The Blackwell Guide to the Philosophy of Religion

Edited by

William E. Mann

Blackwell Publishing

BLACKWELL PUBLISHING
350 Main Street, Malden, MA 02148-5020, USA
108 Cowley Road, Oxford OX4 1JF, UK
550 Swanston Street, Carlton, Victoria 3053, Australia

First published 2005 by Blackwell Publishing Ltd

Library of Congress Cataloging-in-Publication Data

The Blackwell guide to the philosophy of religion / edited by William E. Mann.
p. cm. — (Blackwell philosophy guides ; 17)
Includes bibliographical references and index.
ISBN 0-631-22128-X (hardcover : alk. paper) — ISBN 0-631-22129-8 (pbk. : alk. paper)
1. Christianity—Philosophy. 2. Religion—Philosophy. I. Mann, William Edward, 1947– II. Series.

BR100.B49 2004
210—dc22
2003026163

A catalogue record for this title is available from the British Library.

Set in 10/13pt Galliard
by Graphicraft Limited, Hong Kong
Printed and bound in the United Kingdom
by MPG Books Ltd, Bodmin, Cornwall

The publisher's policy is to use permanent paper from mills that operate a sustainable forestry policy, and which has been manufactured from pulp processed using acid-free and elementary chlorine-free practices. Furthermore, the publisher ensures that the text paper and cover board used have met acceptable environmental accreditation standards.

For further information on
Blackwell Publishing, visit our website:
http://www.blackwellpublishing.com

Chapter 11

Human Destiny

Peter van Inwagen

What is to become of us?

The question is ambiguous. It might be understood to mean, What is to become of us *collectively*: what is to become of the human species? Alternatively, it might be understood to mean, What is to become of us *individually*: what is to become of you and what is to become of me and what is to become of Jack and what is to become of Joan and . . . ? In this essay I shall address only the first of these questions.

The essay has two parts. The first comprises general remarks about the question, What is to become of the human species? These remarks are something of a miscellany, united only by their common theme. The second part is devoted to the "Doomsday Argument" of Brandon Carter and John Leslie. It is my hope that some of the general remarks will justify the proportion of an essay on "human destiny" I have chosen to devote to the Doomsday Argument.

General Remarks

Whether there is an answer to the question, What is to become of the human species? depends upon whether any important aspects of the human future are determined or at least have determinate objective probabilities. That is: consider the set of "possible futures" consistent with both the present state of things and the laws of nature.[1] If all these futures share some feature – the imminent extinction of humanity; a 10,000-year Utopia; the second coming of Christ – there is at least a partial answer to the question, What is to become of the human species? If various subsets of these futures have non-trivial measures that satisfy the usual formal constraints – so that we can meaningfully say things like "In 57 percent of the futures that are consistent with the present and the laws of nature, the human species will become extinct before the year 2200" – then there is an answer to the question, What is to become of the human species? But this answer may be very complicated and may essentially involve probabilities. It might take the form of a set of functions that assign objective probabilities to dates and important "eventualities;" for example, one of these functions might assign to the eventuality "human extinction" and the date "January 1, 2220" the probability 0.57, meaning that there is an objective probability of 0.57 that our species will have become extinct by this date.[2] If no eventuality/date pairs, or none involving eventualities relevant to our hopes and fears concerning the future of humanity, have objective probabilities – if strict determinism is true, then every eventuality/date pair has an objective probability of either 0 or 1 – then there is simply nothing to be said about what is to become of the human species.

If there is an answer to the question, What is to become of the human species? it might nevertheless be idle to ask this question, owing simply to the fact that it is not humanly possible to discover its answer, or even to find any cogent reason to regard any of its possible answers as in any way epistemically preferable to any other (equally specific) answer. Suppose there is an urn in which someone has placed a certain number of black balls and a certain number of white balls, and that we have no way of knowing what these numbers are or of knowing anything non-tautological about the ratio of either to the other. If we know that one ball will be selected at random from the urn – by an ontologically indeterministic mechanism – we know that there is an answer, an informative response, to the question, Will the ball that is drawn be black or white? The answer is "black" if the number of white balls is 0 and the number of black balls is not 0; it is "There is a probability of 0.57 that it will be white and a probability of 0.43 that it will be black" if the number of white balls is 57 and the number of black balls is 43 . . . and so on. But we also know that the question is an idle one, since we can have no reason to accept any of the possible answers.

Some have thought the question answerable, and have given answers to it.

Aristotle and the Hindu religion agree that the world and the human species are eternal. For Aristotle, the past and the future will be pretty much the same as the present: the sun will always shine (and has always shone), there will continue to be the same climatic conditions and the same biological kinds there have always been, and cities and empires and languages will pass away and others will come to be. For the Hindus, although world-history is cyclical, in the long run things will always be pretty much the same: the same cycles are ordained to recur eternally in their given order, just as, in our experience, the same four seasons continually recur in their given order. The Pythagoreans and the ancient atomists allowed that the human race would someday become extinct and that the physical universe itself would dissolve into chaos; but, they held, the physical universe and the human species will be reborn: owing to nothing more than the chance recombination of the basic units of matter, the future of the universe comprises an infinity of deaths and rebirths. Nietzsche, who adopted this thesis, called it "the most scientific of hypotheses,"[3] although he never attempted any careful argument for the scientific necessity of this "eternal return."[4] (The necessity of the eternal return does not follow from the premises he seems to have thought entailed it: a universe that consists of two particles that draw ever nearer for the "first half" of eternity, pass each other like ships in the night, and then draw ever farther apart for the "second half" of eternity, is consistent with those premises. There is, however, an interesting theorem of Poincaré that says roughly this: in a bounded ideal mechanical system in which the elements of the system do not lose energy when they "bounce off the walls" – for example, an ensemble of ideal gas molecules confined forever to an unchanging, ideal container; our "two-particle universe" is not bounded – for all but a set of initial states of the system of 0 measure, the system will return to a state arbitrarily close to its initial state, given enough time.)

The Abrahamic religions[5] – Judaism, Christianity, and Islam – disagree with Aristotle, the Hindus, the Greek atomists, and Nietzsche. These religions maintain, first, that the physical universe came into existence at a certain moment in the past,[6] and, secondly, that at some moment in the future, the physical universe, the earth, and the human species – the theater, the stage, and the actors – will undergo a radical and irreversible transformation. And they hold that human beings (if not the physical universe or the earth) will thereafter always exist in the new state that this transformation will effect. (Each human being will exist eternally and there will be no reproduction.)[7]

Another sort of disagreement with Aristotle et al. can be found in the writings of "historicist" philosophers like Hegel and Marx. Ontologically speaking, historicism is a secular rewriting of Christian eschatology. (Historicism, like Gothic architecture, is an epiphenomenon of Latin Christianity.) The human species (but not the universe or the earth) is to undergo a radical and irreversible transformation. (A spiritual – *geistlich* – or economic transformation, not a biological transformation.) But this transformation will not be brought about by God – as, of course, the Abrahamic religions maintain with respect to the transformation they look forward to – but by the operation of impersonal historical principles. Epistemologically speaking, the coming "historicist eschaton" is said by historicists to be predictable by the exercise of human reason – as opposed to its being revealed by God, divine revelation being, of course, the source of our supposed knowledge of the eschaton according to Jews, Christians, and Muslims. According to historicists, the historical development of reason has reached a level – the "historical moment," which occurred at some point in the nineteenth century – at which reason became able to understand the laws of its own development.

