A Third Route to the Doomsday Argument

preprint

Paul Franceschi

revised May 2003
University of Corsica

p.franceschi@univ-corse.fr
http://www.univ-corse.fr/~franceschi

ABSTRACT In this paper, I present a solution to the Doomsday argument based on a third type of solution, by contrast to on the one hand, the Carter-Leslie view and on the other hand, the Eckhardt et al. analysis. The present line of thought is based on the fact that both aforementioned analyses are based on an inaccurate analogy. After discussing the imperfections of both models, I present then a two-sided model that fits more adequately with the human situation corresponding to DA and encapsulates both Carter-Leslie's and Eckhardt et al.'s models. I argue then that this new analogy also holds when one takes into account the issue of indeterminism and the reference class problem. This leads finally to a novel formulation of the argument that could well be more consensual than the original one.

In this paper, I present a solution to the Doomsday argument (DA, for short) based on a third type of solution, by contrast to on the one hand, the Carter-Leslie view and on the other hand, the Eckhardt et al. analysis. In section 1, I describe the Carter-Leslie view. In section 2, I review the Eckhardt et al. line of reasoning. I point out then in section 3 an atemporal-temporal disanalogy in the Carter-Leslie analogy, which leads to the description of a strengthened variation of this latter model. In section 4, I raise some criticisms against the Eckhardt et al. analogy, thus leading to reformulate this latter analogy more accurately. I present then in section 5 a new two-sided analogy that fits more adequately with the human situation corresponding to DA and encapsulates both Carter-Leslie's and Eckhardt et al.'s models. This leads to a novel formulation of the argument that could well be more consensual than the original one. I argue in section 6 that this last two-sided analogy also holds when one takes into account the issue of indeterminism. Finally, I show in section 7 that the two-sided model is also capable of handling the reference class problem.[1]

1. The Carter-Leslie View

Let us begin by sketching briefly the Doomsday argument. The argument can be described as a reasoning leading to a Bayesian shift, from an analogy between what has been termed the *two-urn case*[2] and the corresponding human situation. Consider, first, the *two-urn case* (slightly adapted from Bostrom 1997):[3]

---

[1] The solution to DA presented here is a somewhat condensed and enhanced version of the ideas expressed in Franceschi (2002), that also discusses at length the problems related to DA: *God's Coin Toss*, the *Sleeping Beauty Problem*, the *Shooting-Room Paradox*, the *Presomptuous Philosopher*.

[2] Cf. Korb & Oliver (1998).

[3] Cf. Bostrom (1997): 'Imagine that two big urns are put in front of you, and you know that one of them contains ten balls and the other a million, but you are ignorant as to which is which. You know the balls in each urn are numbered 1, 2, 3, 4 ... etc. Now you take a ball at random from the left urn, and it is number 7. Clearly, this is a strong indication that that urn contains only ten balls. If originally the odds were fifty-fifty, a swift application of Bayes' theorem gives you the posterior probability that the left urn is the one with only ten balls. (Pposterior (L=10) = 0.999990)'.

The two-urn case An urn[4] is in front of you, and you know that it contains, depending on the flipping at time T_0 of a fair coin, either 10 (tails) or 1000 (heads) numbered balls. The balls are numbered 1, 2, 3, ... At this step, you formulate the H_{few} and H_{many} assumptions with $P(H_{few}) = P(H_{many}) = 0.5$ and you try to evaluate the number of balls which were contained at T_0 in the urn. You know all the above and you randomly draw a ball from the urn at time T_1. Now you get the ball #5 at T_1. You conclude then to an upward Bayesian shift in favour of the H_{few} hypothesis.

The *two-urn case* constitutes an uncontroversial application of Bayes' theorem. It is based on the two following competing hypotheses:

($H1_{few}$) the urn contains 10 balls
($H2_{many}$) the urn contains 1000 balls

and the corresponding prior probabilities: P(H1) = P(H2) = 0.5. Taking into account the fact that E denotes the available evidence that the random ball is #5 and P(E|H1) = 1/10 and P(E|H2) = 1/1000, a Bayesian shift ensues from a straightforward application of Bayes' theorem. As a result, the posterior probability is such that P'(H1) = 0.99.

Let us consider, on the other hand, the *human situation corresponding to DA*. Now being concerned with the final size of the human race, you consider the two following competing hypotheses:

($H3_{few}$) the number of humans having ever lived will reach 10^{11} (*doom soon*)
($H4_{many}$) the number of humans having ever lived will reach 10^{14} (*doom later*)

Now it appears that each human has his own birth rank, and that yours is roughly 60×10^9. Let us assume then, for the sake of simplicity, that the prior probabilities are such that P(H3) = P(H4) = 0.5.[5] Now according to Carter and Leslie, the *human situation corresponding to DA* is analogous to the *two-urn case*.[6] Let us denote by E the fact that your birth rank is 60×10^9. Thus, an application of Bayes' theorem, taking into account the fact that $P(E|H3) = 1/10^{11}$ and $P(E|H4) = 1/10^{14}$, leads to a vigorous Bayesian shift in favor of the hypothesis that Doom will occur soon: P'(H3) = 0.999. For this reason, the Carter-Leslie line of thought can be summarized as follows:

(5) in the *two-urn case*, a Bayesian shift of the prior probability of H_{few} ensues
(6) the situation corresponding to DA is analogous to the *two-urn case*
(7) ∴ in the situation corresponding to DA, a Bayesian shift of the prior probability of H_{few} ensues

From the Carter-Leslie's viewpoint, the analogy with the urn is well-grounded. And this legitimates DA's conclusion according to which a Bayesian shift in favor of doom soon ensues. This last conclusion appears paradoxical or at least counter-intuitive. But the task of diagnosing what is wrong, if any, with the Doomsday Argument proves to be very difficult and remains an open question.

At this point, it is worth mentioning in passing that the reasoning based on the two-urn case does not yield absolute certainty. This last reasoning is probabilistic and as such, it leads to a true conclusion *in most cases*. If the experiment is repeated many times and you bet accordingly, you will win in most cases. But it must be acknowledged that you will sometimes lose. For consider the situation where the coin lands heads and the number of balls in the urn is 1000. In this last case, if you get the ball #5, the reasoning based on the two-urn case leads to the false conclusion that the urn contains only 10 balls. However, this does not preclude us from regarding the corresponding reasoning as sound. For in the long run, it is reliable and yields many more true conclusion than false ones. The following table summarizes this situation:

[4] Bostrom's original description of the two-urn case refers to two urns. For the sake of simplicity, I refer here to one single urn (containing either 10 or 1000 balls) instead of two, since it is equivalent to the original two-urn case.
[5] The reasoning remains unaltered if we consider some alternative prior probabilities (with 0 < P(H3) < 1).
[6] More precisely, Leslie considers an analogy with the *lottery case*.

