The Snake Eyes Paradox

Daniel M. Reeves
manifold.markets/dreev

Wamba Ivanhoe
manifold.markets/ShitakiIntaki

DRAFT 2023-10-30 23:30

1 Problem Statement

Our protagonist, Euclid, is offered a gamble. A pair of six-sided dice are rolled and unless they come up snake eyes he gets a bajillion dollars. If they do come up snake eyes, he's devoured by snakes.

So far it sounds like Euclid has a 1/36 chance of dying, right?

Now the twist. First, we gather up an unlimited number of people willing to play the game, including Euclid. We take 1 person from that pool and let them play. Then we take 2 people and have them play together, where they share a dice roll and either get the bajillion dollars each or both get devoured. Then we do the same with 4 people, and then 8, 16, and so on, doubling each time.

We keep going until one of those groups is devoured by snakes, then the game stops. Is the probability that Euclid will die, given that he's chosen to play, still 1/36?

Argument for NO aka the frequency argument: Due to the doubling, the final group of people that die is slightly bigger than all the surviving groups put together. So if Euclid's chosen to play he has about a 50% chance of dying! 😬 🐍

Argument for YES aka the one-fair-roll argument: The dice rolls are independent and whenever Euclid's chosen, what happened in earlier rounds is irrelevant. His chances of death are the chances of snake eyes on his round: 1/36. 😅

So which is it? What's Euclid's probability of dying, conditional on being chosen to play? If Euclid's mom learns that Euclid was chosen to play in this game and the game is now over, how worried is she? If Euclid wanted to play the one-shot version, does he still want to play the doubling groups version?

Fine print: The game is not adversarial; the dice rolls are independent and fair. Groups are chosen uniformly and without replacement. Void where prohibited. See Bartha and Hitchcock (1999) for details.

2 Solution (With Limits)

We want the probability that Euclid dies given that he is chosen to play, Pr(death | chosen). It seems like we can ignore the 0% chance of rolling not-snake-eyes forever and say that eventually about half the people who are chosen die, but let's Bayes it out carefully:

$$\begin{aligned}\Pr(\text{death} \mid \text{chosen}) &= \frac{\Pr(\text{chosen} \mid \text{death})\Pr(\text{death})}{\Pr(\text{chosen})} \\ &= \frac{1 \cdot \Pr(\text{death})}{\Pr(\text{chosen})}.\end{aligned}$$

But if Euclid's part of an infinite pool, he has a 0% chance of being chosen and a 0% chance of dying. The probability we want is 0/0. **robot-with-smoke-coming-out-of-its-ears-emoji**

Since we can't directly calculate the probability in the infinite case, a natural thing to do is to take a limit.

To get a feel for where we're going, suppose Euclid's one person in a huge but finite pool. Now suppose he is actually chosen. There are two ways that can happen:

1. The pool runs out and everyone survives.
2. The pool doesn't run out and Euclid has about a 50% chance of dying.

But knowing that Euclid is chosen is Bayesian evidence that we had many, many rounds of survival. If an early group died then most of the pool wasn't chosen, so probably Euclid wasn't chosen.

Thinking like a Bayesian means shifting your probability in light of evidence by seeing how surprised you'd be in various universes by that evidence. If an early group died then most people aren't chosen and in that universe Euclid is surprised to be chosen. If *no* group died then everyone was chosen and in that universe Euclid is fully unsurprised that he was chosen. That's the sense in which being chosen is Bayesian evidence that more people survived. In particular it's at least weak evidence that everyone survived.

So even with an absurdly huge pool of people, where there's *essentially* a 0% chance of everyone surviving, if Euclid knows he was chosen (which itself has near zero probability, but, you know, *if*) then that means Euclid is more likely to be in that essentially-0%-probability universe where everyone survives.

Enough hand-waving and appeals to intuition. Let's Bayes it out to see what Pr(death | chosen) is exactly, in a truncated game where we stop after N rounds. Once we have that, we can take the limit as N goes to infinity.