I have given a sketch of three positions according to which important features of the human future are determined by the present state of things (according to the Abrahamic religions, "the present state of things" includes God's intentions as regards the human future), and, according to which, it is possible to know what some of these features will be. These seem to me to be the most important positions of this type that have actually been taken. What can be said for and against them?

The first or "Aristotelian" position may be rejected, for physics and cosmology have shown that the universe does not have an infinite past. (Or have they shown this? The currently "standard" cosmological theory implies that time had a beginning. According to that theory, "Alice lived 20 billion years ago" cannot be true – for a reason analogous to the reason "Alice lives 20,000 miles to the south" cannot be true: the phrase "20 billion years ago," although it is semantically suited for being the name of a time, is not the name of a time, just as "20,000 miles to the south" is not the name of a place although it is semantically suited for being the name of a place. But there are other cosmological theories than the standard theory, theories consistent with the cosmological evidence, according to which the universe does not have a beginning in time. What can be said for certain is this. One cannot properly regard it as having been established, or even as having been shown to be probable, that the age of the universe is infinite; the most one can say is that this is a possibility that hasn't been definitively ruled out.) As of this date the consensus among cosmologists is that the universe has an infinite future, but that only an initial segment of this future (finite, of course: all initial segments are finite) will be at all interesting. Following this initial segment, the physical universe will forever consist entirely of radiation, black-body radiation at a temperature that only respect for the law of the conservation of energy can lead us to distinguish from absolute zero. (There will come a time when the temperature of the radiation-universe in degrees Kelvin will differ from zero only in the seventy-second decimal place; later it will differ only in the eighty-ninth decimal place; later still only in the hundred-and-sixth decimal place, and so on. The radiation-universe of the future has nothing to look forward to but the eternal growth of the initial segment of zeros in the number that measures its temperature.) The heavens will indeed, in the words of Ps. 102, wear out like a garment. This is what cosmology tells us. Geology, paleontology, anthropology, archaeology, and historical scholarship tell us that the history of the earth, the biosphere, the human species, and human culture has been neither static nor cyclical.

As to the possibility of the eschaton expected by Jews, Christians, and Muslims, there is little to be said. I am a Christian and thus "look for the resurrection of the dead, and the life of the world [that is, age] to come." But this is a matter of faith (and it is a part of my faith that it is a matter of faith). If you do not share my faith, I have a great deal to say to you, but you (justifiably) are not reading this essay with the expectation of being evangelized, and neither I nor the editor nor Blackwell wishes to be accused of false advertising.

Historicism, at any rate in its "strong" Hegelian-Marxist form, cannot be taken seriously. Hegel and Marx (or at least Hegel and Marx the prophets) seem merely comic today. But various more modest forms of the thesis that important aspects of the human future are predetermined and knowable by human reason (or have determinate objective probabilities that are knowable by human reason) are worthy of serious consideration. It follows, obviously, from what we have said above about what cosmology tells us that the human species does not have an infinite future. (What we know about what the sun is going to do over the next 5 billion years is also relevant to the question, How long will the human species exist? But it is possible that we shall be able to migrate to other planetary systems. No one, however, is going to organize a migration to another cosmos. At least I don't *think* so.) But little else follows. For all cosmology (or stellar astronomy) can tell us, the human species might have a future that is "imaginatively infinite" – a future that is large in comparison with the human past as intergalactic distances are large in comparison with the walk from the parish church to the post office. (And even the modest conclusion that the human species has only a finite future depends on a premise that, uncontroversial though it may be in some circles, is not accepted by everyone – not, for example, by me and my house. This premise is naturalism, the thesis that the physical universe is "all there is or was or ever will be." If human beings are not, in the words of J. R. R. Tolkien, "confined for ever to the circles of the world," the heat-death of the physical universe does not imply the end of humanity.)

It is possible, then, that, although humanity does not have an infinite future, it has a very long future indeed, a future during which . . . what?

Philosophers of the classical world could suppose, consistently with the astronomical, biological, and historical knowledge of their times (if not with the creation stories they had heard at their mothers' knees), that the universe and the earth had always been much as they were then and always would be much as they were then. A fortiori, they could easily enough have believed in a terrestrial history and future that spanned thousands of millions of years, a history and future of a world at every time much like their own, a world of cities and agriculture, wars and empires. But (as regards the past) they did not know that biological life had a beginning and has a history; their cultural memory did not reach back to the last glaciation or even to the revolutions in agriculture and metallurgy that had created their world of city-states and empires. And (as regards the future) they could not foresee the technological revolution of the second Christian millennium and the population explosion of the nineteenth and twentieth centuries. *We* do not think of history (either cosmic, terrestrial, or human; but I am now thinking primarily of human history) or of the future, including the human future, as they did. We know that there had never in all history been a human world at all like the one we lived in during the second half of the twentieth century, and we are certain that the twenty-second century will be no more like the twentieth century than the twentieth century was like the eighteenth. It seems, therefore, that at least one feature of the human future is now predictable: it will not be like the human present or the human past. So we believe, and I have a hard time seeing how this belief could be wrong. The shape of our lives (to the extent that our biology allows this shape to vary) is to a very large degree a consequence of technology. Technological development has become a self-sustaining, impersonal thing, like a forest fire or a pandemic. (I mean these images to suggest things that are growing and impossible to control and have taken on "a life of their own," not things that are bad; I consider it still an open question whether the Luddites or the technophiles or the proponents of some intermediate view will turn out to be right. I suppose the fact that all the similes that serve my literary purpose are bad things should be a cause for *some* unease in the technophile camp.) Continuing technological innovation is inevitable – unless it should be curtailed by some global disaster (whether a consequence of unbridled technological growth or of some wholly unrelated cause, as in the recently popular spate of “big rock hits the earth” movies). And if such a global disaster were to happen, that, too, would have the result that the future will be very different from the present.

