Public Finance Research Papers

Dipartimento di Studi Giuridici
Filosofici ed Economici

Istituto di Economia e Finanza

Sapienza
Università di Roma

# Public Finance Research Papers

## Is there a fair price in St. Petersburg repeated games? An empirical analysis

Ruggero Paladini

E-PFRP N. 44

2020

Ruggero Paladini

Sapienza Università di Roma e Unitelma Sapienza

ruggero.paladini@uniroma1.it

Please cite as follows:

Ruggero Paladini (2020), “**Is there a fair price in St. Petersburg repeated games? An empirical analysis**”, *Public Finance Research Papers*, Istituto di Economia e Finanza, DSGE. Sapienza University of Rome, n. 44 (http://www.digef.uniroma1.it/pubblicazioni).

E-PFRP N. 44

2020

Ruggero Paladini

# IS THERE A FAIR PRICE IN ST. PETERSBURG REPEATED GAMES? AN EMPIRICAL ANALYSIS


**Abstract**

Can the Foley value (*n+2.53*) be considered a fair price? It is possible to give a positive answer, even if one has to depart to some extent from the classic St. Petersburg game. The only way to put player and dealer in an equal position is to compare the median of an MC simulation of *N* games with the Foley value. The difference should be, apart from exceptional cases, of the order of one monetary unit or less. Of course this difference has to be multiplied by *N* in favour of the winning side, so that higher *N* implies greater win and risk. The departure from the original game of St. Petersburg therefore takes place in two steps: the first step consists in playing not a single game but several ones; the second in repeating the same number of games with the MC technique and checking the median value with respect to the Foley value. In this way the odds of player and dealer would be balanced; the value of the unit stake would depend on their love for risk.



**JEL Classification:.H 25, H 87, K 34**
**Keywords**: **Business taxation, State aid control, Multinational enterprises**

E-PFRP N. 44

2020

# Is there a fair price in St. Petersburg repeated games? An empirical analysis

## 1. *Introduction*[1]

Between the second half of the 16th and the first half of the 17th, Fermat, Pascal, Huygens, and several others, laid the foundations of the theory of probability; among the many debated topics there is that of a game with infinite expected value, but in which no player would bet a sum if not rather modest. In the standard version, the infinite expected value derives from the fact that, by tossing a coin, the payoff consists of one monetary unit if heads come out on the first toss, two units on the second, four on the third and so on. The infinite value is the product of a probability that halves at each toss and an associated win that doubles simultaneously. Therefore the sum of a half times one + a quarter times two + an eighth times four +… ... implies an infinite series of half unities[2], therefore an infinite value.

The attention of scholars is focused on the paradoxical aspect that no one, to whom the game is proposed, is anxious to offer high sums; various hypotheses are put forward, in particular the one formulated by Daniel Bernoulli (1738)[3], according to which it is not the expected value that the player looks at, but the expected utility. Assuming a utility function with respect to income in logarithmic form[4], the infinite series quickly converges to a finite value. The hypothesis of decreasing marginal utility was taken up by Marshall (1890) and gained importance with the emergence of the marginalist approach in economics. When the famous book by von Neumann and Morgenstern (1944) came out, Bernoulli was rightly considered a precursor of the expected utility function.

There is an aspect that remains neglected for a long time: if the expected value is infinite, it seems obvious that it is not possible to define a fair

---

[1] A few words of premise; this work derives from a discussion with my son Andrea, mathematician, on Rodriguez's paper; Andrea told me: don't look at the mean (like Rodriguez), look at the median. At that point I remembered Foley's paper, reconsidering the meaning of Monte Carlo simulations, and arrived at the conclusion that the law of large numbers works, even if imperfectly, and explains the results, paving the way to a different kind of fair price.

[2] In more recent times it is preferred to start with two monetary units at the first toss, thus having an infinite series of 1 + 1 + 1 +….

[3] Since Bernoulli's essay appeared in the Proceedings of the St. Petersburg Academy, where the scholar was at that time, the problem took the name of St. Petersburg Paradox.

