

A Bayesian Analysis of the Doomsday Argument

Jonathan Oliver and Kevin Korb
jono@ultimode.com, korb@cs.monash.edu.au
Department of Computer Science,
Monash University
Clayton, Victoria 3168 Australia

January 27, 1998

Abstract

The doomsday argument purports to show that starting solely from the information of what your birth rank is amongst some population the extinction of that population is near at hand, or at any rate significantly nearer than you had previously supposed. Various authors have suggested that this is a straightforward consequence of Bayesian reasoning. We show that Bayesian reasoning is both more subtle and more plausible than the doomsdayers' understanding of it.

1 Introduction

The doomsday argument (DA) [8, 6, 9] claims that starting solely from the information of what your birth rank¹ is we should prefer the hypothesis that the extinction of humanity is near rather than remote. Intuitively, the argument runs: an inductive inference to a hypothesis should in general make the evidence more probable than do alternative hypotheses; what would make your birth rank more probable than otherwise is the hypothesis that your birth is typical rather than atypical; if homo sapiens endures for millions of years into the future, you would be an atypically early member of the species; therefore, homo sapiens will go extinct relatively soon.

In the form in which the argument has been presented in the literature the doomsday argument requires the following two assumptions:

- we consider *two* competing hypothesis [8, p. 86] [3, p. 78], namely **Quick Doom** (QD) and **Doom Deferred** (DD).

¹A person's birth rank is the position in the birth order, numbered consecutively of all people who will ever be.

- we will make inferences under the *Human Randomness (HR) assumption* [4, p. 248]:

“We can validly consider our birth rank as generated by random or equiprobable sampling from the collection of all persons who ever live.”

Let N be the total number of people who will ever be, and let R be our birth rank. We find that the application of Bayes’ theorem:

$$P(H_{\text{QD}}|R = r) = \frac{P(R = r|H_{\text{QD}})P(H_{\text{QD}})}{P(R = r|H_{\text{QD}})P(H_{\text{QD}}) + P(R = r|H_{\text{DD}})P(H_{\text{DD}})}$$

in combination with the *human randomness (HR) assumption* leads us to prefer the Quick Doom hypothesis over the Doom Deferred hypothesis. In fact, if R is 60 billion, and we are initially indifferent between H_{QD} and H_{DD} , and if we interpret the latter in precise form as the assertion that homo sapiens will go extinct after 100 billion and 1,000 billion humans respectively, then the posterior probability of Quick Doom is the quite threatening value of 0.9091.

Since the argument proposes making a significant change in our assessment of our collective life span on what appears to be, at best, very thin evidence, it is natural to look for assumptions behind the doomsday argument that are vulnerable. We believe that both the identified assumptions are open to serious doubt.

2 A Statistical Analysis of Doomsday

First, we show that the restriction that previous authors have imposed on the doomsday argument by limiting consideration to two hypotheses — Quick Doom and Doom Deferred — is a significant mistake. Whereas those authors have taken this as a simplifying assumption that helps to set out the inferential matters at stake clearly (e.g., [8, p. 86] [3, p. 78]), we believe instead that it *distorts* the problem, leading to serious misunderstandings of it.

A more realistic set of competing hypotheses than simply H_{QD} and H_{DD} would be:

- $H_1 : N = 1$
- $H_2 : N = 2$
- ...
- $H_U : N = U$

where H_i asserts that the population size of homo sapiens will be i and U is some upper limit on the number of people, perhaps based upon subjective beliefs, or perhaps based on considerations such as the estimated mass of the universe,

etc. The question is whether adopting the more realistic set of hypotheses as a priori tenable makes any difference to the conclusion. The answer is that it does, and the difference is enormous.