two-urn case (numbered balls)				
toss outcome	*reference class (numbered balls)*	*#*	*prediction*	*reasoning*
tail (doom soon)	10 numbered balls	#5	true	sound
heads (doom later)	1000 numbered balls	#5	false	sound

2. The Eckhardt et al. Analysis

A line of objection to the Doomsday argument initially raised by William Eckhardt (1993, 1997) and recently echoed by George Sowers (2002) and Elliott Sober (2003) runs as follows. The analogy with the urn at the origin of DA, so the objection goes, is ill-grounded. For in the *two-urn case*, the ball number is randomly chosen. But in the human situation corresponding to DA, our birth rank is not randomly chosen, but rather indexed on the corresponding temporal position. Hence, the analogy is ill-grounded and the whole reasoning is invalid. Eckhardt notably stresses the fact that it is impossible to make a random selection when there exists numerous unborn members in the chosen reference class.[7] Sober (2003) argues along the same lines,[8] by pointing out that no mechanism having the effect of randomly assigning a temporal location to human beings, can be exhibited. Lastly, such a line of objection has been recently revived by Sowers. He emphasizes that the birth rank of each human is not random, because it is indexed on the corresponding temporal position.[9]

In parallel, according to the Eckhardt et al. analysis, the human situation corresponding to DA is not analogous to the *two-urn case*, but rather to an alternative model, the *consecutive token dispenser*. The consecutive token dispenser is a device, initially described by Eckhardt,[10] that expels consecutively numbered balls at a constant rate: '(...) 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'. A similar device - call it the *numbered ball dispenser* - is also mentioned by Sowers:[11]

> There are two urns populated with balls as before, but now the balls are not numbered. Suppose you obtain your sample with the following procedure. You are equipped with a stopwatch and a marker. You first choose one of the urns as your subject. It doesn't matter which urn is chosen. You start the stopwatch. Each minute you reach into the urn and withdraw a ball. The first ball withdrawn you mark with the number one and set aside. The second ball you mark with the number two. In general, the n^{th} ball withdrawn you mark with the number *n*. After an arbitrary amount of time has elapsed, you stop the watch and the experiment. In parallel with the original scenario, suppose the last ball withdrawn is marked with a seven. Will there be a probability shift? An examination of the relative likelihoods reveals no.

Thus, according to the Eckhardt et al. line of thought, the human situation corresponding to DA is not analogous to the *two-urn case*, but rather to the *numbered ball dispenser*. And in this latter model, the conditional probabilities are such that P(E|H1) = P(E|H2) = 1. As a consequence, the prior probabilities of the two alternative hypotheses H_{few} and H_{many} are unchanged. Hence, the corresponding line of reasoning goes as follows:

(8) in the *numbered ball dispenser*, the prior probabilities remain unchanged
(9) the situation corresponding to DA is analogous to the *numbered ball dispenser*

---

7 Cf. (1997, p. 256): 'How is it possible in the selection of a random rank to give the appropriate weight to unborn members of the population?'.

8 Cf. (2003, p. 9): 'But who or what has the propensity to randomly assign me a temporal location in the duration of the human race? There is no such mechanism.'. But Sober is mainly concerned with providing empirical evidence against the hypotheses used in the original version of DA.

9 Cf. (2002, p. 40): 'My claim is that by assigning a rank to each person based on birth order, a time correlation is established (...).' and also (2002, p. 44): 'The doomsday argument has been shown to be fallacious due to the incorrect assumption that you are a random sample from the set of all humans ever to have existed.'.

10 Cf. (1997, p. 251).

11 Cf. (2002, p. 39).

(10) $\therefore$ in the situation corresponding to DA, the prior probabilities remain unchanged

thus yielding $P(H_{few}) = P'(H_{few})$ and $P(H_{many}) = P'(H_{many})$.

3. Strengthening the Carter-Leslie Analogy

As we have seen, according to the Carter-Leslie view, DA is based on an analogy between the human situation corresponding to DA and the *two-urn case*. By contrast, from the Eckhardt et al. standpoint, the analogy associates the human situation corresponding to DA and the *numbered ball dispenser*. In what follows, I shall argue that both analogies suffer from some defects and consequently do not prove to be fully adequate. This leads finally to reformulating the analogy more accurately.

Consider, on the one hand, the analogy with the two-urn case inherent to the Carter-Leslie view. Let us begin with the characteristics of the human situation corresponding to DA. A summary analysis shows indeed that this last situation is *temporal*. In effect, the birth ranks are successively attributed to human beings in function of the temporal position corresponding to their appearance on Earth. Thus, the corresponding situation takes place, say, from T_1 to T_n, 1 and *n* being respectively the rank numbers of the first and of the last human. By contrast, the two-urn case is *atemporal*, for at the moment where the ball is randomly drawn, all balls are already present in the urn.[12] Consequently, the two-urn case takes place at a given time T_0. Thus, the situation corresponding to DA needs to be modeled in a temporal model, while the two-urn case is rendered in an atemporal one. In short, the situation corresponding to DA being temporal, and the two-urn case being atemporal precludes us from regarding the two situations as isomorphic.[13] The importance of the *atemporal-temporal disanalogy* will become clearer later. Roughly, its importance rests on the fact that an atemporal model leads to one single model, while considering a temporal one leads to several competing probabilistic models. In addition, considering a temporal model is best suited for taking into account the issue of indeterminism and the reference class problem in the context of DA. In any case, at this step, it is apparent that the human situation corresponding to DA being temporal should be put in analogy more accurately with a *temporal* experiment.

The atemporal-temporal disanalogy being stated, let us investigate now how this inconvenient could be overcome. Consider then the following experiment, which can be termed the *incremental two-urn case* (let us denote it by *two-urn case*$^{++}$):

> *The synchronic and deterministic two-urn case*$^{++}$ An urn is in front of you, and you know that it contains, depending on the flipping at time T_0 of a fair coin, either 10 (tails) or 1000 (heads) numbered balls. At time T_1, you randomly draw the ball #*e* from the urn and then replace it in the urn. Then a device expels at T_1 the ball #1, at T_2 the ball #2..., at T_n the ball #*n*.[14] Now, according to the outcome of the random drawing performed at T_1, the device stops at T_e when the ball #*e* is expelled. At this step, you formulate the H_{few} and H_{many} assumptions with $P(H_{few}) = P(H_{many}) = 0.5$ and you try to evaluate the number of balls which were contained at T_0 in the urn. Now you know all the above and you get the ball #5 at T_5 when the device stops. You conclude then to an upward Bayesian shift in favour of the H_{few} hypothesis.

The novelty in this variation is that the experiment presents a temporal feature, given that the random selection is made at T_1 and the chosen ball is ultimately expelled at T_5. It is also worth pointing out that in the *synchronic and deterministic two-urn case*$^{++}$, the total number of balls in the urn is definitively fixed at T_0, when the experiment begins. For this reason, the corresponding situation can be termed *deterministic*. An instance of the *synchronic and deterministic two-urn case*$^{++}$ is as follows:

---

[12] It could be pointed out that a small amount of time is necessary to perform the Bayesian shift, once the problem's data are known. But this can be avoided if one considers ideal thinkers, who perform Bayesian shifts at the time when they are informed of the data relevant to the corresponding situation.

