

A E ~ C - D ~ A --
G L ~ L

: I discuss two versions of the doomsday argument. According to “Gott’s Line,” the fact that the human race has existed for 200,000 years licences the prediction that it will last between 5100 and 7.8 million more years. According to “Leslie’s Wedge,” the fact that I currently exist is evidence that increases the plausibility of the hypothesis that the human race will come to an end sooner rather than later. Both arguments rest on substantive assumptions about the sampling process that underlies our observations. These sampling assumptions have testable consequences, and so the sampling assumptions themselves must be regarded as empirical claims. The result of testing some of these consequences is that both doomsday arguments are empirically disconfirmed.

1. G L

Richard Gott (1993, 1997) presents the following version of the doomsday argument:

1. My present temporal position can be treated as if it were the result of random sampling from the times during which S exists.
2. Hence, there is a probability of 0.95 that my present temporal position is in the middle 95% of S’s duration.
3. S began at time t_0 and the present date is t_1 .
4. Hence, there is a probability of 0.95 that S will cease to exist after the passage of a period of time that is greater than $(1/39)(t_1 - t_0)$ and less than $39(t_1 - t_0)$.
5. Hence, we may reasonably predict that S will cease to exist after the passage of a period of time that is greater than $(1/39)(t_1 - t_0)$ and less than $39(t_1 - t_0)$.

Gott used this form of reasoning in 1969 to estimate the durations of the Berlin Wall and the Soviet Union, both of which perished within the intervals he calculated. He also used this argument to predict the duration of the human race; assuming that has so far existed for 200,000 years, Gott’s prediction is that our species will last somewhere between 5100 and 7.8 million more years. I call this argument “Gott’s Line” because it envisages sampling from a time line.

Gott (1997, p. 39) says that his argument requires some restrictions. He points out that if you attend a wedding, you should not use his argument to predict, five minutes after the bride and groom say “I do,” that the marriage will probably end in about five minutes. The reason is that “you are at the wedding precisely to witness its beginning.” He thinks it would be equally wrong to use the doomsday argument to reason that the Universe is about half over. The reason is that “intelligent observers

emerged only long after the Big Bang, and so witness only a subset of its time line.” Gott is pointing to the fact that we sometimes have good reason to think that the first premiss of his argument is false. Our empirical background knowledge tells us that our time at the wedding is a time drawn at random from the full duration of the marriage, nor should it be thought of in accordance with that false assumption. And if it takes a long while for intelligent life to evolve, then our present temporal position should be regarded as a time drawn at random from the full duration of the Universe. Gott’s Line is intended to be a and defeasible argument, one that can be overturned by empirical evidence to the contrary (Bostrom 1997, 2000; Caves 2000). As Gott (1997, p. 39) says, his argument “... is most useful when examining the longevity of something, like the human race, for which there is no actuarial data available. We know only one human race. In predicting your lifetime, you can do better by using statistics on the life spans of people who have died...” It is unclear why the fact that we know only one human race is decisive. We have information about the longevities of many species; why isn’t this relevant to estimating the life expectancy of our own? The same point applies to the Berlin Wall and to the Soviet Union. We are not totally in the dark with respect to the longevities of walls or repressive regimes. In any event, it is clear that Gott intends his argument to apply when (and only when) one lacks frequency data or other empirical evidence from which probabilities can be estimated. If we have no empirical evidence against assigning equal probabilities to the times during S’s duration that we might now be in existence, we are supposed to adopt that flat distribution.

This justification of premiss (1) makes it sound as if Gott’s argument relies on the (notorious) Principle of Indifference. Indeed, Goodman (1994) criticizes the argument in just these terms; he points out that if we are completely uncertain about the number of years in S’s duration, we also are in the dark with respect to the exponents and roots of the number of years. Each of these corresponds to a different probability distribution when the Principle of Indifference is allowed to do its work. Gott (1994) replies that not all versions of the Principle of Indifference are fallacious, and claims that his “Copernican Principle” is one of the principle’s valid versions. This principle says that we regard our temporal position as “special.” From this negative, Gott extracts a positive – that we endorse the sampling assumption described in premiss (1).