2.1 Classical Statistical Estimation

This problem is known to statistics as the *serial number problem* [10, p. 33]. In World War II, the Allies considered the problem of estimating enemy production (in particular, that of tanks) by serial number analysis [10, p. 33] [5, p. 212]. If they had sighted T tanks uniformly sampled (without replacement²) from $1 \dots N$ and R was the maximum serial number (i.e., the highest serial number sighted), then the maximum likelihood estimate of N would be

$$\hat{N}_{\text{MaxLik}} = R$$

An unbiased estimate of N would be

$$\hat{N}_{\text{Unbiased}} = \frac{T+1}{T}R - 1$$

In the doomsday problem we have only sampled one point and hence (at the time of observation)

$$\hat{N}_{\text{MaxLik}} = R, \quad \hat{N}_{\text{Unbiased}} = 2R - 1$$

These estimators were used to good effect by the Allies, in large part because the sample sizes attained were not small. The estimates produced were consistently better than those of a more traditional source, spies. In the meantime, we are told, Allied manufacturers during World War II deliberately left random gaps in their serial numbers.

2.1.1 A Classical Analysis

If the doomsday proponent is prepared to use a classical unbiased estimator, then on the human randomness (HR) assumption, the doomsday scenario is endorsed: If $R = 60 \times 10^9$ (sixty billion) (a possible figure for one of us), then $\hat{N}_{\text{Unbiased}} = 2R - 1 \approx 120 \times 10^9$, which is not so far away. This is not an entirely unreasonable argument, although as with all classical statistics, it totally ignores the prior probability of an early doom. The main difficulty here is that such classical estimators display very high variability with small sample sizes — and here we have a sample size of *one*. In such a case, the standard error of the estimate cannot even be calculated using classical methods. Most proponents of doom, however, have suggested that they are thinking of a specifically *Bayesian* argument for doom. We now assess the Bayesian inference to early extinction.

²We assume that the serial numbers are sampled without replacement since the serial numbers were probably obtained by destruction or capture of the tank.

2.2 Bayesian Estimation

The Bayesian approach to the serial number problem requires us to find the posterior probability distribution over N given R . To do this, we first need a prior probability distribution over N . Following other commentators, we shall adopt a uniform prior distribution over N in the range $[1, U]$. Although this is not a realistic assumption, it will allow us to assess readily whether or not the dramatic shift in the probability distribution, claimed by doomsdayers, is in the offing. Hence,

$$P(N = n) = \frac{1}{U} \text{ for } N = 1 \dots U$$

The likelihood of selecting a particular person $R = r$ given a particular total number of people $N = n$ under the HR assumption is:

$$P(R = r|N = n) = \frac{1}{n} \quad \text{if } r \leq n, \text{ 0 otherwise}$$

We may then calculate the probability of selecting a particular person $R = r$:

$$\begin{aligned} P(R = r) &= \sum_{n=1}^U P(N = n)P(R = r|N = n) \\ &= \frac{1}{U} \sum_{n=r}^U \frac{1}{n} \end{aligned}$$

If U is large, then we may approximate the sum by an integral:

$$\begin{aligned} P(R = r) &\approx \frac{1}{U} \int_r^U \frac{1}{n} dn \\ &\approx \frac{1}{U} (\log_e U - \log_e r) \end{aligned}$$

Given the sample information $R = r$, we now calculate the posterior probability for each hypothesis $P(N = n|R = r)$. We expand the posterior probability using Bayes Theorem:

$$\begin{aligned} P(N = n|R = r) &= P(R = r|N = n) \frac{P(N = n)}{P(R = r)} \\ &\approx \frac{1}{n} \frac{U}{U} \frac{1}{(\log_e U - \log_e r)} \\ &\approx \frac{1}{n(\log_e U - \log_e r)} \end{aligned}$$

In order to estimate of the total number of people who will ever exist we may calculate the expected value of N given the experimental data.

$$E(N|R = r) = \sum_{n=1}^U nP(N = n|R = r)$$

$$\begin{aligned} &\approx \sum_{n=r}^U n \frac{1}{n(\log_e U - \log_e r)} \\ &\approx \frac{U - r}{(\log_e U - \log_e r)} \end{aligned}$$

The expected value is a reasonable Bayesian point estimate of the size of the population, one which summarizes the entire probability distribution³.

