

# A DILEMMA FOR THE DOOMSDAY ARGUMENT

Robert Northcott

## Abstract

I present a new case in which the Doomsday Argument ('DA') runs afoul of epistemic intuition much more strongly than before. This leads to a dilemma: in the new case either DA is committed to unacceptable counterintuitiveness and belief in miracles, or else it is irrelevant. I then explore under what conditions DA can escape this dilemma. The discussion turns on several issues that have not been much emphasised in previous work on DA: a concern that I label *trumping*; the degree of uncertainty about relevant probability estimates; and the exact sequence in which we integrate DA and empirical concerns. I conclude that only given a particular configuration of these factors might DA still be of interest.

## 1. Introduction

The Doomsday Argument was first presented by astrophysicist Brandon Carter more than 30 years ago. Since then it has been elaborated and defended by many philosophers and scientists. It offers a controversial and strikingly pessimistic prediction about the future of humanity based solely on a priori argument.

At the heart of DA's interest is that it gives us reason to modify the predictions that we might derive from empirical considerations alone. Depending on the details, this modification can be counterintuitive – which fact has been used as a criticism. Counterintuitiveness alone, of course, need not derail DA if other considerations tell in its favour. But the verdict of other considerations is much contested, so counterintuitiveness remains potentially decisive.<sup>1</sup> In this paper, I present a new example in which DA's counterintuitiveness is much worse than before.

DA's predictions are probabilistic. In usual presentations, the rival probability derived from empirical considerations alone is not 'sharp'.<sup>2</sup> More formally, if we take the probability value itself to be a random variable that we are seeking to estimate, the probability density function for this variable has a high variance. In other words, we are not very confident about our estimate. As a result, the degree of counterintuitiveness represented by DA's proposed modification is less easily assessed. In this paper's new example, by contrast, the empirical probability is very sharp. It is therefore a much better test case, as it brings matters to a clearer head. In particular, the degree of counterintuitiveness implied by DA modification is laid bare. At the heart of the matter will lie an issue not previously emphasised in the DA literature, which I label *trumping*.

---

<sup>1</sup> For surveys of the lively debate, see Nick Bostrom, *Anthropic Bias: Observation Effects in Science and Philosophy* (New York: Routledge, 2010) and Alasdair Richmond, 'The Doomsday Argument', *Philosophical Books* 47 (2006), pp.129-42. Recent contributions include: Elliott Sober, 'An Empirical Critique of Two Versions of the Doomsday Argument', *Synthese* 135 (2003), pp.415-30; Brian Kierland and Bradley Monton, 'How to Predict Future Duration from Present Age', *Philosophical Quarterly* 56 (2006), pp.16-38; Tom Adams, 'Sorting out the Anti-Doomsday Arguments: a Reply to Sowers', *Mind* 116 (2007), pp.269-73; Nick Bostrom, 'The Doomsday Argument', *Think* 6 (2008), pp.23-8; Alasdair Richmond, 'Doomsday, Bishop Ussher and Simulated Worlds', *Ratio* 21 (2008), pp.201-17; Paul Franceschi, 'A Third Route to the Doomsday Argument', *Journal of Philosophical Research* 34 (2009), pp.263-78; Peter Lewis, 'The Doomsday Argument and the Simulation Argument', *Synthese* 190 (2013), pp.4009-22; and Yann Benétreau-Dupin, 'Blurring out Cosmic Puzzles', *Philosophy of Science* (forthcoming).

This paper remains neutral about other critiques of DA and the replies to them. In effect, it investigates whether, even on the assumption that DA survives these other critiques, still it is vulnerable to a new worry.

<sup>2</sup> Here and throughout, probabilities should be understood as rational credences.

In the new example, DA faces a dilemma: either it indeed insists on modifying the empirically derived probability, at the cost of a counterintuitiveness that now seems unacceptable; or else it can be made consistent with the empirical probability, but only at the cost of thereby rendering itself irrelevant precisely because it no longer implies any modification. Does this dilemma generalise to other cases? When assessing this, what becomes evident is the importance of two further factors: first, the uncertainty mentioned above about the empirical probability estimates – among other things I will argue that such uncertainty is a necessary condition for DA to be of interest. The second further factor is the sequence of analysis – in particular, DA must be interpreted (contrary to what has been the usual recommendation) as setting values for the relevant probabilities prior to any empirical input at all.