2. O G L

If we agree, as I think we should, that the Principle of Indifference is invalid, what verdict does this entail about Gott’s argument? Can we say only that the justification of the argument’s first premiss remains obscure? I want to argue for a stronger criticism – that the first premiss is empirically implausible, at least with respect to the objects (the human race, the Berlin Wall, the USSR) to which the argument has been applied. I should decline to treat my present temporal position as if it were the result of random sampling from the times during which these objects exist because this assumption has implications that are disconfirmed by observations.

To argue this point, let me begin by describing how Gott calculates his estimate – that there is a

95% probability that S will cease to exist between the limits he specifies. First, notice that premiss (1) entails all propositions that have the following form:

(MID) $\Pr(\text{I now am in the middle } p \text{ of } S\text{'s lifetime}) = p.$

The propositions collected together by (MID) say that my confidence that I now am in the “middle” of S’s lifetime increases as I adopt a looser definition of “middling.” Although Gott sometimes uses (MID) with p set at 0.5 to predict how long S will last, he usually uses $p = 0.95$; since the latter value makes for a stronger prediction, I’ll focus on $p = 0.95$ in reconstructing his argument.

Suppose that we now are in the year 2000 and that the S in question has already lasted for 50 years. If we set $p = 0.95$, then (MID) encompasses the two extreme cases depicted in Figure 1, with all the others falling in between. In both scenarios, the present year, 2000, is included in the 95% time interval that lies in the middle of S’s duration. Scenario (A) situates the year 2000 right at the beginning of this middle period, whereas (B) places it at the very end. According to (A), S will last an additional 1950 years; according to (B), S will live another 1.3 years. Thus (MID) says that there is a probability of 0.95 that S will enjoy between 1.3 and 1950 years of additional life.

Figure 1

Although Gott’s Line generates a prediction by using (MID) with p set at a single value, the fact of the matter is that his sampling assumption has many additional consequences. For example, premiss (1) also entails

(FIRST) $\Pr(\text{I now am in the first } p \text{ of } S\text{'s lifetime}) = p$

(LAST) $\Pr(\text{I now am in the last } p \text{ of } S\text{'s lifetime}) = p.$

When these different implications of premiss (1) are , predictions are obtained that are much more specific than the one illustrated in Figure 1. For example, Mid-90 (the proposition obtained from MID by setting $p=90\%$) asserts that there is a 90% chance that S will last between 2.6 and 950 years, Mid-40 says that there is a 40% chance that S will last between 17 and 150 more years, Mid-20 says that there is a 20% chance that S will last between 33 and 75 more years, and Mid-2 entails that there is a 2% chance that S will last between 49 and 51 additional years (all on the assumption that S has so far been around for 50 years). If we put this information together, we obtain the rather more detailed distribution represented in Figure 2.¹

Figure 2

When Gott applies his method to the Soviet Union and the Berlin Wall, he calculates a single temporal interval, which his random sampling assumption entails has a 95% probability of including the

time at which the object in question expires. The temporal interval is wide and the object's duration in these two cases turns out to fall within that wide interval. This is like testing the hypothesis that a coin is fair by seeing if it produces between 1% and 99% heads in a hundred tosses. The hypothesis make that prediction, and the fact that the coin behaves as predicted is of significance. But the prediction is very weak and unspecific, and the confirmation this confers on the hypothesis is correspondingly meager. A better test of the hypothesis would look at the hypothesis' more specific predictions.

So the question we need to ask is this: Is it plausible to apply the specific distribution entailed by Gott's sampling assumption to objects like the Soviet Union, the Berlin Wall, and the human race? I suggest that empirical evidence can be mustered against these applications. Let's begin with the Soviet Union, which was about 50 years old in 1969. Do approximately one in forty of the repressive regimes that have lasted for 50 years go extinct within the next 1.3 years? Do one in forty last for more than 1950 additional years? My bet is that the frequency of the second of these events is far less than one in forty. Similar doubts attach to Gott's discussion of the Berlin Wall. The Wall was eight years old in 1969 when Gott made his prediction. What Gott should have considered is not just the temporal interval generated by setting $p = 0.95$ in (MID), but the full distribution that his sampling assumption entails. One then can see whether observed frequencies conform with this distribution. Some details of that distribution are represented in Figure 3. I would be inclined to assimilate the Berlin Wall not to walls in general (though I bet that that would lead to problems of its own), but to barriers that governments erect to physically separate populations that are part of the same culture. The Berlin Wall was like a law prohibiting travel. Do one in forty of such eight-year old barriers last between 152 and 312 years? Do one in forty last more than 312 years? I doubt that these barriers exhibit the distribution of longevities that Gott's sampling assumption predicts.

