9

The Rev'd Mr Bayes and the Life Everlasting

Peter van Inwagen

Most philosophers will have encountered the "Doomsday Argument," a Bayesian argument for the conclusion that our species will come to an end sooner than we might have expected—that the probability that we should assign to the hypothesis that the end of the human species will come within the next few centuries is significantly higher than the probability we should assign to that hypothesis if we considered only causal factors like thermonuclear or environmental or epidemiological catastrophe—or extraterrestrial disaster scenarios: "big rock hits the earth," "nearby supernova," and the like.

The question has recently been raised whether, if this argument indeed establishes its conclusion, a closely parallel argument might show that it is highly unlikely that there is any such thing as an afterlife—or at any rate a very lengthy one.[1] This question is my subject. I begin with a review of the Doomsday Argument.

1

In order to eliminate some of the messy issues that a satisfactory analysis of the Doomsday Argument as it applies to the real world would involve (such as increasing and decreasing population levels), I'm going to show

[1] See Page (2010: 399–401) and Leslie (2008). In this chapter I attempt to think that question through from first principles and do not directly discuss Page's and Leslie's arguments.

how the argument would work in application not to the end of our species, but to the end of a universe—and not *our* universe, but a very simple universe: the Toy Universe, I'll call it.

Imagine that I am an inhabitant of the Toy Universe. At a certain point in my life, when our story opens, I know very little about the temporal features of the Toy Universe (which, within the fiction that I am one of its inhabitants, I shall refer to as simply the universe). Among the few things I do know is this:

> The universe had a beginning in time: there was no first moment at which it existed, but the set of moments at which it has existed has a greatest lower bound.

I do not, however, know anything about how long it has existed, other than this:

> The universe has existed for less than 1 million years.

(This implies that I know nothing about how long *I* have existed: for example, for all I know, I have existed for only a few seconds. Make this an intelligible assumption in any way you like—one possibility would be to suppose that for all I know, I was created, along with the rest of the universe, a few seconds ago, complete with memories of "a wholly unreal past.") I know that the universe has existed for less than 1 million years because I know that:

> The total span of the existence of the universe will be 1 million years or less.

But I would like to give a rather more precise description of my "initial" knowledge of the temporal features of the universe than this. In order to do this, I will introduce a precise system of dates. Let our unit of temporal measure be the terrestrial sidereal year. Let any real number greater than or equal to 0 be a "date." Then the first year of the existence of the universe comprises those dates greater than 0 and less than or equal to 1, the second year of the existence of the universe comprises those dates greater than 1 and less than or equal to 2, and so on. And the date 23.48592 denotes a period of about 5 minutes that ends approximately 23 years and 15,334,736 seconds after 0. One thing I know is this:

> The universe will come to an end at some point in the interval (0, 1M].

('1M' being my abbreviation for '1,000,000.') And I also know *this*:

> The date of the End is the outcome of some random process—of some ontologically random process that "chose" a single number x that satisfied the condition $0 < x \leq 1\text{M}$ (and might "just as well" have chosen any other number in that interval as the one it did choose). More exactly, the objective probability of the proposition that the universe would end in any given year (not greater than 1M) was (at the moment of creation, so to speak) the same as the objective probability of the proposition that it would end in any other given year: for each year, that probability was literally one in a million. Similarly, the objective probability of the proposition that the universe would end in any given 100-second subinterval of (0, 1M] was the same as the objective probability of the proposition that it would end in any other 100-second subinterval of that interval. Generalizing: For any four dates, m, n, i, and j (all less than or equal to 1M). If $|m-n| = |i-j|$, then the objective probability of the proposition that the End would come between m and n was equal to the objective probability of the proposition that the End would come between i and j: $|m-n|/1\text{M}$ or $|i-j|/1\text{M}$.

Now consider the following hypothesis, which I will call 'K' (for kiloyear).

> The End will fall within (0, 1,000].

What subjective probability should I assign to this hypothesis, given what I know "as our story opens?" Neglecting the "nuisance fact" that, if I consider this question, I must know that, since I have had *time* to consider it, at least a few seconds of (0, 1,000] must already have passed, the answer is evident: the same probability that I should assign to the hypothesis that it will end in any other millennium: one in one thousand, or 0.001.

Now suppose that at some point I acquire the following piece of information (which I will call 'H' for hundred or hectoyear):

> The present date falls within the interval (0, 100].

(That is to say, *this* moment is a moment that falls during the first century of the existence of the universe.) It is important to emphasize that I acquire no more specific information about the present date than that. Imagine, if you like, that the whole episode of my acquiring this new piece of knowledge consisted in God's handing me a slip of paper on which this alone was written: "The present date falls within the interval (0, 100]." So, for all I know, the present year could be 1 or 23 or 87 or 100 or any other year that is not greater than 100.

The Doomsday Argument (applied to the present simple case) is essentially a Bayesian argument for the conclusion that my knowledge that the present moment falls within (0, 100] should lead me to raise the probability I assign to any hypothesis to the effect that the End will come "fairly soon."[2] And not only raise it, but to raise it *significantly*.

It should, for example, lead me to raise—significantly—the probability I assign to K. But why—and by how much?

It is possible to reason one's way to an answer to these questions, but the required piece of reasoning is rather complicated. At any rate, it is complicated if one carries it out from scratch and does not simply begin *in medias res*—by saying something like, "We can calculate the probability I should assign to K, given my new piece of knowledge, H, using Bayes's Theorem, which, as we all know, is . . . " I propose to start from scratch. I will present the complex piece of reasoning I have alluded to in the form of a commentary on some figures—a sort of visual aid—that I will call *probabilistic Venn diagrams* (PVDs).

A PVD is much like an ordinary or "logical" Venn diagram except that:

(i) A PVD contains an outer circle, L, which represents the whole of "logical space;" all the other circles in a PVD are inside this outer circle.

