GOTT'S DOOMSDAY ARGUMENT

by Bradley Monton and Sherrilyn Roush

University of Kentucky and Rice University

bmonton@uky.edu roush@rice.edu

11/20/2001

Acknowledgments

We are equally responsible for this paper; our listed order is the consequence of a coin toss. We thank Nick Bostrom, J. Richard Gott III, Cricky Hitchcock, Brian Kierland, Elliott Sober, and Bas van Fraassen for helpful comments.

GOTT'S DOOMSDAY ARGUMENT

Abstract

Physicist J. Richard Gott uses the Copernican principle that "we are not special" to make predictions about the future lifetime of the human race, based on how long the human race has been in existence so far. We show that the predictions which can be derived from Gott's argument are less strong than one might be inclined to believe, that Gott's argument illegitimately assumes that the human race will not last forever, that certain versions of Gott's argument are incompatible with Bayesian conditionalization, and that Gott's argument is self-refuting.

Introduction

In the prominent science journal *Nature* physicist J. Richard Gott III has given a version of the Doomsday argument, concluding that with 95% confidence the future lifetime of our species is less than 7.8 million years, but more than about 5,100 years (Gott 1993a). Gott's argument has been discussed in such places as *The New York Times* (Lerner 1993, Gott 1993b), *The New Yorker* (Ferris 1999), and *Nature* (Goodman 1994, Mackay 1994, Buch 1994, Gott 1994, Landsberg and Dewynne 1997, Gott 1997), but surprisingly has received scant attention from philosophers. We know of only four references to Gott in the philosophy literature. Three of these (Leslie 1996, 16; Eckhardt 1997, 244; Korb and Oliver 1998, 403) do not explain Gott's argument; they simply mention that Gott's argument is a version of the Carter-Leslie Doomsday argument. Only Bostrom (2002) points out the difference between Gott's argument and the

Carter-Leslie argument. As we will show below, Gott's argument is importantly different from the Carter-Leslie argument – at least it is importantly different from the Bayesian version of the Carter-Leslie argument, which is the version recent literature has focused on (Eckhardt 1997; Bartha and Hitchcock 1999; Korb and Oliver 1998, 1999; Bostrom 1999, 2001, 2002).

There are three reasons we believe that the Gott argument is worthy of independent consideration. First, a refutation of one version of the Doomsday argument does not necessarily refute all versions. We are sympathetic to the refutation of the Carter-Leslie argument suggested by Dieks (1992), and independently developed by Kopf et. al. (1994) and Bartha and Hitchcock (1999). This refutation, however, does not carry over to the Gott argument.[1] Second, Gott's argument usefully formalizes the Copernican principle that "we are not special", and hence enables one to discuss this principle with some precision. Third, Gott's argument is not straightforwardly Bayesian, and this helps to bring out differences between Bayesian and non-Bayesian approaches to statistical inference.

While we believe that Gott's argument is worthy of consideration, we do not believe it is correct. Our attitude, however, is much more sympathetic than that of, for example, Eric Lerner, who in an editorial in *The New York Times* called Gott's argument "pseudo-science" and asked:

> Why would a prestigious journal like Nature publish such astrology and why would a prominent cosmologist, who presumably knows better, write it? (Lerner 1993)

We think that Gott's argument is to be commended for a higher degree of explicitness and precision than those of some other doomsayers. While Gott's argument is mistaken, the reasons it is mistaken are important and illuminating.

Gott's Argument

Gott begins what he calls the 'delta t' argument as follows, where t_{begin} is the time at which the phenomenon whose lifetime we are interested in begins, and t_{end} the time when it ends:

> Assuming that whatever we are measuring can be observed only in the interval between times t_{begin} and t_{end}, if there is nothing special about t_{now} we expect t_{now} to be located randomly in this interval. (Gott 1993a, 315)

This is an application of the so-called Copernican principle that we do not occupy a privileged place in the universe. Gott's argument continues:

> If $r_1 = (t_{\text{now}} - t_{\text{begin}})/(t_{\text{end}} - t_{\text{begin}})$ is a random number uniformly distributed between 0 and 1, there is a probability $P = 0.95$ that $0.025 < r_1 < 0.975$. (Gott 1993a, 315)

Letting $t_{\text{future}} = t_{\text{end}} - t_{\text{now}}$ and $t_{\text{past}} = t_{\text{now}} - t_{\text{begin}}$, it takes just a few lines of math[2] to show that the consequent of the statement above is equivalent to

$$(1/39)\, t_{\text{past}} < t_{\text{future}} < 39\, t_{\text{past}} \text{ (with 95\% confidence).} \qquad (1)$$

Gott applies this equation to situations where we do not have actual data about the longevity of what we are measuring. For example, he first visited the Berlin Wall in 1969, so that $t_{\text{past}} = 8$ years. He assumed that he was a random observer of the Wall, so he expected to be located randomly between t_{begin} and t_{end}.