Detailed (and even rather general) predictions of the future of *homo technologicus*, insofar as they have been predictions of futures that have had time to become the present, have almost always been wrong. Indeed, it seems that the only way for a prediction of the future of a technological civilization to be right is for two people to make predictions that are logical contradictories. (And if two propositions are logical contradictories, one of them, at least, must be very general and abstract.)

Two sorts of people have offered such predictions in print: imaginative writers like Jules Verne and H. G. Wells and George Orwell and Aldous Huxley and Robert Heinlein, on the one hand, and self-described “futurologists” on the other. The former, claiming to be no more than tellers of tales, cannot, perhaps, in the strictest sense, be said have been engaged in the business of “predicting the future.” Still, many of them did the best they could to predict the future (if they had thought some other future more likely than the one presented in their stories, they would have written stories set in that more likely future). Predictions of the future from either source tend to be (but are not invariably) of two types: utopian or dystopian. Wells’s *The Time Machine* cleverly combines both: it “predicts” a utopia of vast duration to be followed, finally, by dystopia and human extinction. Orwell’s *1984* and Huxley’s *Brave New World* are the classic examples of dystopian literary predictions. (Huxley’s forgotten *Ape and Essence*, should, I think, be in the list of classics.) Whether a detailed prediction of the human future is the result of a novelist’s imagination or a futurologist’s computer modeling (garbage in, garbage out, as they say), one thing is certain: it will be wrong. Predictions made in the 1940s or 1950s are now amusing. Those made in the 1920s and 1930s are hilarious. Those made in the nineteenth century are charming. And that is really all there is to be said about predicting the future in detail. There are a vast number of epistemologically possible futures, and any proposal of one of them as the actual future (any prediction) will be based on the author’s hopes and fears, unconstrained speculations, and extrapolations from a minuscule proportion of even the available data: that is, it will be an essentially random choice from among a vast array of possibilities, most of which the chooser will not even have thought of. (Most of which *no one* will have thought of. Most of which no one *could have* thought of. There has been no “historical moment.”)

There is, however, one prediction of an important aspect of the human future that is based not on speculation and the predictor’s prejudices, but on philosophical argument; on one argument, an argument which, although it may be mistaken, is a good deal more worthy of the honorable name "argument" than anything that can be found in the writings of Hegel, Marx, or the twentieth-century futurologists. I refer to the notorious Doomsday Argument.[8] The prediction this argument makes is not specific in the way the predictions of the Club of Rome were specific, but it does not lack interest on that account. The conclusion of the Doomsday Argument is that there is a significant probability that the human species is going to become extinct: and not in a million or even in 10,000 years, but within the next few centuries.

The Doomsday Argument

Imagine that (contrary to your expectations when you went to bed) you awaken one morning in a hotel room. You are informed that you have been drugged and kidnapped and are being held on the island of Antiqua, an island of which you have never heard. (You are very ignorant of geography.) You look out the window – your room appears to be on something like the tenth floor of the hotel – and, a few hundred yards away, you observe the sea (or at least a body of water large enough that its farther shore cannot be seen from your window). Can you infer anything much from this fact about the island of Antiqua? Not obviously, for it is in the nature of an island to be bounded by water. But suppose you discover the following facts about the way hotels are distributed in Antiqua (perhaps you have found a page ripped from a brochure about the Antiqua hotel industry in the wastebasket in your room).

> Antiqua is topographically pretty uniform, and every place in Antiqua is suitable for the construction of a hotel. Partly in consequence of this fact, Antiqua is densely and uniformly populated with hotels: every square mile of Antiqua contains a largish number of hotels and the number of hotels in each square mile is about the same as the number in any other square mile (about 10 or 12, in fact).

And suppose you know (you have overheard one of your captors say this) that, for security reasons, the hotel you are in was chosen at random from among the hotels of Antiqua. That is, the names of every hotel in Antiqua were written on slips of paper, which were put into a hat, and one of the slips was drawn by a blindfolded member of the gang that kidnapped you. You were then taken to the hotel whose name was drawn.

Now suppose that, having this information at your disposal, you reason as follows:

> Suppose Antiqua were a large island, an island the size of Ireland, say, or even larger. [The depths of your geographical ignorance are such that, for all you know, there *is* an island the size of Ireland called "Antiqua."] If Antiqua were of that size, and

> hotels on Antiqua were as numerous and as uniformly distributed as I know them to be, only a very small proportion of the hotels in Antiqua would be this close to the sea. (And the body of water I observe must be the sea or at any rate the body of water that surrounds Antiqua: if there were a huge lake or bay in Antiqua, then Antiqua wouldn't be "topographically uniform," which I know it is.) And I know that this hotel was chosen at random from among all the hotels of Antiqua. If Antiqua were the size of Ireland, therefore, it would be very improbable on what I know that I should be able to observe the sea from my window. I conclude that it is probable that Antiqua is much smaller than Ireland. I can even conclude that Antiqua is probably considerably smaller than, say, Long Island, although the probability I can assign to this conclusion is smaller than the "Ireland" probability. I therefore conclude that Antiqua is a small island.

This reasoning can seem very plausible. It can seem that you'd be justified in using its conclusion ("Antiqua is a small island") in some context in which something of great practical importance hung on its conclusion. Suppose, for example, you had formulated two promising escape plans that differed in only one important respect: plan A would be more likely to succeed on a small island and plan B would be more likely to succeed on a large island. One might suppose that you would be rationally justified in proceeding according to plan A, that it would, in fact, be positively irrational – all other things being equal – for you to prefer plan B to plan A.

But not so fast. There is a lacuna in your reasoning. You have forgotten to take into account the "prior" or "antecedent" or "posterior" probability of Antiqua's being a small island. You, the fictional "you" of the example, are, as I say, very ignorant of geography. Suppose you were to learn the following facts (facts in the fictional world of our example): there are exactly 100 islands in the world, and all but one are about the size of Ireland (and the other is very small); just as your captors chose the hotel in which they would hold you captive at random from among the hotels of Antiqua, so they chose Antiqua as the island of your captivity at random from among the 100 islands of the world. In that case, the prior or antecedent probability of the conclusion of the above reasoning (that this island you are on is a small island) – antecedent, that is, to your observing the sea from your window – is low: it is, in fact, 0.01 or 1 percent. Can you validly conclude that it is highly probable that you are on a small island if you know that the antecedent probability of this hypothesis is low? Can you validly conclude this if you have *no idea* what the antecedent probability of your being on a small island (given that you're on *some* island) is? The answer to the first question is, That depends. The answer to the second is, No. There is, nevertheless, something about your reasoning that is on the right track. It does lead validly to an interesting conclusion, but this conclusion is not that you are probably on a small island; it is that you should now assign to the hypothesis that you are on a small island a *higher*, perhaps a significantly higher, probability than you did before you looked out the window and saw the sea. The new piece of evidence you have acquired, as some people say, "raises the antecedent probability" of the hypothesis that this island you are on, Antiqua, is a small island. (A strictly meaningless phrase, since antecedent probabilities do not and cannot change, but it seems to convey to most people what it is intended to convey.) There is a theorem of the probability calculus, Bayes's theorem, that governs the degree to which antecedent probabilities are raised by new evidence. I will not undertake a technical, or even a non-technical, discussion of this theorem,[9] but I well mention some numbers, just to give a sense of orders of magnitude.