(ii) In an ordinary Venn diagram, each circle represents a class of things. A circle in a PVD, however, represents a proposition or hypothesis (or, if you like, a class: the class of possible worlds in which that proposition is true). In fact, any subregion of L, however shaped or gerrymandered, represents a proposition, whether that subregion has been "outlined" or not. For example, the region common to two or more circles represents a proposition, the conjunction of the propositions represented by those

[2] And, in fact, for any positive number a, $a < 1M$, to raise the probability that I assign to the hypothesis that the End will come in the interval $(0, a]$. As a approaches 1M, however, the prior probability of that hypothesis and the revised probability become closer and closer to each other.

circles. (L, the outer circle, represents the proposition that something is the case.)

(iii) In a PVD, unlike an ordinary Venn diagram, the size (that is, the area) of a region within the outer circle—its area in proportion to the area of the outer circle, the proportion of the outer circle it occupies—is significant. (We will use the area of the outer circle as our—dimensionless—unit of area measure: its area is 1. If, therefore, a region inside the outer circle has an area that is 1/17th the area of the outer circle, its area is 1/17 or about 0.059.) The area of a region represents the *probability* of the hypothesis the region represents. For example, the region within the outer circle mentioned in the parenthesis represents an hypothesis that has one chance in 17 of being right—or has a probability of 0.059.

In the sequel, we will, to simplify our exposition, speak as if a circle in a PVD and the hypothesis it represents were one and the same thing.

Now consider two overlapping—that is, consistent—hypotheses: A and B. We abbreviate "the proportion of B that is A" as "A//B". If, for example, 10 percent of the region B coincides with a part of A (if the area of A ∩ B is one-tenth the area of B), A//B = 0.1. If a point in B is chosen at random—by throwing a dart, perhaps—, there is one chance in ten that that point will also be a point in A. We may therefore think of A//B as the probability of A, given that B—or the conditional probability of A on B.

Now let us draw a PVD that displays our two hypotheses: K and H. (These two hypotheses are obviously consistent and, obviously, neither entails the other. We therefore represent them by circles that overlap, but only partly. And each obviously has a probability of less than unity, so we draw them both as "subcircles" of the outer circle.)

Having done this, let us proceed to "normalize" the region of logical space that comprises all the propositions about the existence of the Toy Universe and its temporal features that I knew before I discovered that the present date fell within (0, 100]. That is, let us agree to ignore the region of logical space in which the conjunction of those propositions is false and assign a new measure of 1 (formerly the measure of the whole of logical space) to the region in which it is true. Call this region T and represent it by a circle—our "new" outer circle. (The circles H and K of course fall within T.) Doing this will have the effect of increasing the probability measure of every subregion of T by the same proportion as that by which the measure of T was increased. Suppose, for example, that T occupies one-seventeenth of logical space. Then, when the measure of T is normalized to 1, the "new" measure of every subregion of T must be seventeen times its "old" measure—since 1 is seventeen times one-seventeenth. If, for example, A is some subregion of T whose old measure was 0.02, the new measure of A is 0.34 (= 17 × 0.02). More generally, suppose that a region of logical space A is normalized and that B is a subregion of A; if A_{old} is the old area of A and if B_{old} is the old area of B, then B_{new}, the new area of B, is equal to $(1/A_{old}) \times B_{old.}$

We may now draw Figure 9.1, a PVD, on the left below.

Now the way I have drawn Figure 9.1 displays more than I have stipulated about the relations between T and H and H ∩ K. Figure 9.1 represents H as occupying a certain determinate proportion of T, and it represents H ∩ K as occupying a certain determinate proportion of H and a certain determinate proportion of K. (K *does* occupy a certain determinate proportion of T: it occupies one one-thousandth of T—obviously the figure is not drawn to scale—, but nothing that has been said has any implications as regards the size of H or of H ∩ K.) In displaying such "extraneous information," the diagram resembles a diagram illustrating a proof of a theorem concerning, say, the interior angles of a triangle in a plane-geometry textbook; all but certain features of the "sample" triangle shown in the diagram are to be ignored: all of them that do not actually figure in the proof of the theorem.

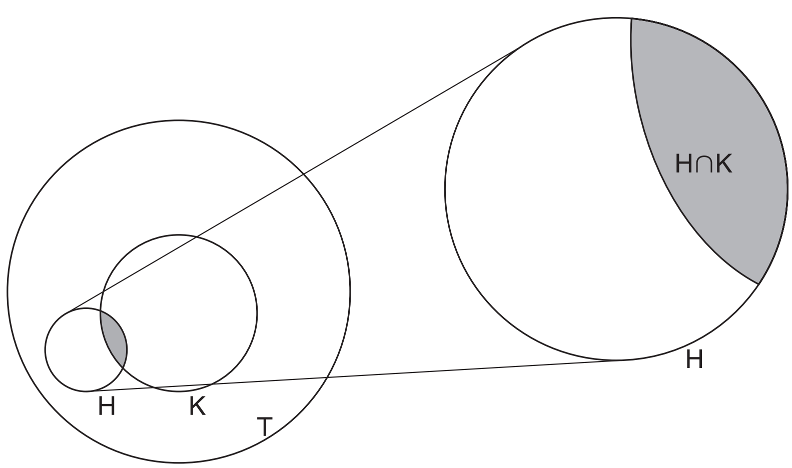

Figure 9.1 **Figure 9.2**

In what follows, I will abandon the first person: *we* shall try to see how to solve the problem that faces the fictional "me," the inhabitant of the Toy Universe: How should my coming to know the truth of H affect the probability I assign to K?

When *we* look at Figure 9.1, therefore, *we* are to attend only to certain of its features: that both circles lie inside T, that they partly overlap, and that the (new or normalized) area of K is 0.001—although this last "feature" cannot be said to be perspicuously displayed in the figure. (A triangle in a geometry-text diagram has to be drawn with *some* determinate interior angles; the spatial regions that represent H and H ∩ K in a PVD have to have *some* sizes.)

Now here is what we want to know: what is the value of K//H? (That is, what probability should we assign to the hypothesis that the End will come in the first thousand years of the existence of the Toy Universe, given that we know that the present falls within the first hundred years of its existence?)

We can determine this value (given a certain assumption that I will come to presently), and the reasoning by which we obtain it will, as I said, be presented in the form of a commentary on a PVD—much as the reasoning that leads to "The sum of the interior angles of a triangle is equal to the sum of two right angles" can be presented in the form of a commentary on a diagram that displays, among other things, a triangle and its interior angles.