[13] I borrow this terminology from Chambers (2001).

[14] From now on, I assume that the intervals of time, i. e. from T_1 to T_n, are regular. Considering alternatively irregular intervals of time would not result in significant differences in the present account.

time	T_0	T_1	T_2	T_3	T_4	T_5
flipping	tails					
range		1-10				
random #		5				
ball #		1	2	3	4	5

At this step, it should be emphasized that anyone who accepts the conclusion of the *two-urn case* would also accept the Bayesian shift resulting from the *incremental two-urn case*.

Furthermore, it appears that other variations of the incremental two-urn case can even be envisaged. For consider the following variant:

> *The diachronic and deterministic two-urn case*$^{++}$ An opaque device contains an urn that has, depending on the flipping at time T_0 of a fair coin, either 10 (tails) or 1000 (heads) numbered balls. At time T_1, a robot inside the device draws a ball[15] at random in the urn (containing the balls #1 to #*n*) and the device expels the ball #1; if the ball #1 has been drawn then the device stops at T_1; else at T_2, the robot draws a ball at random in the urn (now containing the balls #2 to #*n*) and the device expels the ball #2; if the ball #2 has been drawn then the device stops at T_2; ...; else at T_i, the robot draws a ball at random in the urn (now containing the balls #*i* to #*n*) and the device expels the ball #*i*; if the ball #*i* has been drawn then the device stops at T_i; else at T_{i+1}, etc. Now you know all the above and you get the ball #5 at T_5 when the device stops.[16] You formulate the H_{few} and H_{many} assumptions with $P(H_{few}) = P(H_{many}) = 0.5$ and you conclude to an upward Bayesian shift in favor of the H_{few} hypothesis.

In this last case, the random selection is performed gradually and is only made effective when the number of the randomly drawn ball equals the number corresponding to the temporal position, i. e. when the ball #*i* is drawn at T_i. This contrasts with the synchronic version of the experiment, where the random selection is made definitively at time T_1. Nevertheless, in the *diachronic and deterministic two-urn case*$^{++}$, the probability of drawing the ball #*n* at T_n still equals 1/*n*. Let us denote by E the fact of drawing the ball #5 at T_5. It follows that the probability of drawing the ball #5 at T_5 if the urn contains 10 balls is such that P(E) = 9/10 x 8/9 x 7/8 x 6/7 x 5/6 x 1/5 = 1/10. An instance of the *diachronic and deterministic two-urn case*$^{++}$ is as follows:

time	T_0	T_1	T_2	T_3	T_4	T_5
flipping	tails					
range		1-10	2-10	3-10	4-10	5-10
random #		4	7	9	6	5
ball #		1	2	3	4	5

where the random selection is only made effective at time T_5.

At this step, it should be pointed out that the incremental two-urn case (whether synchronic or diachronic) does not face the above-mentioned criticisms concerning the atemporal-temporal disanalogy between the human situation corresponding to DA and the original two-urn case. For it has been shown that the human situation corresponding to DA, being temporal, cannot be put in analogy with the two-urn case, which is atemporal. By contrast, the incremental two-urn case is a temporal experiment. Thus, the incremental two-urn case meets the above mentioned requirements concerning the analogy and can be put legitimately in analogy with the human situation corresponding to DA. In this context, we now face a variation of DA which can be stated explicitly as follows:

(11) in the *incremental two-urn case*, a Bayesian shift of the prior probability of H_{few} ensues

[15] After the ball is drawn, it is replaced in the urn.

[16] This can be equivalently rendered with the following computer algorithm: at T_1, draw randomly a number between 1 and *n*; if 1 is issued then display 1 and stop; else at T_2, draw randomly a number between 2 and *n*; if 2 is issued then display 2 and stop; ...; else at T_i, draw randomly a number between *i* and *n*; if *i* is issued then display *i* and stop.

(12) the situation corresponding to DA is analogous to the *incremental two-urn case*

(13) ∴ in the situation corresponding to DA, a Bayesian shift of the prior probability of H_{few} ensues

And this last variation is not vulnerable to the above objection. The analogy with the urn is now plainly plausible, since both situations are *temporal*.

At this point, it is also worth scrutinizing the consequences of the *incremental two-urn case* (whether synchronic or diachronic) on the Eckhardt et al. analysis. For in the incremental two-urn case, the number of each ball expelled from the device is indexed on the rank of its expulsion. For example, you draw the ball #60000000000. But you also know that the preceding ball was #59999999999 and that the penultimate ball was #59999999998, etc. However, this does not prevent you from reasoning in the same way as in the original two-urn case and from concluding to a Bayesian shift in favor of the H_{few} hypothesis. In this context, the incremental two-urn case has the following consequence: *the fact of being time-indexed does not entail that the ball number is not randomly chosen*. Contrast now with the central claim of the Eckhardt et al. analysis that the birth rank of each human is not randomly chosen, but rather indexed on the corresponding temporal position. Sowers in particular considers that the cause of DA is the time-indexation of the number corresponding to the birth rank.[17] But what the incremental two-urn case and the corresponding analogy demonstrates, is that our birth rank can be time-indexed and nevertheless considered as random for DA purposes. And this point can be regarded as a significant objection to Sowers' analysis. This last remark leads to consider that the concrete analysis presented by Sowers does not prove however sufficient to solve DA. For the problem is revived when one considers the analogy between on the one hand, the human situation corresponding to DA and on the other hand, the incremental two-urn case. One can think that it is this latter analogy which constitutes truly the core of the DA-like reasoning. In this context, Sowers' conclusion according to which his analysis leads to the demise of DA appears far too strong. Echoing Eckhardt, he has certainly provided additional steps leading towards a resolution of DA and clarified significant points, but Sowers' analysis does not address veritably the strongest formulations of DA.

4. Refining the Exchardt et al. Analogy

Let us consider, on the other hand, the analogy with the *numbered ball dispenser*, which is characteristic of the Eckhardt et al. line of thought. As mentioned above, Eckhardt describes the *consecutive token dispenser*, where the tokens are expelled from the urn at constant rates ('one per minute'). Sowers also describes an analogous experiment, where the balls are expelled from the urn and numbered accordingly, at the constant[18] rate of one per minute. In this last experiment, the balls are numbered in the order of their expulsion from the urn.