In the case we have imagined (the "100 islands" case), the antecedent probability that you are on a large (Ireland-sized) island is 99 percent and the antecedent probability that you are on a small island is 1 percent. Suppose that if one were placed on a large island at a place chosen at random, the chance of one's finding oneself near enough to the sea for it to be visible to one (under the conditions we have imagined) would be 0.0002 or $\frac{2}{100}$ of 1 percent. Suppose being placed at a randomly chosen spot on the sole small island would give one a 50/50 chance of being able to observe the sea.[10] This information, given Bayes's theorem, is sufficient for you to be able to calculate the probability you should assign to the hypothesis that you are on a small island, given the new piece of information that has just come your way, to wit, that you can observe the sea. It is over 96 percent.[11]

Let us now move from problems about one's location in space to problems about one's location in time. ("Size" – in this context, area – is a concept involving two dimensions, and time has only one dimension. At the cost of what little realism our "island" example had, however, we could as easily have considered "thin" islands, islands one of whose geographical dimensions could be ignored. We could have considered the effect that finding yourself near one "end" of an island should have on the probability you assigned to the hypothesis that it was a *short* island.)

Let us consider a temporal analogue of our "island" example. Some moment between the present and 1 million AP ("after present") is chosen at random. Call this moment B. Then a moment, A, is chosen at random from the moments between the present and B. I am taken by time machine to the moment A. On arrival, I discover that the year is "only" 2097. Can I conclude anything about how far in the future B is? The question is rather vague. Let us ask a more precise question. Let us call the next 100 years the "very near future," the next 1000 years the "near future," and the period from 1000 AP to 1 million AP the "distant future." Here is a precise question: what is the probability that B lies in the "near future," given that I have found myself in the "very near future"?

To calculate the probability that B lies in the near future (given that I have found myself in the very near future) using Bayes's theorem, we need three numbers: the antecedent probability of B's lying in the near future, the probability of my finding myself in the very near future if B lies somewhere in the near future, and the probability of my finding myself in the very near future if B is somewhere in either the near or the distant future. (The probability we are trying to determine will be the first probability multiplied by the ratio of the second to the third.)

The antecedent probability that B lies in the near future is 0.001, since the near and the distant future are together 1000 times as long as the near future. We can approximate the probability of my finding myself in the very near future given that B is in the near future by noting the following facts: if B were 100 years from now, the probability of my finding myself in the near future would be unity; if B were 200 years from now, this probability would be 1/2; if B were 300 years from now, the probability would be 1/3; . . . ; if B were 1000 years from now, the probability would be 1/10. Intuitively, the probability of my finding myself in the very near future, given that B is *somewhere* in the near future, should be close to the *average* of these 10 probabilities – about 0.3. So we may say that if B is in the near future, the probability of my finding myself in the very near future is approximately 30 percent.[12] An exact calculation (one that substitutes integration for summation) shows that this approximation is pretty good; the actual probability is equal to $\frac{1}{10}$ the sum of 1 and the natural logarithm of 10, or about 0.3302. (The number 10 occurs where it does in this calculation because the near future is 10 times as long as the very near future; natural logarithms come into the picture when one integrates the function 1/x; and where does the function 1/x come from? – well, remember the series of probabilities, 1/2, 1/3, . . . , 1/10.) The probability of my finding myself in the very near future if B is in the near future is thus almost exactly 33 percent. A similar calculation shows that the probability of my finding myself in the very near future if B is randomly chosen from the next million years is $\frac{1}{10{,}000}$ the sum of 1 and the natural logarithm of 10,000. (A million years is 10,000 times as long as the very near future.) This number is about 0.00102. The probability that B lies in the near future given that I have found myself in the very near future is thus (by Bayes's theorem)

$$0.001 \times (0.3302 \,/\, 0.00102) = 0.3234.$$

The bottom line is: if I find myself in the very near future (in the circumstances imagined), I can conclude that the probability that B lies within the next 1000 years is just over 32 percent. (A considerable multiple of the antecedent probability of this hypothesis, which was $\frac{1}{10}$ of 1 percent.)

Now let us map this example on to another example. Suppose that God, as in the book of Genesis, created the human race on a particular day and at a particular hour: at 11 o'clock on that day there were no human beings and at noon there were (fully formed adult language-users). Let us call this event the Creation. But suppose that, departing from the Genesis story, God created several hundred human beings, and that (by his will) the human species was destined to number in the low hundreds and to live clustered together in a single community for the duration of its existence. (Every generation of human beings, moreover, had lifespans of the three-score-and-ten order: there were no Methuselahs.) God told his human creations the following fact immediately after the Creation: they were not (that is, their species was not) to exist forever. At some point (the Omega Point), humanity would come to an end in a natural disaster (always a real possibility for a population confined to a small geographical area). God revealed only this about the time-frame of the extinction of humanity: the Omega Point was *at most* 1 million years after the Creation. The Omega Point in due course arrived and the human species came to an end, like a tale that is told. Many millions of years later, the earth was inhabited by another intelligent species. This race invented time machines and, using these marvelous devices, had covertly observed human history, from the Creation to the Omega Point (the dates of both of which some of them of course knew). For a crime unintelligible to us, a member of this species was sent into temporal exile, to live among the long-extinct human beings. He did not himself know the date of the Omega Point or how long after the Creation it was, but he understood the concept, and he was told this by the authorities: “We are going to place you at a point in human history that we shall choose at random; in effect, we will put all the dates between the Creation and the Omega Point in a hat and draw one of them. (If your ‘arrival time’ turns out to be only a year or two before the Omega Point – well, that’s just too bad.)” After the exile had been placed among the human beings and had learned their language, he was rather surprised to learn that he had “arrived” only 96 years after the Creation.

Could he conclude anything about how long it was till the Omega Point? (We assume the human beings have told him what *they* know: that the Omega Point would occur at most 1 million years after the Creation.) If our reasoning in the previous example is correct, one thing he was in a position to conclude is this: the probability that the Omega Point was within 1000 years of the Creation on the proposition that he has found himself within 100 years of the Creation was about 32 percent. And, of course, any human being to whom he told his story would have been able validly to infer the same conclusion.

