

The end of the world
and life in a
computer simulation

In the reading for today, Leslie introduces a familiar sort of reasoning:

Prima facie, we should prefer theories which make our observations fairly much to be expected, rather than highly extraordinary. Waking up in the night, you form two theories. Each has a half-chance of being right, you estimate. The first, that you left the back door open, gives the chances as 10 per cent that the neighbour's cat is in your bedroom. The second, that you shut the door, puts those chances at 0.01 per cent. You switch on the light and see the cat. You should now much prefer the first theory.

The basic idea here is one which we employ all the time in our ordinary reasoning about the world. It might be summed up as follows:

The principle of confirmation

E is evidence for T1 over T2 if the probability of E given T1 > the probability of E given T2.

Here E is our evidence, and T1 and T2 are two theories between which we are trying to decide. When we talk about the probability of X given Y, we are talking about the probability that X will take place, **if** Y also happens.

The basic idea here is one which we employ all the time in our ordinary reasoning about the world. It might be summed up as follows:

The principle of confirmation

E is evidence for T1 over T2 if the probability of E given T1 > the probability of E given T2.

Here E is our evidence, and T1 and T2 are two theories between which we are trying to decide. When we talk about the probability of X given Y, we are talking about the probability that X will take place, **if** Y also happens.

Now consider the application of that line of reasoning to the examples of balls shot at random from a lottery machine. Suppose that you know that the balls in the machine are numbered sequentially (with no repeats) beginning with 1, but that you don't know how many balls there are in the machine. Now we start the machine, and a ball comes out with "3" on it. You're now asked: do you think that it is more likely that the machine has 10 balls in it, or 10,000 balls in it?

The principle of confirmation suggests - correctly, it seems - that our piece of evidence favors the hypothesis that there are just 10 balls in the machine. This is because the probability that "3" comes out given that balls 1-10 are in the machine is 10%, whereas the probability that this ball comes out given that balls numbered 1-10,000 are in the machine is only 0.01%.

(Note that, whichever hypothesis you endorse, you **could** be wrong; this is not a form of reasoning which delivers results guaranteed to be correct. The question is just which of these hypotheses is most likely, given the evidence.)

This gives us **some** information about how to reason about probabilities, but not very much. It tells us when evidence favors one theory over another, but does not tell us how much. It leaves unanswered questions like: if before I thought that the 10,000 ball hypothesis was 90% likely to be true, should I now think that the 10 ball hypothesis is more than 50% likely to be true? If I assigned each of the two hypotheses prior to the emergence of the "3" ball a probability of 0.5 (50% likely to be true), what probabilities should I assign to the theories **after** the ball comes out?

This gives us **some** information about how to reason about probabilities, but not very much. It tells us when evidence favors one theory over another, but does not tell us how much. It leaves unanswered questions like: if before I thought that the 10,000 ball hypothesis was 90% likely to be true, should I now think that the 10 ball hypothesis is more than 50% likely to be true? If I assigned each of the two hypotheses prior to the emergence of the “3” ball a probability of 0.5 (50% likely to be true), what probabilities should I assign to the theories **after** the ball comes out?

To understand Leslie’s argument, we’ll have to understand how to answer these sorts of questions.



One way to answer these questions employs a widely accepted rule of reasoning called “Bayes’ theorem,” named after Thomas Bayes, an 18th century English mathematician and Presbyterian minister.

To arrive at the theorem, we begin with the following definition of conditional probability, where, as is standard, we abbreviate “the conditional probability of x given y” as “Pr(x | y)”:

$$P(a|b) = \frac{P(a\&b)}{P(b)}$$

This says, in effect, that the probability of a given b is the chance that a and b are **both true**, divided by the chances that b is true.



One way to answer these questions employs a widely accepted rule of reasoning called “Bayes’ theorem,” named after Thomas Bayes, an 18th century English mathematician and Presbyterian minister.

To arrive at the theorem, we begin with the following definition of conditional probability, where, as is standard, we abbreviate “the conditional probability of x given y” as “Pr(x | y)”:

$$P(a|b) = \frac{P(a\&b)}{P(b)}$$

This says, in effect, that the probability of a given b is the chance that a and b are **both true**, divided by the chances that b is true.

Let’s work through an example. Suppose that this is some time before the 2008 election, and let a = ‘Obama wins’, and let b = ‘a man wins.’ Suppose that you think that each of Obama, Hilary, and McCain have a 1/3 chance of winning. Then what is the conditional probability that Obama wins, given that a man wins, using the above formula?