Overall, our understanding of the conditions necessary for DA to be of interest is thus clarified. At the end, I will assess DA’s standing in this new light.

## 2. The Doomsday Argument

Each of us today is roughly the 60 billionth human to be born. Here are two possible figures in turn for the total number that ever will be born: 100 billion; and 10 trillion.<sup>3</sup> Which of these is more likely? DA argues that, conditioning on the fact that we are 60 billionth on the ‘birth list’, and given that we have no a priori reason to think we personally are especially likely to be in any particular position on this list, it follows that the 100 billion figure is much more likely than the 10 trillion one.<sup>4</sup> As it were, we are each a ball drawn randomly from an urn; the fact that our number is a lowly 60 billion tells in favour of the total number of balls in the urn being only 100 billion rather than it being the huge 10 trillion. The punchline is: given likely demographic trends, an eventual total of 100 billion implies an eye-catchingly early end for humanity – say, in just a couple of centuries.

Consider the usual futurology factors: nuclear proliferation, asteroid-detection capability, eventual colonization of space, and so on. Taking all these into account might yield particular probabilities for the 100 billion and 10 trillion totals – by this means, all empirical evidence is incorporated. But, its advocates claim, a priori DA reasoning should then lead us subsequently to *modify* these empirically derived probabilities. In particular, we should shift our credence in the direction of early doom. Humanity, watch out!

DA is explicitly supported by a flourishing recent literature on observational selection effects, i.e. on how, in Nick Bostrom’s words, to correct for biases introduced when our evidence has been filtered by the precondition that a suitably positioned observer exists to have the evidence in the first place. Observational selection effects are important and their recent illumination by epistemologists is welcome.<sup>5</sup> In this paper I will, in effect, examine critically the application of that work to DA.

## 3. The Asteroid example

---

<sup>3</sup> The presentation of DA here roughly follows that of John Leslie, *The End of the World* (New York: Routledge, 1996), which has become standard. The exact figures, and the restriction to two possibilities, do not matter; they are useful merely for illustration.

<sup>4</sup> Let  $H_1$  = the 100 billion hypothesis,  $H_2$  = the 10 trillion hypothesis, and  $O$  = the observation that we are 60 billionth in the birth list. Assume that the prior probabilities of  $H_1$  and  $H_2$  are each 0.5. What is the ratio of their posteriors after updating by  $O$ ? Applying Bayes’s Rule:  $\text{pr}(H_1/O) / \text{pr}(H_2/O) = \text{pr}(H_1) \cdot \text{pr}(O / H_1) \cdot \text{pr}(O) / \text{pr}(H_2) \cdot \text{pr}(O / H_2) \cdot \text{pr}(O) = (0.5) \cdot (1/100\text{bn}) / (0.5) \cdot (1/10\text{tn}) = 10\text{tn} / 100\text{bn} = 100$ . That is, initially  $H_1$  and  $H_2$  were level; but after  $O$ ,  $H_1$  is now 100 times more likely.

<sup>5</sup> See, e.g., Bostrom, *Anthropic Bias*, or Darren Bradley, ‘Four Problems about Self-Locating Belief’, *Philosophical Review* 121 (2012), pp.149-77.

Imagine that a large asteroid, previously unobserved, is discovered to be heading dangerously towards Earth, arriving in just a few days' time. In particular, astronomers calculate that it has a 0.5 probability of colliding with us, the uncertainty being due to measurement imprecision regarding its exact path. This probability estimate is based on well-established science and is not disputed by any expert in the field. Suppose, further, that any collision would be certain to destroy all human life, and moreover that there is no feasible way to avoid it.

In the face of this terrifying news, what is now the rational expectation of humanity's future duration? The empirical answer is clear enough – there is a 0.5 probability that humanity will last just the few days until the big impact ('Doom-Now'), and a 0.5 probability that it will last however long it would have otherwise. What does DA say? Either it is deemed to modify these empirical probabilities; or it is not. The *dilemma* is that either choice tells against it.

#### **4. First horn: the trumping problem**

As per the above, the empirical probability of Doom-Now in Asteroid is 0.5, i.e. just the probability of the asteroid indeed striking Earth. As it were, the 'memory' of our prior credence in Doom-Now is wiped out by the asteroid observation. There is no gradual updating; rather, regardless of our starting point, there is an abrupt shift straight to 0.5.<sup>6</sup>

One might think there is nothing controversial here. However, arguably DA requires that our answer be modified, i.e. that we embrace the first horn of the dilemma. In particular, according to DA, a priori considerations show that the expected duration for humanity is much greater than just a few days. The probability of Doom-Now should accordingly be modified *downwards*.<sup>7</sup> That is, DA urges that in the apocalyptic asteroid scenario the probability of humanity ending in a few days is *less* than 0.5, contrary to the astronomers' calculations. (Below, I will return to whether DA should indeed embrace the dilemma's first horn; regardless, it will prove useful to set out now the implications of doing so.)

