

Some (Vaguely Meaningful) Fun With A Coin Toss Game: The "(St.) Petersburg" Game Paradox

Gary J Duggan
gjdgd@btinternet.com

Abstract

A simple coin toss game, attributed to Nicolaus Bernoulli, in the early 1700s results in a mathematical paradox which still appears to be subject to what might be described as "mainly conceptual" rather than "purely mathematical" solutions. A mathematical solution is given showing that, if the number of games is $2^m - 1$ then the average payout per game for this number of games is $m / (2 - (1/2^{m-1}))$.

The "(St.) Petersburg" Game

There are various summaries of the history and assorted "solutions"^{[1][2]}. To summarise the basics of the game and the paradox:-

A coin is tossed. When a "Head" (H) occurs the game ends and there is a payout. Should a "Tail" (T) occur the coin is tossed until an H occurs.

The payout basis is that the number of tokens paid is equal to 2^{p-1} where p is the number of coin tosses in the game.

The task is to estimate how much the player would be prepared to pay to play the game. Thus, the player needs to calculate the average payout per game and to ensure that they offer significantly less than this to play the game; in this way they will expect to profit.

The standard/normal way of calculating the average payout leads to an infinite value.

Given that this is fun and there is no Fields Medal expected, we shall not be too concerned about a lack of mathematical rigour when discussing infinite values; we shall use popular language. It is recommended just to go with the flow.

Some Additional History

The paradox appears to have originated from the famous Bernoulli family in the early 1700s. Daniel Bernoulli went to St Petersburg in 1725 where his brother Nicolaus was professor of mathematics. Nicolaus had uncovered the paradox prior to Daniel's arrival but their discussions of the problem of the game resulted in Daniel pursuing and publicising it. It first appeared in the *Commentarii* of the Academy and this is the origin of its name; "Petersburg".

The paradox was so shocking to mathematicians of the time and in the absence of any obvious mathematical solution, (this was over 100 years before the birth of Georg Cantor), that Bernoulli was reduced to inventing the concept of "moral expectation" (as opposed to "mathematical expectation") as a means of explaining away the paradox; the first "conceptual" approach to resolving the paradox. Inventing new concepts by way of explanation has become popular over the years, some of which have found application in other areas. For some reason a simple mathematical solution does not seem to have been offered.

At the time, the problem was subject to an experimental mathematical test by Georges-Louis Leclerc (Comte de Buffon) who played a total of 2084 games and paid out a total of 10057 tokens, (Crowns in his case), indicating that the average payout was 4.83. It is worth noting that the formula previewed in the Abstract gives a value of 5.52 which is comfortably close to the experimental value. Probably, if a set of 2084 games were to be played a very large number of times then the average value for the average payout would approach 5.52; or not?

The (Short-Form) Calculation

Consider playing 1 game; it is necessary to calculate the probability of all of the possible outcomes of this game occurring. The probability of the game ending on the first toss (H) is $1/2$. The probability for TH is $1/4$; the probability for TTH is $1/8$ etc. The general probability is $1/2^p$ where p is, as before, the number of coin tosses in the game.

We now have the probability of any particular game occurring and we have the reward for that particular game defined. The table below shows the fractional game outcome probabilities, the reward, their product and the summations of the two infinite series.

Game Outcome Probability	Game Reward	Product (Game Fraction Reward)
1/2	1	1/2
1/4	2	1/2
1/8	4	1/2
1/16	8	1/2
1/32	16	1/2
1/64	32	1/2
1/2 ^p	2 ^{p-1}	1/2
.....		
Column Sum tends to: 1		Column Sum tends to: ∞/2

Thus for 1 game the average reward will be ∞/2. In theory the player should offer an enormous number of tokens for the privilege of being allowed to play the game. Should the player do this they will quickly discover that there is something wrong with the mathematical analysis; something wrong with a mathematical analysis that they have used on many occasions in other circumstances and achieved a fully correct outcome/result.

This is the "(St.) Petersburg Paradox". So where does the problem lie and what is the mathematical resolution to the paradox?

The Root Of The Problem

The game is cleverly constructed so that not only are the individual games not all of the same length (probability of occurrence) but there is the potential for any individual game to be infinite. The potential for a single game to be of infinite duration negates the usual (short-form) methods of calculating probabilities. The problem is therefore linked to infinity and thus a resolution would require a more precise understanding of infinity than was available in the early 1700s.

The (Long-Form) Calculation

Consider a simple 1 toss game (p=1) where 1 token is paid for H and 0 tokens are paid for T. The task, as previously is to calculate the average reward and the usual short-form way of doing this is as follows:-

- i The probability of H is 1/2 and the probability of T is 1/2
- ii The average reward is $(1/2 \times 1) + (1/2 \times 0) = 1/2$
- iii The average reward over n games is $(n \times 1/2) / n (= 1/2)$

In the long-form procedure we shall use the terms "most likely" distribution of results when considering n games and "theoretical" distribution of results for an infinite number of games. For simplicity and clarity we shall also consider only a "convenient" number of games where "convenient" requires that the number of games chosen is capable of producing the required most likely (and ultimately, theoretical) distribution. For this simple 1 toss game we are therefore only considering an even number of games. We start from the same position as for the short-form.