Figure 3

What about Gott's prediction concerning the human race? As far as I know, we don't have much frequency data on species that have existed for a mere 200,000 years. However, Gott's model has testable implications about the process of species extinction; it says that an older species that is around today will probably take longer to go extinct than a younger one will.² It is interesting to compare this claim with the evidence that Van Valen (1973) assembles in support of the hypothesis that, within various taxonomic groups, a genus or family's probability of extinction is of its prior duration. I concede that Van Valen discusses genera and families, not species, and that his comparisons are carried out within taxa, not across them. But surely the point is clear that there can be no model of species extinction.

How should we think about the relationship between a species' prior duration and its expected time until extinction?³ Here is an intuitive picture. Let us imagine that a species is "trying" to track its changing environment. For the sake of a simple example, consider a bear species whose survival depends just on its ability to evolve longer fur as the environment grows colder. If the species is able to

evolve faster than the rate at which the environment changes, the species may be able to improve its chances of surviving as it evolves. However, if the environment changes faster than the species can evolve, the species will lag further and further behind, and its risk of extinction will increase as it evolves. Van Valen's hypothesis is that most species are in between these two cases – adaptive evolution manages to keep species “even” with their changing environments. Van Valen calls this conjecture “the Red Queen Hypothesis,” naming it for the character in

who has to run as fast as possible just to stay in the same place. Gott's Line entails that older species will probably outsurvive younger species; he grants that his argument is and might be over-ruled by empirical evidence. However, in the absence of data, we are told to follow Gott's Line. I'd expect most biologists to say something different – in the absence of data, you should go out and get some.

It is interesting to consider more generally the relationship between a system's prior duration and its expected time until extinction. As just noted, Gott's sampling assumption entails that these two quantities are related. In contrast, the Red Queen Hypothesis says that they are , if the system in question is a biological species. A third possibility is exemplified by the individual organisms that belong to the same species – they'll probably die sooner, the older they are now; within species, an organism's prior duration and its expected number of future years are correlated. Gott's claim is that in the absence of evidence to the contrary, we should assume that the system at hand differs both from Van Valen's species and from what we know about individual organisms. My point is not that we should endorse the Van Valen picture or the organism picture , but that claims about the relationship of prior duration to longevity must be judged empirically.⁴

Although Gott asserts that his sampling assumption is justified by what he calls “the Copernican Principle,” he nonetheless seems to agree implicitly that empirical testing is relevant, in that he goes to the trouble of offering empirical evidence that his sampling assumption generates reliable predictions. The problem we have seen is that the tests he considers are . Stronger tests throw doubt on the sampling assumption when it is applied to the Soviet Union, the Berlin Wall, and the human race.^{5,6}

3. L

John Leslie (1990) presents a different version of the doomsday argument:

One might at first expect the human race to survive, no doubt in evolutionarily much modified form, for millions or even billions of years, perhaps just on Earth but, more plausibly, in huge colonies scattered through the galaxy and maybe even through many galaxies. Contemplating the entire history of the race – future as well as past history – I should in that case see myself as human. I might well be among the first 0.001 per cent to live their lives. But what if the race is instead about to die out? I am then a fairly typical human. Recent population growth has been so rapid that, of all human lives lived so far, anything up to about

30 per cent ... are lives which are being lived at this very moment. Now,

To promote the reasonable aim of making it that I exist where I do in human history, let me therefore assume that the human race will rapidly die out (pp. 65-66, emphasis his).

I call this argument “Leslie’s Wedge” because the argument looks with favor on the hypothesis that the human race started increasing dramatically sometime in the past and will continue to do so until a time in the near future when it abruptly goes extinct. If this growing census size is plotted against time, the history of the human race takes the form of a wedge -- thin on the left and rapidly getting thicker as one moves to the right, until suddenly (and soon!), it drops off to nothing.