> The Wall fell 20 years later giving $t_{\text{future}} = 2.5\, t_{\text{past}}$, within the 95% confidence limits predicted by equation (1). (Gott 1993a, 315)

Gott also applies equation (1) to estimate the longevity of *Homo sapiens*. Our species is roughly 200,000 years old, so letting $t_{\text{past}} = 200{,}000$ years, equation (1) gives (to the nearest year):

$$5{,}128 \text{ years} < t_{\text{future}} < 7{,}800{,}000 \text{ years (with 95\% confidence).}$$

This is the sense in which Gott's argument is a Doomsday argument; you might have assigned probability higher than 0.025 to the proposition that the human race will last longer than another 7.8 million years.[3]

Objection 1: High Confidence Cannot Lead to Belief

For our first objection, we will grant that the premises of Gott's argument are correct, but show that the conclusions that can be drawn from the argument are weaker than they appear. Specifically, we will show that, as long as the confidence level for what interval t_{future} is located in is less than 100%, then on pain of incoherence one cannot *believe* that t_{future} is located in that interval using the grounds that the confidence level for that interval is high.

Suppose for reductio that one does decide to believe on the basis of Gott's argument that $(1/39)\ t_{\mathrm{past}} < t_{\mathrm{future}} < 39\ t_{\mathrm{past}}$, for a particular process under consideration. This requires believing that it is not the case that $0\ ?\ t_{\mathrm{future}}\ ?\ (1/39)\ t_{\mathrm{past}}$, and nor is it the case that $39\ t_{\mathrm{past}}\ ?\ t_{\mathrm{future}}\ ?\ ?$. The incoherence arises because there are *other* intervals which also include t_{future} with 95% confidence. For example, it follows from a Gott-style argument that, with 95% confidence,

$$0\ ?\ t_{\mathrm{future}} < (19/21)\ t_{\mathrm{past}} \text{ or } (21/19)\ t_{\mathrm{past}} < t_{\mathrm{future}}\ ?\ ?. \qquad (2)$$

This simply rules out the middle 5% interval centered around $t_{\mathrm{future}} = t_{\mathrm{past}}$, instead of ruling out the two 2.5% extreme ends as Gott does. But if one chooses to believe $(1/39)\ t_{\mathrm{past}} < t_{\mathrm{future}} < 39\ t_{\mathrm{past}}$ on the basis that the confidence level is 95%, one should also believe (2) on that same basis. Thus, one would have to believe that it is not the case that

$$(19/21)\ t_{\mathrm{past}}\ ?\ t_{\mathrm{future}}\ ?\ (21/19)\ t_{\mathrm{past}}.$$

Similar arguments show that, by this line of reasoning, one will believe it is not the case that

$$(1/39)\ t_{\mathrm{past}}\ ?\ t_{\mathrm{future}}\ ?\ (3/37)\ t_{\mathrm{past}}, \text{ nor } (3/37)\ t_{\mathrm{past}}\ ?\ t_{\mathrm{future}}\ ?\ (5/35)\ t_{\mathrm{past}},$$

$$\text{nor } (5/35)\ t_{\mathrm{past}}\ ?\ t_{\mathrm{future}}\ ?\ (7/33)\ t_{\mathrm{past}}, \ldots, \text{ nor } (37/3)\ t_{\mathrm{past}}\ ?\ t_{\mathrm{future}}\ ?\ 39\ t_{\mathrm{past}}.$$

An incoherence arises because (granting Gott's assumptions) one believes that t_{future} has some finite value, and yet the kind of ground Gott gives us for believing t_{future} is not greater than 7.8 million or less than 5,128 years also gives grounds for believing it is not any *other* particular value. This is formally the same as the lottery paradox, where one believes that some lottery ticket is the winner, but for each ticket one does not believe that that ticket is the winner.

Note that this argument for incoherence of belief applies regardless of what confidence level is involved, as long as the confidence level is not 100%.[4] For example, for a 99.99% confidence level, one will believe that it is not the case that

$0 \; ? \; t_{future} \; ? \; (1/19{,}999)\, t_{past}$, nor $19{,}999\, t_{past} \; ? \; t_{future} \; ? \; ?$.