I submit that, so applied, DA's implications are much more counterintuitive in the asteroid case than in the traditional one. The reason is, as noted, that there is much less uncertainty concerning the empirically derived probability. That is, we are much more confident in the 0.5 estimate of humanity ending in a few days than we were before in the probability estimate of our lasting more than a few centuries. The proposed DA adjustment is accordingly much more clearly implausible. We can say that it is fully *trumped* by the asteroid observations.<sup>8</sup>

Imagine a further update – suppose that after another few hours' fevered observation it was estimated that the asteroid had a 0.75 rather than 0.5 probability of colliding with Earth. The previous empirical estimate would now have been trumped in its turn. Should not DA's proposed modification be trumped similarly? And which modification would we take more seriously – surely the empirically-based pessimistic one, not the DA-based optimistic one?

Further troublesome variations arise. Imagine next *two* large asteroids, one approaching Earth as before plus this time a second approaching Mars. Suppose that each's

---

<sup>6</sup> This 'abrupt shift' is consistent with Bayesian reasoning. Notes: 1) strictly speaking, the shift is to just over 0.5 once we factor in the presumably tiny addition to our credence in Doom-Now from non-asteroid causes; 2) I am assuming plausible background beliefs and that the prior credence in Doom-Now was not 0 or 1.

<sup>7</sup> True, DA reasoning implies that the single most likely total number of humans is the current number, i.e. 60 billion. But although the mode of the distribution is thus 60 billion, the mean is much greater. Thus, Doom-Now is not favoured.

<sup>8</sup> Although I am sure the trumping worry must have occurred to many, the only place I have found it mentioned is in an unpublished paper by Brian Weatherson (<http://brian.weatherson.org/doomsday.pdf>), where it is developed rather differently. Several reviewers of Leslie do raise the broader, but less decisive, issue of how an abstract probabilistic argument could possibly influence concrete causal processes.

collision probability is 0.5. A collision with Mars would presumably have minimal impact on humanity. As a result, according to DA our predictions in the two cases now diverge, for DA modifies our credence that Earth will be impacted but not our credence that Mars will be. So we have two situations which, according to all the physical evidence, have identical collision probabilities – yet if we follow DA our collision predictions for them will be different.

The discomfort can be made yet sharper. Imagine a set-up in which the fate of Earth is made dependent on the outcome of a single indeterministic quantum event. Suppose in particular that this event will yield, say, either an Up or Down result, each with 0.5 probability, and that things are so rigged that Up would trigger Earth's demise while Down would spare us. Therefore, at the moment just before the event, Doom-Now has a probability of 0.5. Nevertheless DA reasoning would still urge us to reduce our credence in Doom-Now, and hence in Up, to less than 0.5, even though that 0.5 value is the direct implication of well-confirmed physical law.<sup>9</sup> Thus our credence here would contradict a law of nature. DA requires us to believe in miracles.

### 5. Second horn: the irrelevance problem

But should DA instead be read as embracing the dilemma's second horn? That is, perhaps it need not be committed to denying that the asteroid has a 0.5 probability of colliding with Earth after all. Here is an argument for this:

Let H = Doom-Now, and O = the observation of the approaching asteroid. Because DA thinks Doom-Now is very unlikely, our prior  $\text{pr}(H)$  is very low. This prior should then be updated in the light of the asteroid observation O:  $\text{pr}(H/O) = 0.5$  (H/asteroid indeed strikes Earth) + 0.5 (H/asteroid misses Earth) =  $0.5(1) + 0.5(\text{prior } \text{pr}(H))$  = just over 0.5. Thus the calculated probability of Doom-Now is now (more or less) what the astronomers endorse. More to the point, DA has been incorporated via the low prior  $\text{pr}(H)$  and yet the calculation remains consistent with assigning a probability of 0.5 to the asteroid striking Earth. So the sharp contradiction with empirical science is avoided.

