

A Bit About the St. Petersburg Paradox

Jesse Albert Garcia

March 20, 2013

Abstract

The St. Petersburg Paradox is based on a simple coin flip game with an infinite expected winnings. The paradox arises by the fact that no rational human would risk a large finite amount to play the game, even though the expected value implies that a rational person should risk any finite amount to play it. Here I describe the St. Petersburg Paradox and give some proposed methods of resolving it.

1 Introduction

The St. Petersburg Paradox is the name now given to the problem first proposed by Daniel Bernoulli in 1738. Bernoulli proposes a coin flip game where one flips until the coin lands tails. The payouts double for each toss that lands heads, and an infinite expected value is obtained. Expected value shows what the player should average for each trial given a large amount of trials. The problem has fascinated mathematicians for centuries. Many have proposed to put a fair price on such a game, but even today there is no one solution.

It is rare that rather simple concepts underlie such a puzzling problem. The main concept to grasp is expected value, which is pretty simple given elementary knowledge in probability. An explanation of expected value is given here, followed by a full explanation of the St. Petersburg Paradox and some attempted resolutions.

2 Expected Value

Expected value is the most concept to grasp in order to learn about the St. Petersburg Paradox. Fortunately, it is pretty simple. The expected value of an event tells you how much you can expect to gain or lose by participating in that event. The expected value of is calculated by taking the amount gained or lost for each possible outcome of the event, multiplying it by the probability of that outcome happening, and then adding them all.

For example, the expected value of rolling a die can be calculated. There are six possible outcomes (which are rolling a 1, 2, 3, 4, 5, or 6). You can view rolling a 1 as "gaining" 1, and there is a $1/6$ chance of rolling a 1 (assuming the die is fair, of course). So you multiply $1(1/6)$. There is also a $1/6$ chance of rolling a 2, a 3, a 4, a 5, and a 6. So the expected value is $1(\frac{1}{6}) + 2(\frac{1}{6}) + 3(\frac{1}{6}) + 4(\frac{1}{6}) + 5(\frac{1}{6}) + 6(\frac{1}{6}) = 3.5$. Now of course you cannot possibly roll a 3.5 on a single roll, but expected value does not imply that you can. It simply means that you should expect that to be your average. For instance, if you roll a die six times and turn up three 4's and three 5's, your average for the six rolls is $\frac{4+4+4+5+5+5}{6} = 4.5$, even though you did not roll a 4.5 on any single roll.

Now let's suppose you were playing a game in which you won \$2.50 if you rolled a 1 or 2, and you lost \$1.50 if you rolled a 3, 4, 5, or a 6. Should you play this game? If your expected value is positive, then you should expect to win money, and if your expected value is negative, you should expect to lose money. Let's calculate the expected value for this game. Here, the amount we gain is not the value of the die that appears as in the last example, but rather the amount in dollars that you actually gain or lose. The two outcomes in this game is winning \$2.50 and losing \$1.50. You will win \$2.50 if the die turns up 1 or 2, which has a probability of $1/3$. You will lose \$1.50 (or equivalently, gain $-\$1.50$) if you roll a 3, 4, 5, or 6, which has a probability of $2/3$. So the expected value is $(\$2.50)(1/3) + (-\$1.50)(2/3) = -\$1/6 \approx -\0.17 . In other words, you can expected to average a loss of about 17 cents per play.

Adding expected values is simple. Say you wanted to know the expected value of the sum of the values of two dice. We know that the expected value of rolling a die is 3.5. So if you roll two dice, and each has an expected value of 3.5, then the expected sum of the two is $3.5 + 3.5 = 7$. In the game above, we obtained an expected loss of about one-sixth of a dollar, or about 17 cents. If we played this game six times, and we expect to lose one-sixth

of a dollar on each game, then we should expect $6(-\frac{1}{6}) = -\$1$, or a total loss of one dollar for the six games.

There is a result related to this called the Law of Large Numbers. A paraphrasing of this statement could read: "As the amount of trials of an event increases, the average of all the trials should tend to the expected value." As you approach an infinite amount of trials, the average of all your trials should get closer and closer to the expected value. The proof of the Law of Large Numbers is unnecessary for the scope of this paper, but it can be obtained in any number of probability texts.

3 St. Petersburg Paradox

The paradox known as the St. Petersburg Paradox is obtained by a simple coin flip game. The rules are simple: keep flipping until you get tails. If your first flip is a tails, then you win \$2; if your first tails is on the second flip, then you win \$4; if your first tails is on the third flip, you win \$8, etc. More concisely, if your first tails is on the k^{th} flip, then you will win $\$2^k$.

Table 1: St. Petersburg Game

Trial where first tails appears	Probability	Payout
1	$\frac{1}{2}$	\$2
2	$\frac{1}{4}$	\$4
3	$\frac{1}{8}$	\$8
4	$\frac{1}{16}$	\$16
...

Let's look at how the probabilities in the table were calculated. The probability of rolling a tails on the first try is simply $\frac{1}{2}$. If your first tails is on the second flip, that means your first flip was heads. The probability of this happening is $(\frac{1}{2})(\frac{1}{2}) = \frac{1}{4}$. If your first tails is on the third flip, then your first two flips must have been heads, and the probability of flipping heads, heads, then tails is $(\frac{1}{2})(\frac{1}{2})(\frac{1}{2}) = \frac{1}{8}$, etc.