The conditional probability is that Obama wins, given that a man wins, is 1/2, since in this case Pr(a&b)=1/3 and Pr(b)=2/3. Intuitively, if you found out only that a man would win, you should then (given the initial probability assignments) think that there is a 0.5 probability that Obama will win.

Using this definition of conditional probability, we can then argue for Bayes’ theorem as follows, assuming that Pr(b)≠0.



Definition of conditional probability

$$P(a|b) = \frac{P(a\&b)}{P(b)}$$

Using this definition of conditional probability, we can then argue for Bayes' theorem as follows, assuming that $\Pr(b) \neq 0$.

Derivation of Bayes' theorem

1. $P(a b) = \frac{P(a\&b)}{P(b)}$	def. of conditional probability
2. $P(b a) = \frac{P(a\&b)}{P(a)}$	def. of conditional probability
3. $P(a b) * P(b) = P(a\&b)$	(1), multiplication by = 's
4. $P(a\&b) = P(b a) * P(a)$	(2), multiplication by = 's
5. $P(a b) * P(b) = P(b a) * P(a)$	(3),(4)
C. $P(a b) = \frac{P(b a)*P(a)}{P(b)}$	(5), division by = 's

The conclusion of this argument is Bayes' theorem. Intuitively, what it says is that if we want to know the probability of some theory given a bit of evidence, what we need to know are three things: (1) the probability of the evidence given the theory (i.e., how likely the evidence is to happen if the theory is true), (2) the prior probability of the theory, and (3) the prior probability of the evidence.

Derivation of Bayes' theorem

1. $P(a b) = \frac{P(a\&b)}{P(b)}$	def. of conditional probability
2. $P(b a) = \frac{P(a\&b)}{P(a)}$	def. of conditional probability
3. $P(a b) * P(b) = P(a\&b)$	(1), multiplication by =’s
4. $P(a\&b) = P(b a) * P(a)$	(2), multiplication by =’s
5. $P(a b) * P(b) = P(b a) * P(a)$	(3),(4)
C. $P(a b) = \frac{P(b a)*P(a)}{P(b)}$	(5), division by =’s

The conclusion of this argument is Bayes' theorem. Intuitively, what it says is that if we want to know the probability of some theory given a bit of evidence, what we need to know are three things: (1) the probability of the evidence given the theory (i.e., how likely the evidence is to happen if the theory is true), (2) the prior probability of the theory, and (3) the prior probability of the evidence.

If we let “h” stand for the hypothesis to be tested and “e” for the relevant evidence, then the theorem can be rewritten as follows:

Bayes' theorem

$$P(h|e) = \frac{P(h)*P(e|h)}{P(e)}$$

This theorem is very useful, since often it is easy to figure out the conditional probability of the evidence given the theory, but very hard to figure out the conditional probability of the theory given the evidence.

A good example of this sort of situation is given by our example of the lottery balls.

Bayes' theorem

$$P(h|e) = \frac{P(h)*P(e|h)}{P(e)}$$

This theorem is very useful, since often it is easy to figure out the conditional probability of the evidence given the theory, but very hard to figure out the conditional probability of the theory given the evidence.

A good example of this sort of situation is given by our example of the lottery balls.

Remember that we have two hypotheses: that there are balls labeled 1-10 in the machine, and that there are balls labeled 1-10,000 in the machine. Let's assume that, at the start, we think that these hypotheses have equal probability: they are each assigned probability 0.5. (Maybe which sort of machine we are given is determined by coin flip.)

Now suppose that the first ball that comes out, again at random, is a "3". Bayes' theorem can tell us, given the foregoing information, how likely it is that the machine before us contains 10 rather than 10,000 balls. According to Bayesians, we figure this out by **conditionalizing on our evidence**, i.e. finding the conditional probability of the theory given the evidence.

Let "h" be the theory that the machine contains 10 balls. Then to compute the relevant conditional probability, we need to know three values.

First, we need to know $\Pr(h)$. From the above, we know that this is 0.5.

Next, we need to know $\Pr(e | h)$. Fairly obviously, this is 0.1.

But we also need to know the unconditional probability of e, $\Pr(e)$. How should we figure this out?