I return to this argument shortly, concluding that it should not be accepted at face value. But before that, I want to re-emphasise a separate but crucial point: even if this argument is accepted and thus the second horn embraced, that would be no salvation for DA. In particular, the original force of DA was that it compels us, by a priori reasoning, to revise probabilities derived from empirical evidence. If it does not do this then it is of little interest, because purely empirical investigation is confirmed to be sufficient after all. Therefore the successful replication by DA of the empirical probability could only be a pyrrhic victory.

### 6. Which horn to embrace?

It will be useful now to discuss the mathematical situation in more detail. In essence, we have two independent estimates of Doom-Now, one given by empirical considerations and one by DA considerations. How should we combine them? The crucial point is that, in general, we should *not* just assign them equal epistemic weight. Rather, we may justifiably have rather more confidence in one estimate than in the other. This particular issue does not seem to have been discussed in the DA literature explicitly but it is crucial here.

To illustrate, imagine we are estimating the probability that a (possibly biased) coin will come up Heads. Suppose that one estimate carries much less uncertainty than another, perhaps because it is based on a larger sample of previous tosses or because the witness to those tosses is more reliable. Then we should put greater weight on the more certain estimate. Similarly, when estimating the true state of public opinion we should put greater weight on an

---

<sup>9</sup> Leslie, *End of the World* does argue that DA applies more weakly if the universe is indeterministic, which many interpretations of quantum theory imply. Nevertheless he remains committed to *some* non-zero probability shift even then.

opinion poll with a low sampling error than on one with a high error, and greater weight on one with a rigorous methodology over one with a more slipshod one.

It is uncontroversial that Bayesian procedure should take into account asymmetric uncertainty. Indeed, one of its virtues is precisely that it can. In particular, the uncertainty of a probability estimate may be represented by the relevant probability density function. A more certain estimate would have a sharp-peaked distribution around the estimate's value, rendering that value relatively insensitive to new evidence. A less certain estimate, by contrast, would have a more diffuse distribution.<sup>10</sup>

Similarly, in statistical meta-analysis there are standard procedures for combining two independent estimates. One is to consider each estimate's variance. In particular, an unbiased combined estimate (of the mean) can be achieved via *inverse-variance weighting*. Roughly, the higher an estimate's variance, the more uncertain that estimate is, and so the less weight we should put on it.<sup>11</sup>

In the DA case, how we balance competing DA and empirical estimates of a probability turns – and must turn – on exactly this issue. Among other things, it determines which horn of the dilemma an advocate of DA should embrace. Some toy numbers will illustrate. By assumption, the empirical estimate of the asteroid collision's probability, and thus of Doom-Now's, is very certain. Suppose that the density function of that estimate is a normal distribution with a mean of 0.5 and, representing the scientists' high degree of certainty, a small standard deviation of 0.001. Next, suppose initially that for DA the equivalent figures are a mean of 0.001 and the same small standard deviation of 0.001. In this case, because the two variances are the same, so an unbiased estimate of the mean would be midway between the two component estimates of it, i.e. midway between 0.5 and 0.001, i.e. approximately 0.25. So now DA would have modified the empirical estimate of Doom-Now's probability significantly, which corresponds to embracing the first horn of the dilemma.<sup>12</sup>

But is such a high degree of certainty about DA's estimate justified? There are reasons to think not. It is true that DA's Bayesian calculations of the *ratios* of posterior probabilities, such as in footnote 4 earlier, are sharp. But that is less true of those probabilities' *absolute* values. These latter depend on many contested factors and assumptions, such as choice of reference class, the rate of future demographic growth, and the possibility of non-human intelligent species. For this reason, even among its advocates there is a consensus that the DA probabilities for the various Doom scenarios are not known exactly. This is implicit in the reluctance, from Leslie on, to give precise numerical values to DA's predictions. It is implicit too in the universal concession that DA's predictions should be modified by empirical considerations: given that in reality, as opposed to the Asteroid example, the relevant empirical probabilities are far from certain, the fact that they are still deemed weighty enough significantly to modify DA's probabilities implies that the estimates of DA's probabilities are themselves far from certain. For these reasons, we may conclude

---

<sup>10</sup> For details of the Bayesian treatment, see Colin Howson and Peter Urbach, *Scientific Reasoning: The Bayesian Approach* (Chicago: Open Court, 1993), pp.401-3.