First, let M be the size of the pool:

$$M = \sum_{i=1}^{N} 2^{i-1} = 2^N - 1.$$

And let p be the probability of snake eyes, 1/36. We can now compute the probability of Euclid being chosen by summing up (1) the probability he's chosen for the first round, $1/M$, plus (2) the probability that the first group survives, $1-p$, and that he's chosen for the 2nd round, $2/M$, plus (3) the probability that the first two groups survive and he's chosen for the 3rd round, etc. Writing that out as an equation gives this:

$$\begin{aligned}\Pr(\text{chosen}) = {} & \tfrac{1}{M} + (1-p)\tfrac{2}{M} \\ & + (1-p)^2\tfrac{4}{M} \\ & + (1-p)^3\tfrac{8}{M} \\ & + \dots \\ & + (1-p)^{N-1}\frac{2^{N-1}}{M} \\ = {} & \sum_{i=1}^{N} \tfrac{1}{M} 2^{i-1}(1-p)^{i-1}.\end{aligned}$$

For Pr(death) the calculation is very similar but every term is multiplied by p. To die, Euclid has to be chosen and then roll snake eyes. This can happen on any round, all of which are mutually exclusive. We can then factor that p out and we have

$$\Pr(\text{death}) = p \cdot \Pr(\text{chosen}).$$

Working out that expression for Pr(chosen) wasn't even necessary! We compute Pr(death | chosen) like so:

$$\begin{aligned}\Pr(\text{death} \mid \text{chosen}) &= \frac{\Pr(\text{death})}{\Pr(\text{chosen})} \\ &= \frac{p \cdot \Pr(\text{chosen})}{\Pr(\text{chosen})} = p.\end{aligned}$$

It doesn't depend on N at all! The limit as N goes to infinity is just... p, or 1/36, the probability of rolling snake eyes. □

3 Can We Roll Not-Snake-Eyes Forever?

What about the argument that, with unlimited people, there will necessarily be a finite round n at which snake eyes is rolled? And for every possible such n, at least half of the chosen players die. After all, the probability of rolling not-snake-eyes forever is zero. (More precisely, in the limit as n goes to infinity, the probability of rolling not-snake-eyes n times in a row goes to zero.)

That's all true but let's work out the probability of rolling not-snake-eyes forever *conditional on Euclid being chosen*. Starting with Pr(snake eyes) as the probability that a game rolls snake eyes—unambiguously 1—we have, by the definition of conditional probability:

$$\Pr(\text{snake eyes} \mid \text{chosen}) = \frac{\Pr(\text{chosen} \wedge \text{snake eyes})}{\Pr(\text{chosen})}.$$

In the infinite setting that's $\frac{0}{0}$ because Euclid has a 0% chance of being chosen from an infinite pool. So let's work it out in the limit with a cap of N rounds and finite pool M as before:

$$\frac{\sum_{i=1}^{N} (1-p)^{i-1} p \cdot \frac{2^i - 1}{M}}{\sum_{i=1}^{N} \frac{1}{M} 2^{i-1} (1-p)^{i-1}}.$$

In the numerator we're summing over every possible round i at which we could roll snake eyes, saying that we need to roll not-snake-eyes $i-1$ times followed by one snake eyes *and* that Euclid is chosen in any round from 1 through i. The denominator, Pr(chosen), is the same as in the previous section.

Now algebra ensues. We multiply the numerator and denominator by M to get rid of the $1/M$ factor, then distribute the $(1-p)^{i-1}p$ over the $2^i - 1$ and split it into two summations:

$$\frac{\left(\sum_{i=1}^{N} 2^i (1-p)^{i-1} p\right) - \left(\sum_{i=1}^{N} (1-p)^{i-1} p\right)}{\sum_{i=1}^{N} 2^{i-1} (1-p)^{i-1}}.$$

These are finite sums so that's kosher. The right side of the numerator is the probability of rolling snake eyes by round N, which is Pr(snake eyes) in the limit as N goes to infinity, so we replace that sum by one:

$$\frac{\left(\sum_{i=1}^{N} 2^i (1-p)^{i-1} p\right) - 1}{\sum_{i=1}^{N} 2^{i-1} (1-p)^{i-1}}.$$

(We can also get the answer of 1 for that summation by pulling out the p to leave a geometric series with common ratio $1-p$, which sums in the limit to $1/p$.)

