused to be a majority or minority member of the final pool by the roll on your scheduled round; specifically you will be a majority member only if the roll is double sixes. By contrast, the final pool frequencies are causally terminal; they influence no other event in the story. Your player status is caused not by your membership in the final pool, but by previous rolls of the dice. Hence it is the dice frequencies and not the final loser frequencies that provide useful guides for action. (This can change, of course, if we expand or modify the context. Sampling randomly from the final pool yields a player with a 90% chance of having lost to double sixes. The final frequencies are in this case a reliable guide because the composition of the final population *influences* the sampling.)

Two peculiarities of the Shooting Room story may be what delude some into supposing it reasonable to use causally terminal frequencies to govern earlier behavior. *i*) The final frequencies are known beforehand with unusual precision; *ii*) the round on which you are caused to belong to the final pool comes before the roll of the dice on which you bet. These combine to make it appear that the determinate statistics of the final pool are already operative when the dice are rolled. A crucial point, and one that bristles with significance for the Doomsday argument, is that it is an error to consider yourself a random or typical player until you lose your bet. Before then you have only about a 3% chance of belonging to the 90% majority. (What lends an air of paradox is that an *unlikely* event makes a group of players *typical*.) This means *you should not consider yourself random until the game series is over*. We can highlight the analogy with Doomsday reasoning by the following consideration: if a player about to bet were truly random among all players, then he or she would have better than a 90% chance of losing, which would mean that with high probability the game series would be over sooner than expected.

We return to the Doomsday argument. The HR assumption has its effect because it stipulates a quantitative relationship between the probability of having your birth rank and the number of people who come after you. For the argument to be valid, this crucial probability cannot be based on your causal antecedents, but has instead to be based not only on how many were born before you but also how many are to be born after you. How it is possible in the selection of a random rank to give the appropriate weight to unborn members of the population? In presuming that unborn populations have somehow been factored into the current selection procedure the Doomsday reasoners tacitly presuppose retroactive causal effects. Suppose some crucial event prevents a holocaust in the year A.D. 2150 and doom is thereby delayed a thousand years. By Doomsday reasoning, the probability of having a current birth rank is therefore lower than it would have been were this crucial event not to happen in A.D. 2150. This is true reverse causality, not some correlation artifact. One can perhaps see this most easily in terms of a counterfactual formulation of causality; a counterfactual alteration to Doomsday, for any reason whatever, is allegedly reflected in *current* probabilities. It is probably this tacit but largely unacknowledged reverse causation that is at the bottom of all the thrashing about over determinism. Unless the determinate numerosity of future generations somehow exerts the right influence on one's birth rank selection the argument is endangered; indeterminism threatens to render fuzzy and indistinct the needed retroactive future causes. Such retroactive influence is not inconceivable; however, its use further compromises the already tattered plausibility of the Doomsday argument.

We summarize the major parallels between Doomsday reasoning and the fallacious outcome based approach to the Shooting Room. Both are claimed to depend on determinism, yet in either case what is truly needed is not a determined future, but a retroactively influential one. In both the assumption of the user's randomness is crucial and in both it has the consequence that the "game" is likely to be over sooner than otherwise expected. In the Shooting Room this kind of thinking can be made to yield absurd and contradictory results; in the case of the Doomsday argument, the elusiveness and ambiguity of human destiny help to conceal the argument's invalidity.

6) Random Identity

Examples such as the Consecutive Token Machine illustrate the impossibility of choosing a random rank without knowledge of how many ranks exist overall (or without being able to interact with the population as in drawing from an urn). I believe the Doomsdayer response to this would be that we may be unable to *select* random humans, but this does not prevent us from *being* random humans.

Doomsdayers precede rather freely from the an item's membership in a class to its randomness in that class. However, the mere fact of being a likely candidate for random member in one class often precludes a similar status in some other class. A random human is surely an exceptional vertebrate; etc. Leslie proposes that a single human be taken as a random sample from the class of all humans *and* from the class of all mammals [Leslie (1993 p. 491]. However, a random drawing from the class of all mammals has only a minuscule probability of producing a human being; in fact, to have drawn a human from this class is excellent evidence that the drawing was *not* random. The instructions to select something that is both a random human and a random mammal are incoherent; if the item is to be a random mammal, we cannot also stipulate that it be human.