Leslie’s Wedge differs from Gott’s Line in several respects. First, Leslie specifies no minimum or maximum; he merely provides a qualitative comparison of the hypothesis “doomsday soon” and the hypothesis “doomsday in the distant future.” Second, Gott wants to situate the present moment in the middle of S’s career, whereas Leslie’s favored hypothesis locates it near the end. Third, the evidence on which Gott bases his prediction concerns how long the human race has lasted to date; for Leslie, the relevant observation is simply that I am now alive (as we will see, the past is irrelevant to Leslie’s argument). But perhaps the most fundamental difference is this -- Leslie’s is not an argument about what will happen. It is an argument about .

Here I use the term “likelihood” in the technical sense introduced by R.A. Fisher. The likelihood of a hypothesis H , relative to observations O , is the probability that H confers on O , not the probability that O confers on H . H ’s likelihood is $\Pr(O * H)$, not $\Pr(H * O)$. The fundamental role that likelihood plays in evaluating the testimony of evidence is encapsulated in the Likelihood Principle (Hacking 1965, Edwards 1972, Royall 1997, Sober 1999, Forster and Sober 2003):

Observation O supports hypothesis H_1 more than O supports hypothesis H_2 (i.e., O H_1 over H_2) if and only if $\Pr(O * H_1) > \Pr(O * H_2)$.

Within a Bayesian setting, likelihood is the vehicle whereby new evidence is used in updating one’s prior assignment of probabilities. It follows from Bayes’ theorem that the posterior probabilities, likelihoods, and priors of the competing hypotheses H_1 and H_2 are related as follows:

$$\frac{\Pr(H_1 * O)}{\Pr(H_2 * O)} = \frac{\Pr(O * H_1) \Pr(H_1)}{\Pr(O * H_2) \Pr(H_2)} .$$

If H_1 has a higher likelihood than H_2 , then the ratio of posterior probabilities will be larger than the ratio of priors.

One reason I interpret Leslie's Wedge as a claim about likelihoods derives from the explanation he provides of his argument. He says that his reasoning is the same as the reasoning that figures in other examples that he describes, including the following two "urn stories" (pp. 68-69):

You draw a ball from an urn containing many balls, replace the ball in the urn, and shake the urn before drawing again. After several such draws, you notice that the same ball has been drawn each time. This result favors the hypothesis that trickery is at work or that the drawn ball is especially light-weight over the hypothesis that the balls are drawn at random.

One of the balls in an urn has your name on it. Balls are drawn without replacement and the seventh ball drawn is 'yours.' This result favors the hypothesis that the urn contained just twenty balls over the hypothesis that it contained a thousand.

My other reason for thinking that Leslie's Wedge draws a conclusion about likelihoods is his comment that his argument, "strictly speaking, ... is only for a in any estimate of the risk of our race's imminent extinction (p. 70)." Leslie (1996) is perfectly clear on this point. Whatever probability you assigned to the doomsday soon hypothesis before you considered Leslie's argument, you should revise that probability (see also Bartha and Hitchcock 1998).

If Leslie's Wedge is a likelihood argument, how does he manage to show that

- (I) $\Pr(I \text{ exist in the year 2000} * \text{doomsday soon}) > \Pr(I \text{ exist in the year 2000} * \text{doomsday in the distant future})?$

Leslie obtains this inequality by making the following sampling assumption:

- (A) My now existing in the year 2000 can be treated as if it were the result of a random sample drawn from the set of all human beings who ever exist.

It is useful to visualize Leslie's idea in terms of the time line shown in Figure 4. The probability of my being one of the first n people, if doomsday comes soon and (A) is correct, is $n/(n + m_1)$. The probability of my being one of the first n people, if doomsday occurs far in the future and (A) is true, is $n/(n + m_1 + m_2)$. The likelihood advantage that the hypothesis of "doomsday soon" enjoys would remain in place if, instead of calculating my probability of being one of the first n people, I considered my probability of having precisely the birth order position I have in the sequence of people who constitute the human race. Here the relevant inequality is $1/(n + m_1) > 1/(n + m_1 + m_2)$. Notice that the human population's past history is not relevant to these inequalities. Regardless of the past pattern of population growth – whether it is a wedge or a brick – proposition (I) must be true if (A) is. In fact, Leslie's Wedge leads to a stronger conclusion. Although Leslie usually compares "doomsday soon" with "doomsday in the distant future," his argument entails