A natural thought is that since we know that each of the two hypotheses - 10 ball and 10,000 ball - have a probability of 0.5, $\Pr(e)$ should be the mean of the conditional probabilities of e given those two hypotheses. In this case:

$$P(e) = 0.5(0.1 + 0.0001) = 0.05005$$

Bayes' theorem

$$P(h|e) = \frac{P(h)*P(e|h)}{P(e)}$$

Now suppose that the first ball that comes out, again at random, is a “3”. Bayes' theorem can tell us, given the foregoing information, how likely it is that the machine before us contains 10 rather than 10,000 balls. According to Bayesians, we figure this out by **conditionalizing on our evidence**, i.e. finding the conditional probability of the theory given the evidence.

Let “h” be the theory that the machine contains 10 balls. Then to compute the relevant conditional probability, we need to know three values.

First, we need to know $\Pr(h)$. From the above, we know that this is 0.5.

Next, we need to know $\Pr(e | h)$. Fairly obviously, this is 0.1.

But we also need to know the unconditional probability of e, $\Pr(e)$. How should we figure this out?

A natural thought is that since we know that each of the two hypotheses - 10 ball and 10,000 ball - have a probability of 0.5, $\Pr(e)$ should be the mean of the conditional probabilities of e given those two hypotheses. In this case:

$$P(e) = 0.5(0.1 + 0.0001) = 0.05005$$

Plugging these numbers into Bayes' theorem, we get:

$$P(h|e) = \frac{0.5*0.1}{0.05005} = 0.999.$$

Hence, after you see the “3” ball come out of the lottery machine, you should revise the probability you assign to the 10-ball hypothesis from 0.5 to .999 - that is, you should switch from thinking that the machine has a 50% chance of being a 10-ball machine to thinking that it has a 99.9% chance of being a 10-ball machine.

Bayes' theorem

$$P(h|e) = \frac{P(h)*P(e|h)}{P(e)}$$

Hence, after you see the “3” ball come out of the lottery machine, you should revise the probability you assign to the 10-ball hypothesis from 0.5 to .999 - that is, you should switch from thinking that the machine has a 50% chance of being a 10-ball machine to thinking that it has a 99.9% chance of being a 10-ball machine.

An intuitive way to think about what this all means is in terms of what bets you would be willing to accept. If you think that something has a 50% chance of happening, then you should be willing to accept all bets which give you better than even odds, and willing to reject all bets which give you worse odds. Analogous remarks apply to different probability assignments.

This link between probability assignments and bets is one way to bring out a strength of the Bayesian approach to belief formation. Following the Bayesian rule of conditionalization is the only way to avoid being subject to a **Dutch book**.

A Dutch book is a combination of bets which, no matter what the outcome, is sure to lose. For example, suppose that we have a 10 horse race, and you have the following views about the probabilities that certain horses will win.

- #3 has at least a 40% chance of winning
- #6 has a slightly better than even chance of winning
- #8 has a better than 1 in 3 chance of winning

If these are your probability assignments, then you should be willing to make the following three bets:

- \$2 on #3 to win, at 3-2 odds
- \$3 on #6 to win, at even odds
- \$1.50 on #8 to win, at 3-1 odds

Bayes' theorem

$$P(h|e) = \frac{P(h)*P(e|h)}{P(e)}$$

This link between probability assignments and bets is one way to bring out a strength of the Bayesian approach to belief formation. Following the Bayesian rule of conditionalization is the only way to avoid being subject to a **Dutch book**.

A Dutch book is a combination of bets which, no matter what the outcome, is sure to lose. For example, suppose that we have a 10 horse race, and you have the following views about the probabilities that certain horses will win.

- #3 has at least a 40% chance of winning
- #6 has a slightly better than even chance of winning
- #8 has a better than 1 in 3 chance of winning

If these are your probability assignments, then you should be willing to make the following three bets:

- \$2 on #3 to win, at 3-2 odds
- \$3 on #6 to win, at even odds
- \$1.50 on #8 to win, at 3-1 odds

But this is a Dutch book. Can you see why?

One of the main arguments given in favor of updating your beliefs using Bayesian conditionalization, using Bayes' theorem, is that other ways of updating your beliefs will leave you, in the above sense, subject to a Dutch book. And being subject to a Dutch book seems like a paradigm of irrationality. It seems sort of like the probabilistic version of having inconsistent beliefs. Hence it is a strength of Bayesianism that it avoids this situation.