Almost there! Pull a $2p$ out of the remaining sum in the numerator to get this:

$$\frac{2p\left(\sum_{i=1}^{N} 2^{i-1} (1-p)^{i-1}\right) - 1}{\sum_{i=1}^{N} 2^{i-1} (1-p)^{i-1}}.$$

Notice that the sums in the numerator and denominator are now identical. We distribute the denominator,

$$2p - \frac{1}{\sum_{i=1}^{N} 2^{i-1} (1-p)^{i-1}},$$

and combine the terms in the sum,

$$2p - \frac{1}{\sum_{i=1}^{N} \left(2(1-p)\right)^{i-1}},$$

to see that the denominator is a geometric series with common ratio $2(1-p)$. As long as the common ratio is greater than or equal to 1, the denominator diverges and the above approaches $2p$ in the limit as N goes to infinity. How do we know $2(1-p) \geq 1$? Because we can rearrange it as $p \leq 1/2$ and that's true for us, namely $p = 1/36$.[1]

In conclusion, the probability of eventually rolling snake eyes, conditional on Euclid being chosen to play, approaches $2p = 1/18$ in the limit. Which is to say that the conditional probability of rolling not-snake-eyes literally forever is 17/18. 🤯

(Or to say it less sensationally: For any finite N, the conditional probability of taking more than N rolls to hit snake eyes is greater than 17/18.)

This vindicates our initial intuitive argument that being chosen is Bayesian evidence—strong Bayesian evidence, it turns out!—of never rolling snake eyes. And it invalidates the intuition that we can safely condition on snake eyes being rolled just because it definitely will be rolled (unconditionally). Another version of that intuition is that any event with probability 1, such as rolling snake eyes eventually, must be independent of any other event. But if being chosen and rolling snake eyes were independent then, by definition of independence, $\Pr(\text{chosen} \wedge \text{snake eyes}) = \Pr(\text{chosen}) \cdot \Pr(\text{snake eyes})$. And if that were true, we'd conclude from the above derivation of $\Pr(\text{snake eyes} \mid \text{chosen})$ that

$$\begin{aligned}
&\Pr(\text{snake eyes}) \\
=&\frac{\Pr(\text{chosen})\Pr(\text{snake eyes})}{\Pr(\text{chosen})} \\
=&\frac{\Pr(\text{chosen} \wedge \text{snake eyes})}{\Pr(\text{chosen})} \\
=&\Pr(\text{snake eyes} \mid \text{chosen}) \\
=&1/18.
\end{aligned}$$

Which contradicts $\Pr(\text{snake eyes}) = 1$. The temptation to treat $\Pr(X)$ as $\Pr(X \mid \text{snake eyes})$ since $\Pr(\text{snake eyes}) = 1$ leads us astray!

4 To Infinity And Beyond (With A Nonuniform Prior)

What if we reject the whole idea of defining a truncated version of Snake Eyes to take a limit of? Can we math out an answer for the infinite game directly? Yes! The only monkey wrench is that we can't have a uniform prior over an infinite set.[2] So let's just say we don't *quite* have a uniform prior. Maybe Euclid thinks he's equally likely to be any of the first trillion people chosen to play and that it gradually becomes less likely after that. We can make that "trillion" as high as we like.

As long as the probability of being chosen isn't exactly zero, there's no 0/0 problem like before.

Is that fair though, to reject the stipulation in the problem statement that Euclid is chosen uniformly? Well, it's arguably less of a leap than we made before in defining a truncated version of the game where it's possible for no one to die. We're just saying Euclid is not quite chosen uniformly because he *can't* be and have any probability of being chosen at all. But we can get arbitrarily close to uniform! We can even consider the limit as the distribution approaches uniform. Great, let's get to it!