Leslie's expanded prescription [Leslie (1996) pp. 256-263] that a selected item count as random in diverse classes to which it belongs leads to incorrect results; this can be readily seen by enlisting the former in an attempt to beat the Consecutive Token Machine. Suppose the sampler always selects a token from among the first fifty according to the principle that such a token is random in both smaller and larger pools. These fail as random choices for cases of late doom since, for that subset all selected tokens are numbered 50 or lower, whereas randomly this should be true of only about half of them.

Whether or not we are likely candidates for the status of random human, we cannot be counted random mammals, random vertebrates or random solid objects; we are too atypical as members of the last three classes.

Even though randomness in a class cannot be considered an automatic consequence of membership, there remains the possibility that randomness might arise from some ineffable element of human consciousness or human identity. Perhaps the best way to undermine this is to examine how the allegedly random user of the Doomsday argument is elected. In order to derive significance from the value of the birth rank selected, we need to know something of what statisticians would call the *sampling density* of the selection procedure; the only conceivable way of accomplishing this is to examine how the sample is drawn. This is so obvious as scarcely to need mentioning, but the Doomsdayers seem reluctant to take this step. In spite of their silence on what triggers selection, it seems clear that such sampling is undertaken only after the Doomsday argument, or its central concept that we are randomly selected humans is discovered; in other words, the sampling density is that of waiting for a discovery. (Just as the randomness of selection from among all mammals is compromised by the precondition that the mammal be human, so the randomness of selection from among all humans who ever live is compromised by the precondition that the human be in possession of Doomsday reasoning.) If the mean time until discovery of the Doomsday argument is significantly shorter than the mean time until extinction of the race, then we would expect Doomsday argument self-samples to be *early* members of the complete human population. In this case the relative lowness of our birth-rank loses its implication of early doom. To lend support to Doomsday reasoning we need a contrary assumption that the average time it takes to discover the argument is long compared to the average time it takes for the race to become extinct. The fact that the Doomsday argument was independently discovered at least three times in recent years certainly detracts from the plausibility of this last assumption, but without it the strongest conclusion we can hope to draw is that we are likely to be of average rank among potential first discovers of the Doomsday argument, which themselves may be biased toward earlier segments of humanity's complete history. In this case discovery of the argument and the resultant sample rank imply nothing about the likelihood of early doom. In fact the only conclusion about the timing of Doomsday we can validly draw from our birth rank is the trivial one that Doomsday does not occur until after the selection [Eckhardt, 1993, p. 487].

We may have a variety of reasons to fear that the human race is doomed, but lowness of our own birth ranks cannot be counted among them.

Appendix: A general formula for Doomsday arguments.

For greatest generality, we break the Doomsday A_n scenarios into the finest possible alternatives by interpreting A_n to mean "doom will occur when the total cumulative human populations reaches n." (We then know $P(A_k) = 0$ for all k corresponding to people already born, since we know doom has not befallen us yet.) These A_n's can then be grouped and approximated, but it is an essential feature that "time left until doom" be reckoned though additions to the population; the HR assumption imputes randomness to individual, not to time periods.

We let R stand for the fact that one's birth rank is r. By Bayes' Theorem

$$P(A_n|R) = \frac{P(R|A_n)P(A_n)}{\sum_{j=0}^{j=\infty} P(R|A_j)}$$

The crucial probability in this expression is $P(R|A_n)$, the probability of having birth rank r *given* that doom will occur at cumulative population n. If, quite reasonably, one supposes that the probability of having birth rank r remains the same as one runs through various *later* doom scenarios, the argument stops in its tracks; for then $P(A_n|R)$ is just $P(A_n)$ and the prior probabilities are left unrevised. To secure Doomsdayer conclusions, we need the HR assumption from which it follows that $P(R|A_n) = \frac{1}{n}$ and that for any n,

$$P(A_n|R) = \frac{\frac{1}{n}P(A_n)}{\sum_{j=0}^{j=\infty}\frac{1}{j}P(A_j)}.$$

Prior probabilities $\{P(A_n)\}_n$ shift to $\left\{\frac{1}{n}P(A_n) \Big/ \sum_j \frac{1}{j}P(A_j)\right\}_n$ under the impact of the HR assumption.