(I*) $\Pr(I \text{ exist now} * \text{doomsday occurs } x \text{ years from now}) > \Pr(I \text{ exist now} * \text{doomsday occurs } y \text{ years from now}), \text{ for all } x < y.$

The sooner doomsday comes the better, as far as likelihood is concerned.⁷

Figure 4

Given this reconstruction of Leslie's Wedge, it is clear that his sampling assumption (A) is incompatible with Gott's sampling assumption (1), unless it happens that people's birthdates are uniformly distributed over the duration of the human race. This difference should alert us to the danger already noted of using the Principle of Indifference. If our temporal location is to be thought of as if it were the result of a random draw from a uniform distribution, uniform distribution should we use? We could draw from time intervals of equal duration, assuming that they are equiprobable, or we could draw from the list of birthdates occupied by human beings past present and future, on the assumption that these are equiprobable. I set this point aside now; since I don't think that Gott's sampling assumption is plausible, I can hardly hold it against Leslie that he makes a different assumption.

Two observations help place Leslie's claim about likelihoods in its proper context. The first is that thoroughly preposterous hypotheses can have high likelihoods. For example, if I hear noises in my attic, the hypothesis that there are gremlins bowling up there has a likelihood of unity, but few of us would say that this hypothesis is very probable. The second is that the probability shift justified by Leslie's argument may be tiny. However, these two do not undermine Leslie's thesis – that doomsday soon is more likely than doomsday in the distant future, given the observation that I now exist. This is equivalent to the claim that our prior probability for doomsday's coming soon should be updated; the likelihood argument entails that we should increase the probability we assign to that hypothesis.

4. O L

As the reader will no doubt expect from my discussion of Gott's Line, I think that Leslie's sampling assumption (A) is false. My present temporal position is the result of random sampling from the temporal locations of all human beings. And if this sampling assumption is false, why should I assign probabilities as if it were true? If the Principle of Indifference is offered to you as a reason,

Leslie (1996, p. 209) addresses a similar objection, which he formulates as follows: "... urn analogies are inappropriate. We weren't given our birth times by a deity who pulled our souls from an urn at successive seconds and put them into human bodies." Here is Leslie's reply:

Urn analogies are relevant to many statistical calculations. For example, (1) Jim and Mike drive cars equally frequently in the same city. Jim has had twenty accidents and Mike not a single one. Are they equally good drivers? Consider an urn filled with two balls, one marked 'Jim' and the other 'Mike.' Balls are drawn again and again. In each case the drawn ball is put back in the urn, which is then shaken. Would it be likely that the 'Jim' ball would be drawn every single time? (2) You are hit by an arrow while walking around on a small island. Was this bad luck, or was the arrow aimed at you? If only luck was involved then this would be (as a rough approximation) as if the grid references of every square foot of the island had been put on slips of paper in an urn, your name being written on just one of them, etc.

Leslie's reply allows me to clarify my objection. I have nothing against urn models, In fact, some of my best friends are urn models. Rather, my claim is that they are appropriate when and only when there is an underlying chance process that one is modeling. This requirement is satisfied in Leslie's two examples. We imagine that Jim's propensity to have accidents is stronger than Mike's. These two propensities correspond to different probabilities. Each man's probability characterizes the process whereby his driving a car puts him at risk of having an accident. In the island example, we imagine an archer who has one of two propensities (shoot at random or aim at me), and we are trying to figure out which. These two possible states of the archer correspond to two different probability models of the process by which he or she sends arrows into the air. Jim and Mike have their propensities, and archers have theirs. 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.

I now come to my second objection to Leslie's Wedge. Not only do I think that Leslie has no justification for the probability assignment he uses. In addition, I think that this probability assignment is in fact false, and that my present temporal location has no evidential relevance to estimating when doomsday will occur. Of course, it is not to be denied that my existing in the year 2000 and the fact that I am a human being together entail that the human race cannot go extinct 2000. But notice that Leslie's likelihood inequalities (I) and (I*) go beyond that modest point. It is these more substantive claims that I think are empirically disconfirmed.