However, the numbered ball dispenser can be criticised on the grounds that its protocol seems inaccurately defined. This inaccuracy concerns in particular the mechanism that expels a given ball #*n* at T_n. What makes the device stops at T_n after the ball #*n* has been expelled? The numbered ball dispenser seems to be designed for whatever way of choosing a given ball. So, could it be said, any mechanism for choosing the ball #*n* would be acceptable. But this won't do as a response, I think. For consider a deterministic situation, where the total number of balls in the urn is already fixed before the experiment begins. And suppose that a device chooses a ball at random at T_1 in the urn, say #5, and expels then accordingly the ball #5 at T_5. But it appears now that the corresponding situation is fully

[17] Cf. Sowers (2002, p. 40): 'My claim is that by assigning a rank to each person based on birth order, a time correlation is established in essentially the same way that the stopwatch process established a correlation with the balls.'.

[18] It could be pointed out that both Eckhardt's and Sowers' experiments do not exactly correspond to the human situation corresponding to DA. For in this latter situation, the humans appear on Earth at variable intervals of time, while Eckhardt and Sowers consider constant rates. However, this last disanalogy can be regarded as a minor qualm. For both Eckhardt's and Sowers' experiments could be eventually restated with items which are expelled at irregular rates instead of constant ones. In this context, a constant rate numbered ball dispenser can even be regarded as a useful simplification, for our present purposes of modeling the human situation corresponding to DA.

isomorphic with the *synchronic and deterministic two-urn case*$^{++}$. Thus, at least on one particular interpretation, the numbered ball dispenser proves to be identical to the *synchronic and deterministic two-urn case*$^{++}$. But as we have seen, this latter experiment leads to a straightforward Bayesian shift in favour of the H_{few} hypothesis, in opposition with the numbered ball dispenser where the prior probabilities remain unchanged. However, such interpretation of the numbered ball dispenser should be arguably discarded, on the grounds that it is at the opposite of the Eckhardt et al. viewpoint. But this shows that the protocol of the numbered ball dispenser stands in need of refinement and must be defined more accurately. This urges us to search another interpretation of the protocol of the numbered ball dispenser that fits more adequately with the spirit of the Eckhardt et al. line of thought.

Let us consider, second, another interpretation. Such interpretation arises from Sowers' description of the numbered ball dispenser. Sowers mentions in effect that the last ball is #7 ('In parallel with the original scenario, suppose the last ball withdrawn is marked with a seven'). Now let us repeat the experiment many times. In the long run, the numbered ball dispenser will always yield the ball #7 (or alternatively, a small number). Under this interpretation, the repeatability of the experiment show that the numbered ball dispenser has a *bias* towards #7. Let us call it the *biased numbered ball dispenser*. Although it should be acknowledged that this is also one *possible* interpretation of the numbered ball dispenser, I don't think neither that it fits adequately with what Sowers' has in mind. For it seems that Sowers' is concerned with a last ball expelled which is marked with whatever number (recall: 'In general, the *n*th ball withdrawn you mark with the number *n*'). For that reason, this second interpretation should also be rejected.

Let us consider then a third alternative interpretation. For it seems that an adequate interpretation of the numbered ball dispenser must do justice to Eckhardt' idea that *it is impossible to make a random selection when there exists numerous unborn members in the reference class*. Both previous interpretations of the numbered ball dispenser fail to incorporate this last idea. But consider now the following variation of the numbered ball dispenser:

> *The synchronic and deterministic numbered ball dispenser* An opaque device contains an urn that has 10 balls at T_0, but will ultimately have either 10 or 1000 numbered balls. The final number of balls in the urn will be determined by the flipping of a fair coin at T_0. If heads, it will add 990 numbered balls (#11 to #1000) in the urn a given time T_i ($1 \leq i < 10$), say at T_6. If tails, it will do nothing at T_6. At time T_1, you randomly draw the ball #*e* from the urn and then replace it in the urn. Then a device expels at T_1 the ball #1, at T_2 the ball #2..., at T_n the ball #*n*. Now, according to the outcome of the random drawing performed at T_1, the device stops at T_e when the ball #*e* is expelled. At this step, you formulate the H_{few} and H_{many} assumptions with $P(H_{few}) = P(H_{many}) = 0.5$ and you try to evaluate the number of balls which were contained at T_0 in the urn. You formulate the H_{few} and H_{many} assumptions relating to the total number of balls in the urn at T_6 with $P(H_{few}) = P(H_{many}) = 0.5$. Now you know all the above and you get the ball #5 at T_5 when the device stops. You conclude then that the prior probabilities remain unchanged.

An instance of the *synchronic and deterministic numbered ball dispenser* is then as follows:

time	T_0	T_1	T_2	T_3	T_4	T_5
flipping	tails					
range		1-10	2-10	3-10	4-10	5-10
random #		8	3	9	9	5
ball #		1	2	3	4	5

The novelty in this variation is that the urn contains 10 balls at the beginning but 990 other balls are eventually added later, say at T_6, depending on the outcome of a coin's toss. If the coin lands tails, nothing is done at T_6 and the urn remains with only 10 balls, the ball being drawn continuously in the range [1, 10]. If the coin lands heads, 990 balls are added in the urn at T_6. In this last case, the ball is drawn in the range [1, 10] until T_6, but from T_6 onwards, the ball is drawn in the range [1, 1000]. The protocol of this experiment can be described more generally in the following terms: if the urn contains only 10 balls, a ball is drawn randomly in the range [1, 10]. But if the urn contains 1000 balls, a ball is drawn randomly in the range [1, 1000]. Thus, the ball is drawn randomly, according to the *actual* number of balls in the urn. At this step, it should be apparent that this last protocol does do justice to Eckhardt's idea that it is impossible to make a random selection when there exists numerous unborn members in the reference class. In the present experiment the 990 balls that are added at T_6 represent those unborn members and the random process operates in the range [1, 10] until T_6, even in the case where the reference class will ultimately contain 1000 balls. Under these conditions, the *synchronic and deterministic numbered ball dispenser* appears well as a more robust variation of the numbered ball dispenser, that also vindicates Eckhardt's insights. This last variation does not face the above criticism of inaccuracy in its protocol. In addition, it incorporates an element of randomness, which is in line with our intuition that we are in some sense random humans. In this latter situation, it would be plainly erroneous to conclude to a Bayesian shift in favor of the H_{few} hypothesis. What is rational to infer in this situation, rather, is that the prior probabilities remain unchanged.

At this step, it is worth pointing out that a diachronic variation of the preceding experiment can even be envisaged. For consider the following variant of the numbered ball dispenser:

> *The diachronic and deterministic numbered ball dispenser* An opaque device contains an urn that has 10 balls at T_0, but will ultimately have either 10 or 1000 numbered balls. The final number of balls in the urn will be determined by the flipping of a fair coin at T_0. If heads, it will add 990 numbered balls (#11 to #1000) in the urn a given time T_i ($1 \le i < 10$), say at T_6. If tails, it will do nothing at T_6. At time T_1, a robot inside the device draws a ball[19] at random in the urn (containing the balls #1 to #10) and the device expels the ball #1; if the ball #1 has been drawn then the device stops at T_1; else at T_2, the robot draws a ball at random in the urn (now containing the balls #2 to #10) and the device expels the ball #2; if the ball #2 has been drawn then the device stops at T_2; ...; else at T_i, the robot draws a ball at random in the urn (now containing the balls #*i* to #*n*) and the device expels the ball #*i*; if the ball #*i* has been drawn then the device stops at T_i; else at T_{i+1}, etc. You formulate the H_{few} and H_{many} assumptions relating to the total number of balls at T_6 with $P(H_{few}) = P(H_{many}) = 0.5$. Now you know all the above and you get the ball #5 at T_5 when the device stops. You conclude then that the prior probabilities remain unchanged.