By the above reasoning, one will also believe that it is not the case that

$(1/19{,}999)\, t_{past} \; ? \; t_{future} \; ? \; (3/19{,}997)\, t_{past}$,

... , nor $(9{,}999/10{,}001)\, t_{past} \; ? \; t_{future} \; ? \; (10{,}001/9{,}999)\, t_{past}$,

... , nor $(19{,}997/3)\, t_{past} \; ? \; t_{future} \; ? \; 19{,}999\, t_{past}$,

and thus the same incoherence arises.

We should point out that, as far as we know, Gott never says that one should *believe* that t_{future} is located in the interval given by his argument. Nevertheless, he does not point out that his argument applies equally to all the various types of intervals considered in this section. The objection of this section shows that Gott's argument is weaker than one might otherwise be tempted to suppose.

Objection 2: The Human Race Need Not End

It is reasonable to assign a non-zero probability to the proposition that the human race will last forever. For one (admittedly speculative) proposal for how the human race could last forever, see Tipler 1994; for general discussion see Cirkovic and Bostrom 2000. If you are unwilling to bet your life that the human race will some day end in return for a dollar, then this suggests that you assign a non-zero probability to the proposition that the human race lasts forever.

Gott's argument entails that the probability that the human race will last forever is zero (or infinitesimal, if one uses non-standard measure theory). Gott's argument allows one to increase the confidence level that t_{future} is in the given range by increasing the interval that contains r_1. For example, there is a probability $P = 0.99$ that $0.005 < r_1 < 0.995$. If the human race lasts forever, then $r_1 = 0$. As long as P ? 1, P can be made arbitrarily large without having the interval which contains r_1 include 0. Thus, Gott's argument entails that the probability that the human race lasts forever is zero (or an infinitesimal). Since Gott's argument can also be used for intelligent life in general, his argument also entails that the probability that intelligent life lasts forever is zero (or an infinitesimal).

One might think that this means Gott's argument gives us reason to consider revising our assignment of a non-zero probability to the proposition that the human race will last forever. However, this impression is unwarranted, because the assignment of zero or infinitesimal probability to the proposition that the human race lasts forever is actually a *premise* of the argument, which does not follow from the Copernican principle and is, in our view, unjustified. The way Gott makes definite the thesis that our temporal location is not special is by maintaining that r_1 is randomly distributed between 0 and 1.[5] This means that he is treating the $r_1 = 0$ case no differently than the case where r_1 equals some other number; the proposition that the human race lasts forever is treated no differently than the proposition that the human race will end exactly 17,000 seconds from now. The claim that one should assign equal probability to these two cases implicitly involves a partitioning of the possibility space that could be done otherwise, and to different effect. Consider the following alternative partition: assign equal probabilities to the possibility that our species lasts forever and the possibility that it has an end in time. Then the probability that our species lasts forever is 0.5 rather than zero, and the probability that its demise occurs between 5,128 and 7.8 million years from now is (0.95)(0.5) = 0.475 rather than Gott's 0.95. It is well known that in calculating probabilities by applying a principle of indifference, as Gott does, one's results can depend sensitively on how the possibility space is divided into equally probable possibilities. Admittedly, variant partitions can be generated at will, but many of these will be perverse or unmotivated. Our variant partition is simple, and is motivated by the fact that finite and infinite age are two exclusive and exhaustive cases which, it could be argued, should be treated the same.

To be fair, Gott (1993a) does give arguments for the claim that the probability that our species will last forever is zero, but these arguments occur after and independently of the delta t argument discussed above, and he never acknowledges that the delta t argument depends on these arguments. We will now show that these latter arguments are unsuccessful.

Gott's first defense of the claim that the human species will not last forever is an extension of an argument he gives against the thesis that our species lasts for a long finite lifetime. That argument is as follows. Supposing that the human species is subject to some unknown extinction rate, if that rate is very tiny then the longevity of the species is very great and the total population of human beings who will ever have lived is enormous. Only a very small fraction of that total will find themselves in the first 200,000 years of the existence of the species, meaning that finding ourselves there, as we do, is unlikely. This tells against the extinction rate being so low. Gott then says:

> In the limit where we expect our species to live forever, [the extinction rate] goes to 0 and $P(t_p$?? 200,000 years) goes to 0, (Gott 1993a, 316)

where t_p is t_{past} for our species. Gott concludes that the extinction rate cannot be zero, and thus the human species cannot last forever.