Bayes' theorem

$$P(h|e) = \frac{P(h)*P(e|h)}{P(e)}$$

With this Bayesian apparatus in hand, let's return to Leslie's argument.

Leslie asks us to consider two hypotheses about the future course of human civilization:

Doom Soon. The human race will go extinct by 2150, with the total humans born by the time of such extinction being 500 billion.

Doom Delayed. The human race will go on for several thousand centuries, with the total humans born before the race goes extinct being 50 thousand billion.

Leslie thinks that we can use a certain kind of evidence we have, along with Bayesian conditionalization, to show how likely these two hypotheses are. To do this, though, we'll first have to think about how likely we think these two hypotheses are to begin with.

Doom Soon is pretty grim; it means that the human race will go extinct during the lives of your grandchildren. Let's suppose that we think that this is pretty unlikely - maybe that it has a 1% chance of happening. Let's suppose (we'll relax this assumption later) that Doom Delayed is the only other possibility, so that it has a 99% chance of happening.

What evidence could we possibly have now to help us decide between these hypotheses now? Some obvious candidates spring to mind: the proliferation of nuclear weapons; astronomical calculations of the probabilities of large asteroids colliding with the earth; prophecies involving the end of the Mayan calendar; etc. But the evidence that Leslie has in mind is of a different sort.

Doom Soon. The human race will go extinct by 2150, with the total humans born by the time of such extinction being 500 billion.

Doom Delayed. The human race will go on for several thousand centuries, with the total humans born before the race goes extinct being 50 thousand billion.

Bayes' theorem

$$P(h|e) = \frac{P(h)*P(e|h)}{P(e)}$$

Doom Soon is pretty grim; it means that the human race will go extinct during the lives of your grandchildren. Let's suppose that we think that this is pretty unlikely - maybe that it has a 1% chance of happening. Let's suppose (we'll relax this assumption later) that Doom Delayed is the only other possibility, so that it has a 99% chance of happening.

What evidence could we possibly have now to help us decide between these hypotheses now? Some obvious candidates spring to mind: the proliferation of nuclear weapons; astronomical calculations of the probabilities of large asteroids colliding with the earth; prophecies involving the end of the Mayan calendar; etc. But the evidence that Leslie has in mind is of a different sort.

That evidence is: each of us is one of the first 50 billion human beings born. Let's now ask, using the Bayesian method, what probability, in light of this evidence, we should assign to Doom Soon (DS) and Doom Delayed (DD).

First, what is the probability of our evidence, conditional on the truth of DS? It seems that it should be 0.1; if there will be 500 billion total human beings, then I have a 1 in 10 chance of being among the first 50 billion born.

Second, what is $\Pr(e | DD)$? By analogous reasoning, it seems to be 1 in 1000, or .001.

But what is the probability of e , by itself? On the current assumption that either DS or DD is true, and that DD is 99 times more likely to be true, it seems that $\Pr(e)$ should be equal to:

$$[99 * \Pr(e | DD) + 1 * \Pr(e | DS)] / 100 = .00199$$

Doom Soon. The human race will go extinct by 2150, with the total humans born by the time of such extinction being 500 billion.

Doom Delayed. The human race will go on for several thousand centuries, with the total humans born before the race goes extinct being 50 thousand billion.

Bayes' theorem

$$P(h|e) = \frac{P(h)*P(e|h)}{P(e)}$$

That evidence is: each of us is one of the first 50 billion human beings born. Let's now ask, using the Bayesian method, what probability, in light of this evidence, we should assign to Doom Soon (DS) and Doom Delayed (DD).

So we have the following:

$$\begin{aligned} Pr(DS) &= 0.01 & Pr(e) &= 0.00199 \\ Pr(DD) &= 0.99 & Pr(e|DS) &= 0.1 \\ & & Pr(e|DD) &= 0.001 \end{aligned}$$

And this is all we need to figure out the probabilities of DS and DD conditional on the evidence that we are among the first 50 billion human beings born.

$$P(DS|e) = \frac{P(DS) * P(e|DS)}{P(e)} = \frac{.01 * .1}{.00199} = .503$$

$$P(DD|e) = \frac{P(DD) * P(e|DD)}{P(e)} = \frac{.99 * .001}{.00199} = .497$$

So, even if we begin by thinking that the probability of Doom Soon is only 1%, reflection on the simple fact that we are born among the first 50 billion humans shows that we should think that there is a greater than 50% chance that the human race will be extinct in the next 150 years.