We know that H is true. Let us therefore normalize the measure of H to 1. The result of this operation is shown in Figure 9.2, given earlier. The shaded region, of course, is H ∩ K, or the conjunction of H and K.

As I have said, to normalize a region of logical space A is to multiply the area of all its subregions by the reciprocal of the area of A—that is, the area of A before we "redefined" the measure of that area to be 1. If, for example, the "old" area of A was 0.2, the "new" area of A is 0.2 × 1/0.2, which is, of course, 1. And if B was a subregion of A with area 0.013, the new area of B is 0.013 × 1/0.2 or 0.065.

Now consider the shaded region in Figure 9.2, H ∩ K. Now that we are in effect treating H as the whole of logical space—now that its area has been normalized to 1—, the new area of H ∩ K is simply the proportion of H that is K. (If, for example, the new area of H ∩ K is 0.3, then three-tenths of H is occupied by a part of K.) That is to say, now, after normalization, the area of H ∩ K = K//H. (Note that the values of proportions like K//H and H//K do not change when we normalize the area of H to 1, since the areas of H and K increase *proportionately*.)

And we can say something about the present area of $H \cap K$ (now equal to K//H) in terms of H//K (the proportion of K that is H) and the old areas of H and K and $H \cap K$. The old area of $H \cap K$ was equal to H//K times the old area of K. (If a plane region is one-tenth red, the area of its red part is one-tenth of its total area.) As we did on an earlier occasion, let us use subscripts to keep all these areas straight: 'H_{old}' will denote the old area of H, and similarly for 'H_{new},' 'K_{old},' 'K_{new},' '$(H \cap K)_{old}$,' and '$(H \cap K)_{new}$.'

We have:

$$(H \cap K)_{old} = H//K \times K_{old}$$

And we know that $(H \cap K)_{new}$ is equal both to K//H and $(H \cap K)_{old} \times 1/H_{old}$. We therefore have

$$K//H = (H \cap K)_{new} = H//K \times K_{old} \times 1/H_{old},$$

or

$$K//H = H//K \times (K_{old}/H_{old}),$$

an expression some of you may recognize despite my non-standard notation.[3] And a way of calculating K//H—that is, the probability I should assign to the hypothesis that the End of the Toy Universe will come within its first thousand years (K), given my knowledge that the present moment lies within the first century (H), is just what I promised to deliver. Of course, to make that calculation, I still need to determine three numbers: the prior probability that the Toy Universe will end in its first thousand years (the probability I assigned to that hypothesis prior to my discovery that the present moment lay in the first century), the prior probability of the present moment's falling within the first century, and the probability of the present moment's falling in the first century given that the Toy Universe will end at some point in its first millennium.

The prior probability that the Toy Universe will end at some point in its first thousand years is easy to come by: it is, as we have seen, 0.001 or one one-thousandth. The other two numbers are more difficult. Let us

[3] In standard notation: $P(K|H) = P(H|K) \times (P(K)/P(H))$.

begin with the third, H//K or the proportion of K that is H or the probability that the present falls in the first century, given that the End falls in the first millennium. I can deliver this number if you allow me a certain assumption.

> *The Crucial Assumption*
> I may regard the present moment as a moment chosen at random from the interval (0, the End].

This, I believe, is the assumption on which the Doomsday Argument turns. I will use it in the calculation that follows, and in one other calculation: the calculation of the prior probability of H.

Let P be the present moment and let E be the End. Let us assume that E falls within the first millennium and see what this assumption can tell us about the probability we should assign to H.

Notice that if E fell within the first century, the probability of H would be unity, and that if E fell in the interval (100, 1,000] it would be less than unity. If the date of E were exactly 200, the probability of H would be 1/2—owing to the fact that in that case P is a moment chosen at random from the interval (0, 200] and hence has a 50–50 chance of falling within the interval (0, 100]. (Note that the Crucial Assumption was a premise of that little argument.) And, similarly, if the date of E were exactly 300, the probability of H would be 1/3—and so on. Suppose we knew that the date of E was *one* of the "century" dates (100, 200, 300, . . . , 900, 1,000)—and knew the probability of the date of E's being any one of these was equal to the probability of its date being any one of the others—and knew nothing else about its date. Then, intuitively, we should have to assign to H a probability that was the average value of the numbers 1, 1/2, 1/3, 1/4, 1/5, 1/6, 1/7, 1/8, 1/9, and 1/10—or roughly 0.2929.

This intuitive calculation suggests that, if we measure time in centuries rather than years, the probability of H on K should be the *average* value on the interval (0, 10]—that is, the interval that runs from 0 centuries to 10 centuries—of a function *f* defined on that interval such that

$$f(x) = 1 \text{ if } 0 < x \leq 1$$
$$f(x) = 1/x \text{ if } 1 < x \leq 10.$$

The arguments of *f* are dates greater than 0 centuries and less than or equal to 10 centuries. The values of *f* are probabilities. For example,

$f(5.34)$ is the probability that P falls within the first century given that the date of E is the final moment of the year 534—about 0.187.

The average value of a function on an interval is given by the Mean Value Theorem, which is highly intuitive, at least in the first quadrant. Since the interval we are considering ((0, 10]) has 0 as its greatest lower bound, we need consider only such intervals. The following diagram (Figure 9.3) illustrates the Mean Value Theorem for this special case.