5. A Third Route

Given the above developments, we are now in a position to evaluate the adequacy of the analogy inherent to DA. We now face two competing analogies with the human situation corresponding to DA. At this step, the question that arises is the following: is the human situation corresponding to DA analogous to the *two-urn case*$^{++}$ or to the *numbered ball dispenser*. For we are now in presence of two alternative models for the analogy with the human situation corresponding to DA: (i) the - *synchronic* or *diachronic* - *deterministic two-urn case*$^{++}$ and (ii) the - *synchronic* or *diachronic* - *deterministic numbered ball dispenser*. As we have seen, these two variations are strong variations of their respective models, since they are not vulnerable to several objections that can be raised against the original ones.

Consider, to begin with, the part of the experiment that takes place from to T_1 to T_n. Consider, first, the analogy between the human situation corresponding to DA and the *incremental two-urn case*. As we have seen, this latter version constitutes a strong variation of the argument in the sense that it is not open, first, to the charge of putting in correspondence an atemporal model with a temporal situation. This last variation is not vulnerable, second, to the objection that arises from the Eckhardt et al. analysis, according to which our birth rank is not random because it is time-indexed. Let us question now whether the human situation corresponding to DA is analogous or not to the part of the experiment that takes place from to T_1 to T_n in the *incremental two-urn case*. At this step, it appears that both situations are temporal and relate to numerous objects (or individuals) whose number is that of their expulsion (or birth) rank. Thus, the part of the experiment that takes place from to T_1 to T_n proves to be fully analogous to the human situation corresponding to DA. In this sense, the analogy proves to be strongly established.

[19] After the ball is drawn, it is replaced in the urn.

Let us turn now to the analogy between the human situation corresponding to DA and the - *synchronic* or *diachronic* - *deterministic numbered ball dispenser*. Just as with the incremental two-urn case, it appears that the part of the experiment that takes place from to T_1 to T_n is entirely analogous to the human situation corresponding to DA. This should not be surprising because the *external* parts of both the incremental two-urn case and numbered ball dispenser are identical. For an external observer, there is no difference between the two experiments. The upshot is that for what concerns the part of the experiment that takes place from to T_1 to T_n, there is no difference between the incremental two-urn case and the deterministic numbered ball dispenser. Hence, both models incorporate adequately the corresponding features of the situation corresponding to DA.

Let us turn now to the random process. In the synchronic and deterministic two-urn case^{++}, the random drawing takes place at T_0. In the same way, this random drawing also occurs at T_0 in the synchronic and deterministic numbered ball dispenser. For in both *synchronic* experiments, the random selection of a numbered ball is made at that very moment. Now does an analogous random selection take place with the same degree of certainty at the eve of the beginning of humankind? Do we have any evidence that such random selection has occurred just before the birth of the first human? No. For we lack any proof that the birth rank of future humans is determined by a random selection having occurred just before the beginning of humankind. We currently lack evidence of any such random process. In the case of the human situation corresponding to DA, the occurrence of such random selection remains fully hypothetical. However, this latter disanalogy does not appear to be a fundamental one. In effect, it appears that both synchronic models admit of a diachronic variation. For in both the *diachronic and deterministic two-urn case*$^{++}$ and the *diachronic and deterministic numbered ball dispenser*, the random process is performed diachronically and is made gradually from to T_1 to T_e. Now it appears that both models are capable of modelling adequately the random process[20] that determines the birth rank of each human. Whether such random drawing has been made before the beginning of humankind or perhaps more plausibly, during the course of the existence of the human race does not matter. For both the *deterministic two-urn case*$^{++}$ and the *deterministic numbered ball dispenser* admit of variations that model adequately these two situations.

What precedes casts light on the crucial point that the experiment analogous to the human situation corresponding to DA should reflect this last property, namely that both the *two-urn case*$^{++}$ and the *numbered ball dispenser* are capable of modelling properly the situation corresponding to DA. Now the question that arises is the following: Does there exist an objective criterion allowing the preferential choice of one of the two competing models? Is there any clue that allows to prefer the incremental two-urn case or the numbered ball dispenser? The answer is no. We currently lack any objective criterion allowing to choose between the two alternative analogies. In the lack of the relevant evidence, it is then wise to retain both models as roughly equiprobable. What remains in force, thus, is an indeterminate situation. We currently lack any objective criterion allowing to choose on rational grounds between the two-urn case^{++} and the numbered ball dispenser. In consequence, the preferential choice of either the two-urn case^{++} or the numbered ball dispenser appears well as a one-sided attitude. It appears now that an adequate model should reflect the fundamental property of being *two-sided*. At this point, we are in a position to describe a new model for the human situation corresponding to DA which is not open to the charge of being one-sided:

> *The two-sided model* With a probability P such that 0 < P < 1, a device performs (whether synchronically or diachronically) either a *numbered ball dispenser* or a *two-urn case*$^{++}$. Now you get the ball #*e* at T_e when the device stops. Given the H_{few} and H_{many} assumptions and $P(H_{few}) = P(H_{many}) = 0.5$, how would you update your prior probabilities?

Now, whatever the chosen random number, the device expels the ball #1 at time T_1, the ball #2 at time T_2, ..., the ball #*e* at time T_e. This last experiment presents the following property: there is no *external* difference whether the performed experiment is a *numbered ball dispenser* or a *two-urn case*$^{++}$. Let us analyze then the *two-sided model* in more detail. On the one hand, the situation that takes place from

[20] To put it metaphorically. Needless to say, such random process need not be taken at face value.

T_1 to T_n does reflect the part of the human situation corresponding to DA that takes place from the beginning of humankind[21] to our current birth. On the other hand, the characteristics of the random process that takes place at T_0 (synchronically) or from T_1 to T_e (diachronically) now correspond adequately to our current situation. Finally, the corresponding line of reasoning is as follows:

(14) the situation corresponding to DA is either analogous to the incremental two-urn case or to the numbered ball dispenser
(15) in the incremental two-urn case, a Bayesian shift of the prior probability of H_{few} ensues
(16) in the numbered ball dispenser, the prior probability remains unchanged
(17) $\therefore$ in the situation corresponding to DA, either a Bayesian shift of the prior probability of H_{few} ensues or the prior probability remains unchanged

in replacement of the steps (5)-(7) of the Carter-Leslie line of reasoning and of the steps (8)-(10) of the Eckhardt et al. line of argument.