Doom Soon. The human race will go extinct by 2150, with the total humans born by the time of such extinction being 500 billion.

Doom Delayed. The human race will go on for several thousand centuries, with the total humans born before the race goes extinct being 50 thousand billion.

Bayes' theorem

$$P(h|e) = \frac{P(h) * P(e|h)}{P(e)}$$

$$Pr(DS) = 0.01$$

$$Pr(e) = 0.00199$$

$$Pr(DD) = 0.99$$

$$Pr(e|DS) = 0.1$$

$$Pr(e|DD) = 0.001$$

$$P(DS|e) = \frac{P(DS) * P(e|DS)}{P(e)} = \frac{.01 * .1}{.00199} = .503$$

$$P(DD|e) = \frac{P(DD) * P(e|DD)}{P(e)} = \frac{.99 * .001}{.00199} = .497$$

So, even if we begin by thinking that the probability of Doom Soon is only 1%, reflection on the simple fact that we are born among the first 50 billion humans shows that we should think that there is a greater than 50% chance that the human race will be extinct in the next 150 years.

Varying our initial assumptions changes the outcome dramatically. Suppose that reflection upon the risk of climate change, nuclear war, etc. makes us think that we should assign DS and DD **roughly equal** initial probabilities - suppose we think, before considering the fact that we were born in the first 50 billion people, that the probability of each is approximately 0.5. On this assumption, the probability of Doom Soon **after** conditionalizing on our evidence is **just over 99%**.

Alternatively, we might think (as Leslie says, p. 202) that if the human race survives past 2050, then it will likely colonize other planets, making it quite likely that the population of humans will ultimately grow to be at least 50 million billion (50 quadrillion). On this assumption, even if we begin by assigning Doom Soon a chance of only 1%, conditionalizing on the evidence that we were among the first 50 billion born gives Doom Soon a probability of 99.9%.

Doom Soon. The human race will go extinct by 2150, with the total humans born by the time of such extinction being 500 billion.

Doom Delayed. The human race will go on for several thousand centuries, with the total humans born before the race goes extinct being 50 thousand billion.

Bayes' theorem

$$P(h|e) = \frac{P(h)*P(e|h)}{P(e)}$$

So, even if we begin by thinking that the probability of Doom Soon is only 1%, reflection on the simple fact that we are born among the first 50 billion humans shows that we should think that there is a greater than 50% chance that the human race will be extinct in the next 150 years.

Varying our initial assumptions changes the outcome dramatically. Suppose that reflection upon the risk of climate change, nuclear war, etc. makes us think that we should assign DS and DD **roughly equal** initial probabilities - suppose we think, before considering the fact that we were born in the first 50 billion people, that the probability of each is approximately 0.5. On this assumption, the probability of Doom Soon **after** conditionalizing on our evidence is **just over 99%**.

Alternatively, we might think (as Leslie says, p. 202) that if the human race survives past 2050, then it will likely colonize other planets, making it quite likely that the population of humans will ultimately grow to be at least 50 million billion (50 quadrillion). On this assumption, even if we begin by assigning Doom Soon a chance of only 1%, conditionalizing on the evidence that we were among the first 50 billion born gives Doom Soon a probability of 99.9%.

The numbers are surprising. But more surprising than the numbers is the way we arrived at them. One thinks of revising one's view about the likelihood of Doom Soon based on empirical claims about nuclear weapons, climate change, asteroids, etc. It seems crazy that such dramatic changes in our view about the extinction of the human race should result from mere reflection on how many human beings have been born.

This is why this argument deserves to be considered a paradox. We have a plausible argument from Leslie that we should radically change our view of the future based on the number of human beings who have lived, but it seems clear that it is unreasonable to change our views on this topic for this reason.

This is why this argument deserves to be considered a paradox. We have a plausible argument from Leslie that we should radically change our view of the future based on the number of human beings who have lived, but it seems clear that it is unreasonable to change our views on this topic for this reason.

We'll move on shortly to consider some objections to Leslie's argument. But first let's consider another, even more surprising, application of the same sort of reasoning, due to Nick Bostrom (in "Are you living in a computer simulation?").

Suppose we met some Martians, or other beings from another planet, who were made of entirely different materials from us - suppose, for example, that they were composed almost entirely of silicon. Could such beings be conscious?

