

This is our question:

Scenario 7: Bank v Casino

A bank plays the following game against a casino, for which the bank pays the casino a price \$x to start playing.

A coin is tossed repeatedly until it lands heads up. If heads appears on the first toss, the casino pays the bank \$2. If heads first appears on the second toss, the casino pays the bank \$4. If heads first appears on the third toss, the casino pays \$8; on the fourth toss, \$16; on the fifth toss, \$32; and so on.

How much should the bank be willing to pay the casino to start playing this game?

Proposition A: The bank should play no matter how high the price \$x is.

Proposition B: The bank should only play if the price to start, \$x, is less than the finite amount \$ ____ .

Let's get the expected value:

$$E = \frac{1}{2} * 2 + \frac{1}{4} * 4 + \dots + \frac{1}{2^k} * 2^k * \dots = 1 + 1 + 1 \dots = \infty$$

Since we can actually have all of these values.

However, are you actually willing to cash out any amount of money for this infinite prize pool? Not really...

In fact, this is a common problem that's been studied for hundreds of years.

But first let's show that infinite dollars arrive for us.

So, suppose we have a finite number of turns. Let's say three to begin with. We will have an expected value, by playing the game, of $2 * 1/2 + 4 * 1/4 + 8 * 1/8 = 1 + 1 + 1 = 3$. If you do this for n times, we have an expected value of \$n.

$$\lim_{n \rightarrow \infty} n = \infty$$

So of course we have an infinite expected value.

*The profits we make will be $2^n - c$ where n is the number of flips in a row until heads and c is the initial price paid to play the game. The bank makes $c - 2^n$ dollars.

But let's test our game out. Let's have a play.

[Now we will view a Python program I made to flip coins for the problem.]

So, it doesn't seem like we should actually cash out a lot of money into this game per round, especially if we're getting such low amounts! You would call someone crazy to even bet \$1000 into the game. (We'll get into the maximum or expected amount they should pay soon – the results are more nuanced)

So, it was definitely hard to even earn a bit of money, right?

Well, let's see the type of reasoning that would make someone believe the flip doesn't have infinite value, based on their intuition (and that's the problem – humans find logic hard, mathematics harder, and probability even more challenging (common(?) saying by mathematicians)).

This is a seemingly “good” argument, until you realise you actually make way more money than you think from this game:

1st: \$2

2nd: \$4

3rd: \$8

...

(n-1): $\$2^{(n-1)}$

n: $\$2^n$

But each consecutive probability is halved, so the increase in money is cancelled out with likelihood of the event occurring.

$$E_k = \frac{1}{2^k} * 2^k = 1$$

Now if we go delta more steps to reach a head, finally, from this value, the expected value of that head is:

$$E_{(k+\Delta)} = \frac{1}{2^{(k+\Delta)}} * 2^{(k+\Delta)} = 1$$

Now we just sum these two expected values to get 2.

Now, the problem is that it fails to realise that we should add up the expected values because if the game continues without heads we can keep adding up the profits. With equal expected values for each separate point where we reach heads, we can notice that rather than reaching a specific value once (like a game where you need to only get say u things in a row), we can reach any of them and make money. Thus it adds up without bound, reaching positive infinity.

So, how much are you willing to pay to play the game?

Well, if you were a robot with purely “mathematical” reasoning, then you would place your life savings on a single game and be okay with that.

But well, no rational human is going to do that? Why?

Well, although you can earn a million, a million+1, a trillion, a googol amount of money, these high values have such low probabilities that we don’t even think about them. I mean, in a finite world even 100 flips in a row is $2^{100} = 10^{(100 \log(2))} = 10^{30}$, which is a nonillionth chance of that happening. However, you do earn a nonillion dollars if you do!

Also, if you were to bet say \$1000, you should be scared, especially if you’re not some billionaire that can have a few tries at the game to earn their money back or give the law of large numbers a test. What if you only get away with a heads on the first try, earning back \$2? Us humans are scared of loss even more than we want profits, oftentimes. I mean, to even get \$1000 you need about 10 in a row, and that’s already tough.

So, then, how have mathematicians worked with this problem to find out how much a human would pay?

I didn't realise this problem was actually one that was interesting to mathematicians until recently while browsing economics articles on Wikipedia. It turns out this problem is called the "St Petersburg Paradox". There's no actual paradox (mathematically), but the paradox is that from anybody surveyed to play this game, people would rarely spend large amounts of money. Many said they would've only paid around \$5 for a game, with a maximum of \$25.

Well, Daniel Bernoulli, who popularised the problem something like three centuries ago, used a new methodology related to the wealth of the player, and how much they would be willing to spend. His work on utility theory, which basically tells you the worth or value of something, is still used in economics today.

Let's get into the formulation of one of Bernoulli's utility functions.

A utility function models risk-taking behaviour such that:

- If someone has more wealth, she will be much comfortable to take more risks, if the rewards are high. (i.e. a rich gambler)
- But, if someone has less wealth, she will be more concerned about the worse case, and therefore, she will think twice before taking a risk of losing, even though, the reward can be high.

Bernoulli wants a function that allows us to analyse this risk, and he uses wealth as the variable for the function. He proposes that marginal utility is inversely proportional to wealth. That's because marginal utility diminishes. What I mean here is that something like \$1000 might seem like a lot to you, but to Bill Gates it's nothing. That's why we have this inverse correlation. Note that marginal utility is the derivative of utility.

So,

$$\frac{d}{dw} U(w) = \frac{a}{w}$$

Solving for $U(w)$, we just integrate both sides with respect to w :

$$U(w) = a * \ln(w) + c$$

Now let's use this function in our problem.

Instead of the expected value, we use "expected" utility.

For each possible event, the change in utility, which is $\ln(\text{wealth after the event}) - \ln(\text{wealth before the event})$, will be weighted by the probability of the event occurring (which we have as our powers of $\frac{1}{2}$). Let c = cost charged to enter a game.

Our wealth after the event is:

$w + 2^k - c$, since we start off with \$ w , gain \$ 2^k from the game, and subtract the \$ c we paid to play.

Our wealth before the event is of course just w .

And the probability for each possible point is $1/(2^k)$.

If we add these up we get the total utility.


Combining all of the information we have, the expected incremental utility of the lottery now converges to a finite value:

$$\Delta E(U) = \sum_{k=1}^{\infty} \frac{1}{2^k} (\ln(w + 2^k - c) - \ln(w)) < \infty$$

And from this formula, we can work out c (how much the player is expected/willing to pay) if we input some value of w (if we know how wealthy the player is).

We can put this function into Desmos, leave our single variable as c by setting some w, and by looking at the x-intersect of this summative logarithmic function we get the price the player is willing to pay to play.

Let's do this for say somebody with a worth of \$2. I tried doing higher values, but Desmos couldn't graph the series precisely, likely due to computational problems.



Interestingly, and a result which doesn't happen once enough wealth is achieved, is that the c-value here is \$3.35, which is more money than the player has. This suggests they should borrow \$1.35 and gamble with this money for profits. The reason I say "gamble", is because there's still risk involved, and the expected value shows what happens after infinite turns in this case, in terms of simulation.

*Note that this is just one proposed idea to predict how much somebody's willing to pay to play.

Many mathematicians have proposed their own solutions to work out how much a human is willing to play, however Bernoulli's function, which he created centuries ago, is one of the most popular and one of the most intuitive.

PROPOSITION A IS TECHNICALLY CORRECT, BUT PROPOSITION B APPLIES TO HUMANS. IN A REAL-WORLD ENVIRONMENT, THERE COULD BE NO "INFINITELY-LONG" GAME WITH "INFINITE" REWARD.