This extension of the argument to the infinite case is incorrect, however, since in the case where the lifetime of our species is infinite *any* particular finite interval in which to find the age of the species will have the same zero probability that the interval 0 to 200,000 years has. That our species is less than 200,000 years old is no more unlikely than that it is less than 700 billion years old, for the case where the total lifetime of the species is infinite. Thus, even though $P(t_p$?? 200,000 years) goes to zero as the extinction rate goes to zero, that does not give evidence that the extinction rate is not zero.

Gott's second defense of the claim that the human species will not last forever consists in submitting that, were the human species to last forever, there would be three telltale signs which do not in fact obtain:

> What we would expect to observe in the limit as [the extinction rate] goes to 0 is that the value of t_0 we observe goes to infinity, t_p ??t_0 and $(t_0 - t_p)$?? t_0, (Gott 1993a, 316)

where t_0 is the present age of the universe.

Focussing on the first telltale sign, Gott is implying that if the extinction rate were zero, so that the human race lasts forever, then we would observe t_0 to be infinite. But this claim is unjustified. The lifetime of our species may be infinite in the forward direction without anything following logically about whether the lifetime of our species or of the universe is infinite in the backward direction. Focussing on the second and third telltale signs, it is true that t_p is far from t_0, relative to how close they will be in the distant future (assuming that the human race lasts that long). But this claim can be made regardless of how close t_p and t_0 are now. Even if the difference between t_p and t_0 were only 0.1%, say, this difference still appears large compared to a time in the distant future when the difference is only 10^{-100}%, say. Thus, the appeal to the second and third telltale signs does not provide a severe test of the hypothesis that the human race will last forever. If the human race will last forever, no matter how large t_p is we would still observe t_p far from t_0, relative to how close they will be in the distant future. The general idea we are appealing to here is that there is no random or 'non-special' place on a continuum which is infinite in one direction but not the other. Every point is a particular finite distance from the starting point, and every point is an infinite distance from the endpoint. Every point is both unique on the one hand and entirely typical on the other.

Objection 3: Gott's Argument is Incompatible with Conditionalization

Gott (1993a) does not rely on Bayesian conditionalization. We will now show that Gott's argument is actually incompatible with Conditionalization – the rule that one should always update one's subjective probabilities using Bayes' Theorem, or the generalization proposed by Jeffrey (1983). Indeed, we will show that Gott's argument is incompatible with the Reflection Principle (van Fraassen 1984, 1995), which is entailed by Conditionalization but does not entail it (van Fraassen 1995, 17).

The Reflection Principle, for situations where people have precise numerical probabilities for their opinions, is (van Fraassen 1995, 19):

$P(A|\, p_t(A) = x) = x$ when defined.

Here P denotes one's current probability function, p_t denotes one's probability function at later time t, and "$p_t(A) = x$" denotes the proposition that at later time t, the probability one assigns to proposition A is x. In the situation where one believes that $p_t(A) = x$, then by Conditionalization one's posterior probability for A, $P^*(A)$, is x. In the situation where one does not believe that $p_t(A) = x$, but instead believes that either $p_t(A) = x$ or $p_t(A) = y$, then

$$P^*(A) = P(A|\, p_t(A) = x)\, P(p_t(A) = x) + P(A|\, p_t(A) = y)\, P(p_t(A) = y)$$

$$= x\, P(p_t(A) = x) + y\, P(p_t(A) = y). \tag{3}$$

Now, suppose as before that t_{past} for the human race is 200,000 years. Gott's argument entails that, to the nearest year,

$$5{,}128 \text{ years} < t_{\text{future}} < 7{,}800{,}000 \text{ years (with 95\% confidence).}$$

Suppose you believe that, unless the human race has ended, five years from now you will believe Gott's argument. Five years from now, t_{past} will be 200,005 years, so Gott's argument will entail that

$$5{,}128 \text{ years} < t_{\text{future}} < 7{,}800{,}195 \text{ years (with 95\% confidence).}$$

Let A be the proposition that the human race will end between the year 7,133 (that is, 2005 + 5,128) and the year 7,802,200 (that is, 2005 + 7,800,195). Five years from now, unless the human race has ended, you will assign probability 0.95 to A. If the human race has already ended five years from now, then $p_5(A) = 0$: if you were, per impossible, around in 2005, you would assign probability 0 to the human race surviving past 2005, since it has already ended. It follows from equation (3) that

$$P^*(A) = 0.95\, P(p_5(A) = 0.95) + 0$$

But $P(p_5(A) = 0.95)$ is just the probability that the human race will not end within five years from now. This probability can be calculated using Gott's argument:

$$P(5 \text{ years} < t_{\text{future}} < 8 \text{ years}) = P((1/40{,}000)\, t_{\text{past}} < t_{\text{future}} < 8\ \ t_{\text{past}})$$

$$= P(0 < r_1 < 0.999\,975\,001)$$

$$= 0.999\ 975\ 001$$

Thus, $P^*(A) = 0.949\ 976\ 251$. It follows that Gott's argument applied five years from now along with the Reflection Principle assigns probability 0.949 976 251 to the following interval considered from now:

$$5{,}133 \text{ years} < t_{\text{future}} < 7{,}800{,}200 \text{ years.} \quad (4)$$