- iv The probability of H is 1/2 and the probability of T is 1/2.
- v As the number of games played increases the distribution of the number of games between the two outcomes will approach closer and closer to the most likely distribution for the number of games (n) played.
- vi When the number of games is infinite *with respect to the number of possible outcomes of a game* the calculated theoretical distribution is achieved.
- vii When the conditions of vi are met we may write $((\infty/2 \times 1) + (\infty/2 \times 0)) / \infty$ as the expression for the average reward over an infinite number of games. Permitting the ∞s to cancel we obtain 1/2.
- viii To find the average reward for n games ∞ is replaced by n and obviously 1/2 is again obtained.

It may be said that the long-form of the procedure is unnecessarily pedantic and it is, excessively so, for finite games. Where a game is potentially infinite the pedantry is not unnecessary.

For the (St.) Petersburg game, using the short-form procedure we have shown, (by a procedure equivalent to ii above), that the average reward per game is $\infty/2$ and we can say the average reward over n games is $(n \times \infty/2) / n$ which is also $\infty/2$.

Using the long-form procedure differences occur and in particular the basis of v_i (*emphasis*) becomes important.

The individual probabilities for the fractional game outcomes obviously remain the same; $1/2, 1/4, 1/8$ etc. We should now consider a "convenient" number of games and this number needs to take into account the fractional outcome probabilities. As n increases the most likely distribution of the outcomes will follow these fractional probabilities; there will be twice as many H outcomes as TH outcomes and twice as many TH outcomes as TTH outcomes etc. So, for example, for 15 games the most likely distribution would be 8, 4, 2, 1 for $1/2, 1/4, 1/8, 1/16$. Generalising this distribution, a convenient number for n will therefore take the form $2^m - 1$ (which is the (St.) Petersburg game equivalent of demanding even numbers for the simple $p=1$ H/T game described previously), where m is an integer >1 .

It is now necessary to consider the condition (*emphasis*) given in v_i . We require to establish the number of games which is both infinite in respect of the (∞) number of outcomes and which is also a convenient number. A number of games satisfying both conditions is therefore $2^\infty - 1$ (which it will be noted is a higher order of infinity than the number of outcomes).

By analogy with v_{ii} we can calculate the total reward and divide it by the number of games to generate the average reward:-

$$\begin{aligned} \text{Total Reward} & \quad (2^{\infty-1} \times 1) + (2^{\infty-2} \times 2) + \dots (2^{\infty-r} \times 2^{r-1}) \dots (2^0 \times 2^{\infty-1}) = \infty \times 2^{\infty-1} \\ \text{Average Reward} & \quad (\infty \times 2^{\infty-1}) / (2^\infty - 1) \end{aligned}$$

Dividing top and bottom by $2^{\infty-1}$ permits the Average Reward to be rewritten as $\infty / (2 - (1/2^{\infty-1}))$

It is worth noting that this reduces to $\infty/2$ (the result from the short-form method) when allowing for the different orders of infinity.

Finally, ∞ can be replaced by m to produce $m / (2 - (1/2^{m-1}))$ and we see that the Average Reward has a dependency on the number of games played; which is what was intuitively clear from the initial analysis in respect of how many tokens should be offered for the privilege of playing the game. It is worth emphasising that n is the number of completed games, not the number of coin tosses. The table shows the Average Reward for some values of m .

m	n	Average Reward
2	3	1.3333333
4	15	2.1333333
6	63	3.0476190
8	255	4.0156863
10	1023	5.0048876
12	4095	6.0014652
14	16383	7.0004273
16	65535	8.0001221
18	262143	9.0000343
20	1048575	10.0000095

The average reward tends to $m/2$ as m increases and hence when m is infinite the average reward is $\infty/2$ which agrees with the short-form result. However, the average reward per game (n) increases only very slowly with n . In simple terms the short-form approach to the Petersburg game picks the incorrect order of infinity for the game; the long-form picks the correct order of infinity for the game.

Generalisation

The procedure may be generalised for a game with any number of outcomes; b . The general form of the short-form procedure is shown below. The probability of one defined outcome, of the b possible outcomes, occurring is $1/b$

and the probability of this one defined outcome not occurring prior to this is $(1-1/b)^q$ where q is the number of undefined outcomes which occur prior to the defined outcome occurring.

Game Occurrence Probability	Game Reward	Product (Game Fraction Reward)
$(1-1/b)^0 \times 1/b$	$1/(1-1/b)^0$	$1/b$
$(1-1/b)^1 \times 1/b$	$1/(1-1/b)^1$	$1/b$
$(1-1/b)^2 \times 1/b$	$1/(1-1/b)^2$	$1/b$
$(1-1/b)^3 \times 1/b$	$1/(1-1/b)^3$	$1/b$
$(1-1/b)^4 \times 1/b$	$1/(1-1/b)^4$	$1/b$
$(1-1/b)^5 \times 1/b$	$1/(1-1/b)^5$	$1/b$
$(1-1/b)^q \times 1/b$	$1/(1-1/b)^q$	$1/b$
.....		
Column Sum tends to: 1		Column Sum tends to: ∞/b

Thus for 1 game the average reward will be ∞/b . For $b=2$ the generalised table reproduces the binary outcome table given previously. A similar generalisation may be applied to the long-form if required.