[4] This implies that when the income doubles, the marginal utility (the derivative) halves.

price between player and dealer. Until William Feller (1945) proves that there is indeed a "fair price", not for a single game, but for a repeated number of games (*Appendix 1*). When the number of games ($N = 2^n$) approaches infinity, the probability that the price per game differs from $n$[5] tends to zero. However, Feller puts the expression fair price in quotes, as it warns that, at a given (unit) price of the game, it is possible to incur heavy losses. As we will see this is especially true for the dealer[6].
Feller's analysis was followed by a series of theoretical contributions that highlighted various aspects of the theme of repeated games. The empirical studies appeared later, with the advent of personal computers. We can divide the empirical tests into two groups; those analysing the average payoffs for groups of games with increasing numbers, and those analysing the Monte Carlo (MC) iterations, i.e. the repetition, for a high number of times, of a given number of games. The formers are of limited use; one can note that the payoffs tend to grow with the number of games, but in a rather irregular way. Instead the iterations of MC brought to light more interesting aspects; in particular, the fact that the payoff distribution has similar characteristics both with iterations of a million times or with much less numerous ones. That is, even one hundred MC iterations of a number $N$ of games are enough to obtain a structure similar to that of ten thousand, one hundred thousand or one million iterations.

## 2. *Simulations*

The empirical simulations carried out were conducted to verify the distribution to the various launches of twenty[7] groups of one hundred games, from 16 to 65,536, each repeated fifteen times with the MC program. A first program made it possible to study in detail the distribution of the games released at the various launches. In the next fourteen MC simulations a second program directly provided the one hundred payoffs for the twenty groups of games[8].
Now generally MC simulations are far more numerous repetitions than 100 games. However, it is interesting to note that the twenty distributions

[5] Feller indicates the fair price as *n/2* because he starts with the payoff of a monetary unit at the first toss.

[6] In the case of a dealer that plays with a plurality of players, the statement is also true in the case of only one game (Samuelson 1977).

[7] Thirteen game groups are $2^n$ ones with *n* from 4 to 16, and seven are in the form $(2^n+2^{n+1})/2$. In this second case the theoretical distribution to the different throws ends with 12,6,3. At the next roll, since we cannot have 1.5 games, the distribution closest to the theoretical one is 2 at the nth roll and 1 at the n+1 roll. In this case the payoff is n + 2.66. The programs were created with Python by Andrea Paladini.

[8] However, given a payoff, it is always possible to accurately reconstruct the distribution of the games released in the individual launches. In some cases, the distribution, giving the same payoff, can be more than one, but the most probable is the one closest to the theoretical distribution, for the reasons explained in the text.

all have a similar distribution pattern, with the payoff concentrated between *n-1* and *n+1* about 33%. A much smaller percentage (9%) has payoff values below *n-1*; therefore, more than half of the payoffs have values greater than *n+1*. This means that about half the time, particularly from launches *n+1* onwards, there are one or more rolls in which no game comes up. Sometimes there is only one roll with zero games, more rarely they can be two.
It is also noted that the distribution structure, in the number of games released at the various launches, although it is almost never in total agreement with the theoretical distribution, does not differ excessively from it either. With the increase in the number of games (i.e. from $2^4$ to $2^{16}$) the percentage differences in the first launches are reduced to below 1%; on subsequent throws the percentage increases slightly, and only in the last throws more consistent differences may be noted (but, in some cases, it may happen that there is a coincidence with the theoretical distribution, see *Appendix 2*).
Thus, for example, the reason for starting with a group of 16 games derives from the fact that less numerous game sets, such as 4 or 8 games, have a distribution that is more irregular than the theoretical one. Even in the case of 16 games, 11 times out of 100 the games released on the first launch were less than those released on the second. In the case of 32 games, on the other hand, most of the time the games released in the first three launches follow a descending order; only in 3 games, out of the 100 simulated, those released on the second launch are greater than the first, and 5 times those of the third launch greater than the second. Furthermore, the sum of the games released in the first three throws, most of the time, coincides with the theoretical ones, that is 28 (16 + 8 + 4), or they differ by only one unit. The reason is easy to understand: if 18 games come out instead of 16 on the first roll, there are 14 games remained on the second roll, so there is a high probability that 7 will come out. If on the third roll out of the remaining 7 there are 3, overall we have 28 games which went out in the first three launches[9].
Moving on to more numerous groups of games (64, 128, 256….), one can note that the phenomenon of games that are more numerous at a certain launch compared to the previous one moves progressively; already at 64 no group of games at the second throw is greater than the first. In only one