6. The Issue of Indeterminism

Let us examine now whether the two-sided model is affected or not by the issue of indeterminism in the context of DA. This last issue is an important one, since it has notably motivated Eckhardt's above-mentioned point that it is impossible to make a random selection when there exists numerous unborn members in the reference class. At this step, it is worth noting that the preceding experiments suggests that some variations are capable of handling an *indeterministic* situation, namely where the total number of balls in the urn is unknown at the time where the experiment begins and is only determined with certainty during the course of the experiment. As an example, the following variation takes into account an indeterministic situation:

> *The diachronic and indeterministic two-urn case*$^{++}$ An opaque device contains an urn that has 10 balls at T_0, but will ultimately have either 10 or 1000 numbered balls. The final number of balls in the urn will remain undetermined until an internal mechanism will flip a fair coin at a given time T_i ($1 \leq i < 10$). If heads, it will add 990 numbered balls (#11 to #1000) in the urn at T_i. If tails, it will do nothing. At time T_1, a random generator inside the device issues a number in the range [1, 1000] and the device expels the ball #1; if the number 1 has been issued then the device stops at T_1; else at T_2, the random generator issues a number in the range [2, 1000] and the device expels the ball #2; if the number 2 has been issued then the device stops at T_2; ...; else at T_{i-1}, the random generator issues a number in the range [*i*-1, 1000] and the device expels the ball #*i*-1; if the number *i*-1 has been issued then the device stops at T_{i-1}; else at T_i ($1 \leq i < 10$), the random generator issues a number in the range [*i*, *n*] (the total number of balls in the urn after the flipping of the coin is *n*) and the device expels the ball #*i*; if the number *i* has been issued then the device stops at T_i; else at T_{i+1}, etc. Now you know all the above and you get the ball #*e* at T_e when the device stops. You formulate the H_{few} and H_{many} assumptions relating to the total number of balls in the urn after the flipping of the coin with $P(H_{few}) = P(H_{many}) = 0.5$. Now you know all the above and you get the ball #5 at T_5 when the device stops. You conclude then to an upward Bayesian shift in favor of the H_{few} hypothesis.

An instance of the *diachronic and indeterministic two-urn case*$^{++}$ is then as follows:

time	T_0	T_1	T_2	T_3	T_4	T_5
flipping				tails		
range		1-1000	2-1000	3-1000	4-10	5-10
random #		857	326	92	9	5
ball #		1	2	3	4	5

[21] To be accurate: from the beginning of the chosen *reference class*.

The novelty in this variation is that it takes into account an indeterministic situation. In effect, the number of balls present in the urn is unknown at the time where the first ball is expelled from the device and is only determined at T_i. Such a variation shows that a random selection can even be made when the number of balls in the urn is unknown at the time where the random process begins. And this appears as a counter-example to Eckhardt's attack against the random sampling assumption in DA, based on the impossibility of making a random selection when there exists many unborn members in the given reference class. But the *diachronic and indeterministic two-urn case*$^{++}$ shows that a random selection can even be made, under certain indeterministic circumstances.

However, it should be acknowledged that this only partly undermines Eckhardt's point and that this latter experiment does not handle every type of indeterministic situation. For it could be retorted that Eckhardt could provide a counterpart of the *diachronic and indeterministic two-urn case*$^{++}$. In effect, Eckhardt could reply that what he has in mind is an experiment of the following type:

> *The diachronic and indeterministic numbered ball dispenser* An opaque device contains an urn that has 10 balls at T_0, but will ultimately have either 10 or 1000 numbered balls. The final number of balls in the urn will remain undetermined until an internal mechanism will flip a fair coin at a given time T_i ($1 \leq i < 10$). If heads, it will add 990 numbered balls (#11 to #1000) in the urn at T_i. If tails, it will do nothing. At time T_1 a robot inside the device draws a ball at random in the urn (containing the balls #1 to #10) and the device expels the ball #1; if the number 1 has been issued then the device stops at T_1; else at T_2, a robot inside the device draws a ball at random in the urn (containing the balls #2 to #10) and the device expels the ball #2; if the number 2 has been issued then the device stops at T_2; ...; else at T_{i-1}, a robot inside the device draws a ball at random in the urn (containing the balls # *i*-1 to #10) and the device expels the ball #*i*-1; if the number *i*-1 has been issued then the device stops at T_{i-1}; else at T_i, a robot inside the device draws a ball at random in the urn (containing the balls #*i* to # *n*) (the total number of balls in the urn after the flipping of the coin is *n*) and the device expels the ball #*i*; if the number *i* has been issued then the device stops at T_i; else at T_{i+1}, etc. Now you know all the above and you get the ball #5 at T_5 when the device stops. You formulate the H_{few} and H_{many} assumptions relating to the total number of balls in the urn after the flipping of the coin, with $P(H_{few}) = P(H_{many}) = 0.5$. You conclude then that the prior probabilities remain unchanged.

An instance of the *diachronic and indeterministic numbered ball dispenser* is as follows:

time	T_0	T_1	T_2	T_3	T_4	T_5
flipping				tails		
range		1-10	2-10	3-10	4-10	5-10
random #		7	6	10	7	5
ball #		1	2	3	4	5

To the difference of the preceding case, the random drawing of the ball is made here in the range [1, 10] until the flipping of the coin. This increases the probability that the device stops, says at T_5, even if the urn will ultimately contain 1000 balls, after the coin has eventually landed heads. In such a case, the drawing of the ball #5 at random gives us no grounds for concluding to a Bayesian shift in favor of the H_{few} assumption. In effect, in this last situation, it is very probable to draw a number in the range [1, 10], even if the coin lands heads at T_i.

The upshot is that taking into account the issue of indeterminism gives us no clue to decide whether the incremental two-urn case or the numbered ball dispenser is an adequate model for the human situation corresponding to DA. For both models admit of a variation which is capable of modeling the human situation corresponding to DA, even if it is of an indeterministic nature. Hence, whether the situation is inderministic or not, we still find ourselves on the third route.

7. The Reference Class Problem

At this stage, it is worth recalling the *reference class problem*.[22] Roughly, it is the problem of how to define 'humans'. More accurately, it can be stated as follows: How can the reference class be objectively defined for DA-purposes? For an extensive or restrictive definition of the reference class can be given. An extensively defined reference class would include for example the somewhat exotic future evolutions of humankind, for example with an average I.Q. of 200 or with backward causation abilities. Conversely, a restrictively designed reference class would only include those humans who correspond accurately to the characteristics of, say, *homo sapiens sapiens*, thus excluding the past *homo sapiens neandertalensis* and the future *homo sapiens supersapiens*. To put it more in adequation with our current taxonomy, the reference class can be defined at different levels which correspond respectively to the supergenus *superhomo*, the *homo* genus, the *homo sapiens* species, the *homo sapiens sapiens* subspecies, etc. At this step, it appears that we lack an objective criterion to choose the corresponding level non-arbitrarily.