As an aside; it is worth noting that a single game with infinite outcomes of known probabilities can be a convenient way of generating infinite series where the sum is (by definition) equal to 1 and where it would not be the simplest task to establish the value of the summation by other means. The generalised occurrence probability column summation is the infinite series:-

$$\sum_{q=0}^{q=\infty} \frac{1}{b} \left(1 - \frac{1}{b}\right)^q$$

where b is an integer. The summation is 1 by definition.

Numerical Checking

A simple programme, replacing the Herculean efforts of Leclerc, can be used to check the predictions of the long-form procedure. The elements are:-

- i A random number generator produces either a 0 or 1
- ii For efficiency, two games may be played in parallel; one game ending with a 0 the second game ending with a 1. The number of games is stored (for each of the two parallel games).
- iii The reward is calculated for each (parallel) game and the cumulative reward is stored.
- iv When the number of games (for either the "0" game or the "1" game) equals 2^m-1 (where m is an integer), the average reward per game is calculated and stored.
- v Results are outputted in graphical form where the y axis is the average reward and the x axis is m

When the author first did this using the built-in BASIC programming language on a Computers Lynx 48k computer^[3] (with output to cassette tape; chormium dioxide of course) in the early 1980s it took quite a long time to reach $m=15$; but reach it it did.

Commonly, after the initial wild fluctuations settle down, a curve is generated which is not as smooth as might be expected. Oscillations may be observed, steps up/down may occur, or some otherwise disjointed curve is produced; this indicates that the random number generator built into the programme is not actually generating truly random numbers.

The randomness of the random number generator can be improved by, for example, using it to fill, say, an 8x8 array with 0s and 1s and then using it to choose a random row (by generating a random number 1-8) and a random column (again by generating a random number 1-8) and using the 0 or 1 found from the row/column co-ordinates. This simple "trick" of introducing a second type/level of randomisation is usually sufficient to ensure that, as m increases, the curves for both the 0 and 1 parallel games approach smoothly ever closer to the value predicted by the equation.

The same procedure can be adapted for the general case where there are b outcomes. In particular the matrix used to improve the randomisation will require (greatly) extending.

A More Serious Suggestion

With due regard to the comments in Numerical Checking concerning the limitations of random number generators it is possible to turn the analysis of the (St.) Petersburg game round. Having established the relationship between the average reward and the number of games it would be possible, when dealing with very long sequences of numbers, to use the degree of compliance achieved to the predictions of the equation as a measure of randomness in the sequence. The sequence could be binary or any base via the generalisation given. A failure to achieve, ultimately, a smooth approach to the theoretical curve could, for example, imply that there was a pattern in the data stream and there may be applications where this would be a valuable thing to establish.

Finally; "Or Not?"

"Probably, if a set of 2084 games were to be played a very large number of times then the average value for the average payout would approach 5.52; or not?"

We may imagine Leclerc, being a scientist of some standing, deciding that a single experimental mathematical test was not sufficiently reliable. He resolves to repeat the experiment a number of times, 8 to be precise, but being a Count he decides to use some of his staff to assist. 8 workers are given a tent each and are required to play the game and to record the H/T sequence until each one has played 2084 games. In this way Leclerc has 8 identical experiments and hence 8 independent results to average. When all of the staff indicate that the task has been completed the results are passed to Leclerc; he is somewhat unnerved to find that the results are quite scattered and that his reliability situation has not been greatly improved.

Also, unbeknown to Leclerc, Daniel Bernoulli still has a great interest in this paradox and has collected the sequence from each tent and produced one longer sequence of 16672 games. Neither 2084 nor 16672 are "convenient" numbers but we can still apply the formula:-

For 2084 games	$m = 11.02444712$	Average reward = 5.5148711
For 16672 games	$m = 14.02505303$	Average reward = 7.0129472

Leclerc's ambition to obtain a more accurate value for his 2084 game series would appear to be doomed. This is not unexpected where the return is dependent on the number of games played. It is often said that "the wheel has no memory" in the context of playing roulette and this statement is equally applicable to the coin. There is, however, an assumption built into these truisms which is that the game is finite. More fully, the truism should be "the wheel/coin has no memory when playing a game that has no memory"; and when dealing with a game which has an infinite number of outcomes and a return which has a dependence on the number of game played then the second condition is not met.

I trust it has been fun.

References

- [1] https://en.wikipedia.org/wiki/St._Petersburg_paradox
- [2] <https://plato.stanford.edu/entries/paradox-stpetersburg/>
- [3] https://en.wikipedia.org/wiki/Camputers_Lynx