First, think of every player in the pool being lined up in an infinite queue. This ordering is established once, before any dice rolls, and is independent of the dice rolls. Now we can let QR_c be the event that Euclid is positioned in the queue such that, if the game gets far enough, he'll be in round c. (QR for "queued for round".) Let SE_s be the event that snake

[1] What would happen if we had $p > 1/2$? In that case, by the preceding derivation, $\Pr(\text{snake eyes} \mid \text{chosen}) = 1$ so no chance of everyone surviving. That makes sense because the nature of the paradox changes if $p > 1/2$: The probability of dying in the one-shot version is already greater than the fraction of people who die when the game ends in snake eyes. The frequency argument and the one-fair-roll argument aren't necessarily in conflict when $p > 1/2$.

[2] Not in standard analysis anyway. If infinitely many things are all equally likely then they all have zero probability. Or to be slightly more formal, there's an elegant proof by contradiction: First, the sum of the probabilities of each element of the set must be 1. That's part of what it means to have a prior over a set of possibilities. Now suppose every element in your infinite set has equal probability ϵ. That's what we mean by a uniform prior. Further suppose that $\epsilon = 0$. Then the sum of the probabilities is 0. So that's no good; we must have $\epsilon > 0$. But the sum of an infinite number of positive ϵ's is infinity. So that's no good either. $\rightarrow\leftarrow$

eyes is rolled in round s. Again, with QR defined purely in terms of Euclid's position in the pool, it's independent of SE.

Now define

$$p_{cs} = \Pr(\mathrm{QR}_c \wedge \mathrm{SE}_s)$$

as the probability of a game where Euclid is positioned to be chosen in round c and snake eyes is rolled in round s. (So $c > s$ is possible, just that it means a game where Euclid doesn't end up chosen because snake eyes was rolled before we got to him.) Summing p_{cs} over every possible c and s—every possible game—necessarily gives us 1:

$$\sum_{s=1}^{\infty}\sum_{c=1}^{\infty} p_{cs} = 1.$$

The independence of QR_c and SE_s gives us the following:

$$\begin{aligned} p_{cs} &= \Pr(\mathrm{QR}_c \wedge \mathrm{SE}_s) \\ &= \Pr(\mathrm{QR}_c) \cdot \Pr(\mathrm{SE}_s) \\ &= \Pr(\mathrm{QR}_c) \cdot (1-p)^{s-1} \cdot p. \end{aligned} \tag{1}$$

That final line is because the only way to get snake eyes on round s is by rolling not-snake-eyes $s-1$ times in a row followed by one snake eyes.

We can write the unconditional probability of death like this:

$$\Pr(\text{death}) = \sum_{i=1}^{\infty} p_{ii}. \tag{2}$$

That's just summing up all the infinite ways Euclid can be chosen on the same round that snake eyes is rolled.

For the unconditional probability of being chosen to play, we can get it two ways:

$$\Pr(\text{chosen}) = \sum_{s=1}^{\infty}\sum_{c=1}^{s} p_{cs} = \sum_{c=1}^{\infty}\sum_{s=c}^{\infty} p_{cs}. \tag{3}$$

In the first double sum, the outer sum iterates over every round s on which we might roll snake eyes and the inner sum covers all the cases where Euclid is chosen on or before s. In the second double sum, the outer sum iterates over every round c in which Euclid can be chosen and the inner sum covers all the cases where snake eyes is rolled on or after c.

Eventually we want to find the probability of death given that Euclid is chosen. As we saw in the original derivation, Bayes' Law tells us that this is $\Pr(\text{death})/\Pr(\text{chosen})$. But first let's compute $\Pr(\text{death} \mid \text{chosen} \wedge \mathrm{QR}_c)$, Euclid's probability of death given that he is chosen on a particular round c. We expect that probability to be $p = 1/36$ because it amounts to the one-shot scenario: a specific round c when Euclid is chosen means there's exactly one way for him to die, namely, rolling snake eyes on that specific round. To be totally sure, and to sanity-check our p_{cs} definition, let's now compute it rigorously. We start with the definition of conditional probability:

$$\Pr(\text{death} \mid \text{chosen} \wedge \mathrm{QR}_c) = \frac{\Pr(\text{death} \wedge \text{chosen} \wedge \mathrm{QR}_c)}{\Pr(\text{chosen} \wedge \mathrm{QR}_c)}.$$