Leslie's treatment of the reference class problem is exposed in the response made to Eckhardt (1993) and in Leslie (1996).[23] Leslie's response to the reference class problem is as follows. According to Leslie, one can choose the reference class more or less as one wishes, i.e. at a somewhat arbitrary level. Once this choice is performed, it suffices to adjust the prior probabilities accordingly to get the argument moving. Leslie's sole condition is that the reference class should not be chosen at an extreme level of extension or of restriction.[24] In addition, Leslie addresses the resulting fact that each human belongs to several different classes, restrictively or extensively defined. However, this is not a problem from Leslie's standpoint, since the argument works for all these classes. In effect, a Bayesian shift ensues for whatever reference class arbitrarily chosen, at a somewhat reasonable level of extension or of restriction. Leslie illustrates this point with an urn analogy. To the difference of the two-urn case, he considers an urn that contains balls of different colors, say red and green. A red ball is drawn from the urn. In this context, from a restrictive viewpoint, the ball is a random red ball and there is no difference in this case with the classical two-urn case. But from a more extensive viewpoint, it is also a random red-or-green ball.[25] According to Leslie, although the prior probabilities are different in each case, a Bayesian shift ensues in both cases.[26] In sum, on Leslie's view, the reference class problem can be overcome because the argument works for *all* reference classes. For that reason, Leslie's account can be termed an *undifferentiated* account of the reference class problem.

At this step, it appears that the incremental two-urn case and the numbered ball dispenser can be easily adapted, in order to incorporate the elements of the reference class problem. In both models, it suffices to consider that the 10 first balls are red and that the 990 remaining balls are green ones. Now the urn is filled with red-or-green balls and a given red ball (or a green ball) can also be considered as a red-or-green ball. In particular, this provides an adequate framework for modeling the paradigm case of the neandertalian. Leslie addresses in effect the case of a neandertalian who would have implemented a DA-like reasoning:[27]

---

[22] The *reference class problem* in probability theory is notably exposed in Hájek (2002, s. 3.3). For a treatment of the reference class problem in the context of DA, see notably Eckhardt (1993, 1997), Bostrom (1997, 2002, ch. 4 pp. 69-72 and ch. 5), Franceschi (1998, 1999). The point of Franceschi (1999) is that the reference class problem also arises in confirmation theory.

[23] In the part entitled 'Just who should count as being human?' (pp. 256-63).

[24] Cf. 1996, p. 260: 'Widenings of reference class can easily be taken too far.' and p. 261: 'Again, some ways of *narrowing* a reference class might perhaps seem inappropriate.'.

[25] Cf. Leslie (1996, p. 259): 'Suppose all the balls in the urn are numbered. A ball is drawn. It turns out to be bright red. Note that it is not only a bright red ball whose number has been drawn *at random* from the numbers of all the *bright red* balls in the urn, but also a red-or-reddish ball whose number has been drawn *at random* from the numbers of all the *red-or-reddish* balls in the urn.'.

[26] Cf. Leslie (1996, pp. 258-9): 'The thing to note is that the red ball can be treated either just as a red ball or else as a red-or-green ball. Bayes's Rule applies in both cases. When we're interested in how many red balls there are in the urn, we need to treat the ball just as a red ball. The 'prior probabilities' entering into our Bayesian calculation are then probabilities for such and such numbers of red balls. When, in contrast, what interests us is how many red-or-green balls the urn contains, then we have to treat the red ball as red-or-green. Correspondingly, the prior probabilities entering into the calculation are the prior probabilities of various numbers of balls in the red-or-green-ball class. [...] All this evidently continues to apply to when being-red-or-green is replaced by being-red-or-pink, or being-red-or-reddish.'.

[27] Cf. (1992, pp. 527-8).

> Consider the protest that any Stone Age man who had used the argument would have been led to the erroneous conclusion that the human race would soon die out. A first reply is: So what? It is not a defect in any merely probabilistic argument if it leads someone improbably situated - someone very early in time, maybe, or someone who has thrown a dozen dice with eyes shut and expects (mistakenly, in view of what is actually on the table) not to see a dozen sixes upon opening them - to an erroneous conclusion.

From the above excerpt and the fact that Leslie considers the neandertalian's conclusion as erroneous, it is implicit that the corresponding reference class is the somewhat extensively defined *homo sapiens* species. Now the *homo sapiens* species (red-or-green balls) has historically included the *homo sapiens neandertalensis* subspecies (red balls) and then the *homo sapiens sapiens* subspecies (green balls). Leslie acknowledges the fact that a neandertalian who would have implemented a DA-like reasoning relating to the *homo sapiens* reference class would have been led to a false conclusion. But as Leslie explains, this is due to the above-mentioned[28] fact that the reasoning based on the two-urn case does not yield absolute certainty. In the neandertalian case, the *doom later* hypothesis (heads) has been finally confirmed and the neandertalian had concluded falsely to a nearest extinction. But her anthropic reasoning were nevertheless sound. The following table summarizes the corresponding situation:

incremental two-urn case (red-or-green balls - homo sapiens)					
toss outcome	*reference class (balls)*	*reference class (humans)*	*#*	*prediction*	*reasoning*
tail (doom soon)	10 red-or-green balls	10^7 *homo sapiens*	#5	true	sound
heads (doom later)	1000 red-or-green balls (10 red balls and 990 green balls)	10^{10} *homo sapiens* (10^7 *homo sapiens neandertalensis* and 10^{10} - 10^7 *homo sapiens sapiens*)	#5	false	sound

At this step, it appears that Leslie's treatment of the neandertalian case is fully in adequation with the incremental two-urn case.

But let us plug alternatively this situation into the *numbered ball dispenser* model. Consider, as an example, the *diachronic and indeterministic numbered ball dispenser*.[29] Now it appears that the neandertalian who would have based her reasoning on this latter model would have been right if she has reasoned as follows: It is well possible that some unborn members are not currently taken into account and hence, my birth rank cannot be considered as random in the corresponding reference class. Consequently, I better leave my prior probabilities unchanged. Now it should be apparent that the *diachronic and indeterministic numbered ball dispenser* also fits adequately with the neandertalian case. The upshot is that if the adequate model is a numbered ball dispenser, then the just-mentioned reasoning is entirely in accordance with this latter model.

Under these circumstances, it appears that the situation arising from the reference class problem parallels that of the issue of indeterminism in the context of DA. In effect, it appears that the paradigm case of the neandertalian can be modeled in two different ways. Just as in the issue of indeterminism, the neandertalian case can be rendered either by an *incremental two-urn case*, or by a *numbered ball dispenser* (whatever the chosen variations, i.e. synchronic or diachronic, deterministic or indeterministic). It follows that both models are capable of rendering the corresponding analogy. However, this has significant consequences, since the former leads to a Bayesian shift, while the latter leaves the prior probabilities unchanged.