In this diagram, g is the function whose average value on some interval (0, a] we wish to determine. That interval is shown in the diagram as the interval between the $g(x)$-axis and the vertical line AB—an interval whose length is a. The line CD represents the average value of g on that interval—that is, its "height" above the x-axis ('h' in the diagram) is the average value of g on that interval. The Mean Value Theorem tells us that the area of the rectangle with base a and height h is equal to the area of the shaded region—the area of the region under the graph of g whose width is the same as the rectangle with base a. And that's

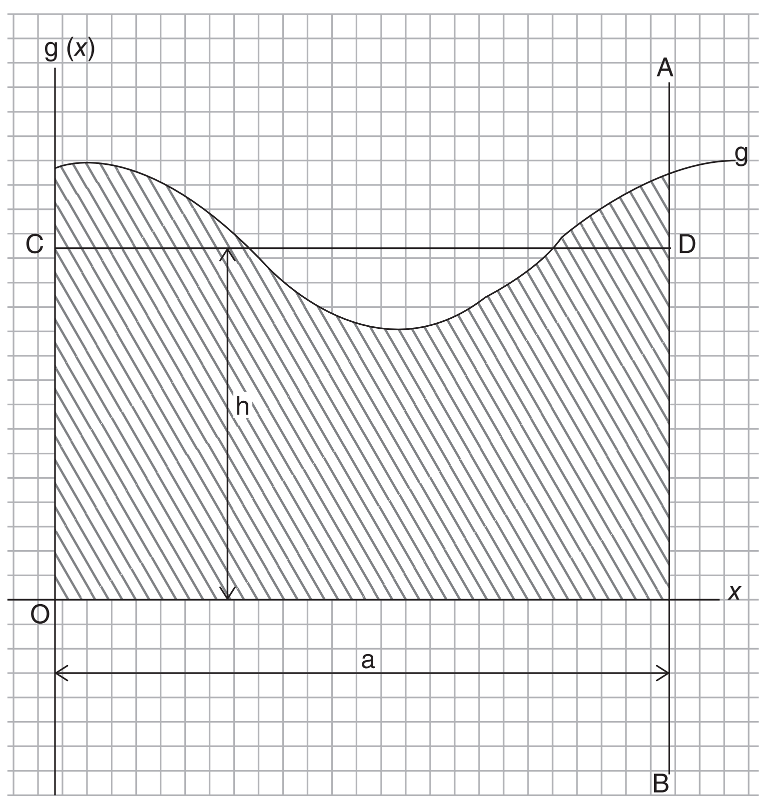

Figure 9.3

so intuitive an idea that in some developments of real analysis it's the basis of the *definition* of the average value of a function on an interval. Incidentally, since we're in that part of mathematics, real analysis, we treat both the areas of regions and the lengths of intervals as simple numbers (the measures of the regions and the intervals)—that is, we needn't worry about units and square units as we do when we make calculations about distances and areas in the physical world.

So armed, let us return to our function f and calculate its average value v on the interval (0,10]. The hard part of this problem is to determine the area A under f in that interval. If we know what A is, then, by the Mean Value Theorem, $A = 10v$—that is, A is equal to the area of a rectangle whose base is 10 and whose height is v. And since $A = 10v$, $v = A/10$. Here is a diagram (Figure 9.4) that represents the relations between A and v.

Figure 9.4 represents a consequence of the Mean Value Theorem: that the average value of f on the interval (0, 10] is equal to the height of the rectangle drawn in the figure, the rectangle whose area is equal to the area (A) of the shaded region.

And what is A? What is the area under f on the interval (0, 10]?—or since the area is the same whether the interval is closed or open at either

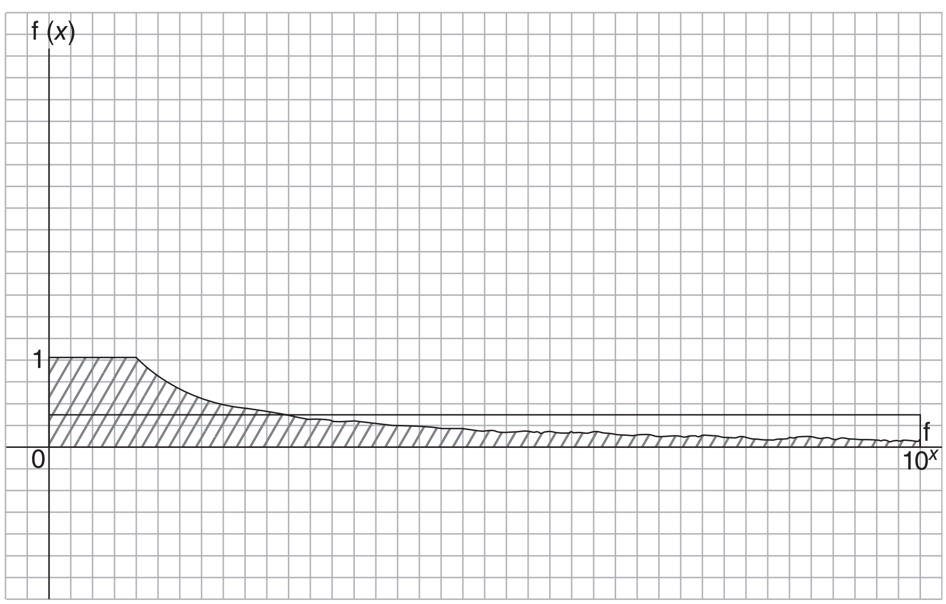

Figure 9.4

end, we can simply ask what the area under *f* is in the region that starts at 0 and ends at 10.

The function *f* is defined partwise, so the question divides. The area of the region (a square) under the "flat" part of the function is 1×1 or 1. The area under the curved part of the function, the "1/*x*" part, is

$$\int_1^{10} \frac{1}{x}\,dx$$

And that definite integral, if you remember your first-year calculus, is equal to the natural logarithm of 10 minus the natural logarithm of 1. Any sort of logarithm of 1 is equal to 0, so the area under the curved part is equal to the natural logarithm of 10, and the area under the flat part *plus* the curved part is 1 plus the natural logarithm of 10 (usually written "1 + *ln*10"). And, since, as we have seen, the average value of *f* between 0 and 10 is equal to one-tenth the area under its graph on that interval, the average value of *f* between 0 and 10 is one-tenth the sum of 1 and the natural logarithm of 10, or (1+ *ln*10)/10. You can look up the values of logarithms in tables, or, these days, on your smart phone (which will also do the arithmetic for you), and, having consulted my iPhone, I'm in a position to tell you that the result of our calculations is roughly 0.3302—not *too* different from the figure (0.2929) we got by averaging the ten fractions 1/1, 1/2, 1/3, . . . 1/9, 1/10.

And that—0.3302—is the probability that P lies in the interval (0, 1], given that E lies in the interval (0, 10]. (Or, measuring temporal intervals in years rather than in centuries, the probability that P lies in the interval (0, 100], given that E lies in the interval (0, 1,000].)