Suppose A_x and A_y are two doom scenarios with nonzero prior probability and such that upon application of the transformation A_x gains in likelihood and A_y loses; then

$$\frac{\frac{1}{x}P(A_x)}{N} - P(A_x) > 0 > \frac{\frac{1}{y}P(A_y)}{N} - P(A_y)$$

where $N = \sum_j \frac{1}{j}P(A_j)$. Since $P(A_x), P(A_y)$ are positive $\frac{1}{xN} - 1 > 0 > \frac{1}{yN} - 1$

Then $\frac{1}{xN} > \frac{1}{yN}$ and $x < y$. Therefore, it is always earlier scenarios that gain likelihood from later scenarios.

To undermine the claim that the Doomsday argument can only revise pre-existing risk assessments, we employ a “prior” that renounces all such preconceptions in a thorough manner. (At this point, but only at this point, do we invoke techniques that are Bayesian in the controversial sense.) A *noninformative prior* can express *complete ignorance* as to e.g., the size of something [Berger (1985) pp. 82-90]. As applied to the magnitude of the final cumulative human population such complete ignorance encapsulates not only freedom from risk assessments, but a thoroughgoing lack of hope, fear, prejudice or opinion of any kind as to the size of this population (except for the belief that it will be finite.)

Let r be the rank of an infant born now. For $n > r$ assign weight $w(A_n) = \frac{1}{n}$ to A_n.

(These weights are not probabilities since they necessarily sum to infinity.) Applying the above transformation to these weights we obtain,

$$P(A_n) = \frac{\frac{1}{n^2}}{\sum_{j=r}^{j=\infty} \frac{1}{j^2}}$$

Unlike the original weights, these are true probabilities; note that the probability of doom scenario A_n depends exclusively on r and n.

REFERENCES

Berger, James O. (1985) *Statistical Decision Theory and Bayesian Analysis*, Springer, New York.

Eckhardt, W. (1993) "Probability Theory and the Doomsday Argument", *Mind, 102*, 407, pp. 483-488.

Epstein, Richard A. (1977) *The Theory of Gambling and Statistical Logic*, Academic Press, New York.

Feller, William (1968) *An Introduction to Probability Theory and Its Applications*, Vol. 1, Wiley, New York.

Gott, J.R. (1993) "Implications of the Copernican Principle for Our Future Prospects', *Nature, 363*, pp. 315-319.

Leslie, John (1990) "Is the End of the World Nigh?", *The Philosophical Quarterly, 40*, 158, pp. 65-72.

____ (1992) "Time and the Anthropic Principle", *Mind, 101*, 403, pp. 521-540.

____ (1993) "Doom and Probabilities", *Mind, 102*, 407, pp. 489-491.

____ (1996) *The End of the World*, Routledge, New York.

Nielsen, H.B. (1989) "Random Dynamics and Relations Between the Number of Fermion Generations and the Fine Structure Constants" *Acta Physica Polonica, B20*, 5, pp. 427-468.

THE FOLLOWING WAS CIRCLED, BUT NOT INSERTED

Doomsday reasoners predict the equipability of *arg*-sampling on the inchoate idea that one could have been any of the members of a given reference class, such as all humans who ever live, or all users of radio in the history of the universe, etc. They ignore the question of what it is that occasions self-selection But what can it be other than the discovery of anthropic or Doomsday reasoning? This constitutes the sole explanation as to why some of our contemporaries consider themselves to be so selected whereas previous generations did not.

Gott (1993, p. 317) casts our possession of the Doomsday argument in a different role. He questions how, if human civilization is to continue for a long time, we are so atypical as to be extraordinarily early among all who possess Doomsday reasoning.

If a random discovery is selected from a lengthy time interval, then one is *not* likely to be one of the earliest possessors of *that* discovery. However some discoveries are likely to take place during one's lifetime, and one is quite likely to be among the earliest possessors of those discoveries. On expects to be among the earliest possessors of any *recent* discovery.

For those to whom Bayes' Theorem remains irredeemably opaque, the power of the HR assumption to lower expectations of human survival can be illustrated in terms of the familiar procedure of sampling to determine an average value. The average of the numbers 1 through d is $\frac{1}{2}d$. Via the HR assumption the Doomsday argument provides us with a single random rank. It may seem statistically flimsy to reason from a single datum, but under these assumptions it is the only such datum available and it would be a mistake not to derive as much from this unique opportunity as possible. The sample average is then *r* (a single datum, averages to itself) and it provides an estimate of the population average ½*d*, in other words, the best estimated of *d* is *2r*. At current population growth it would require only a few generations to reach a cumulative population figure of *2r*, or so the argument goes. Of course, without the HR assumption, there is no absolutely reason to believe that *r* is an estimate of *½d.*

The issue can be clearly framed: what difference should the knowledge of our own birth rank make to our estimates of the cumulative size of future populations?