I have expressed my own view on the *reference class problem* in Franceschi (1998, 1999). By contrast to Leslie's viewpoint, it can be characterized as a *differential* account of the reference class problem. Let us examine how it handles the neandertalian case.[30] It is worth pointing out first that Leslie's account of the neandertalian case is only concerned with the *homo sapiens* reference class. But

[28] Cf. §1.
[29] The same goes for the other variations of the numbered ball dispenser.
[30] Adapted from Franceschi (1998, p. 243).

my point is that the neandertalian can consider himself at a somewhat extensive level, as a member of the *homo sapiens* species (as Leslie emphasizes) or at a slightly more restrictive level, as a member of the *homo sapiens neandertalensis* subspecies. And the consequences of this choice are not negligible. In effect, had the neandertalian identified the reference class with the *homo sapiens neandertalensis* subspecies, her anthropic prediction would have been then successful. But the corresponding anthropic prediction would have failed if she had chosen, more extensively, the *homo sapiens* species as the relevant reference class. At this step, it should be pointed out that the history of humankind is such that the extinction of the early *homo sapiens neandertalensis* subspecies has been followed by the appearance of our current *homo sapiens sapiens* subspecies. Thus, the doom later hypothesis for the *homo sapiens* reference class has been historically confirmed. But the point is that a doom later hypothesis from the *homo sapiens* class viewpoint is also a doom soon hypothesis from the *homo sapiens neandertalensis* class standpoint. More generally, it appears that a *doom later* situation for a given reference class is also a *doom soon* situation for a more restrictive reference class. Furthermore, such type of situation among evolutionary species is very widespread. For the extinction of our current *homo sapiens sapiens* subspecies could well be followed by the survival of the *homo sapiens supersapiens* subspecies. Consider also, at a greater level of restriction, a reference class consisting of all *homo sapiens sapiens having not known of the computer*. Doesn't there exist serious grounds for considering that this last reference class is promised to a nearest extinction? More generally, for whatever chosen reference class, I can still choose a slightly more extensive class that will survive. And it should be pointed out that this ambivalent effect has the effect of depriving the original argument from its initial terror. At this step, it should be apparent that the consideration of a *differential treatment of the reference class problem* renders DA innocuous. Finally, this gives a way of accepting its conclusion by rendering the argument less counterintuitive than in its original formulation.

Now the preceding remarks concerning the reference class problem can be combined with the conclusion of the preceding developments relating to the analogy with the urn. Let us recall the conclusion of the amended DA based on the two-sided analogy: for a given reference class, a Bayesian shift *possibly* ensues. In effect, it has been shown that DA only possibly works, for a given class. And this leads to a novel formulation of the argument. For if there existed a given reference class for which the argument were conclusive, there could well exist a more extensive class for which the argument would fail. This vindicates the differential treatment of the reference class problem and finally renders the argument innocuous, by depriving it of its initially associated terror. At the same time, this leaves room for the argument to be successful for a given reference class, but without its counterintuitive consequences.

To sum up now. What results from the foregoing developments is that the Doomsday Argument must be weakened in two ways. First, the analogy underlying the argument must be defined more accurately. As we have seen, this can be done with the help of the two-sided model, which does justice to the respective insights of both the incremental two-urn case and numbered ball dispenser. Within this new two-sided model, the Bayesian shift associated with the Doomsday Argument now appears just as a *possible* inference from the premises, and not as an absolutely certain consequence. Second, the reference class problem must be taken into account, thus leading to the conclusion that the Doomsday Argument could work but without its originally associated terror. This has the effect of rendering the conclusion of the argument less counter-intuitive than in its original formulation. Given these two sidesteps, it seems that the resulting novel formulation of the argument could be more consensual than the original one.

Lastly, what precedes casts light on an essential facet of the Doomsday Argument. For on a narrow sense, it is an argument about the fate of humankind. But on a broad sense (the one I have been concerned with) it emphasizes the difficulty of applying probabilistic models to real-life situations,[31] a difficulty which is usually largely underestimated. This opens a path to a whole field of practical

[31] This important underpinning of the argument is also underlined in Delahaye (1996). This is also the main point of Sober (2003).

interest, consisting of a taxonomy of probabilistic models, whose philosophical importance would have been unravelled without John Leslie's robust and courageous defence of the argument.[32]

References

Bostrom, N. (1997) 'Investigations into the Doomsday argument', *Preprint* at http://www.anthropic-principle.com/preprints/inv/investigations.html

Bostrom, N. (2002) *Anthropic Bias: Observation Selection Effects in Science and Philosophy,* New York, Routledge

Chambers, T. (2001) 'Do Doomsday's Proponents Think We Were Born Yesterday?', *Philosophy*, 76, 443-50

Delahaye, J-P. (1996) 'Recherche de modèles pour l'argument de l'apocalypse de Carter-Leslie', *unpublished manuscript*

Eckhardt, W. (1993) 'Probability Theory and the Doomsday Argument', *Mind*, 102, 483-88

Eckhardt, W. (1997) 'A Shooting-Room view of Doomsday', *Journal of Philosophy*, 94, 244-259

Franceschi, P. (1998) 'Une solution pour l'argument de l'apocalypse', *Canadian Journal of Philosophy*, 28, 227-46

Franceschi, P. (1999) 'Comment l'urne de Carter et Leslie se déverse dans celle de Hempel', *Canadian Journal of Philosophy*, 29, 139-56, English translation at http://www.anthropic-principle.com/preprints

Franceschi, P. (2002) *Une application des n-univers à l'argument de l'Apocalypse et au paradoxe de Goodman*, doctoral dissertation

Hájek, A. (2002) 'Interpretations of Probability', *The Stanford Encyclopedia of Philosophy* (Winter 2002 Edition), E. N. Zalta (ed.), http://plato.stanford.edu/archives/win2002/entries/probability-interpret

Korb, K. & Oliver, J. (1998) 'A Refutation of the Doomsday Argument', *Mind*, 107, 403–10

Leslie, J. (1992) 'Time and the Anthropic Principle', *Mind*, 101, 521-40

Leslie, J. (1993) 'Doom and Probabilities', *Mind*, 102, 489-91

Leslie, J. (1996) *The End of the World: the science and ethics of human extinction*, London, Routledge

Sober, E. (2003) 'An Empirical Critique of Two Versions of the Doomsday Argument - Gott's Line and Leslie's Wedge', *Synthese*, 135-3, 415-30

Sowers, G. F. (2002) 'The Demise of the Doomsday Argument', *Mind*, 111, 37-45

[32] I am very grateful to Claude Panaccio and Daniel Andler for comments on an ancestor version of this paper. I especially thank Jean-Paul Delahaye for very useful comments and discussion. I am also indebted to John Leslie and Elliott Sober for comments on an earlier draft. I thank Nick Bostrom for very helpful discussion on the reference class problem.