U	$E(N)$	$E(N R = 60 \times 10^9)$
100×10^9	50×10^9	78.3×10^9
$1,000 \times 10^9$	500×10^9	334×10^9
$10,000 \times 10^9$	$5,000 \times 10^9$	$1,942 \times 10^9$
$100,000 \times 10^9$	$50,000 \times 10^9$	$13,471 \times 10^9$

Table 1: Expected Population Sizes

Letting $R = 60 \times 10^9$ as before, then Table 1 gives the prior expected value $E(N)$ and the posterior expected value $E(N|R = 60 \times 10^9)$ for various values of U . Graphically, Figures 1 and 2 illustrate the prior and posterior probability distributions in the doomsday argument when R is 60 billion and U is 100 billion and 1,000 billion, respectively. Although the posterior distributions in both cases peak at 60 billion, and are above the prior probability for 60 billion, it is important to note that the actual posterior probability at this peak is low, at around 0.033 and 0.0059, respectively — compared to the alarmist 0.9091 we computed for the first version of doomsday reasoning.

³“Fundamentalist” Bayesians reject attempts to summarize a posterior density by any point estimate as being unsound. We, however, consider it a reasonable practice in many circumstances.

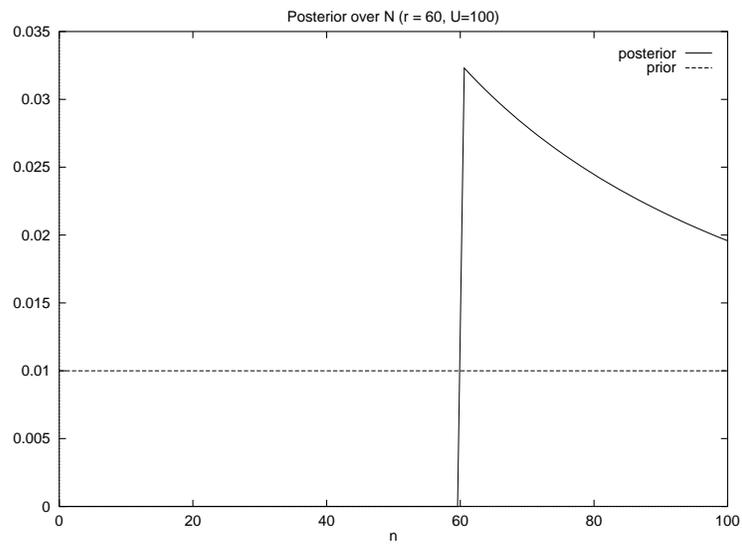


Figure 1: DA Probabilities for R=60 and U=100

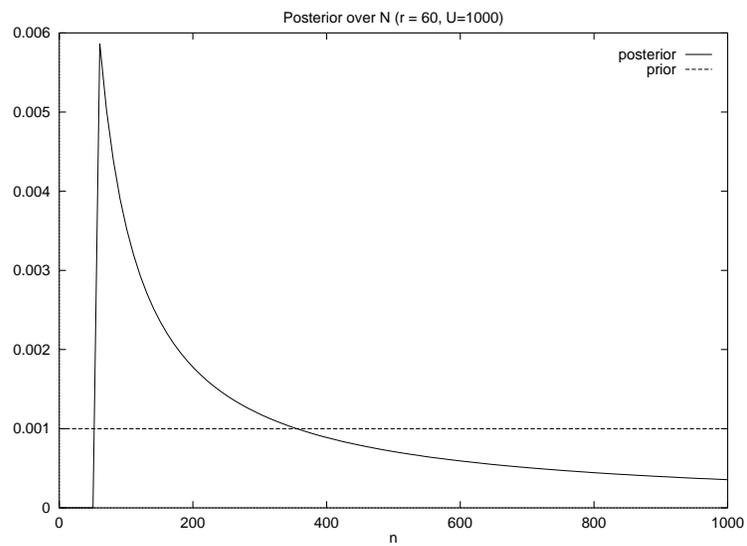


Figure 2: DA Probabilities for R=60 and U=1000

More important than looking at the peaks of the posterior distributions however, is considering the whole. Figure 3 gives the prior and posterior expected population size for $R = 60$ billion and varying U from 60 billion to 1,000 billion. We see that the expected population size changes only slightly between prior and posterior distribution.