(There is no tendency for a single equiprobable sample to lie near the middle; a larger sample tends to <u>average</u> nearer the middle, but that is another matter.)

Appendix: The obstacle in the Doomsday argument is not that of sampling randomly from a population that is indeterminate but rather that of sampling randomly from a process of unknown termination. We illustrate this with an urnlike example, but one that does not beg the central question of the HR assumption in that randomness of sampling is the issue in doubt, just as in the Doomsday argument.

Consider a machine equipped with a large store of balls that, when started; expels either 50 balls (early doom) or 100 balls (late doom) at the rate of one per minute according to a procedure outlined below. The rank of a ball is a number from 1 to 50 or to 100 that describes the ball's place in its series. In each trial, the player's task is to select a ball <u>as it is expelled</u> that is of random rank for that trial. We consider two cases.

i) On the nth trial, the machine preloads into an antechamber either 50 or 100 balls based on whether the nth digit of $\sqrt{23}$ is even or odd. Thus, it is not only a definite fact before play begins how many balls are to be expelled, but this fact is decided by a purely deterministic procedure. Suppose the player knows all this but does not know the particular number, $\sqrt{23}$, being used, so the player has no way of knowing on a given trial whether there are to be 50 or 100 balls. These conditions are allegedly favorable to Doomsday reasoning.

A moment's reflection should show that the task set for the player is impossible. In a long series of trials, the player would have to contrive to spread the selections from 1 to 50 in cases for which the antechamber begins with 50 and from 1 to 100 in the other cases, without knowing which were which. Unless the Player were psychic or extraordinarily lucky, he or she would succumb to the tendency to select too late with the series of 50's or too early with the series of 100's or both. A procedure that would secure that kind of randomness is inconceivable. The insurmountable obstacle is precisely that of sampling equiprobably from a sequence of unknown termination. With knowledge of the termination, on the other hand, the task becomes easy.

ii) In this case, only after the 50th ball has been expelled does the machine decide whether or not it is going to expel 50 more. This decision is made with even chances on the basis of an indeterministic quantum device. Whether the population of balls is ultimately to be 50 or 100 is indeterminate as the trial begins. The player would find this second version exactly as difficult as the first; in either case, he or she would be equally ignorant of the termination point. When we make a prediction, as in the Doomsday argument, uncertainty owing to ignorance is as detrimental as equal amounts of uncertainty owing to indeterminism. Would you, e.g., play a hand of bridge differently if the deck has been randomized by consulting a quantum device than if it had been randomized by consulting the decimal expansion of $\sqrt{23}$?

We are in the same position with regard to Doomsday as the player in either of these cases; we cannot randomly sample from the collection of all who ever live without knowledge of the collection's termination. It is small wonder that making the gratuitous assumption that we have a truly random rank has grave consequences for the nearness of Doomsday.

What are, I believe, downright errors in Leslie's treatment of the more explicit Shooting Room story closely parallel highly problematic features of his treatment of the Doomsday Argument. It is allegedly determinism that unleashes the full force of Doomsday reasoning but also the incorrect conclusions that the D-series games are unfavorable and should be declined. More importantly, the error the declining D-series player makes is that of assuming that he or she is a *typical* or *random* player (and hence ought to decline since typical players lose in this story). In fact, the player has no good reason to believe he or she is "typical" in this sense except in the event that the player actually loses. To emphasize the analogy, suppose before playing, a player learns his or her rank, i.e., how many have played before. This may tell the player how many previous rounds have gone by without rolling a double six, but surely that does not affect the chances for a double six this round. The player ought to draw the conclusion that, no matter how high his or her rank is, the final player population has 35/36 probability of being more than ten times larger. Accordingly, the player ought to bet no matter how high the learned rank. The fact that at least 90% of the population of players will lose is relevant prospectively only if the player knows, in addition to his or her own rank, *the final pool size*, in which case being beyond the tenth percentile is excellent evidence the player is going to lose the favorable bet.