It seems clear that they could; beings could be conscious even if made of radically different materials than us. But what makes such beings conscious? Plausibly, what makes them conscious is not what they are made of, but how what they are made of is organized; even something made of silicon could be conscious if its silicon parts were arranged in a structure as complex as a human brain.

If this is true, then it seems clear that computers could be conscious. Indeed, given reasonable assumptions about future increases in the computing power available to us, it seems quite plausible that in the relatively near future computers will be able to run simulations of human life which contain many, many conscious beings.

With continuing increases in computing power, there will presumably come a time in which a vast, vast majority of all conscious beings to have existed will have existed only in computer simulations.

Are **you** living in a computer simulation right now? It seems that reasoning analogous to that used in the doomsday argument suggests that it is reasonable for you to believe that you are. After all, if the hypothesis that the vast number of conscious beings to have existed exist only in a computer simulation is true, then the probability that you do not exist in a computer simulation is extremely low.

Of course, there is another option: human beings could become extinct before the relevant advances in computing power are realized. So, it seems, it is overwhelmingly likely that either you are living in a computer simulation, or human beings are soon to become extinct.

How might one respond to Leslie's argument? (We'll leave it an open question whether the responses we consider will also apply to the "computer simulation" argument.)

One could of course respond that the Bayesian apparatus on which it depends is faulty. But that would seem hasty; let's consider some other possibilities.

One possibility is that the problem begins with the obviously false assumption that Doom Soon and Doom Delayed are the only relevant possibilities. Would Leslie's argument still work if we relaxed this assumption, and took into account the fact that there are many possible futures for the human species? What would the analogous situation be with the example of the numbered balls and the lottery machine?

Let's consider a different line of objection:

Look, we know that something must be wrong with this way of arguing. After all, couldn't someone in ancient Rome have used this reasoning to show that the end of the world would come before 500 AD? And wouldn't they have been wrong? So mustn't there also be something wrong with **our** using this reasoning?

Is this a good objection?

A more promising line of objection focuses on the apparent assumption that one's location in the birth order of human beings is random. Leslie asks you to, in effect, assimilate this case to the lottery ball example. But why think that the fact that I am born in the first 50 billion people is relevantly analogous to the number "50 billion" coming out of the lottery machine?

To develop this objection, let's think more closely about exactly which assumptions are involved in the lottery machine example. (This follows the discussion in Mark Greenberg's "Apocalypse not just yet.")

Doom Soon. The human race will go extinct by 2150, with the total humans born by the time of such extinction being 500 billion.

Doom Delayed. The human race will go on for several thousand centuries, with the total humans born before the race goes extinct being 50 thousand billion.

A more promising line of objection focuses on the apparent assumption that one's location in the birth order of human beings is random. Leslie asks you to, in effect, assimilate this case to the lottery ball example. But why think that the fact that I am born in the first 50 billion people is relevantly analogous to the number "50 billion" coming out of the lottery machine?

To develop this objection, let's think more closely about exactly which assumptions are involved in the lottery machine example. (This follows the discussion in Mark Greenberg's "Apocalypse not just yet.")

An analogous case would be a lottery machine which we know to contain either 500 billion or 50 thousand billion balls. To make the numbers smaller, let's think of a pair of lottery machines with, respectively, 500 and 50,000 balls. Suppose you are confronted with a lottery machine which you know to be of one of the types just described.

Suppose now that the lottery machine spits out a ball which reads "50." (This is supposed to be analogous to finding that you are among the first 50 billion people born.) Isn't this evidence, as Leslie says, that the machine has 500 rather than 50,000 balls in it?

It is - **but only if we make two assumptions about the machines.**

First, we must assume that the ball which comes out is randomly selected by the lottery machine. If, for example, balls with numbers higher than "100" are slightly larger and never come out of the machine, then obviously the fact that a ball labeled "50" came out of the machine would be **no evidence at all** about how many balls are in the machine.

Second, and just as important, we must assume that the ball with the label "50" on it **would still be in the urn if it contained 500 balls.**

Doom Soon. The human race will go extinct by 2150, with the total humans born by the time of such extinction being 500 billion.

Doom Delayed. The human race will go on for several thousand centuries, with the total humans born before the race goes extinct being 50 thousand billion.

Suppose now that the lottery machine spits out a ball which reads “50.” (This is supposed to be analogous to finding that you are among the first 50 billion people born.) Isn’t this evidence, as Leslie says, that the machine has 500 rather than 50,000 balls in it?