To develop this point, let's consider what the sampling assumption (A) entails about the time at which I'll exist if doomsday occurs soon, assuming that I exist during S's lifetime. Since my temporal location is to be assigned randomly, it could be anywhere. However, its is in the middle of S's duration; this is represented in Figure 5 by E_s .⁸ In contrast, if doomsday occurs far in the future, my expected temporal location is E_f , also represented in Figure 5. Notice that E_s is earlier than E_f .

Figure 5

This means that Leslie's sampling assumption is subject to the following test. What we need is an ensemble of species. All of them exist right now⁹ and have been around for about

200,000 years. They perhaps vary in their ecological circumstances, now and in the future, and they also may vary with respect to their present and future phenotypic and genetic characteristics. Let us suppose that I have belonged to each of these species at some time or other in the past. The point represented in Figure 5 is that Leslie's sampling assumption predicts that the doomsdays of these different species should be with the times at which I happen to belong in them. The species I belong to sooner should (on average) go extinct earlier.

Of course, I am a member of just one species, and there is just one of me, so this test cannot be conducted. However, there are analogous situations in which analogous tests can be run. Although I belong to just one species, I have joined different organizations at different times in my life. In childhood, I joined the American Numismatic Association. As a teenager, I was in an orchestra. When I first came to Wisconsin, I joined a sailing club. Now, in the wisdom of my maturity, I am a card-carrying member of the American Philosophical Association. Suppose these organizations all began at the same time and that each is still around now. When will each exit from the scene? Leslie's likelihood inequality, applied to these organizations, predicts a correlation. These organizations should go extinct sooner, the earlier in my life I joined them. If Leslie is right, this is bad news for the ANA, but good news for the APA.

We don't have to wait for the demise of the ANA or the APA to assess what the outcome of this experiment will be. When people join one organization before they join another, and both exist now, is there a tendency for the first organization joined to be the first to go extinct? There is no reason to restrict this question to joining organizations. It applies equally to businesses that one patronizes, cities that one visits, and magazines that one reads. Surely the general pattern is that if I bear relation R to x before I bear R to y, where x and y both exist now and are equally old, there is no general tendency for x to go extinct before y.

Although I think the general pattern is as just described, this does not mean that correlations of the kind we are considering are impossible. For example if my joining an organization (or visiting a city, or patronizing a business, or ...) caused a time bomb to start ticking that eventually goes off and blasts that organization to oblivion (or if my joining and the planting of the time bomb had a common cause), then a correlation of the sort predicted by Leslie's Wedge would exist.¹⁰ It doesn't matter whether the bomb's time until detonation is determined or is subject to stochastic variation around a mean value. As long as different organizations are subject to time bombs of the same expected duration, the order in which I join provides evidence about the order in which they will go extinct. Correlations induced by the triggering of time bombs are not just conceivable; they actually exist. Individuals with AIDS start time bombs ticking in others when they have unprotected sex with them. If an infected individual has unprotected sex with x before he or she has unprotected sex with y, this is evidence (weak and defeasible though it may be) that x will die before y does.

Leslie's as-if sampling assumption predicts that there should be a correlation between when I exist and when the human race will go extinct. This correlation is to be expected if my birth sets a time

bomb ticking that eventually blasts the human race to oblivion, or if the birth and the extinction are joint effects of a common cause. There is, however, no evidence at all for either of these arrangements. A far more plausible hypothesis is that the date of doomsday is _____ of when I exist (except for the obvious fact that my existing now places a minimum date on the human race's day of doom).

5. C C

Gott and Leslie both use assumptions about random sampling that they think are reasonable on grounds. Gott uses his sampling assumption to associate _____ with the different longevities that the human race (and other objects as well) might have. Leslie uses his sampling assumption to compare the _____ of different hypotheses about when the human race will draw to a close. I have argued that both sampling assumptions make testable predictions. For example, Gott's assumption predicts that there should be a correlation between a species' prior duration and its number of future years until extinction; Leslie's predicts that there should be a correlation between when I joined an organization (or visited a city, or subscribed to a magazine ...) that still exists and when in the future that organization (or city or magazine ...) goes extinct. These and other empirical predictions that the models generate show that the sampling assumptions are not justifiable _____. In addition, there appears to be substantial evidence that the models are false, at least for the examples considered here. This does not rule out the possibility that the sampling assumptions might fare better when applied to other examples. But whether this is so is an empirical question, and that is the main point.