<sup>11</sup> Formally, given a sequence of independent observations  $y_i$  with variances  $\sigma_i^2$ , the inverse-variance weighted average is given by:

$$\hat{y} = \frac{\sum_i y_i / \sigma_i^2}{\sum_i 1 / \sigma_i^2}$$

For details, see, e.g., Joachim Hartung, Guido Knapp, and Bimal Sinha, *Statistical Meta-Analysis with Applications* (Hoboken, NJ: Wiley, 2008).

<sup>12</sup> In this scenario, in the calculation in section 5 the coefficient on the term (H/asteroid indeed strikes Earth) would take a value less than 0.5, and the other coefficient a value of greater than 0.5. In other words, the coefficient values would be modified from those implied by the empirical evidence, and thus the argument for embracing the second horn would fail.

that the uncertainty of DA's estimates is surely less than that of the scientists' estimates in Asteroid.

Accordingly, we can calculate a second scenario, with new toy numbers. So, suppose as before that in Asteroid the density function of the empirical estimate is a normal distribution with a mean of 0.5 and a small standard deviation of 0.001. And this time, suppose that for DA the equivalent figures are still a mean of 0.001 but now, say, a standard deviation of 0.1, i.e. a hundredfold larger than before. Now the ratio of variances is therefore 10,000:1, and accordingly in the inverse-variance weighting calculation the empirical estimate is given 10,000 times the weight of the DA one. Running through the numbers, this time the unbiased estimate for the mean is therefore 0.4999501 – in other words, almost indistinguishable from the scientists' estimate of 0.5. This corresponds to DA embracing the second horn of the dilemma, in which it (almost) entirely fails to modify the empirical probability at all.

### **7. From numbers to philosophy**

The above illustrates the mathematics of how DA and empirical probabilities should be combined. However, it is crucial to understand correctly such calculations' philosophical significance – or otherwise. How do they fit into the overall argument?

The main thing to appreciate is that the calculations concern which horn of the dilemma should be embraced *by a DA advocate*. They do not necessarily capture the combined estimate of the probability that *we* should endorse. In fact, the heart of this paper's philosophical argument is precisely that they do not. Rather, the trumping objection to DA applies regardless of the exact variance of DA's estimate. Thus, in Asteroid, the correct probability estimate remains exactly 0.5 in all cases.

Below, I will argue that rationally we should ignore DA's modification of the empirical probability, however big or small that modification is. In other words, no DA modification should be conceded, whether via the inverse-variance weighting rule or any other rule. To act on the preceding section's calculations would therefore be a philosophical mistake, namely the mistake of thinking that the empirical estimates should be modified at all. The true importance of the certainty of the scientists' estimate in Asteroid is not its impact on inverse-variance calculations but rather its role as an intuition pump. That is, the dialectical point of Asteroid is not that the certainty of its empirical estimate is what justifies this estimate's supremacy; rather, it is that its certainty is *what makes intuitively evident* its supremacy. The trumping argument would still hold regardless.

Overall, the trumping argument thus holds regardless of the certainty of *either* the empirical or DA estimates. This is the sense in which the calculations in section 6 are philosophically beside the point.

Nevertheless, finally, there is a potential sting in the tail. For when the incorporation of DA is conceived of in a different way, there can arise a combination of circumstances in which the calculations above do become relevant again after all. At the end (section 10), we will see exactly how that can be, and how the overall dialectical landscape allows for this relevance sometimes even while denying it otherwise.

### **8. Is the dilemma general?**

So Asteroid presents DA with a dilemma, but is this dilemma peculiar only to Asteroid or does it generalise? In particular, does the trumping argument generalise? Although modification of the empirically derived probability seems unacceptably counterintuitive in Asteroid, might it seem rather less so in cases where that probability is less sharp? And in practice DA is indeed usually applied to a case much less sharp than Asteroid – for, given

that no asteroid is actually coming, we are in fact rather uncertain of the empirical probability of humanity coming to extinction at any particular date.

To investigate, begin by tweaking Asteroid into a less sharp version. In particular, suppose that this time scientists are unsure of the probability that the asteroid will strike Earth, in other words that the relevant probability density function has a higher variance. Suppose nevertheless that the mean of this function is still 0.5. The key question is: would DA considerations now, unlike before, shift our rational estimate of  $\text{pr}(H)$  to below 0.5? Because the true probability of Doom-Now has become hazier to us, perhaps the tension between the DA modification and the empirical findings might seem correspondingly less serious.