Now, what is the *prior* probability of H—the hypothesis that P falls in the first century? We can answer this question very simply by the method we have just used if we note that it is equivalent to the question: What is the probability of the hypothesis that P falls in (0,1] given that E falls *somewhere*, at some randomly chosen point, in (0, 10,000]? (One million years is ten thousand centuries.) Our previous answer was (1 + *ln*10)/10 $\approx$ 0.3302. To get the number we now require, we need only substitute 10,000 for 10: the prior probability of H is (1 + *ln*10,000)/10,000 or (the heavy-duty calculator in my iPhone tells me) approximately 0.00102. And now, finally, we have the numbers we need to use the equation

$$\mathrm{K}//\mathrm{H} = \mathrm{H}//\mathrm{K} \times (\mathrm{K}_{\mathrm{old}}/\mathrm{H}_{\mathrm{old}})$$

The probability that the Toy Universe will end at some point in its first thousand years, given that the present moment falls in the first century of the existence of the Toy Universe (and given the Crucial Assumption—that I may treat the present moment as a moment drawn at random from all the moments at which the Toy Universe exists[4]), is, therefore,

$$0.3302 \times (0.001/0.00102)$$

You don't need a heavy-duty calculator for that one: any old calculator will do.

And the answer is . . .

$$0.3234$$

Before I realized that the present moment was in the first century—when I had no idea what the present date was other than that it was not greater than 1,000,000—, I assigned a probability of 0.001 to the hypothesis that the Toy Universe would end in the first millennium, the same probability I assigned to the hypothesis that it would end at some point in the 576th millennium or at some point in any other given millennium among the first 1,000 millennia. Now, it seems, having gained that piece of knowledge, I must assign that hypothesis a probability more than 300 times greater than 0.001—certainly a case of a significant upward revision of a prior probability!

And, of course, if I revise some probabilities upward, I must revise others downward. The prior probability of the hypothesis that the Toy Universe will end at some point in the interval (1,000, 1M] was 0.999. The revised probability is 1 − 0.3234 or 0.6733: a near certainty has become a probability of about 2/3. The prior probability of the hypothesis that the Toy Universe would last at least half a million years was 0.5. The revised probability is close to 0.07 or 7 percent: an even chance has become a very bad bet.

I will not, in this chapter, examine the Crucial Assumption or question the cogency of the Doomsday Argument in any other way. I want instead

[4] Given the Crucial Assumption, the problem of determining K//H becomes essentially a problem of pure mathematics. Stated very abstractly, it is this:

> Let a number B be chosen at random from the interval (0, 1M]. Once that choice is made, let a number A be chosen at random from the interval (0, B]. What is the probability that B falls within the interval (0, 1,000], given that A falls within the interval (0, 100]?

The solution is, of course, 0.3234.

to raise the following question: Assuming that the kind of reasoning illustrated here is correct, can parallel or analogous reasoning be used by each of us human beings to show that he or she should assign a significantly lower probability to the hypothesis that he or she will exist 1,000 years from now than the probability that that person would otherwise have assigned to that hypothesis?

You may wish to remind me that many people already assign a probability of 0 (or a probability so close to 0 that the difference is of no more than abstract, theoretical interest) to this hypothesis. And that is certainly true. Still, not everyone does. Its probability, conditional on the truth of the teachings of more than one religion, is 1, and many people who practice a religion assign (in effect, perhaps) a probability of 1 to the hypothesis (if Kierkegaard will forgive my use of that word in this context) that the teachings of their religion are true—or, if not 1, then certainly some probability significantly higher than 0.5. My question could be put this way: Can Bayesian reasoning parallel to or analogous to or in some way suggested by the "Doomsday" reasoning I have presented be used to show that these probability assignments are wrong or irrational or in some other respect objectionable?

2

Let us construct an argument parallel to the Toy Universe argument, an argument that turns on the question how long I shall exist.

Let us suppose that God has revealed to me that my existence will end at some randomly chosen point within 1 million years of my conception (or whenever it was that I began to exist). Let us suppose I know that I am now 100 years old or less, and that I know nothing more specific about my age: for all I know I am 1 year old or 23 years old or 99 years old. Or, again refining our system of chronology, for all I know I am 0.783 years old (to the nearest thousandth of a year) or I am 78.662 years old or . . .

What probability should I assign to the hypothesis (K) that I shall exist for 1,000 years or less if I know (H) that my present age is 100 years or less?

I can answer this question if—and only if—I am granted a certain assumption: the "personal analogue" of the Crucial Assumption of the Doomsday Argument.

> *The Crucial Assumption (Personal)*
> I may regard the present moment as a moment chosen at random from the interval (0, the final moment of my existence[5]].

(Here, of course, '0' represents the greatest lower bound of the set of my "possible ages" and my "possible durations": whatever my present age may be, it is greater than 0 years; whatever the length of the total span of my existence may be, it is greater than 0 years.)

Given this assumption, we have already, in effect, calculated this probability: it is 0.3234. That is, my knowledge that I am less than 100 years old raises the probability of K from 1 in a thousand to about 1/3. And, therefore, it *lowers* the probability that I shall exist for *more* than a thousand years from 0.999 to about 2/3—and lowers the probability that I shall exist for at least 500,000 years from 1/2 to about 7 percent.

It seems, therefore, that Doomsday-style reasoning can be adapted to show that even if I can expect to continue to exist after my death, I cannot reasonably expect to exist for *very long*—not as the religions that preach a post-mortem existence measure "very long," at any rate. Not, for example, "so long as the sun and the moon endure." If we say that the sun will endure till it leaves the main sequence (it will destroy the moon when it does, if the moon is still around), that period is about 5 billion years. Suppose God has told me that I myself shall "endure" for some randomly chosen number of years between now and 10 billion years from now and that I know nothing more specific about the duration of my existence than that. Then the prior probability that I shall endure for as long as the sun and the moon is 0.5. But the probability conditional on my being less than 100 years old (given the Personal version of the Crucial Assumption) is 0.0361—less than 4 percent.[6]

[5] I assume for the sake of simplicity that there will be a moment that is the last moment of my existence—rather than assuming only that there will be a moment that is either the last moment of my existence or the first moment of my non-existence. Incidentally, I doubt whether such mathematical niceties have any application to the real world. If I am one day destined to cease to exist, I doubt whether there will be even a final nanosecond of my existence, never mind a mathematical point in time that is the least upper bound of the set of points in time at which I exist.