The reason this causes a problem for Gott's argument is that one will make incompatible probability assignments by using both Gott's argument applied now and Gott's argument applied five years from now along with the Reflection Principle. Applying Gott's argument now, there is a probability $P = 0.\ 949\ 976\ 251$ that

$0.025\ 000\ 959 < r_1 < 0.974\ 977\ 210$, or equivalently

$$5{,}133 \text{ years} < t_{\text{future}} < 7{,}799{,}693 \text{ years.} \quad (5)$$

If one were to decrease the upper bound of (4), that would lower the probability assigned to the interval. It follows that Gott's argument applied now assigns a lower probability to the interval (5) than Gott's argument applied five years from now along with the Reflection Principle; the probability assignments are incompatible. Thus, Gott's argument is incompatible with the Reflection Principle, and since Conditionalization entails the Reflection Principle, Gott's argument is incompatible with Conditionalization.

At this stage, of course, there are two options: one could reject Gott's argument, or one could reject Conditionalization. We leave this to the reader's preference. It is worth pointing out, ad hominem, that Gott (1994) uses Bayesian conditionalization, and so would presumably be surprised that his argument is incompatible with Conditionalization.

Before moving on to the next objection, we should point out that there are some versions of Gott's argument to which the objection presented in this section does not apply. Thus, the most that the objection of this section can show is that Gott's argument has restricted applicability. For example, consider the birth rank version of Gott's argument, where instead of considering the number of years the human race has existed, you consider the number of humans who have been born before you. Five years from now your birth rank will be unchanged, so if the predictions for t_{future} are based on birth rank then Gott's argument applied now and Gott's argument applied five years from now make the same probability assignments.

Objection 4: Self-Reference Leads to Refutation

P. T. Landsberg and J. N. Dewynne (1997) attempted to derive a Russell-type paradox by applying Gott's theory of life-span prediction to the theory itself (where by 'Gott's theory' we mean the theory that Gott's argument gives correct results). We will show that the situation is not a paradox but a refutation of Gott's theory.

In 1997, when Gott's theory had gone unrefuted for 51 months, Landsberg and Dewynne derived that with 95% confidence the remaining life-span of Gott's theory would be no more than 165.75 years, according to that theory applied to itself. In other words, with 95% confidence there are (as of the year 2001) 161 years left before Gott's theory must be refuted in order to be verified. Landsberg and Dewynne thought this situation was paradoxical. They granted that Gott was saved from the paradox they announced by the fact that his predictions are always probabilistic, and any outcome is strictly consistent with a probabilistic prediction. However, we submit that even if Gott's predictions were deterministic there would be no paradox, and the probabilistic nature of his predictions does not save his theory from refutation.

There are two possibilities: Gott's theory is refuted within the next 161 years or it is not.

If it is refuted, then one of the predictions of the theory is fulfilled. If it is not refuted, then one of the predictions of the theory is not fulfilled. Both possible outcomes are consistent with the theory being false, and neither is consistent with it being true (except in the second case by dint of the predictions being probabilistic). In the first case the theory is refuted, and this is so regardless of its having gotten it right that it would be. (A false theory can yield true predictions.) In the second case the theory is empirically falsified, with 95% confidence. (A true theory cannot yield false predictions, ignoring the probabilistic nature of the predictions.) The situation has the air but not the structure of the Russell paradox, since here the analogues of true and false—verified and refuted—are not contradictory. A theory may be both refuted (hence false) and verified in some, or many, instances. Here there is a stable solution to the question of what happens when the theory is applied to itself, and it is that the theory is false or probably false.