Now: remove the “temporal exile” from the story. Or turn him into a fiction within the story: a science-fiction writer among the human beings of the First Century tells the story of the appearance in their century of an alien temporal exile from the far future, and of the unsettling piece of reasoning that his appearance in the First Century occasions. A reader of this tale smiles at its cleverness – and then a disquieting line of argument occurs to him:

> Am *I* not in a position very like that of the Temporal Exile? True, I “arrived” in the First Century not by being brought from another era in a time machine to a point in human history that was chosen at random, but by being born in it. Nevertheless, I simply *find myself* in the First Century – just as the exile in the story did. Now maybe the metaphysical theory of the essentiality of origins is correct, and *I* couldn’t have been born at any other point in human history. Maybe the probability of my being born in any other century is precisely 0 – whereas the Temporal Exile was as likely to find himself at any point in human history as any other. Still, this consideration, if true, seems somehow irrelevant to my worries. Suppose God were to reveal to us human beings that we don’t come into existence in our mothers’ wombs. Suppose he told us that we had some sort of pre-existence in a Platonic heaven, and that the year of each person’s physical birth was selected by angels throwing celestial dice.

> Would that revelation change the way I should regard the logical force of the unpleasant argument that has occurred to me? I am imagining that I have learned that it is literally true that for any two calendar years in human history there was an equal and non-zero probability of my being born in either year. Having learned this, should I say, "*Now* I see that the argument that troubles me is a good one – but if I hadn't learned the strange truth about my pre-existence and how the year of my birth was selected, I ought to have said that the argument had no force"? (There can be no doubt that if the imaginary revelation were true, my epistemic position would be exactly that of the Temporal Exile.) For the life of me, I can't see that the imaginary revelation would make a difference to the force of the argument. Whatever I should conclude on the assumption of pre-existence-and-random-selection-of-year-of-birth, I should conclude the same thing on the assumption that I came into existence in my mother's womb and could not possibly have come into existence at any other point in history. I very much fear that my reasoning is right: I must conclude that there is a probability of 32 percent that the human species will come to an end by the thousandth year after Creation. If I were a bookmaker (an ideally rational bookmaker), and if someone came to me and wanted to bet that the Omega Point would come by the thousandth year after Creation, then (assuming there was some way to settle the bet) I shouldn't be able to offer him very attractive odds. I should have to offer to pay him some amount less than $2.12 on the dollar if he won.[13]

This reasoning seems to me (me, the author of this paper, not me, the fictional reasoner) to be pretty good. If it has a weak point, that weak point has, I believe, nothing to do with the applications of Bayes's theorem it contains or with the determination of this or that probability. It is, rather, a premise of this reasoning that could be put a little more explicitly in these words:

> Let n be the number of years humanity exists – from beginning to end. If a human being (who does not know what number n is) knows he was born m years after the beginning of humanity, he may, for the purposes of probabilistic reasoning like that illustrated in our examples, treat m as a number chosen at random from among the numbers 1, 2, 3, . . . , n–1, n.

I am really not too sure how plausible this premise is. I had my fictional reasoner defend it by a clever mixture of rhetoric and picture-thinking. But how plausible is it – really? Again: I'm not sure. I'll leave it to the reader to decide. In aid of this task, I leave the reader with two more imaginary cases, cases that are intended to serve as "intuition pumps." The first (these are my own subjective reactions; the reader's may differ) works against the principle and the second works in its favor. (I am at a loss to explain my differing reactions to these two cases.)

In the first case, suppose you are a passenger on a fully automated "generation ship," which is crossing some vast intersidereal or even intergalactic gulf at a very small fraction of the speed of light. As generally happens in science-fiction stories about generation ships, the passengers have, after a generation or two, somehow forgotten all the essential information about their voyage, including its destination and length (except this one fact: they know that the voyage will not be longer than 1 million years). You have no reason to regard any hypothesis about the length of the voyage as preferable to any other equally specific hypothesis, provided both hypotheses respect the 1-million-year upper bound and the lower bound established by the currently elapsed time-in-voyage. And you know that there will be passengers on the ship for the whole voyage – the automated systems that care for the passengers will see to it that the generations do not fail. At a certain moment, as you stare at the Great Chronometer (which registers elapsed time-in-voyage), it strikes you that it has been 96 years since the voyage began. You reason as follows: "Since human beings will inhabit the ship for the whole voyage, I may treat the number 96 (which corresponds to the point in the voyage at which the population of the ship currently 'finds itself') as a number chosen at random from among the numbers 1, 2, 3, . . . , n–1, n, where n is the length of the voyage in years. I can therefore calculate that there is a 32 percent chance that the length of the voyage is 1000 years or less."

In the second case, suppose you are a member of a tribe that lives on the banks of the Great River – and your tribe has always lived there, since the beginning of time. The ancient stories of your tribe (which cannot be doubted) tell you that the gods who made the Great River were constrained to make it 1 million miles long or less. Casting the sacred knuckle-bones, the gods chose at random a number n between 1 and 1,000,000 and made the River n miles long. (What the number n is, only the gods know.) Owing to the bounty of the gods, the River is densely populated by tribes much like your own along its entire length. The rules forbid members of your tribe ever to venture downstream, but great heros of the old days have explored the River upstream and have found (as expected) other tribes every few miles – *and* they have found that your tribal village lies 96 miles from the source of the River. You reason as follows: "Since the River is densely and uniformly populated along its entire length, I may treat the number 96 (which corresponds to the location at which *my* tribe 'finds itself') as a number chosen at random from among the numbers 1, 2, 3, . . . , n–1, n, where n is the length of the River in miles. I can therefore calculate that there is a 32 percent chance that the length of the River is 1000 miles or less."

As I say, I am not sure how plausible the premise I have called attention to is. There are, however, many objections to the kind of reasoning exemplified by the foregoing arguments, that are unrelated to the question of the plausibility of this premise. None of them seems very convincing to me. I will mention only one, the first to occur to almost everyone (including me) on his or her first encounter with the Doomsday Argument. Applied to our last example, this objection could be put as follows. "Look, given that the banks of the Great River are populated in the way you have imagined, *some* tribes have to find themselves within 100 miles of the source of the River. To simplify the picture, assume that the River is divided into 100-mile 'segments,' and that exactly one tribe lives in each segment. Your imaginary reasoner's tribe just happens to be the one that does 'find itself' in segment 1. They can't *infer* anything from this fact. Suppose there were 10,000 rivers, one of them one segment long, one two segments long, and so on . . . up to the one that is 10,000 segments or 1 million miles long, all of them having one tribe living in each of its constituent segments. On each of these rivers, there will be a tribe that lives in its first segment. But only 10 of these tribes, $\frac{1}{10}$ of 1 percent of them, will live on the banks of a river 1000 miles long or less. A member of one of the first-segment tribes should therefore believe that there's only one chance in 1000 that 'his' river is 1000 miles long or less."