But it seems to me that, on reflection, to endorse the DA adjustment here would be to go against the scientists' recommendations just as unacceptably as in the original version of Asteroid. By assumption, the density function's mean of 0.5 is empirically well founded. The only disanalogy with the original version is that its variance has increased. However, why should that make it any more rational to depart from its mean? Analogously, suppose we were estimating the probability that a coin will come up Heads, and doing so now on the basis of 10 previous tosses rather than on 100 previous tosses. The uncertainty of our estimate is thus increased, but why does it not remain equally irrational to deviate from whatever estimate the evidence yields us? The trumping problem still applies. Embracing the first horn of the dilemma remains as hard to justify in the less sharp as in the original version of Asteroid.

This conclusion is supported by running through for the new case the thought-examples from the old one. Thus, imagine again identical physical situations with respect to asteroid collisions for Mars and Earth. Any DA modification would again imply that we nonetheless assign divergent collision probabilities, and this still seems unacceptable even when the empirically based estimates of these probabilities are subject to greater uncertainty. And in the quantum example, our mean guess would still be that the fatal outcome has a probability of 0.5. Given that this guess is the result of standard scientific methods supported by relevant background knowledge, any DA-inspired deviation from it would still imply deviation from well-supported empirical findings and ultimately therefore again a belief in miracles.

### **9. The traditional case revisited**

Return now to the traditional presentation of DA ('Traditional'), which focuses on the probabilities of Doom-Soon and Doom-Late rather than on that of Doom-Now. Should we accept that, influenced by DA, our estimates of these can reasonably deviate from those derived from (empirical) futurology alone? And thus, should we accept that there is in this respect a disanalogy between Traditional and Asteroid? If so, then in Traditional the dilemma's first horn, i.e. trumping, would lose its force.

Initially, this escape route for DA might seem tempting. However, I do not think it is viable. The reason is that it relies on a dubious discontinuity claim. We saw above that the transition between the two versions of Asteroid – from sharp to less sharp – did not mitigate the unacceptability of DA-inspired modification of the empirical probability. Why should such a modification gain extra support just from the mere fact that the implications of empirical rationality are now a further degree more difficult to assess with certainty? To use DA to modify the empirical probabilities in Traditional is to go against scientists' recommendations no less than it is in Asteroid. There is no extra space for DA to slip into. Here too, any deviation from our best empirical prediction, even when that prediction is uncertain, is implicitly to believe in miracles. To deny this continuity claim, defenders of DA would need to say where and why in the continuum of increasing uncertainty between

Asteroid and Traditional a deviation from empirical rationality, and thus belief in miracles, somehow becomes acceptable.

### 10. DA priors – a different way out

There remains a different line of defence for DA though. Suppose we imagine that DA sets our prior probability for some particular date of Doom, and that now we interpret ‘prior’ to mean before consideration of any empirical evidence at all.<sup>13</sup> It is true that, numerically speaking, incorporating DA right at the start of our sequence of calculations would yield the same results as incorporating it at any other stage. Nevertheless philosophically speaking the difference is significant, this numerical equivalence notwithstanding. In particular, the a priori adjustment of empirical estimates falls foul of the trumping problem, but the a priori setting of priors does not. Thus, this (and only this) method of incorporating DA mathematics offers an escape route from our dilemma. In particular, it enables DA to embrace the dilemma’s first horn without any longer falling foul of trumping.<sup>14</sup> Whether DA actually does so, turns on the mathematics from section 6, i.e. on the relative size of the empirical and DA uncertainties. Unlike earlier, those calculations are thus now philosophically relevant again. Let us go through them in the various cases.

First, in Asteroid, the certainty of the empirical probability of Doom-Now means that DA’s prior value for this probability is trumped by the new empirical evidence of the asteroid observation, just as earlier empirical observations are trumped too. So DA is still impotent. Next, what of the less sharp version of Asteroid? Now the trumping is no longer complete because the uncertainty about the empirical probability of Doom-Now opens the door for the DA-informed prior to retain some influence. Thus, in principle the posterior probability of Doom-Now could now take a value below 0.5 legitimately. As per section 6, how much it actually does so depends on the relative uncertainties.

