

# St. Petersburg Paradox\*

Graeme Taylor

August 04, 2004

Ariels'<sup>1</sup> writeup above gives a description of a Martingale System to guarantee winning a fixed amount; but the original St. Petersburg paradox is in fact concerned with a slightly different setup- a game of chance seemingly offering infinite potential winnings- yet one which you would be hard pressed to find someone willing to offer a £100 to play.

The St Petersburg Paradox was first described by Daniel Bernoulli (one of many mathematically-notable Bernoullis) in a paper published by St. Petersburg Academy in 1738 (hence the name; Bernoulli had spent time in St. Petersburg as a mathematics professor and during that time profitably worked with Euler.); although the focus of the paper, entitled "*Exposition of a new theory of the measurement of risk*" was on probability and expectation in general and not just the so-called paradox.

The scenario is as follows- a game is to be played which involves tossing a coin until a head is obtained. If on the first attempt a head is obtained, you win 1 ducat. If not, then the coin is tossed again- with a stake of 2 ducats. Each time a tail is obtained, the stake is doubled and the process repeated, until, eventually, the coin lands heads up.

Assuming a fair coin, how much would you be prepared to pay to play this game?

The standard mathematical approach to such a question (at least as far as undergraduate stats level) is to calculate the expected return- an average of all possible outcomes, weighted according to the odds of them occurring. For a much simpler example, if I offered you £1 for an outcome of heads and £2 for an outcome of tails, then you'd expect for a fair coin to win on average £1.50; so you'd gladly play if I was charging less than this as (at least in the long run) you should come out ahead. This value hasn't simply been intuitively grabbed out of thin air; it's  $\frac{1}{2} \times £1 + \frac{1}{2} \times £2$ . To see if you understand what's going on, see if you can verify that the expected (in a mathematical sense) score for one roll of a die is 3.5.

---

\*First appeared on Everything2, at [http://www.everything2.com/index.pl?node\\_id=1668974](http://www.everything2.com/index.pl?node_id=1668974)

<sup>1</sup>An Everything2 user, <http://www.everything2.com/?node=ariels>

What happens if we apply an expectation calculation to the St. Petersburg paradox?

The probability of winning 1 ducat is  $\frac{1}{2}$  (the coin lands on heads for the first throw). Failing that, we win 2 ducats if the second toss yields a head, but the previous one was tails, an event of probability  $\frac{1}{4}$ . Continuing in this vein, we see that the odds of winning  $2^n$  ducats is  $\frac{1}{2^{n+1}}$  (as we are specifying the outcome of  $n + 1$  events of probability  $\frac{1}{2}$ ). So our sum looks like this-

$$\begin{aligned}\text{Expectation} &= \frac{1}{2} \times 1 + \frac{1}{4} \times 2 + \frac{1}{8} \times 4 + \dots \\ &= \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \dots\end{aligned}$$

Or more precisely

$$\sum_{n=1}^{\infty} \frac{1}{2^n} \times 2^{n-1} = \frac{1}{2} \sum_{n=1}^{\infty} \frac{1}{2^{n-1}} \times 2^{n-1} = \frac{1}{2} \sum_{n=1}^{\infty} 1 = \infty$$

In other words, if we really intend the game to last until a head is obtained, our expected payout is infinite! But before you try such a gamble down the pub, consider this- three quarters of the time, your winnings will be just 2 ducats or less (well, more probably 2 quid/dollars, depending on location); and your chance of winning over 100 ducats depends on about a four in a thousand chance. (Specifically,  $\frac{255}{256}$ ths of the time you get 128 ducats or less.)

Bernoulli's solution revived ideas that would prove revolutionary for economics (I say revived since the Swiss mathematician Cramer independently came up with the same solution some ten years before Bernoulli). The idea was to introduce the concepts of expected utility and diminishing marginal utility. Within this framework, utility from wealth is not linearly related to wealth, but suffers from a decreasing rate of increase relative to wealth. Thus although the prizes increase in size at an impressive enough rate to ensure an infinite expectation, what we should really be considering is the utility of each payout: some lesser amount that takes into account the diminishing marginal utility.

For instance, with a utility model that claims that utility is simply the (base 10) logarithm of the raw cash amount, we get the following sum-

$$\frac{1}{2} \times \ln(1) + \frac{1}{4} \times \ln(2) + \frac{1}{8} \times \ln(4) + \dots = \frac{1}{2} \times 0 + \frac{1}{2} \times 0.301 + \frac{1}{8} \times 0.602 + \dots$$

Where the numerators (utility) grow more slowly than the denominators (growing exponentially in base 2) and thus our sequence turns out to have a non-infinite limit. Various means of damping the value to provide a utility measure have been suggested which achieve this goal, from the simple (Cramer suggests simply cutting off at a certain value, that is having a maximum effective payout) to the more complex- square roots or logarithms as used above. In Bernoulli's formulation, the

value of the game depended on the wealth you had prior to playing, with only an infinitely-wealthy man expecting infinite utility (and even in that case, if you can consider ratios of infinities, the payout was a tiny proportion of the original wealth).

Somewhat akin to constructing diagonal arguments, it turns out that these solutions are only of merit for the stated problem (or equivalently scaled weightings). For if we set the payouts based not on exponential growth but rather the growth necessary to offset the decline in utility, whatever the model, we can effectively engineer a new game that now has both infinite expectation and infinite expected utility. We could repeat the process with a more powerful model of utility decline, but that would always be open to an even more outrageously priced version of the game.

The standard criticism of this game is that no-one can truthfully offer it to you, due to the chance of an unbounded payout. Since no bank could cover that, they'd have to demand an infinite down-payment from you first to cover your later winnings: that is, the only fair price for an infinite expectation game is itself an infinite amount. If you have that to hand, why bother playing? Furthermore, the amount of money in the world (and, more trivially, the length of life that can be dedicated to flipping coins) is bounded above, so the scenario is perhaps just too unrealistic (especially in modified forms to account for utility).

So the St. Petersburg paradox is an interesting insight into the conflicts between mathematical theory and the practical realities of finance; and can also lead into ideas such as risk aversion (have you decided how much you'd be prepared to pay yet? Would it make a difference if you had more money?). It's even been linked to the growth of stocks- by changing the probability structure, the expectation series takes on the form of a discounted series of dividend payments.

## References

For the original (save for translation) introduction to the paradox by Bernoulli, see <http://www.math.fau.edu/Richman/Ideas/petersburg.htm>, which itself quotes from the full text which was published in *Econometrica* 22 (1954) 123-136

“*St Petersburg Paradox and growth stocks*” - *Journal of Finance*; Sep57, Vol. 12 Issue 3, p348-364