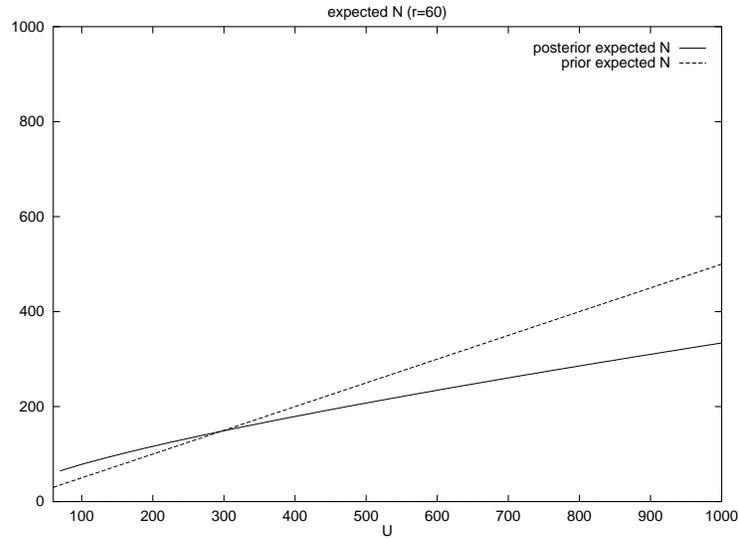


Figure 3: Prior and Posterior Expectation ($R=60$)

2.2.1 A Bayesian Analysis

From these computational results we can draw the following conclusions:

- If you believe prior to learning the current birth rank of some species that the species will be long lived, then learning that the current birth rank is “small” will only moderately change your beliefs about the expected life span of the species. (In all cases above, the change was less than an order of magnitude.)
- If you believe prior to learning the current birth rank of some species that the species will be short lived, then learning that the current birth rank is “small” moderately confirms your prior beliefs about the expected lifespan of the species.

In general, learning the relevant birth rank only has a moderate effect on one’s expectation — given that the hypothesis space contains some plausible range of alternative hypotheses, rather than the doomsdayers’ preferred forced choice between two alternatives.

Other authors have made much stronger claims on the basis of the doomsday argument. For example, Leslie [8, pp. 86-7] claims that

“[the doomsday calculations] suggest rather strongly that the risk of Quick Doom is usually under-estimated and that we at least have little ground for thinking that a long-lasting human race is probable.”

In a similar vein, Bostrom [1] claims that

“by Bayes’ theorem, you should update your beliefs about mankind’s prospects and realise that an impending doomsday is much more probable than you have hitherto thought.”

However, Table 1 and Figure 3 show that if there has been an underestimation of Quick Doom, it has been a minor one. If we believed prior to knowing anyone’s birth rank that it was equiprobable that there would be between one and one hundred thousand billion people, and we were informed that our birth rank was 60 billion, then we should still expect the species to survive into the tens of thousands of billions. Since the doomsday argument cannot provide guidance about which hypothesis space (i.e., which value for U) to adopt in the first place, it does not suggest “rather strongly” that the risk of Quick Doom is high. Indeed, without specific argument constraining the size of our hypothesis space, Figure 3 shows that we can expand the expected size of the population arbitrarily, simply by adopting larger hypothesis spaces.

Again, Eckhardt [4, p. 248] asserts (in an article attempting to rebut the doomsday argument):

“When correct allowance is made for the randomness assumption, doomsday reasoning has an unalterable bias towards earlier doom. 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.”

Furthermore, Eckhardt produces a mathematical proof (see his footnote 8): that under the evidential circumstances of the doomsday argument, if $P(N = n | R = r) > P(N = n)$ and $P(N = n' | R = r) < P(N = n')$, then $n < n'$. This result would appear to endorse the doomsdayers’ theme. However, Eckhardt’s *interpretation* of his proof goes too far. Whereas the posterior probabilities of individual (unrefuted) hypotheses nearer to the sample value necessarily increase (simply as a result of their higher likelihoods on the evidence), the *expected*

value of the posterior probability distribution is by no means guaranteed to shrink. Indeed, Figure 3 illustrates over 200 billion counterexamples to that interpretation: every case where U is less than the cross-over point between prior and posterior expectation (around 300 billion) is one where the posterior expected population size is greater than the prior expectation.