Answer. Well, yes – so he should, if he knows he is on the bank of one of 10,000 rivers having the lengths and population-distributions you have imagined. But the situation is different for a tribal reasoner who knows that there's just one river and, antecedently to his observation that his tribe lives in its first segment, has no reason to prefer any of the hypotheses "The River is 1 segment long," . . . , "The River is 10,000 segments long" to any other. Both these "states of knowledge," and the reasoning each authorizes, can be represented by "pre-existence and random placement" analogies, but the appropriate analogies are different.

Here's the appropriate pre-existence-and-random-placement analogy for the case in which the tribal reasoner knows there are 10,000 rivers. There *are* 10,000 rivers, and they have the lengths you have imagined. Simple arithmetic shows that these rivers together contain a total of 50,005,000 segments.[14] Suppose that these segments have been numbered in the obvious way, and that each of 50,005,000 people (myself among them) is assigned a number from 1 to 50,005,000 at random and is then placed in the segment whose number he has been assigned. When all this has been accomplished, I am surprised to find myself in the initial segment of one of the rivers. (Surprised because there was only about one chance in 5000 that I'd find myself so placed. Still, I might reflect, it had to happen to someone; in fact it had to happen to 10,000 people.) Now what is the probability that the river on which I have been placed is at most 10 segments (1000 miles) long? Why, just the proportion of the rivers that are at most 10 segments long: $\frac{1}{10}$ of 1 percent. Note that the antecedent unlikelihood of my finding myself in the initial segment of one of the rivers does not figure in this calculation. I'd make the same calculation if I knew that I had been set down in a segment randomly chosen from the initial segments.

Here's the appropriate pre-existence-and-random-placement analogy for the case in which the reasoner knows there is one river. Again, there are 10,000 rivers of the different lengths you have imagined. One of the rivers is chosen at random. Then one of that river's segments is chosen at random, and I am placed in it. I find myself in its first segment. Now what is the probability that the river on which I have been placed is 10 segments long or less? As our earlier calculations show, it is about 30 percent.[15] In this analogy, it is the initial choice of one river that corresponds to there *being* one river; once that choice has been made, the other rivers might as well not exist. The chance that this one, chosen river will be 10 segments long or less is of course only $\frac{1}{10}$ of 1 percent. This is the antecedent probability of the river I find myself on being 10 segments long or less. But this antecedent probability is raised by the fact that I have found myself in its first segment, for the probability of this outcome is about 0.3 on the hypothesis that the river is 10 segments long or less, and only about 0.001 on the hypothesis that the river is 10,000 segments long or less.

I conclude that the reasoning that led our First Century human being to assign a probability of 32 percent to the hypothesis that the Omega Point would occur within 1000 years of the Creation is at the least very plausible.

Now what is the application of this reasoning to *our* situation, the actual human situation? Someone might protest that it has none, because (even if it contains no mistake) we *don't* "find ourselves" at a point in time that is at all close to the origin of our species. After all, however we define "human being," there have been human beings for a least 100,000 years. If there have been human beings for 100,000 years, it seems, reasoning in the style of the Doomsday Argument may convince us that it is likely that the human species will come to an end in the next million years, but not that it is likely that it will come to an end in the next 1000 years.

But note. Suppose humanity lasts another 10,000 years and that population levels remain at least what they are now for the remainder of our existence. (The assumption that the human population will be at least as large as it is now for the remainder of our existence as a species is very plausible. It is not, of course, certain: it is possible that some natural or man-made disaster – or some revolution in human reproductive ethics – will reduce the human population to a few millions, and that human beings will thereafter gradually decrease in numbers till the species finally flickers out of existence 10,000 years from now.) In that case, we present-day human beings do not find ourselves close to the point of human origin if we measure "closeness" in terms of simple duration: we find ourselves perhaps 90 percent of the way along the road from the origin to the extinction of our species. But suppose we measure "closeness to origin" another way: in terms of the number of human beings who have preceded us and the number of human beings who will follow us. By that measure, we are very close to the point of human origin indeed, since only a very small proportion of the whole set of human beings – past, present, and future – precedes us; almost every member of this set is yet to be born. This is, of course, because of the extraordinary population explosion of the last 200 years: for most of human history, the number of human beings has been only a very small fraction of what it is now and (given our assumption about the future) of what it will be at any time in the next 10,000 years.

If the human population were always about the same (or so the Doomsday Argument attempts to convince us), one would not expect to find oneself in the first $\frac{1}{100}$ of its span of existence; that is to say, one would not expect to find oneself among the first 1 percent of the human beings who ever live. But is it not reasonable to suppose that this second way of describing what human beings would not expect would apply no matter what a graph plotting the human population against time looked like? If the human species ends a few hundred years from now (and if its numbers remain high till shortly before their final, precipitous drop to 0), you and I shall be "average" human beings in this sense: the number of human beings who lived before us and the number of human beings who will live after us are roughly comparable. But if humanity lasts another 10,000 years at current or higher population levels, the number of human beings who will live after us is vastly greater than the number who lived before us. It seems that reasoning similar to the "doomsday" reasoning we have considered – but without the "constant population" assumption – suggests that there is a significant probability that the human species will come to an end soon. Again, the reader is invited to consider spatial analogies. Suppose one is an inhabitant of an island (its size is unknown: it may be anywhere from 200 miles across to the size of Australia) that is very sparsely populated near its coasts but becomes increasingly densely populated along any line drawn from a point on its coast to its center – or at least this increase displays itself on the first 100 miles of any such line. You find yourself at a point 100 miles from the coast. The interior of the island is somehow hidden from you: you are able to look only "outward," toward the coast. A powerful telescope shows you that the island is very sparsely populated within 90 miles of the coast (at the coast, the population density is about one person per 1000 square miles, a figure that gradually increases to about one person per 700 square miles 90 miles inland). Between 90 and 98 miles inland, the population density increases rapidly and between 98 and 100 miles inland, it increases explosively: at 98 miles inland there are four people per square mile, and, at 100 miles inland (where you live) there are 40 people per square mile. Would it be more reasonable for you to believe that the island's center is just a few miles inland from you or that it is hundreds of miles inland from you? If you had always unreflectively believed that your island was a huge island, an island the size of Australia, should the discovery of the facts about the increase in population density between the coast and the point at which you find yourself lead you to revise the probability of the "huge island" hypothesis downward? Significantly downward? In considering these questions, it might be useful to ask yourself whether someone in your situation can properly think of himself as living very close to the shores of a "population island," an island that is, as it were, made of people, and to ask yourself whether you can properly phrase the "huge island" question this way: Do I live near the shores of a large or of a small population island? (If so, the question seems closely analogous to the question, "Am I on a large or a small island?" as it presents itself to the central character in the "hotel room" case.)