[6] "The upper bounds in your two calculations—a million years, ten billion years—are entirely arbitrary. It would be more interesting if you told us what the probability was that you would exist for (say) five billion years if you knew that you were less than 100 years old and knew that you would cease to exist at some moment chosen at random *from the whole*

These calculations, however, ignore the fact that the afterlife that the religions of the world promise, those that do promise an afterlife, is supposed to be everlasting, to go on for ever. No religion I know of teaches that a supernatural authority has decreed that we human beings are to exist post-mortem, but only for some finite period.

Suppose, then, that we address the following question: Where F is the hypothesis that I shall exist *for ever*, what is F//H? According to the Rev'd Mr Bayes, the answer is given by the equation

$$F//H = H//F \times (F_{old}/H_{old})$$

(Or so we are often given to understand. I have never read Bayes's "Essay towards Solving a Problem in the Doctrine of Chances," but I have often been told that it does not actually contain "Bayes's Theorem.") Let us see whether we can indeed use Bayes's Theorem to calculate F//H. We begin by considering H//F.

We must first ask how H//F is to be defined. We can't define it this way (again we shall be measuring intervals of time in centuries):

> Chose a date (that is, a real number) at random from the set of all dates (from the set of all real numbers greater than 0). H//F is the probability that that randomly chosen date is less than or equal to 1.

We cannot define H//F this way because there is no such thing as choosing a real number at random from that set—or, indeed, choosing at random a number from any set of numbers that has a lower bound and no upper bound (cf. footnote 6).

I would suggest a "limit" definition. We have in effect seen that if the Crucial Assumption (Personal) is true, then for any date x greater than 1, if x is the final moment of my existence, the probability that the present moment lies in (0,1] is equal to $1/x$. I propose, therefore, that we define H//F as the limit of the reciprocal function as its argument increases without bound. On this definition, and no other definition suggests itself, the value of H//F is 0, since the function $f(x) = 1/x$ is asymptotic to the x-axis.

infinite future." I will be happy to tell the Interlocutor what this probability is—but first she must tell me what it means to choose a moment at random from the whole infinite future (or to choose a number at random from the set of all real numbers greater than 0, or from any set of numbers with a lower bound but no upper bound).

Given this value for H//F, it is tempting to reason as follows:

F//H is 0, since it is equal to the product of 0 and the ratio or fraction F_{old} /H_{old}.

But this reasoning is correct[7] only if the ratio F_{old}/H_{old} exists. And it does not exist if H_{old} is equal to 0. Is it?

Before we can answer this question, we must ask what we mean by 'H_{old}.' What is the prior probability of H, the probability I should assign to H before I consider the "effect" on that probability of my knowledge that I am 100 years old or less, and when I know *nothing* about the total span of my existence—when, for all I know, I shall (a) exist for some unknown finite period of time, or (b) exist for ever. I can think of only one way to define this 'H_{old}.' We have in effect seen that if the Crucial Assumption (Personal) is true, then for any date x, the probability that the present moment falls in (0,1], given that the final moment of my existence is *some* date in (0, x], is equal to $(1 + ln\ x)/x$. I propose, therefore, that we define H_{old} as the limit of $f(x) = (1 + ln\ x)/x$ as x increases without bound. That limit exists, and it is equal to 0.[8] This thesis is hardly unintuitive. The natural logarithm of x is the power to which the number e must be raised to obtain x, e being an irrational number approximately equal to 2.718. To "obtain" 10,000, for example, we must raise e to a power very close to 9.21. ($2.718^{9.21}$ is about 9,987.) That is to say, the natural logarithm of 10,000 is approximately 9.21. It should be obvious that although $1 + ln\ x$ increases as x increases, x increases "much faster" than $1 + ln\ x$—with the consequence that as x increases, $(1 + ln\ x)/x$ is going to get very small very fast. And it is intuitively evident that for every real number y greater than 0, a sufficiently large choice of x will result in $(1 + ln\ x)/x$ having a value greater than 0 and smaller than y.

[7] Correct or not, it ought to arouse our suspicions immediately, because it implies that the probability of the hypothesis that I shall exist forever conditional on the hypothesis that I have not "already" existed for an infinite period of time is 0.

[8] Here is the sketch of a proof: $(1 + ln\ x)/x = (1/x) + (ln\ x)/x$. So, by a well-known fact about limits, the limit of $f(x) = (1 + ln\ x)/x$ as x increases without bound is equal to the sum of the limits of $f(x) = 1/x$ and $f(x) = (ln\ x)/x$ as x increases without bound. We have seen that the limit of $f(x) = 1/x$ as x increases without bound is 0. That the limit of $f(x) = (ln\ x)/x$ as x increases without bound is 0 can be shown by an application of L'Hôpital's rule.

If, therefore, this definition of H_{old} is adopted, the ratio F_{old}/H_{old} does not exist and the product of F_{old}/H_{old} and H//F does not exist. (For no x does $0 \times (x/0)$ exist.) And, therefore, Bayesian reasoning cannot be used to calculate the value of F//H.[9]

3

In closing, I will say something more about "vast but finite afterlife" cases. These are interesting for their own sake, but considering them will also afford an opportunity to examine and evaluate the Crucial Assumption (Personal). (If the position defended by Lara Buchak—see footnote 9—is correct, this examination will have important ramifications for the "infinite afterlife" case.)

Imagine a species I will call the Centurions, since, without exception, they all live for exactly 100 years. The Centurions inhabit our own planet a million or so years in the future. Some Centurions are followers of a religion called Eonism, whose principal teaching is that when Centurions die, they continue to exist—but not eternally: the total span of their existence (including their century of pre-mortem existence) will be 5 billion years. Although they will not exist forever, they will exist as long as the sun and the moon endure. (Or till a' the seas gang dry and the rocks melt wi' the sun. The two periods are more or less equal, since the sun's departure from the main sequence will do for the sun in its present form and the moon and the seas—if the moon and the seas still exist—in one fell swoop. And, as for the rocks, they will quite literally melt wi' the sun.)