[9] The distribution of the 28 games in the first three launches has a much lower weight in the overall result of the average payoff than the remaining four, whose distribution is instead decisive in determining whether the price approaches or moves away from 5. Taking some examples from the simulations carried out , with 18-7-3 the contribution is 3.14 instead of 3.43. But the other four come out 1-2-0-1, with a contribution of 6.5 (208/32), so the overall prize is 9.25 instead of 5. Another example: at 128 games a distribution of the first four throws is 64-32-15-9, that is 120 like the theoretical ones, but the remaining eight are distributed like this: 3-1-1-2-1. While the first four lead to 4.33, very close to the theoretical 4.27, the remainder lead to 10.25, for an overall payoff 14.25 more than double than the theoretical 7.

case (always out of 100) they are equal; in one case the number of games released on the third throw are greater than in the second, and in two cases they are the same.
The phenomenon therefore seems to be general: the great majority of games, from 128 upwards, are distributed in a way that, even if it is not equal to the theoretical distribution, it is not either far from it. The differences often tend to compensate each other in such a way that the contribution to the overall average premium does not differ much from that which would occur with a theoretical[10] distribution. For example, in the case of 512 games, a payoff of 9 (that is *n*) was obtained with the following empirical distribution:

| Launches | empirical distribution | theoretical distribution | deviation% |
|---|---|---|---|
| I | 245 | 257 | -4,7 |
| II | 142 | 128 | +10,9 |
| III | 56 | 64 | -12,5 |
| IV | 36 | 32 | +12,5 |
| V | 19 | 16 | +18,8 |
| VI | 10 | 8 | +25 |
| VII | 2 | 4 | -50 |
| VIII | 0 | 2 | -100 |
| VIII | 2 | 1 | +100 |

In the other distributions of the 512 games on the first and second throw, deviations were all lower in percentage terms, but the payoffs were different from 9 (generally, as mentioned, higher).
For smaller sets of games, a distribution close to the theoretical one cannot have *n* as payoff, but *n+1* or *n+2*; for example, with $2^4$ the distribution 8,4,2,2 gives a payoff of 5 while 8,4,2,1,1 gives a payoff of 6. In the simulation carried out with 100 repetitions, 4 was obtained with some distributions, all quite distant from the theoretical one (10,3,2,1 - 6,9,0,1 - 6,7,3). With $2^5$ games, the payoff of 5 was obtained with 16,10,3,2,1 in which the second and third throw differ from the theoretical distribution (8,4). Even the payoff of 6 can be obtained from a distribution very close to the theoretical one (16,8,4,2,2), but also from 20,6,2,1,3 which instead diverges clearly from it[11].

[10] The only change is that one more game is attributed to the first throw.
[11] These payoffs were obtained in the 100 launches performed by Python.

It is the last games, and often the very last (*Appendix 2*), that determine the result, which may be similar to the theoretical one, that is *n*, or more or less distant (especially much higher).

## 3. *Law of large numbers and the nth game*

Suppose we play *N* games ($N=2^n$) where *n* is an integer. The theoretical distribution foresees that *N/2* games come out at the first launch; there are other *N/2* games to be played. The games that come out on the second launch are $N/2^2$ and there are still $N/2^2$ left, and so on. When one arrives at the games that come out at the nth launch, we will have $N/2^n/N/2^n = 1$, but it still remains the last game (Pianca 2007). But obviously, at the launch *n+1*, one cannot get half a game; so it will be *1* or *0*. The same at the launch *n+2* and so on. The meaning of all this is that, always respecting the theoretical distribution, there is 50% probability that the game will come out on throw *n+1*, 25% that it will come out on throw *n+2*, and so on. That is, if the empirical distribution always follows the theoretical one, repeating the exercise several times with a MC simulation, we should find 50% of the time a game released at launch *n+1*, 25% at launch *n+2* and so on.

The problem of the nth game, therefore, is completely general every time we play a number of games equal to $2^n$ where *n* is an integer. If *n* is not an integer, there is inevitably a deviation, however minimal, from the theoretical distribution. For example, if we play 100 games (approximately $2^{6.64}$) we could have either a sequence 50-25-13-6-3-2-1 with an average payoff of 7.5 or 50-25-12-6-3-2-1-1 with an average payoff of 10, both payoffs higher than 6.64.