In short, the proponents of early doom, in simplifying the doomsday problem to employ only two hypotheses, have simplified the relevant probabilistic calculations to the point where they no longer tell us anything useful about the real world. Examining more realistic hypothesis spaces reveals that, while the probability shift they have identified is real enough, it is so minor an effect that it must be swamped by our uncertainties about what prior hypotheses to consider.

3 Human Randomness

Although we believe the above probabilistic considerations take all of the sting out of the doomsday argument, the second major assumption of that argument can also be questioned — and indeed it has been questioned by both Eckhardt [4] and us [7]. First, it is clear that there is in fact no physically uniform selection process yielding the rank evidence at hand. *Birth* is a causal process and might be construed as a selection process. However, it is an elementary consequence of evolutionary theory and the facts at hand that the next human birth does *not* have an equal probability of reproducing Ishi, the last native Californian, as did Ishi’s actual birth, contrary to HR. What is a more plausible interpretation of the doomsdayer’s position is that whatever *non-uniform* selection process yields our rank, we have no specific knowledge that it is predisposed to produce for us a birth rank that is either early or late, so on a principle of indifference we should treat it as uniformly random. The doomsday argument thus rests upon an epistemic direct inference: the physical probability of our selection for our particular birth ranks is unknown, but, lacking contradictory or biasing evidence of early or late selection, our subjective probability of our particular birth ranks should be treated as uniform.

The fact that HR depends upon epistemic direct inference raises some previously unconsidered difficulties for the argument. In his treatment of the doomsday argument, John Leslie repeatedly asserts that the point of the doomsday argument is not specifically that Doom is near at hand, but rather that whatever the probability of Quick Doom was prior to considering the argument, that probability is substantially greater *after* considering the argument [9, p. 204 and p. 213]. The point of the argument is to induce a probability *shift* in favor of quick extinction, a shift which operates regardless of the available evidence concerning the lifetime of the population (and hence is a priori). In order to support the *further* claim, not merely that Doom is nearer than previously expected, but also that Doom is in fact near to hand, Leslie also brings forth numerous

a posteriori arguments in support of Quick Doom, such as our problems with military technology, overpopulation, and environmental pollution. However, it is generally accepted that an epistemic direct inference is defeated if biasing prior knowledge is available (see, for example, [11]). Unfortunately for Leslie's version of the doomsday argument, all his empirical evidence of a near-term Doom tends to support the assertion that our birth ranks are biased towards being *late* in the total ordering of births. In addition, other arguments might be advanced to the effect that we are *early* in the likely lifespan of homo sapiens — including such things as our growing understanding of the risks of extinction, the prospects of space colonization, etc. The plausible result is that we know of biasing factors in the selection of our birth ranks, but we do not know what the cumulative effects are of these biases. Hence, these considerations of Leslie's, and the others, tend primarily just to show the inapplicability of the principle of indifference operative in epistemic direct inference.

Setting selection biases aside, there is an additional difficulty, or at least oddity, with the use of HR in the doomsday argument. As Eckhardt points out [4, p. 257], a physically random uniform selection process over one class (say, humans), *cannot* simultaneously be a physically random uniform selection process over a distinct class (say, mammals), since the probability of uniform selections in two distinct classes cannot both sum to one.⁴ Nor is this point restricted to physical, as opposed to epistemic, uniform selection. Still, John Leslie's version of the doomsday argument appears to avoid any trouble here. Leslie emphasizes the first-person perspective in his version of the Argument. After all, the doomsday argument historically developed out of Brandon Carter's anthropic principle [2], and much of the point there was that reflexive observations about ourselves observing the universe have interesting inferential implications. The DA came from an attempt to apply such reflexive considerations to our collective lifespan. But, the only beings we know who can apply reflexive considerations of any kind are just homo sapiens. Hence, the object of such considerations must be all and only homo sapiens. We believe such a response to be mere evasion, for two reasons: first, we have no good reason to believe that homo sapiens exhausts the class of sapient (indeed, there is at least good reason to believe that other homos were sapient despite our implicit denial in naming conventions); second, despite its origins, the DA lives or dies by the (probabilistic) logic it employs, and that logic has nothing to do with who is doing the arguing. The first renders obscure both what population is being argued about and also what birth rank is correct; the second throws the issue of the right reference class wide open, since the argument may be applied to any number of populations.