We can also illustrate this possibility qualitatively, in terms of the Mars-Earth and quantum cases. First, to repeat, DA can be relevant only when there is uncertainty regarding the relevant empirical probability estimates. So, let us assume an otherwise analogous set of observations in the Mars and Earth asteroid cases, save that now these yield only rather uncertain estimates of the collision probabilities. As noted, in such circumstances a DA prior implies a lower probability of Doom-Now, and so a lower probability of an asteroid collision for Earth than for Mars. DA considerations thus make it more likely that the two situations are *not* physically identical, the similar observations notwithstanding. The impact of adding empirical uncertainty to the case is that the DA deviation created by our different priors for Mars and Earth is no longer automatically fully trumped. Similarly, in the quantum case, if (but only if) there is uncertainty regarding our 0.5 probability estimate, then that again creates room for DA – via the lingering influence of its priors – legitimately to shift our best guess about it to below 0.5 without implying any belief in miracles.

This line of defence of DA does still face a challenge though. It must stipulate, remember, that DA enters before *any* empirical evidence – otherwise the considerations in sections 8 and 9 kick in again. How does this impact, finally, on our analysis of Traditional? There, the DA-informed priors must by now have been modified by an enormous body of empirical evidence, indeed by every bit of evidence relevant to humanity’s future duration that has ever been gathered: all of our accumulated knowledge about humanity’s evolutionary history, the likely astronomical fate of Earth, the frequency of mass extinction events such as

---

<sup>13</sup> This would therefore be to reverse the sequence of application endorsed by, for instance, Leslie in *End of the World*.

<sup>14</sup> In Bayesian terms, empirical evidence serves merely to modify our prior beliefs, so why not let those prior beliefs be informed by a priori considerations? This need imply no contradiction with best empirical practice – but only if DA is interpreted as motivating a prior probability value before any empirical input.



supervolcanoes and asteroid collisions, the possibility of human self-extermination via war or some other cause, the feasibility of future colonization of other planets, and so on. A vast range of observations has contributed to this knowledge. This inevitably raises the suspicion that, numerically speaking, whatever priors DA bequeathed us have long since been overwhelmed, simply because the variance of an empirical estimate drawn from so many sources might be very low.<sup>15</sup> If so, then DA would be rendered irrelevant and the second horn of the dilemma would apply again. Of course though, such a suspicion is difficult to prove – or equally to disprove – given that the relevant probability calculation is so intricate and enormous that it inevitably remains somewhat hypothetical.<sup>16</sup> The point is that the fate of DA depends critically on the answer.

## 11. Conclusion

One's opinion of the wider debate around DA will dictate how seriously one is willing to take it. But to the extent that one is so willing, this paper has shown how DA's applicability is constrained – or, more positively, under what conditions it might be endorsed. In particular, it is illuminating to apply DA to the case of Doom-Now rather than to Doom-Soon versus Doom-Later, as this directs our attention to the two epistemically crucial factors: first, trumping; and second, the degree of certainty of our probability estimates.

In summary, the asteroid example creates a dilemma for DA: either it is committed to unacceptable counterintuitiveness and belief in miracles, or else it is irrelevant. This dilemma extends to a version of the asteroid example with uncertainty about the relevant probabilities, and thus by a continuity argument it extends too to the usual case to which DA is applied, namely the actual expected future duration of humanity. The only line of defence is to take DA as informing our initial beliefs regarding the date of humanity's doom, prior to any empirical evidence on the matter at all. This defence is potentially sound in cases where the empirical probability is uncertain – but only in such cases. (Thus it already concedes DA's helplessness in the original Asteroid example, for instance.) And even then, it requires in addition that the vast mass of empirical evidence to hand does not swamp the DA priors into irrelevance.<sup>17</sup>

*Department of Philosophy  
Birkbeck College  
Malet Street  
London WC1E 7HX  
r.northcott@bbk.ac.uk*

---

<sup>15</sup> As noted in section 6, the literature seems implicitly to accept that, by contrast, the uncertainty of DA's probability estimates is non-negligible, which is what opens the door to empirical swamping.

<sup>16</sup> Leslie, *End of the World*, pp.201-2 calculates that DA should modify our credence in Doom-Early from 0.01 to 0.5 or even to 0.99. But these and similar calculations in the literature are irrelevant here. The reason is that in effect they assume that the relevant estimates for the empirical and DA probabilities have equal uncertainties, and thus that DA can greatly impact our overall estimate. But given the vast number of observations underpinning the empirical estimate, this equality of uncertainty cannot just be assumed and indeed is precisely what is at issue.

<sup>17</sup> I would like to thank anonymous referees at this and another journal for helpful comments, and also to thank the audience at Birkbeck to whom I presented a very early version of this paper.