Returning to $N=2^n$ games with *n* integer, if we perform a series of throws, and if the distribution corresponds to the theoretical one, then the average value of *N-1* games ($N/2^1$ come out at the first throw, $N/2^2$ at the second ... $N/2^n$, i.e. 1 at the nth toss) is given by the general formula of the average value $AV= n2^n/(2^n-1)$, for all repeated games. As one can see, *AV* converges rather quickly to *n*. But, as we have just said, the last game[12] should come out 50% of the time on roll *n+1*, 25% on roll *n+2*, 12.5% on roll *n+3* and so on, thus bringing the *AV* at *n+2, n+4, n+8* and so on. The distribution would follow the theoretical one for *n* rolls, so it can follow 1 at launch *n+1*, or 0 and then 1 at launch *n+2*, or *0, 0,* and then 1 at launch *n+3*. Obviously, as it is easy to guess, and as we will see later, in reality almost never, especially if the number of games is high enough (in fact from $2^6$ onwards), the empirical distribution coincides with the theoretical one, but the deviations are small. On the other hand, we find, in about half of the simulations, the existence of one empty launch (no game comes out) and in a few cases a high number of empty launches, which obviously implies a

[12] Last in the sense of the game that is missing to get *N*.

high payoff, sometimes very high. In such a way that, when the last game comes out, it determines most of the game's payoff.
So by examining the results of the twenty simulated games, each with a hundred MC iterations, and repeating fifteen times (therefore for a complex of 30,000 games), we arrive at the following results: taking the payoffs released, with values *n, n+2, n+4, n+8*..., within a range of ± 0.5% we obtain on average the following results:

*n 12.8%*
*n + 2 12.5%*
*n + 4 9.6%*
*n + 8 5.8%*
*n + 16 3.4%*
*n + 32 1.9%*

The set of the highest payoffs (*n+64, n+128* ...) adds up to 0.9%. Therefore we have that, on the whole, while the mode is found at *n*, about one third of the set of games have high payoffs; in which the payoff is determined by games with a distribution very close to the theoretical one, since it conforms to the law of large numbers that a game can come out after a few empty launches[13].
If we take into consideration the payoffs with values *n ± 10%*, the percentages tend to double on average. Therefore, for payoffs greater than *n*, we can note thickening from one hand, and voids from the other; phenomenon that finds an explanation precisely in the action (even if imperfect) of the law of large numbers (*Appendix 3*). Simulating a group of games by MC, about half of the game groups give payoffs with some launches (particularly from roll *n+1* onwards) empty.
Furthermore, the fact that the empirical distributions are very close to the theoretical ones, particularly in the first launches, explains why the payoffs with values lower than *n-1* are a percentage lower than 10%. In fact, in the case of four games ($2^2$) the chance that all four will be released at the first launch is 6.25%, but at sixteen games ($2^4$) we go down to 1.56% and at sixty-four ($2^6$) games we go down to 0.39%. Overall, of all the simulations, the lowest payoff has always values ranging between 60% and 80% of *n*.

## *4. Foley and the relation between payoff and probability*

We have seen that the distribution of MC payoffs shows a mode (12.8%) around *n*, but with a very strong variability. In tests carried out with ten thousand MC iterations (Rodriguez 2005, Paladini 2017) or with one

---

[13] To give just one example, in one of the fifteen simulations of the 65,536 games ($2^{16}$) it was obtained a payoff of 528.34. This payoff was achieved with a game released on 25th throw, contributing 512 to the overall payoff. It can be deduced that all the other games have given a contribution of 16.34, which can be had when the distribution is very close, even if not identical, to the theoretical one.

million iterations (Klyve and Lauren 2011, Olivero 2016), the mode always ranks around 12.6%, with a strong asymmetry on the right. The probability that a group of $2^n$ games will have an average payoff greater than *n* is therefore very high. Foley (2015) has built an interesting model of the relationship between *n*, as the average price per game, and *p*, the probability for the player to gain or not suffer losses (payoff greater than or equal to *n*, *1-p* being the reciprocal probability for the dealer). He starts from the observation that as the number of games increases, the average payoff tends to grow, but "that the mean one arrives at even after a large sample is not enough to predict with any reasonable accuracy the mean one will get if one repeats the experiment again". Foley shifts his attention to the median of an MC distribution of various sets of games, noting that it tends to grow steadily as the number of simulated games increases.
The model he arrives at is built on the basis of several tens of billions of $2^n$ games, with *n* going from 1 to 15, and aims to identify a relationship between *n* and *p*, the probability of winning (or breaking even). The relationship is as follows:

$$(1) \qquad f(n, p) = (2{,}53+n)^{(\pi + 1 - 1/\mu)}$$

where $\pi = [(p/(1-p))^{(0{,}2+0.01n)}]/\mu;\ \mu = 1/(0{,}65+0{,}115n)^2$

Foley's comment is: "For example, we can take a look at the famous case of 2048 trials by Compte de Buffon in which he arrived at \$9.82 as the mean gain from his experiment. According to the model that we have found, that value would be only the 14th percentile of the distribution. If a casino had done a similar experiment to Compte de Buffon[14] and set the entry fee at \$9.82 then there would be an 86% chance of loss. The 86th percentile, an event of similar likelihood on the higher end of the distribution, would be \$28.83".
We can ask ourselves what probability Buffon (as a dealer) would have had by setting the unit price of the game to 11. The answer is that he would have been in the 25th percentile, thus leaving to a player three-quarters of the probability. By calculating the relationships between *n* and *p*, we see an increasing relationship, and with $2^{30}$, that is, over a billion games, a price of 30 increases the percentage to the 28th percentile[15]. By increasing the number of games, however, the odds no longer increase, as they begin to decrease slightly. It should also be considered that, as the

---

[14] On the Buffon's experiment, see *Appendix 2*.

[15] Passing from *n = 4 to n = 16*, the Foley formula indicates probabilities ranging from 16.7% to 26.2%, with decreasing increases. In the simulations carried out, it is observed that the fifteen iterations, taking the central values as before, have a correlation index (with the Foley probability) of 75.2%. A correlation therefore still exists, but clearly weaker, which is matched by the fact that the standard deviations for the various groups of games (i.e. the fifteen values of p for each *n*) were almost ten times higher (on average around 5) . There remains the problem of how to interpret this difference between the values found in the case of the medians (*p = 0.5* for all *n*) and those of the different probabilities for the different *n*s.

number of games increases, even a small difference from *n* can inflict colossal losses on the dealer (more likely) or the player (less likely).
According to the Foley formula, what would be the fair price in the case of an equal probability between player and dealer? Observing the formulas, it can be seen that in the case of $p/(1-p) = 1$ function *(1)* is reduced to

$$(2) \qquad f(n, 2^{-1}) = n+2{,}53$$

a price that is quite close to *n+2*, that is to the average payoff in the event that the distribution is equal to the theoretical one and in which the last game comes out at throw *n+1*.

*5. A test of Foley's median*

I carried out, for the twenty groups of games, one thousand five hundred MC iterations; thus obtaining a total of three hundred medians. The standard deviations (mean square error) of the twenty groups are very low; between the highest and lowest values, the difference is on average of two units, so the standard deviations are around 0.5. Let's see the relationship between the central value and that of Foley (i.e. *n+2.53*). The central value is obtained by taking the average of the three central values for each group of games:

| Number of games | Empirical median | Foley's median |
|---|---|---|
| $2^{4}$ | 6,69 | 6,53 |
| $2^{5}$ | 7,47 | 7,53 |
| $2^{6}$ | 8,27 | 8,53 |
| $2^{6,59}$ | 9,48 | 9,12 |
| $2^{7}$ | 9,5 | 9,53 |
| $2^{8}$ | 10,42 | 10,53 |
| $2^{9}$ | 11,49 | 11,53 |
| $2^{9,59}$ | 12,02 | 12,12 |
| $2^{10}$ | 12,7 | 12,53 |
| $2^{10,59}$ | 13,01 | 13,12 |
| $2^{11}$ | 13,89 | 13,53 |
| $2^{11,59}$ | 14,3 | 14,12 |
| $2^{12}$ | 14,55 | 14,53 |
| $2^{12,59}$ | 15,36 | 15,12 |
| $2^{13}$ | 15,81 | 15,53 |
| $2^{13,59}$ | 15,98 | 16,12 |
| $2^{14}$ | 16,66 | 16,53 |
| $2^{14,59}$ | 17,3 | 17,12 |
| $2^{15}$ | 17,57 | 17,53 |
| $2^{16}$ | 18,52 | 18,53 |

As can be seen, the two series have very similar values and trends; the correlation index (Pearson) between the two series is 99.8%. The analogous index for each of the single median fifteen (of the twenty games) compared to that of Foley ranged between 98.4% and 98.6%. Moving to the central value, the index therefore rose by more than one point[16]. In truth, by increasing the Foley values by three hundredths of a point (i.e. *n+2.56*), the correlation index goes to 99,9%. But it is plausible to think that by repeating the iterations there could always be variations of a few cents more or less.
It seems therefore to be a confirmation, in the case of MC simulations of *N* St. Petersburg games, of the Foley value as regards to the median of the simulations (i.e. when the probability is even for both player and dealer), with the simple relation: *n+2.53*.