Gott argues [6] that all reference classes are right — that our birth rank relative to the distinct populations in which we participate can be used to make inferences about the lifespans of all those diverse populations. Since the birth

⁴We ignore cases where the difference set has measure zero, since we are talking about classes of finite size.

ranks we have are relative to each population, it is possible that each is (relatively) uniformly randomly selected, even by a single selection process. Applying DA to these diverse populations, in each case we are being asked to view ourselves as typical members of the population with regard to birth rank: typical homo sapiens, typical homo, typical mammal, typical vertebrate, typical organism, typical thing. Such a view has the decidedly atypical consequence that recently evolved populations shall die out soon and older populations will endure: in other words, that evolution is symmetrical with devolution, with homo sapiens fading into the night, followed by homo (who, then, would be left?), followed by mammals generally, and eventually with the same kind of primordial ooze lying about out of which we all originally came. One might consider that such an outlandish story illustrates what happens when we push too hard to have our theories fit the data: by rendering every data point (birth rank) as likely as possible, we have rendered our theory as a priori unlikely as possible.

4 Conclusion

The doomsday argument is, we believe, an example of probabilistic reasoning gone astray. The preconditions of proper Bayesian reasoning are not simply that one have the probabilities that are explicitly asked for in Bayes' theorem. In addition, one must satisfy at least the additional conditions of:

- Having a reasonable set of hypotheses under consideration. Having the wrong set almost necessarily distorts the prior probability distribution and, so, the posterior probability distribution.
- Satisfying the total evidence condition. If the evidence used in Bayes' theorem does not report all of the relevant evidence — for example, by omitting biasing selection factors which are *known* to be relevant, then operating Bayes' theorem will produce posterior probabilities which can deviate arbitrarily from the correct values.

The doomsdayers have neglected these preconditions for proper Bayesian reasoning and have suffered the consequence of promoting exaggerated and unsupported claims about the impact of a single observation of one's birth rank.

Acknowledgments

We would like to thank Patrick Wilken for introducing us to the debate, Alan Croke for some valuable discussions, and to Malcolm Clarke for giving us references to the classical statistics literature.

References

- [1] N. Bostrom. The perplexing Carter-Leslie Doomsday Argument. Web page at <http://www.hedweb.com/nickb/doomsday/doomsday.htm>, Dept of Philosophy, London School of Economics, 1997.
- [2] B. Carter. The anthropic principle and its implications for biological evolution. *Philosophical Transactions of the Royal Society of London A*, 310:346–363, 1983.
- [3] D. Dieks. Doomsday – or: The dangers of statistics. *The Philosophical Quarterly*, 42:78–84, 1992.
- [4] W. Eckhardt. A shooting-room view of doomsday. *The Journal of Philosophy*, 59:244–259, 1997.
- [5] W. Feller. *An introduction to probability theory and its applications*. Wiley, New York, 2nd edition, 1957.
- [6] J.R. Gott III. Implications of the Copernican principle for our future prospects. *Nature*, 363:315–319, 1993.
- [7] K.B. Korb and J.J. Oliver. A refutation of the doomsday argument. Forthcoming in *Mind*, 1998.
- [8] J. Leslie. Doomsday revisited. *The Philosophical Quarterly*, 42:85–89, 1992.
- [9] J. Leslie. *The end of the world*. Routledge, New York, 1996.
- [10] G.E. Noether. *Introduction to statistics: the nonparametric way*. Springer-Verlag, New York, 1991.
- [11] W.C. Salmon. *Statistical explanation and statistical relevance*. University of Pittsburgh, 1971.