[9] When this chapter was read at a conference at Purdue University in honor of Richard Swinburne's 80th birthday, the commentator was Lara Buchak. In her brilliant comments on the chapter, Buchak presented an alternative way of calculating the value of H//F. If the value of H//F is calculated in that way, it is equal to 0. While I have questions and reservations about Buchak's way of calculating H//F, the issues involved in evaluating it are very subtle, and I have to concede that she may be right. (She is certainly a much better mathematician than I am.) If I were to attempt to evaluate her proposal, however, this chapter would have to be much longer than it is. I will content myself with pointing out that her way of calculating H//F depends—as does mine—on a sort of "limit analogue" of the Crucial Assumption (Personal), and that in the sequel I shall present reasons for rejecting the Crucial Assumption (Personal).

Can Bayesian reasoning be used to prove to the Eonists (during their earthly lives) that it is highly improbable that this central teaching of their religion is correct?

We have already seen that (always granting the Crucial Assumption (Personal)) if one knows that one's existence will end at some randomly chosen point in the next 10 billion years, the discovery that one has existed for at most 100 years lowers the probability of the hypothesis that one will exist for at least 5 billion years from 0.5 to about 0.036—lowers it to about one-fourteenth of its original value. On this pass, however, I'll consider a different sort of Bayesian Argument, one that contains hardly any mathematics (no doubt to the reader's considerable relief).

Among the Centurions who are not Eonists, there are, so to call them, Eonic agnostics. Eonic agnostics hold that either Centurions enjoy no post-mortem existence *or* that the Eonists are right and that all Centurions exist for 5 billion years—and Eonic agnostics accept *neither* of the disjuncts of that proposition. Eonic agnostics, moreover, do not assign a negligible probability to either disjunct. One of them may assign a lowish probability to one of the disjuncts (1 percent, say), but none of them assigns to either disjunct a probability on the order of the probability that I assign to the proposition that I shall die at 11:00 a.m. on June 17, 2019 as the result of having been gored by a water buffalo while crossing Times Square. Can Bayesian reasoning provide them with grounds for lowering the probability that each of them assigns to the 5-billion-year hypothesis to some fraction, perhaps some minute fraction, of the probability he or she had hitherto assigned to it?

Let us consider a particular Eonic Agnostic, Agnes. Every day, Agnes devotes some serious thought to the question: Which is true, "No post-mortem existence" (NPE) or "Five billion years" (FBY)? One day, while she is thinking about this question, the Crucial Assumption (Personal) forces itself upon her mind as true (to borrow a phrase of Gödel's). And that assumption, we remember, is:

> I may regard the present moment as a moment chosen at random from the interval (0, the final moment of my existence].

"Let's see," Agnes muses. "I'm 43 years old. I may therefore regard the number 43 as a number drawn at random from the number of years from the beginning to the end of my existence." Reflection on this fact—she supposes it to be a fact—leads her to construct a carefully formulated analogy:

> Suppose that there are two enormous urns, A and B, standing side by side. Urn A contains, as an urn should, only empty space within its walls. Urn B, however, could almost be described as a fake urn or a mock urn. It contains only a little empty space at the top; if you looked down into it, you'd see a sort of floor not very far below you: beneath that level surface lies only solid urn-stuff. The empty space in each urn is at some point filled with tennis balls. A then contains 5 billion balls and B only a hundred. A demon or some such agency proceeds to number the balls in each urn consecutively, starting with 1, and writing its number on each ball. The balls in A are, of course, numbered 1 to 5 billion, and the balls in B 1 to 100. The positions of the balls in each urn are then thoroughly scrambled or randomized, by some process I'll leave to your imagination. Ernest, a Centurion who knows all this stuff but does not know which of the two urns before him is A and which is B, chooses one of the two by tossing a coin, reaches into the chosen urn, and withdraws a ball. (The distance between the opening at the top of B and its solid "floor" is considerably greater than the length of Ernest's arm.) He examines the ball and discovers that it is numbered 43. He concludes that the urn he drew the ball from is almost certainly B—since he drew the ball from one or the other, and the probability of his having drawn it from B is much, much higher on the evidence he has (to wit, that the ball he drew was numbered 43) than is the probability of his having drawn it from A on that evidence.[10]

Agnes, after careful consideration of her own case and this imaginary case, reasons as follows: "My case and the case I have imagined are exactly parallel. Just as Ernest rightly concluded that the urn from which he drew the ball was almost certainly the urn that contained

[10] Exercise: where A is the hypothesis that Ernest drew the ball from A, B the hypothesis that he drew the ball from B, and H the proposition that the ball he drew bore a number less than or equal to 100, calculate $A//H$ and $B//H$. Assume that the prior probability of both A and B is 0.5. (After all, Ernest chose the urn he drew a ball from by a coin-toss.)

only 100 balls, so I may conclude that the span of my existence is almost certainly 100 years—and that NPE is therefore almost certainly true and FBY almost certainly false."

Has Agnes reasoned rightly? Well, she was certainly right when she said that Ernest's reasoning was correct. And *if* the Crucial Assumption (Personal) is true, she was certainly right when she said that her case was exactly parallel to his.

But *is* the Crucial Assumption (Personal) true? I think that this thesis or principle or assumption, or whatever you want to call it, is at best extremely doubtful. I will explain why someone like you and me, a present-day analytical philosopher, should find it extremely doubtful "in his or her own case": (remember that it is a first-person principle). (What I will say about why I should find it extremely doubtful and why you should find it extremely doubtful could easily enough be adapted to Agnes's case.)