The numerator can also be written $\Pr(\mathrm{SE}_c \wedge \mathrm{QR}_c)$ or p_{cc}, the probability that Euclid is both chosen in round c and that snake eyes is rolled on round c. And we can write the denominator in terms of p_{cs} by summing over all the ways Euclid can be chosen in round c, namely by snake eyes being rolled on or after round c:

$$\frac{p_{cc}}{\sum_{s=c}^{\infty} p_{cs}}. \tag{4}$$

Now we use (1) to expand that to

$$\frac{\Pr(\mathrm{QR}_c) \cdot (1-p)^{c-1} \cdot p}{\sum_{s=c}^{\infty} \Pr(\mathrm{QR}_c)(1-p)^{s-1}p}$$

and cancel common factors (notice we're summing over s, not c) to get this:

$$\frac{(1-p)^{c-1}}{\sum_{s=c}^{\infty}(1-p)^{s-1}}.$$

Because the denominator is a geometric series starting at $(1-p)^{c-1}$ and with common ratio $1-p$ we can replace it with its closed form and simplify the above to this:

$$\frac{(1-p)^{c-1}}{\frac{(1-p)^{c-1}}{p}}.$$

And that simplifies to p. Phew!

Knowing that (4) equals p implies that

$$\sum_{s=c}^{\infty} p_{cs} = \frac{p_{cc}}{p}. \tag{5}$$

Finally we have everything we need to work out Euclid's chances of dying if he's chosen to play. Recall that

$$\begin{aligned}\Pr(\text{death} \mid \text{chosen}) &= \frac{\Pr(\text{chosen} \mid \text{death}) \Pr(\text{death})}{\Pr(\text{chosen})} \\ &= \frac{\Pr(\text{death})}{\Pr(\text{chosen})}.\end{aligned}$$

By (2) and (3), that becomes

$$\frac{\sum_{i=1}^{\infty} p_{ii}}{\sum_{c=1}^{\infty} \sum_{s=c}^{\infty} p_{cs}}.$$

Coup de grâce coming up. The inner sum in the denominator is the left-hand side of (5) so we can substitute that in like so:

$$\frac{\sum_{i=1}^{\infty} p_{ii}}{\sum_{c=1}^{\infty} \frac{p_{cc}}{p}}.$$

And we're home free. Factor out the $1/p$ and the sums are the same sum:

$$\frac{\sum_{i=1}^{\infty} p_{ii}}{\frac{1}{p} \sum_{c=1}^{\infty} p_{cc}}.$$

They cancel and the $1/p$ flips to the top as p and we're done! □

Amazingly, we didn't need to define a finite version of the game. We just need a valid prior on when Euclid is chosen. And even more amazingly, the answer is completely independent of what that prior is. For example, say it's uniform for the first N possible values of where Euclid is in the queue of people in the pool. Now compute Pr(death | chosen) in terms of N. The answer, as we just saw, is p. No N in sight. So in the limit as our prior approaches uniform? Still p.

Or maybe you don't like that the above prior has a finite cutoff. No problem. Here's a prior that's both arbitrarily close to uniform and puts positive probability on all infinitely many future rounds in which Euclid could be picked:

- Euclid's probability of being chosen first is 1 in a million
- Euclid's probability of being nth in the queue is 99.9999% as much as his probability of being $n-1$st.

In the limit as that "million" goes to infinity (and the $99.9999\% = 1 - 1/10^6$ correspondingly goes to 1) we again have a uniform prior. Paradox: resolved and double-resolved.

Bibliography

Bartha, P. and Hitchcock, C. (1999). The shooting-room paradox and conditionalizing on measurably challenged sets. *Synthese*, 118(3):403–437.