Here I leave and commend study of the Doomsday Argument to you.

I close with two observations.

First, as we have seen, the cogency of the Doomsday Argument depends partly on the antecedent probability of the hypothesis that the human species will come to an end soon. Bayesian reasoning can suggest to us that we ought to revise our antecedent estimate of this probability upward, perhaps significantly upward, perhaps to multiply it by, say, 100. But an insignificant probability multiplied by 100 may still be an insignificant probability. Is there any good reason, any good reason independent of the considerations put forward in the Doomsday Argument, for thinking that the human species may come to an end soon? Well, there are all sorts of perfectly respectable scenarios (scenarios put forward by recognized authorities as representing real possibilities) according to which humanity will come an end within a few hundred years.[16] There are various all-too-real possibilities that any thinking person must recognize: thermonuclear or ecological or epidemiological catastrophe. And there are a large number of scenarios that are (at a reasonable guess) individually of small probability but which must each be added into the aggregate reckoning: a comet or asteroid may hit the earth in the near future; a nearby (astronomically speaking) supernova may irradiate us; an ill-advised experiment in high-energy physics may tip us catastrophically out of the metastable false vacuum in which, for all we know, we exist into the yawning abyss of the true vacuum; our computers and automated systems may eventually achieve intelligence, or a reasonable facsimile thereof, and decide to dispense with us. About 30 or 40 such "small probability" extinction scenarios have been suggested by respectable scientists; no doubt even their aggregate probability is not very large, but we can't be sure that a few of them don't have a much greater objective probability than we suppose. And, of course, there is always the unknown: our species may soon enough face dangers that we can no more conceive of than Jules Verne could have conceived of hydrogen bombs.[17] In sum, we are not in a position to say that the antecedent probability (antecedent to our consideration of the Doomsday Argument) of humanity's coming to an end in the next few hundred years is negligible. Any argument, therefore, that tells us that this probability is significantly greater than we could have guessed simply by contemplating "doomsday scenarios" and trying to estimate their aggregate probability is of real interest. (For all sorts of reasons: for one thing, if it were generally accepted, it might lead us to be more cautious in our military and political and scientific undertakings. The conclusion of the Doomsday Argument is not, and has never been presented as, a *sentence* of doom; the argument concerns probabilities, not certainties.)

Secondly, there is a theological point to be made. Let us return to the eschaton expected by the Abrahamic religions. Suppose one believes, as I do, that there will someday be such an eschaton. If so, there will be no doomsday – not in the sense of the extinction of the human species,[18] for humanity has, we believers suppose, an eternal future. Is the Doomsday Argument therefore an argument that should lead us believers to revise downward the probability we assign to the expected eschaton (and therefore to the whole set of our religious beliefs, for the eschaton is an essential and inseparable component of those beliefs)? By no means. Doomsday, in both the popular sense and the strict theological sense (the eschaton) implies the end of human life as it has always been: the end of the cycle of reproduction and death, the end of the addition of new members to the human species. Proponents of metaphysical naturalism naturally take "the end of human life as it has always been" to be just exactly the end of human life. Those who believe in a coming eschaton, I would suggest, should regard the Doomsday

Argument this way: whatever probability they assign to the eschaton's coming in the next few centuries, they should regard the Doomsday Argument, or Bayesian reasoning in the doomsday style, as significantly raising this probability. If naturalists who find the Doomsday Argument cogent doubt whether Jews, Christians, and Muslims can so comfortably accommodate doomsday reasoning to their belief in an eschaton, I suggest they consider the following story. Suppose the search for extraterrestrial intelligence is finally successful, and we receive a message from an ancient species elsewhere in the galaxy. This information is contained in the message: there are many intelligent and technologically able species in the universe, hundreds of which the originators of the message have been in contact with, and *every* such species sooner or later achieves physical immortality, and thereafter ceases to reproduce, continuing in existence for geologically vast periods of time simply in virtue of the immortality of its individual members – those fortunate enough (or in my view unfortunate enough; but that's beside the point) to have been alive when their species achieved immortality. Let us call a species' transition to this state its secular eschaton. After this information has been given us, it is reasonable to for us to believe that the human species will one day experience a secular eschaton. Should not anyone who, before this information was made known to us, regarded the Doomsday Argument as significantly raising the probability that humanity would become extinct soon, now regard the argument, or the style of reasoning it embodies, as significantly raising the probability that humanity will achieve its secular eschaton soon?

Notes

1 Strictly speaking, I should have said something like "the laws of reality." The laws of nature coincide with the laws of reality only if nature coincides with reality: that is, if naturalism is true. If naturalism is false – because, say, there exists a God who has plans for humanity that are not constrained by the laws of nature – it may be that something awaits us in the future that belongs to none of the possible futures that are consistent with the present and the laws of nature. If God exists, "the laws of reality" will be just that set of propositions that supervene on the divine nature. If God does not exist but naturalism is nevertheless false – if naturalism is false for some other reason than that there is a God – the laws of reality would supervene (or at least partly depend) on the nature or natures of whatever beings it is whose existence is not a part of the natural order. (For I suppose that naturalism is false if and only if there are beings – not abstract entities, but beings with causal powers – whose existence is not a part of the natural order.) In the text, I shall, for the sake of simplicity, use the phrase "the laws of nature" to mean "the laws of reality."

2 For present purposes, "There is a probability of 0.57 that it will be the case that p" may be understood to mean, "In 57 percent of the physically possible futures that are continuous with the present, it will be the case that p." This conception of the probability of future events would not do for most purposes. It has the consequence that if the world is deterministic, then, for any p, the probability that it will be the

case that p is either 0 or 1; and it has the consequence that only in certain special cases (the case of a proposition asserting that an eclipse will take place at a given moment, for example) would it be possible to know or even make any reasonable guess about whether the probability of a proposition about the future was 0 or 1. This consequence of our conception of probability of a proposition about the future, therefore, shows that this conception can hardly be the conception that is employed by insurance companies. It will do for our highly abstract and theoretical purposes, however. In fact, it will not only do, but is exactly the conception our purposes require.

A remark for conceptual puritans: "In 57 percent of the possible futures, p" is a colloquial way of saying, "The measure of the set of possible futures in which p is 0.57 times the measure of the whole set of possible futures."

3 For citations of Nietzsche's scattered remarks on *die ewige Wiederkunft*, see Arthur C. Danto, *Nietzsche as Philosopher* (New York: Macmillan, 1965), chapter 7.