This argument itself appears to be sufficient to refute Gott's theory. It would clearly be sufficient if the theory were deterministic, since in that case both of the possible outcomes of self-applying the theory would yield the theory's falsehood. The first possible outcome is one in which the theory is refuted, implying that it is false. The second possible outcome, where the theory is not refuted, would yield the falsehood of the theory by having produced a counterexample to its predictions. For a deterministic theory one counterexample is enough for refutation, ceteris paribus. However, for a probabilistic theory such as Gott's theory actually is, one counterexample does not have the same force. We should judge a theory that makes predictions at a 95% confidence level, such as Gott's does, by asking whether it makes the correct predictions 95% of the time. But we can see how the self-application argument can be adapted to the probabilistic context: we can imagine people applying the Gott formula to itself an indefinite number of times in the future, all yielding different predictions about when Gott's theory will be refuted due to having applied the formula at different times in its lifetime. Despite the different predictions, though, all the predictions will yield the two possibilities described above, one in which the theory is definitively false, the other in which a counterexample has been produced. That is, since there are an infinite number of times in the future there are also an infinite number of counterexamples that can be produced to Gott's theory, drowning out the effect of the successful predictions he has made with it thus far. Self-application refutes Gott's theory.[6]

Gott has insisted that his theory was not meant to apply to itself. If his theory does not apply to itself then the argument we just made has no force. However, we think the reason Gott gives for refraining from self-application is ineffectual. He rightly points out that you should not apply his formula to predict the longevity of a particular marriage when you are at the wedding of those people. This is because it is plain that at the wedding you are observing the marriage at a special time, so it does not qualify as a random time. It is appropriate to apply Gott's theory only when you have no reason to think the time you apply it is special in the phenomenon's lifetime. Gott thinks a similar point holds when we apply his theory to itself. He says:

> My paper and papers written by people who were present in 1993, like guests at a wedding, are located by definition at a special place in the history of when my formula will be known. (Gott 2001, 220)

The year 1993, when the formula was written down, does mark the beginning of the 'life' of Gott's theory, and that is, by definition, a special time in its history. However, it is not 1993 anymore, and no time after 1993 is special *by definition*. Let us be clear that the question of whether we may apply Gott's formula depends on whether the time of *application* of the theory is special, not whether the time of inception is. The time of inception will always be special, but if we apply the theory some random time after that, the time of application is not special, and Gott's formula may be used. (The authors first encountered Gott's formula in the year 2000, so we are not even like wedding guests who waited around for a while to apply the formula, a strategy which would arguably be dishonest.)

It is plausible that Gott's theory does have a restricted range of application, for consider that it would predict of a mathematical theorem published ten years ago that it will be refuted with 95% confidence between approximately three months from now and 390 years from now. We tend to think published mathematical results are more reliable than this. On the other hand, this estimate seems plausible if the theory in question is one of natural science. Perhaps Gott would count his theory as mathematical (or a priori), and hence escape the objection. Even so, the burden would be on Gott to articulate what the criteria for a theory's being a priori or empirical are.

The range of applicability of his theory, and his relative silence on the matter, give Gott trouble in another way. Consider Caves's objection (Caves 2000, 145) that according to Gott's argument, when one randomly meets a person and discovers that she is 50 years old, one can conclude that there is a 1/3 probability that the woman will survive to be 150 years old. Gott might reply that the delta *t* argument cannot be applied here or in any case where we have information relevant to the longevity of the phenomenon in question. (Information about the longevity of other human beings gives us information about this woman's longevity.) This would be because that information sullies the application of the principle of indifference. Yet Gott may be hard-pressed to tell us when we do or do not have relevant information in a consistent way. Consider that he is happy to cite an order of magnitude coincidence between the result of the delta *t* argument for our species's lifetime and the record of the longevities of other species, particularly species like ours (mammals), suggesting that the statistics on other species are relevant (Gott 1993a, 316). Nevertheless, he considers the application of the principle of indifference to our species justified because there is a further property (intelligence) which distinguishes us from all those species on which there is data about extinction (Gott 1994, 108). The trouble is that the same could be done with the woman, since there will be properties that distinguish her from every human being who has ever lived, and so make her a class for which no statistics exist. Gott owes us criteria for distinguishing cases where his argument can be used from those where it cannot, and it is not clear how he can give them in a way consistent with both his aims and the facts.

Leslie's Doomsday Argument

Sometimes John Leslie presents the Carter-Leslie Doomsday argument in a way that makes it just a less precise version of Gott's argument. For example, Leslie (1990, 66) says that the underlying principle behind the Doomsday argument is that "one should, all else being equal, take one's position to be fairly typical rather than very untypical". Leslie (1990, 67) says that being very early in human history would be very untypical, and he suggests (1990, 68) that being very late in human history would be very untypical too. This is akin to Gott's selecting out the first 2.5% and the last 2.5% of the interval over which r_1 is randomly distributed.