A

I am grateful to Nick Bostrom, Branden Fitelson, John Leslie, Brad Monton, Norma Sober and the referees of this journal for useful comments, and I am _____ grateful to Ellery Eells, Martin Barrett, and Christopher Hitchcock for their detailed comments on previous drafts.

Bartha, P. and Hitchcock, C.: 1998, 'No One Knows the Date or the Hour – an Unorthodox Application of Rev. Bayes's Theorem', _____ (Proceedings) **68**, S339-S353.

Bostrom, N.: 1997, 'Investigations into the Doomsday Argument', available at the following URL: <http://www.anthropics-principles.com/preprints/inv/investigations.html>.

Bostrom, N.: 2000, _____ . Ph.D. Dissertation, London School of Economics and Political Science, available at the following URL: <http://www.anthropics-principles.com/phd/>.

Caves, C.: 2000, 'Predicting Future Duration from Present Age – A Critical Assessment',

(January 24, 2000)

- Edwards, A.: 1972, *Bayesian Epistemology*, Cambridge University Press, Cambridge.
- Forster, M. and Sober, E.: 1994, 'How to Tell Which Theory is More Unified, or Less Unified, or Less Accurate', *Philosophy of Science* **61**, 1-36.
- Forster, M. and Sober, E.: 2003, 'Which Theory is More Unified?', in J. A. Cover and S. Lee (eds.), *Philosophy of Science*, Cambridge University Press.
- Goldman, S.: 1994, 'Future Theories', *Philosophy of Science* **61**, 106-107.
- Gold, R.: 1993, 'Implications of the Copernican Principle for our Future Prospects', *Philosophy of Science* **60**, 315-329.
- Gold, R.: 1994, 'Theories', *Philosophy of Science* **61**, 106-107.
- Gold, R.: 1996, 'Our Future in the Universe', In V. Trimbur and A. Reisenegger (eds.), *The Future of the Universe*, Springer-Verlag, New York.
- Gold, R.: 1997, 'The Copernican Principle', *Philosophy of Science* **64**, 140-151.
- Gold, R.: 1999, 'A Grim Reckoning', *Philosophy of Science* **66**, 21-39.
- Hacking, I.: 1965, *Logic of Statistical Inference*, Cambridge University Press, Cambridge.
- Leslie, J.: 1997, 'Is the End of the World Nigh?', *Philosophy of Science* **64**, 65-72.
- Leslie, J.: 1998, *The End of the World Nigh*, Routledge, London.
- Reichenbach, H.: 1938, *Experience and Prediction*, University of California Press, Berkeley.
- Royal, R.: 1907, *The Philosophy of Science*, Chapman and Hall, London.
- Sober, E.: 1975, 'The Principle of the Common Cause', In J. Fetzer (ed.) *Philosophy of Science*, Reidel, Dordrecht, pp. 211-28. Reprinted in *Philosophy of Science*, Cambridge University Press, Cambridge.

Sober, E.: 1999, 'The Principle of the Common Cause', available at the following URL:

Sober, E.: 1999, 'The Principle of the Common Cause', *Philosophy of Science* **66**, 1-16.

Sober, E.: 1999, 'A New Evolutionary Law', *Philosophy of Science* **66**, 1-30.

Philosophy
Madison, Wisconsin 53706
USA

Figure 1: If S is now fifty years old and the present year (assumed to be the year 2000) falls within the middle 95% of S's lifetime, then S will last somewhere between (A) 1950 and (B) 1.3 more years.

Figure 2: Gott's model entails the probabilities shown here for the different possible numbers of additional years that S might enjoy, if S began 50 years ago.

Figure 3: Gott's model entails the probabilities shown here for the different possible numbers of additional years that S might enjoy, if S began 8 years ago.