It is - **but only if we make two assumptions about the machines.**

First, we must assume that the ball which comes out is randomly selected by the lottery machine. If, for example, balls with numbers higher than “100” are slightly larger and never come out of the machine, then obviously the fact that a ball labeled “50” came out of the machine would be **no evidence at all** about how many balls are in the machine.

Second, and just as important, we must assume that the ball with the label “50” on it **would still be in the urn if it contained 500 balls.**

This second condition is easy to miss, since we of course assume that a machine with N balls in it contains those balls labeled 1-N. But of course things don’t have to work this way. Suppose that the 500 balls in the 500-ball machine were randomly selected from the balls in the 50,000 ball machine. Then the “50” ball might not be in the 500-ball machine. In this case, would the emergence of the “50” ball from the machine count in favor of the hypothesis that the machine before us contains 500 rather than 50,000 balls?

The problem is that any way of understanding the analogy between the Doomsday argument and the lottery machine example seems to violate one of the two assumptions needed to legitimate the reasoning used in the lottery machine case.

To see this, think about the question: what is the analogue in the Doomsday argument of the numbers written on the balls in the lottery machine case?

Doom Soon. The human race will go extinct by 2150, with the total humans born by the time of such extinction being 500 billion.

Doom Delayed. The human race will go on for several thousand centuries, with the total humans born before the race goes extinct being 50 thousand billion.

Suppose now that the lottery machine spits out a ball which reads “50.” (This is supposed to be analogous to finding that you are among the first 50 billion people born.) Isn’t this evidence, as Leslie says, that the machine has 500 rather than 50,000 balls in it?

It is - **but only if we make two assumptions about the machines.**

First, we must assume that the ball which comes out is randomly selected by the lottery machine. If, for example, balls with numbers higher than “100” are slightly larger and never come out of the machine, then obviously the fact that a ball labeled “50” came out of the machine would be **no evidence at all** about how many balls are in the machine.

Second, and just as important, we must assume that the ball with the label “50” on it **would still be in the urn if it contained 500 balls.**

To see this, think about the question: what is the analogue in the Doomsday argument of the numbers written on the balls in the lottery machine case?

Here’s one idea: the number “written on you” is the number you happen to be in the birth order of the human species. So (let’s suppose) my number is “50 billion” because I just happened to be the 50 billionth person born.

On this idea, people don’t have “built in” numbers: rather, they are assigned the numbers they get based on when they are born.

But now imagine that the lottery machine worked this way. On this view, there are an undisclosed number of balls in the machine, none of which have numbers written on them. We write the numbers on the balls as they come out of the machine, beginning with “1.” Now suppose we get to “50.” Is the fact that a ball with such a low number came out evidence that we have a 500-ball machine before us? Clearly not, because the first assumption - the assumption of random selection - is violated.

Doom Soon. The human race will go extinct by 2150, with the total humans born by the time of such extinction being 500 billion.

Doom Delayed. The human race will go on for several thousand centuries, with the total humans born before the race goes extinct being 50 thousand billion.

Suppose now that the lottery machine spits out a ball which reads “50.” (This is supposed to be analogous to finding that you are among the first 50 billion people born.) Isn’t this evidence, as Leslie says, that the machine has 500 rather than 50,000 balls in it?

It is - **but only if we make two assumptions about the machines.**

First, we must assume that the ball which comes out is randomly selected by the lottery machine. If, for example, balls with numbers higher than “100” are slightly larger and never come out of the machine, then obviously the fact that a ball labeled “50” came out of the machine would be **no evidence at all** about how many balls are in the machine.

Second, and just as important, we must assume that the ball with the label “50” on it **would still be in the urn if it contained 500 balls.**

So let’s suppose instead that every human being comes with a “built in” number, just like the lottery balls have numbers written on them prior to their emergence from the lottery machine.

There are two problems with this suggestion. First, how do we know what anyone’s number is, on this way of viewing the analogy? What does it even mean to say that people have a certain number?

Second, if we can come up with a way of assigning numbers to all the people that could have existed - we violate the second assumption. **For there is no guarantee that, if N people exist, they will have the numbers 1-N.** This is just like the case in which 500 balls are randomly selected from the 50,000 ball machine - in this case, even if we somehow knew that my number was 50 billion - this would provide us no evidence at all about how many human beings will eventually exist.