The more formal version of Leslie's argument, however, is incompatible with Gott's argument. As we mentioned in the Introduction, it is this version which has received the most attention in the literature. For example, following Leslie (1990, 69), suppose that your name is written on a ball and put in an urn with other balls, and there are two possibilities: the urn contains 20 balls, or the urn contains 1000 balls. (These are meant to be analogous to the two possibilities of the human race containing a total of, say, 100 billion people or 5 trillion.) Balls are drawn out of the urn one by one. If you discover that the seventh ball withdrawn is yours, then you should increase your probability for the hypothesis that the urn contains just 20 balls. Note that this probability shift would take place in the same way as long as your ball was one of the first 20 withdrawn. Thus, on this version of Leslie's argument, the Bayesian shift in probability assignments is always in favor of doom sooner (Eckhardt 1997, 249).

There are two important ways in which this version of Leslie's argument is different from Gott's argument. First, suppose that, before applying Leslie's argument, you believe that there is a 50% chance the human race will end tomorrow, and a 50% chance the human race will end at least 5,128 years from now. Leslie's argument dictates a Bayesian shift which will increase your probability that the human race will end tomorrow. Gott's argument, in contrast, suggests that the probability you assign to the proposition that the human race will end tomorrow is too high; you should believe that there is a 97.5% chance that the human race will end at least 5,128 years from now. Leslie's argument predicts that doom will come sooner than you had thought, while Gott's prediction for doom is independent of your personal prior probabilities.

This leads to the second difference between the two arguments: Leslie's argument requires an input of personal prior probabilities, while Gott's argument does not. Gott's argument specifies probabilities which are independent of any personal prior probabilities you might have had, by assuming that r_1 is randomly distributed between 0 and 1. (This is made clear in Gott (1994), where Gott gives a Bayesian version of his argument and specifies flat prior probabilities.) Thus, even though Gott's argument can be rewritten in a Bayesian form, the argument consists not merely of a certain conditionalization but also of assuming a particular (flat) assignment of prior probabilities, leaving no room for personal priors. This is why, although Gott's argument can be written in Bayesian form, it can still turn out that applying it at multiple times can be inconsistent with Conditionalization, as argued above.

Conclusion

In summary, we find first that it would be paradoxical to believe that the demise of our species will fall within the 95% confidence interval on the basis of Gott's argument, since a delta *t* style argument can be constructed that rules out *any* 5% within the 95% confidence interval. Second, Gott unjustifiably partitions the possibility space so that the probability that our species lasts forever is zero. Third, Conditionalization imposes a coherence constraint on a person's beliefs over time that one would violate if one applied Gott's delta *t* argument about our species at different times in one's life. This makes use of Gott's argument incompatible with Conditionalization. Fourth, Gott's argument appears to be refuted by consideration of its application to itself. At least, Gott needs to make explicit the conditions under which his argument can and cannot be applied, and it seems that it will be difficult to do this consistently. Finally, Gott's argument is distinct from Leslie's Bayesian Doomsday argument since even when the delta *t* argument is put in Bayesian form it gives the same result regardless of one's personal prior probabilities, while Leslie's argument always revises one's prior in favor of an earlier demise. This is because the Bayesian form of Gott's argument is not merely an application of Conditionalization, but depends essentially on a flat prior probability distribution.

References

Bartha, P. and Hitchcock, C.: 1999, 'No One Knows the Date or the Hour: An Unorthodox Application of Rev. Bayes's Theorem', *Philosophy of Science*, 66 (Proceedings), S339-S353.

Bostrom, N.: 1999, 'The Doomsday Argument is Alive and Kicking', *Mind* 108, 539-50.

Bostrom, N.: 2001, 'The Doomsday Argument, Adam & Eve, UN++, and Quantum Joe', *Synthese* 127, 359-387.

Bostrom, N.: 2002, *Anthropic Bias: Observational Selection Effects in Science and Philosophy*, New York: Routledge, forthcoming.

Buch, P.: 1994, 'Future Prospects Discussed', *Nature* 36, 107-8.

Caves, C.: 2000, 'Predicting Future Duration from Present Age: A Critical Assessment', *Contemporary Physics* 41, 143-53.

Cirkovic, M. and N. Bostrom: 2000, 'Cosmological Constant and the Final Anthropic Hypothesis'. *Astrophysics & Space Science* 279, 675-87.

Dieks, D.: 1992, 'Doomsday—Or: The Dangers of Statistics', *Philosophical Quarterly* 42, 78-84.

