

FiveThirtyEight's May 1, 2020 Riddler

Emma Knight

May 5, 2020

The riddler this week is about flipping coins in private:

Question 1. *You are locked in the dungeon of a faraway castle with three fellow prisoners (i.e., there are four prisoners in total), each in a separate cell with no means of communication. But it just so happens that all of you are logicians (of course).*

To entertain themselves, the guards have decided to give you all a single chance for immediate release. Each prisoner will be given a fair coin, which can either be fairly flipped one time or returned to the guards without being flipped. If all flipped coins come up heads, you will all be set free! But if any of the flipped coins comes up tails, or if no one chooses to flip a coin, you will all be doomed to spend the rest of your lives in the castle's dungeon.

The only tools you and your fellow prisoners have to aid you are random number generators, which will give each prisoner a random number, uniformly and independently chosen between zero and one.

What are your chances of being released?

Extra credit: Instead of four prisoners, suppose there are N prisoners. Now what are your chances of being released?

First off, a direct answer to the question. It is fairly clear (at least to the author) that the optimal strategy is that each person chooses a number p_N (which will be written as p in blatant abuse of notation), decides to not flip if their random number is less than p , and flips otherwise. Moreover, everyone chooses the same number as each person is unable to distinguish themselves from everyone

else. One then easily computes the odds of success:

$$\begin{aligned}
 P(\text{exactly } k \text{ people flip a coin, none get tails}) &= \binom{N}{k} p^{N-k} \left(\frac{1-p}{2}\right)^k \\
 P(\text{nobody gets tails}) &= \sum_{k=0}^N \binom{N}{k} p^{N-k} \left(\frac{1-p}{2}\right)^k \\
 &= \left(p + \frac{1-p}{2}\right)^N \\
 &= \left(\frac{1+p}{2}\right)^N \\
 P(\text{nobody gets tails and someone flips a coin}) &= \left(\frac{1+p}{2}\right)^N - p^N
 \end{aligned}$$

An alternative derivation of that formula: the odds that one logician doesn't flip tails is $p + \frac{1-p}{2} = \frac{p+1}{2}$, so the odds that nobody flips tails is $\left(\frac{p+1}{2}\right)^N$. Since the odds that nobody flips a coin is p^N , one gets the desired formula.

This is now a simple optimization problem: let $f(p) = \left(\frac{1+p}{2}\right)^N - p^N$, and we want to find the maximum of f . Setting $f'(p) = 0$, one gets $\frac{N}{2} \left(\frac{p+1}{2}\right)^{N-1} = Np^{N-1}$. Letting $\alpha = \sqrt[N-1]{1/2}$, one can take $(N-1)^{\text{st}}$ roots, solve the linear equation, and get $p = \frac{\alpha}{2-\alpha}$. Plugging this in, one gets $f\left(\frac{\alpha}{2-\alpha}\right) = \frac{1-\alpha^N}{(2-\alpha)^N}$, which is the probability of success. When $N = 4$, one gets a 28.48...% chance of success.

But there's more to this story: notice that $\ln(\alpha) = \frac{-\ln(2)}{N}$, so, for N large, $\alpha \approx 1 - \frac{\ln(2)}{N}$. Additionally, $2 - \alpha \approx \alpha^{-1}$, so the probability that the logicians choose is roughly α^2 . Additionally, one also sees that $\frac{\alpha^2+1}{2} \approx \alpha$, so the probability of success is roughly $\alpha^N - \alpha^{2N} \approx \frac{1}{2} - \frac{1}{4} = \frac{1}{4}$. Interestingly (and surprisingly if this is the first time you are seeing this type of phenomenon) this is bounded away from zero; this says that no matter how many prisoners you have, there is a probability of roughly one in four that they can get everyone to flip tails. Is there a nice way to see this at a glance?

The following is the best I have. Let's say that you choose your probability p such that the odds that nobody flips is $\frac{1}{x}$. Then, as discussed, $p \approx 1 - \frac{\ln(x)}{N}$. To approximate the odds that k people flip a coin, one sees

$$\begin{aligned}
 P(\text{exactly } k \text{ people flip a coin}) &= \binom{N}{k} p^{N-k} (1-p)^k \\
 &\approx \left(\frac{N^k}{k!}\right) \left(\frac{1}{x}\right) \left(\frac{\ln(x)}{N}\right)^k \\
 &= \frac{(\ln(x))^k}{k!x}
 \end{aligned}$$

and one similarly gets that the odds that k people flip a coin and none get tails is roughly $\frac{(\ln(x)/2)^k}{k!x}$.

Summing this up, one sees that the probability that nobody gets tails is roughly $\frac{\exp(\ln(x)/2)}{x} = \sqrt{x}$, and so the odds of success are $\sqrt{x} - x$, which is maximized at $x = \frac{1}{4}$.