Let's calculate our expected value for this game. We have a $\frac{1}{2}$ chance of winning \$2, a $\frac{1}{4}$ chance of winning \$4, a $\frac{1}{8}$ chance of winning \$8, etc. So our expected value is $(\$2)(\frac{1}{2}) + (\$4)(\frac{1}{4}) + (\$8)(\frac{1}{8}) + \dots$, or $\sum_{k=1}^{\infty} (2^k)(\frac{1}{2^k}) = 1 + 1 + 1 + 1 + \dots = \infty$. The expected value for this game is infinite! Because

of this, it can be said a player should accept any finite price to play this game because it will be less than the expected winnings.

But how much should one actually pay to play this game? Most people probably would not pay more than \$20 to play this game because the odds of making a profit are so low (about 6.25%). (Casinos usually offer games at prices that are just above the expected value of the winnings so that in the end, they are expected to make a profit; for each game, the casino wins your entry fee and loses the prize amount they have to pay you). If this game was offered for \$1000 dollars, you could spend your whole life trying to find someone who would play this game for that much. There have been many attempts to try and find a resolution to the paradox, or in other words, to try to find a fair price to offer this game for.

4 Resolutions of the St. Petersburg Paradox

There are resolutions of the St. Petersburg Paradox in which people try to put a fair price on this game. They are not exactly resolutions in the traditional sense, in that even today there is not an exact figure on a fair price for this game, and new articles pertaining to a resolution for the St. Petersburg Paradox still appear periodically.¹ Some particularly well-known resolutions have been given.

4.1 Resolution by experiment

The simplest way to try and resolve the paradox to find a fair price for the game was done by Georges-Louis Leclerc, Comte de Buffon, a French mathematician. He simply ran the experiment himself a total of 2,048 times. These were his results:⁵

We can calculate the expected winnings based off this experiment. Be careful when calculating the probabilities. For example, the probability of rolling a tails on the first flip **based off this experiment** (the experimental probability) is $\frac{1061}{2048}$, not $\frac{1}{2}$, because Buffon got this result 1,061 out of the 2,048 trials he did. So the expected winnings based off this experiment is $(\$2)(\frac{1061}{2048}) + (\$4)(\frac{494}{2048}) + (\$8)(\frac{232}{2048}) + (\$16)(\frac{137}{2048}) + (\$32)(\frac{56}{2048}) + (\$64)(\frac{29}{2048}) + (\$128)(\frac{25}{2048}) + (\$256)(\frac{8}{2048}) + (\$512)(\frac{6}{2048}) \approx \9.82 . One should expect to pay about \$10 to play this game, based off this experiment.

Table 2: Buffon's St. Petersburg Experiment

Trial (k) where first tails appears	Frequency	Payout (2^k)
1	1061	\$2
2	494	\$4
3	232	\$8
4	137	\$16
5	56	\$32
6	29	\$64
7	25	\$128
8	8	\$256
9	6	\$512

4.2 Resolutions based on utility

Some resolutions use the concept of utility to derive a fair price for this game. Utility can be thought of as how much practical use you gained. Surely going from \$0 to \$1,000,000 is more dramatic than going from \$1,000,000 to \$2,000,000, even though they are both increases of a million dollars. In the former, you went from nothing to millionaire, and in the latter, you essentially went from rich to rich again. So it can be said that if you won \$2,000,000 then you would gain more utility in your first million than you would gain in your second million.

Daniel Bernoulli used the equation $\text{Utility} = \log_{10}(w)$, where w is the winnings in dollars (Bernoulli actually used a different currency than U.S. dollars, but the currency is arbitrary)⁴. So if you win \$2, then by this equation, you have won $\log_{10}(2) \approx 0.301$ in utility. A new table can be constructed for the St. Petersburg game now:

Table 3: St. Petersburg Game

Trial (k) where first tails appears	Probability	Payout (2^k)	Utility ($\log_{10}(2^k)$)
1	$\frac{1}{2}$	\$2	.301
2	$\frac{1}{4}$	\$4	.602
3	$\frac{1}{8}$	\$8	.903
4	$\frac{1}{16}$	\$16	1.204
...

Now instead of expected winnings in dollars, we can calculate the ex-

pected utility, which turns out to be about 0.602, or about \$4.

This resolution is not satisfactory however. All one must do is change the payouts to $\$10^{2^k}$ and the result is an infinite expected value in dollars and an infinite expected utility, so the paradox returns. In fact, it turns out that any unbounded utility function would prove to be unsatisfactory in the same way. Based off of articles in past centuries mentioning the St. Petersburg Paradox, it took quite a while for the need for a bounded utility function to be recognized.¹

References

- [1] Gilbert W. Bassett, Jr., The St. Petersburg Paradox and Bounded Utility, *History of Political Economy* **19:4** (1987), pp. 517-523
- [2] C. Ariel Pinto and Paul R. Garvey. *Advanced Risk Analysis in Engineering Enterprise Systems*. Boca Raton, FL: CRC Press, 2012. pp. 415-419.
- [3] Kate Crowley (Ed.), Paul A. Samuelson. "St. Petersburg Paradoxes: Defanged, Dissected, and Historically Described." *The Collected Scientific Papers of Paul Samuelson* Vol. 5, Massachusetts Institute of Technology, 1986. pp. 24-54
- [4] Martin, Robert, "The St. Petersburg Paradox", *The Stanford Encyclopedia of Philosophy (Winter 2011 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2011/entries/paradox-stpetersburg/>
- [5] Donald Richards. "The St. Petersburg Paradox and the Quantification of Irrational Exuberance." PowerPoint Presentation. Penn State University. URL = <http://sites.stat.psu.edu/richards/bgsu/bgtalk1.pdf>