Eckhardt, W.: 1997, 'A Shooting-Room View of Doomsday'. *Journal of Philosophy* 44, 244-59.

Ferris, T.: 1999 'How to Predict Everything', *The New Yorker*, July 12, 35-9.

Goodman, S.: 1994, 'Future Prospects Discussed', *Nature* 368, 106-7.

Gott, J. R.: 1993a, 'Implications of the Copernican Principle for our Future Prospects', *Nature* 363, 315-9.

Gott, J. R.: 1993b, 'The Chances Are Good You're Random', *The New York Times,* July 27.

Gott, J. R.: 1994, 'Future Prospects Discussed', *Nature* 368, 108.

Gott, J. R.: 1997, ‘Tories Out on Cue’, *Nature* 387, 842.

Gott, J. R.: 2001, *Time Travel in Einstein’s Universe*, Boston: Houghton Mifflin.

Jeffrey, R.: 1983, *The Logic of Decision*, second edition, Chicago: University of Chicago Press.

Kopf, T., Krtous, P., and Page, D.: 1994, ‘Too Soon for Doom Gloom?’, *Physics Preprint Archive* gr-qc9407002, available at http://xxx.lanl.gov.

Korb, K. and Oliver, J.: 1998, ‘A Refutation of the Doomsday Argument’, *Mind* 107, 403-10.

Korb, K. and Oliver, J.: 1999, ‘Comment on Nick Bostrom’s “The Doomsday Argument is Alive and Kicking”’, *Mind* 108, 551-3.

Landsberg, P. T. and Dewynne, J. N.: 1997, ‘A Probable Paradox’, *Nature* 389, 779.

Lerner, E.: 1993, ‘Horoscopes for Humanity?’, *The New York Times*, July 14.

Leslie, J.: 1990, ‘Is the End of the World Nigh?’, *American Philosophical Quarterly* 40, 65-72.

Leslie, J.: 1996, *The End of the World: The Science and Ethics of Human Extinction*, New York: Routledge.

Mackay, A.: 1994, ‘Future Prospects Discussed’, *Nature* 368, 107.

Tipler, F. J.: 1994, *The Physics of Immortality*, New York: Doubleday.

van Fraassen, B.: 1984, ‘Belief and the Will’, *Journal of Philosophy* 81, 235-56.

van Fraassen, B.: 1995, ‘Belief and the Problem of Ulysses and the Sirens’, *Philosophical Studies* 77, 7-37.

[1]Kopf et. al. (1994, 5) suggest that their objection does apply to the Bayesian version of Gott's argument, but the most that their objection can show is that the time interval for 95% confidence must be larger than what is specified by Gott. Their objection doesn't address the fundamental issues which arise in the Gott argument.

[2]Rewriting the inequality for r_1, we get

$$39/40 > t_{\text{past}}/(t_{\text{past}} + t_{\text{future}}) > 1/40,$$

and by taking the inverse, multiplying by t_{past}, and then subtracting t_{past}, we get equation (1). Surprisingly, Caves (2000, 144-145) calls this step an "error" in the argument, saying that it "has no justification in probability theory", and "is sufficient to invalidate" Gott's argument. Perhaps it has no justification in probability theory, but that's because it is justified instead in arithmetic.

[3] Gott can be taken to be counting the possibility where our species evolves into a different species as survival rather than demise, as long as the descendent species is intelligent, since a delta *t* argument with the same consequences can be constructed for that case (Gott 1993a, 316).

[4]In fact, the reasoning will apply for certain intervals where the confidence level is 100%, but there is no need to go into that complication here.

[5]Gott does not specify whether this distribution is inclusive or exclusive. If the distribution is exclusive then he is simply ignoring the possibility that the human race lasts forever; the charitable reading is to take the distribution as inclusive.

[6] Gott has objected (personal correspondence) that a theory may become dead (as predicted)

through being forgotten, leaving the possibility that it is still true. Gott's formula does not involve restrictions on what 'dead' may be taken to mean, though, as long as the sense of 'dead' in our conclusions matches the sense of 'dead' in our assumptions, and has the appropriate structural properties with respect to time. The sense of 'dead' as refuted qualifies according to both criteria, and is moreover the only interesting meaning of 'dead' to apply to theories. Forgotten theories are frequently revived, whereas extinct species and dead individual living organisms are generally not. In any case, that there may be other consistent ways of interpreting 'dead' with respect to theories does not change the fact that predictions can be made with the sense of 'dead' as refuted, and those predictions yield trouble for the theory.