4 Nietzsche believed that he himself would return eternally, and embraced this idea joyfully. The Greek atomists found this a horrible possibility and maintained that the "future duplicates" of Nietzsche (or whomever) that the cosmos produced would not be Nietzsche but rather people, as we should say today, qualitatively but not numerically identical with him. It seems plausible to suppose that if the Greek atomists had read Kripke on the essentiality of species-origins, they would have maintained that the human-like beings of future cosmic reorderings would not be human beings but would rather be members of infinitely many numerically distinct species with the same anatomical and physiological characteristics as human beings.

5 And possibly the Old Norse or "Odinic" religion. I say "possibly" because the Old Norse eschatology cannot be said with certainty to be historically independent of Christian eschatology.

6 Indeed, most, if not all, pagan mythologies contain stories about how the physical world came to be.

7 Apparently something like the Greek-atomist position (minus, perhaps, intermittent cosmic dissolutions and reorderings) was not uncommon during the so-called ages of faith – although of course it is not well represented in the written records of the period, which were mostly the works of clerics and monks. The following words were written around the year 1200 by one Peter of Cornwall, prior of Holy Trinity, Aldgate: "There are many people who do not believe that God exists . . . They consider that the universe has always been as it is now and is ruled by chance rather than by Providence." Peter's manuscript is unpublished. Apparently he gives no explanation of what he means by "many people." These two sentences from his manuscript are quoted in Robert Bartlett, *England under the Norman and Angevin Kings, 1075–1225* (Oxford: Clarendon Press, 2000). I have taken them from a review of that book by John Gillingham, which appeared in *The Times Literary Supplement* (May 5, 2000), p. 26. I suppose the translation to be Bartlett's.

8 The Doomsday Argument is the work of several thinkers. The physicist Brandon Carter has the best claim to being its inventor. The most systematic exposition and defense of the argument is to be found in John Leslie, *The End of the World: The Science and Ethics of Human Extinction* (London and New York: Routledge, 1996). The presentation of the Doomsday Argument in the present essay is not based on this book, which I deliberately did not read till after I had finished a first draft of this

essay. I believe that nothing I say is inconsistent with Leslie's treatment of the Doomsday Argument. My presentation and defense of the argument (which is, of course, very abbreviated) contains nothing inconsistent with Leslie's. The main difference seems to me to be that I have made considerably more use of spatial analogies than Leslie does; but spatial analogies figure prominently in his treatment of the argument.

Like "the Big Bang," "the Doomsday Argument" is a dyslogistic name, coined by the intellectual opponents of its referent, and ultimately accepted by its proponents. As Leslie says, "[the more accurate name] 'anthropic argument suggesting that we have systematically underestimated the risk that the human race will end fairly shortly' would have been far too lengthy a label" (ibid., p. 194).

9 Suppose we have some evidence that is relevant to a thesis called Hypothesis. Call this Old Evidence. Suppose we acquire some relevant new evidence: New Evidence. Old Evidence and New Evidence together compose Total Evidence. How should we revise the "old" or "prior" or "antecedent" probability of Hypothesis (the probability of Hypothesis on Old Evidence) in the light of New Evidence? – that is, how are we to calculate the probability of Hypothesis on Total Evidence, given that we know the probability of Hypothesis on Old Evidence? Bayes's theorem is an answer to this question. Very roughly speaking, it tells us that the new probability of Hypothesis is a function of its antecedent probability and the degree to which adding Hypothesis to Old Evidence raises the probability of New Evidence. More exactly, Bayes's theorem tells us that the probability of Hypothesis on Total Evidence is equal to the probability of Hypothesis on Old Evidence multiplied by a certain ratio, the ratio of

The probability of New Evidence on (Old Evidence plus Hypothesis)

to

The probability of New Evidence on Old Evidence alone.

10 In these "hotel room" examples, I have glossed the fact that the sea may be invisible from a hotel-room window (for at least two obvious reasons) even if the hotel is close to the sea. No point of principle is affected by this fudge.

11 $0.01 \times 0.5 / [(0.99 \times 0.0002) + (0.01 \times 0.5)] = 0.9615384$.

12 Does this seem intuitively to be too high? Try the following game with playing cards (or do something equivalent with your computer or a table of random digits). Take 10 cards, ace to 10. Draw one at random and discard the higher cards. Then draw a card at random from the remaining cards. Do this a large number of times and record your results. You will find that you "end up with" the ace about 30 percent of the time. (And the deuce about 20 percent of the time, the trey about 15 percent of the time, the four about 10 percent of the time . . . and the 10 about 1 percent of the time. The sum of the 10 percentages is of course 100 (or would be but for my rounding off).

13 You can see that the odds my reasoner has calculated are *about* right if you consider the case in which the probability of the expected event is exactly $\frac{1}{3}$. Suppose someone wanted to bet that a thrown die would land either three or six. It's two to one that a thrown die won't land three or six: "won't" is twice as likely as "will." If you want

to bet that a die will land either three or six, a rational bookie will agree to give you any odds less than two to one: he'll agree to pay you $1.99 on each dollar you bet (or any smaller sum) if you win. Agreeing to pay a bettor $2.00 on the dollar in a bet that a die will land three or six would be a waste of a bookie's time, for, in an eternity of such bets, he'd simply break even. (The 12-cent difference – $2.12 versus $2.00 – is due to the difference between the probabilities 0.3333 . . . and 0.32000 . . .)

14 The sum of the first n integers (1, 2, . . . , n) is n(n + 1)/2.

15 30 percent rather than 32 percent because of the "graininess" of the present example. In this example, I am placed within a randomly chosen 100-mile segment of a river, and not at a randomly chosen *point* on the shore of a river.

16 For an extended discussion and evaluation of these scenarios (I can think of none he doesn't consider), see chapters 1 and 2 of Leslie's *The End of the World*.

17 One of Verne's characters – I think in *From the Earth to the Moon* – does speculate, in a rather jocular fashion, that the world may end when an enormous boiler, heated to a pressure of 10,000 atmospheres, explodes. That wouldn't actually be possible, or not without fantastic innovations in materials engineering: a boiler the size Verne's character imagines, and made of any material *we* could imagine, would instantly collapse under its own weight. But there was an "unknown" possibility, a possibility now actual, the possibility of a device just as apt to destruction as the impossible device Verne put into his character's mouth.

18 In modern English, "doom" is hardly more than a romantic word for "death" and our understanding of "doomsday" reflects this sense of "doom." The original meaning of "doom," however, was something like "judgment," and the original meaning of "doomsday" – as in "Domesday Book" – was "day of judgment."

Suggested Further Reading

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

Edwards, Paul (ed.) (1992) *Immortality*. New York: Macmillan.

van Inwagen, Peter (1997) *The Possibility of Resurrection and Other Essays in Christian Apologetics*. Boulder, Colo.: Westview Press.

Zimmerman, Dean (1999) The Compatibility of Materialism and Survival: The "Falling Elevator" Model. *Faith and Philosophy*, 16, 194–212.