Let us call a piece of reasoning that turns on or essentially involves the Crucial Assumption (Personal) a CAPA—for "Crucial Assumption (Personal) Argument." Suppose that one is considering a CAPA at a certain moment *t*, and that this occasion is one of the first few occasions on which one has considered a CAPA. Can one regard *t* as a moment chosen at random from the whole span of one's existence (including one's post-mortem existence, if such there be)? Or, more cautiously, can one regard *t* as a moment chosen at random from the whole span of one's existence exclusive of the initial segment of one's existence during which one was not yet intellectually capable of considering a CAPA? It seems to me that one cannot—not, at least, if one is an analytical philosopher or a cosmologist who is alive at the present time (a category in which you and I and John Leslie and Don Page all fall). It seems to me quite likely that one will encounter a CAPA during the first century of one's existence (as you in fact have and I in fact have and Leslie and Page in fact have), however long the whole span of one's existence may be. It was quite likely that Page would devise a CAPA during the first 100 years of his existence (given his interests, his cognitive capacities, his education, and the ideas available to him in the intellectual environment he had inhabited during his life up to the point at which he devised his argument). And, moreover, it was quite likely, given the way ideas or "memes" are transmitted in our culture, that this CAPA would find its way to most of us during the first century of our existences.

Now, if I knew that God had chosen a moment t' at random from the set of moments at which I shall exist (a supposition that is possible only on the assumption that I shall exist only for a finite period of time) and had so arranged matters—by whatever means Providence arranges matters—as to ensure that I should first encounter a CAPA at t', then I should indeed conclude that the fact that t' falls within the first century of my existence significantly lowers the probability that I shall exist for a number of years that is a large multiple of 100. But nothing suggests that this is so—or suggests that any other factor has ensured that the moment at which I first consider a CAPA (or any other moment at which I consider a CAPA) can be regarded as a moment chosen at random from the whole span of my existence. As we have seen, moreover, there are reasons for regarding it as doubtful whether the moment at which I first consider a CAPA is a moment chosen at random from that span. And, of course, when I am considering a CAPA, *that moment* is a moment at which I am considering a CAPA. There is, therefore, no reason to suppose that when I am considering a CAPA, I have any right to regard *that moment*—the moment the CAPA refers to as "the present moment"—as a moment chosen at random from the set of moments at which I exist.

I will supplement these arguments with an "intuition pump." Consider a person I'll call Alice because her real name is unpronounceable by human beings. Alice was a member of a long-extinct species that flourished about 7 billion years ago. Nevertheless, she was a colleague of ours, a philosopher of our sort, and like most of us present-day human philosophers, she eased into that condition over some rather fuzzy stretch of time during her twenties. She worked at philosophy for pretty much the rest of her life till she died at the age of 89. But that was not the end of her: like all members of her species, she had a post-mortem existence of 5 billion years (less 89 years in her case). The obliging God of the Philosophers has chosen for us *at random* a moment t from within the 5 billion years of Alice's existence. At some point during that 5 billion years, moreover, Alice first encountered a CAPA. Now, which do you think probably came first: t or the moment at which she first encountered a CAPA?

My conclusion is that when *you* consider an argument for the conclusion that you are almost certainly not going to have a long but finite post-mortem existence, an argument that depends on the principle that you may regard the present moment as a moment chosen at random from the total span of your existence, the moment you then call "the present moment" *cannot* be regarded as a moment chosen at random from that span: that moment was in effect chosen by the fact that it is a moment at which you are considering a CAPA. You know, moreover, that this is one of the first few times that you have considered a CAPA. And you have no reason to regard that moment, the moment you call the present moment or *this* moment when you are considering a CAPA for the first time or one of the first few times, as a moment chosen at random from the set of moments at which you will exist. And you have at least some reason to suppose that that moment was likely to do what it has in fact done: to fall within the first century of your existence, and to have done that even if you are going to exist for a thousand years or for 5 billion years—or for ever.

I will end with a parable.

Christian set out one day to make his journey to the Celestial City. He passed through the Gate which is called Baptism and he began to walk along the Road to the Celestial City. That Road is as straight as any arrow, but no one who has not traveled its whole length can know that length—which for all the traveler knows when he commences his journey may be only a few short miles or it may be many thousands of miles. Christian had walked but a little way, not so much as three miles, when he came to a stone column on which some words were inscribed. "I will read these words," said Christian. And the words he read were these:

> Reason, O traveler, as follows: "I may regard my present position on the Road as a point chosen at random from its whole extent. I have traveled only a few miles. It is therefore likely that I have only a few more miles to travel."

As Christian was pondering these words, he looked up and saw on the Road an old man leaning on a staff, and the man's name was Right Reason. Right Reason said to Christian, "I see, O Christian, that thou hast read the Words of False Comfort." "Certes they be words of comfort," Christian replied, "but why dost thou call their comfort false?"

And Right Reason answered, "Dost *thou* not see that the place which the Words entreat thee to call 'my present position on the Road'—and in that they do not deceive thee, for it is indeed thy present position on the Road—is no more than the place at which thou hast encountered them? Art thou so blind that thou canst not see that the column on which those Words are written hath been put in such a place as to stand very near that Gate whereat thou beganest thy journey—and that this is the place at which thou must encounter them, howsoever long the Road may be? Didst thou know that whatever agency had set the column in this place had chosen this place at random from the whole extent of the Road, then thou shouldest indeed reason as the Words entreat thee to reason. But think on this, Christian: thou dost *not* know this—nor knowest thou of any reason which argues that it is so."[11]

References

Leslie, John. (2008). "Infinitely Long Afterlives and the Doomsday Argument," *Philosophy*, 83: 519–24.

Page, Don N. (2010). "Scientific and Philosophical Challenges to Theism," in Melville Y. Stewart (ed.), *Science and Religion in Dialogue*, Vol. I. Oxford: Wiley-Blackwell, 396–410. (An online version of this paper is available at <http://arxiv.org/abs/0801.0247>.)

van Inwagen, Peter. (2005). "Human Destiny," in William E. Mann (ed.), *The Blackwell Guide to the Philosophy of Religion*. Oxford: Blackwell, 245–65.

[11] Here are two important questions I have not considered: (1) Can the arguments I have given to show that the Crucial Assumption (Personal) is doubtful be adapted to show that the original Crucial Assumption is doubtful? (I had thought not, but conversation with Brian Cutter has convinced me that my reasons for thinking this were mistaken.) (2) If the Crucial Assumption does stand, does the conclusion of the original Doomsday Argument, which depends on the Crucial Assumption (full stop) but not on the Crucial Assumption (Personal), imply that there is no personal immortality?—or no post-mortem existence that is very, very long in comparison with threescore years and ten? I have very briefly addressed this question in van Inwagen (2005: 261–2).