Figure 4: If n individuals have lived until now, then there will be a total of $(n+m_1)$ human beings if doomsday comes soon, and a total of $(n+m_1+m_2)$ if doomsday arrives much later.

Figure 5: If doomsday occurs soon, the expected value of my temporal location is E_s ; if doomsday occurs far in the future, the expected value is E_f .

Figure 1


(A)

D	1950	2000	3950
L	9	9	9
	[--2.5%--][-----95%-----][--2.5%--]		

(B)

D	1950	2000	2001.3
L	9	9	9
	[--2.5%--][-----95%-----][--2.5%--]		

Figure 2



N

1. Martin Barrett (pers. comm.) has derived the full distribution that Gott's sampling assumption entails. Suppose that S begins at time 0 and ends at time D (thus D is both the date of doomsday and the number of years that S exists), and that N is the present date (where $0 < N < D$). Let u be the percentage that marks the beginning of a percentage interval of S 's lifetime and v be the percentage that marks the end of that interval, thus:

0%.... u v100%

The corresponding interval of S 's life measured in years is $[uD, vD]$. For example, with $p = 0.95$, (First) is equivalent to setting $u=0$ and $v = 0.95$, (Mid) is equivalent to setting $u = 0.025$ and $v = 0.975$, and (Last) is equivalent to $u = 0.05$ and $v = 1.0$. In each case, $p=v-u$. Gott's sampling assumption is given by

$$\Pr(uD < N < vD) = p = v-u \quad (\text{for all } v > u).$$

This is equivalent to

$$\Pr(N/v < D < N/u) = p = v-u \quad (\text{for all } v > u),$$

which provides the full distribution.

2. Notice that in both Figure 2 and Figure 3, the median number of additional years is identical with the system's prior duration.

3. Here and in what follows, my talk of the "expected value" of a quantity refers to the mathematical expectation – the average value that would arise under repeated trials.

4. In other scientific problems, the default assumption is often that two quantities, whose relationship is unknown, are independent. The rationale for this policy is worth pondering, and the statistical literature on model selection is relevant. The Akaike information criterion describes a general context in which a model that says that two quantities are independent will have a higher estimated degree of

than a model that says that they are dependent, if the two models fit the data about equally well. See Forster and Sober (1994) for discussion.

5. Are there systems in which prior duration and subsequent longevity are positively correlated?

Gott (1996) enumerates the 44 plays and musicals that were on Broadway at the time of his 1993 paper in , applying Mid-90s to each. An anonymous referee has pointed out to me that 38 of

these had closed by the start of 2002. We therefore can compute how often Gott's time of observation, in 1993, fell in the first 5% of a play's run, in the second 5%, and so on. The referee provides the following count:

3,1,2,0,3,1,2,2,3,3,2,1,2,0,2,2,3,2,4,0.

Perhaps Broadway shows, unlike Van Valen's species or conspecific organisms, are such that the longer they've been around, the longer they are apt to last.

6. Martin Barrett (pers. comm.) has pointed out to me that Gott's calculations impose no upper bound on how long the system in question might last. What happens if we assume that the system has, say, at most another trillion years to go? One suggestion would be to give equal probabilities to equal time intervals. However, this has the result that our present temporal location is special – we are assuming that we live at the very beginning of the system's lifetime. Alternatively, we might assume that, in expectation, we are in the middle of the system's lifetime. This would lead to a specific distribution different from the ones depicted in Figures 2 and 3.

7. It is worth noting that Leslie's argument assumes that the more people there are in the future, the further away doomsday probably is. Although there is some plausibility to this assumption, it is easy to imagine possible counterexamples. For the sake of argument, I grant the assumption in what follows.

8. Recall that the expected value is the average value that would be exhibited if my temporal location were chosen repeatedly. It is not the specific value we should expect; the randomness assumption entails that each time has the same probability as every other.

9. The reason I require that all these species exist now is that I want to control for the obvious fact mentioned above – that if I belong to a species at time t , the species can't go extinct until after t . Since

Leslie's likelihood inequalities (I) and (I*) go beyond this triviality, I want the experiment to control for its effect.

10. Here I am using ideas consonant with Reichenbach's (1956)

See

Sober (1987, 2001) for discussion.