*6. Conclusion*

Can the Foley value (*n+2.53*) be considered a fair price? It is possible to give a positive answer, even if one has to depart to some extent from the classic St. Petersburg game. The original logic consists in giving a price for a single game. In Feller's version we have a price for the average payoff of several games; the difference between the unit price and the payoff, multiplied by the number of games, determines the player or dealer win. But in this way, even using the Foley value instead of *n*, the dealer would be exposed to much higher losses than the player[17], given the asymmetry on the right in the payoff distribution, highlighted by all the MC simulations. That is, although there is an equal probability of winning or losing between the two sides, the possible losses of the dealer would be much greater than those of the player.
The only way to put player and dealer in an equal position is to compare the median of an MC simulation of *N* games with the Foley value. The difference should be, apart from exceptional cases, of the order of one monetary unit or less. Of course this difference has to be multiplied by *N* in favour of the winning side[18], so that higher *N* implies greater win and risk.
The departure from the original game of St. Petersburg therefore takes place in two steps: the first step consists in playing not a single game but several ones; the second in repeating the same number of games with the MC technique and checking the median value with respect to the Foley value. In this way the odds of player and dealer would be balanced; the value of the unit stake would depend on their love for risk.
It is worthwhile to note the difference, with respect to the proposed solution of a fair price for repeated games, from the case discussed in

[16] Godex and Dulaney (2012) provide nine medians with iterations of 100 MC (games from 10 to $10^9$); the correlation with the Foley's median is 82.6%.
[17] The last case happens when the payoff is less than *n+53*.
[18] The choice of a more or less large N therefore affects only the amount of the stake.

Samuelson (1963), where he proves that if someone declines a single gamble, he should decline any number of repeated gambles, if preferences conform to expected utility theory. Samuelson considers the opposite view as a fallacious application of the law of large numbers. Chew and Epstein (1988) instead obtain a positive answer, using non expected utility. Here things are different, since what is implied is a reference to the median of a MC simulation of a given number of St. Petersburg games. As we have shown, considering each single payoff[19] of a MC simulation would imply an almost certain loss for the dealer.

*Appendix 1*

*Feller's approach*

In 1945 the statistician William Feller approached the problem of the St. Petersburg paradox by a different point of view from that on which scholars had concentrated for two centuries. He argued that even if the single game has an infinite expected value, due to the fact that the prices increase in the same way as the probabilities decrease, it was however possible to define the "fair" price of a number $N$ of repeated games. He showed that the average "fair" price P* tends in probability to $log_2N$ when $N$ tends to infinity (Feller 1945 and 1957). In other words, if $N=2^n$, the "fair" price P* will be the closer to $n$ the higher the latter.

Therefore, even if the expected value of the single game is infinite, in the case of a large number of games, a fair price can be defined, that is $n$[20]. He warned, however, that for a given $N$ the price $n$ could result in significant losses for the players; this is the reason why he used the term "fair" in quotation marks.

To clarify the logic of Feller's reasoning, if the number of throws follows the distribution given by the theoretical probabilities (i.e. that heads come out at the first throw with probability 50%, at the second with probability 25%, etc.) the total price Pt is given by

$$Pt = (2^n/2)2 + (2^n/4)4 + .... + (2^n/2^n)2^n,\ \text{that is } n \text{ times } 2^n;\ \text{hence } P^* = n.$$

Obviously the result is necessarily approximate if $N$ is not a multiple of 2, or if $n$ is not an integer, because in this case the games cannot come out, at any number of launches, a fractional number of times.

For example, if we suppose that we have 128 ($2^7$) games, and that these are distributed exactly according to the a priori probabilities: 64 at the first throw, 32 at the second, 16 at the third ... + 1 at the seventh; then the total payoff will be two times 64 plus four times 32 plus … arriving at 896. Dividing the total payoff by 128 we get exactly 7. It should be borne in mind, however, as mentioned in the text, that only 127 games have been

[19] Or, it is the same, the average payoff multiplied by $N$.

[20] If, as Feller assumed, the prize at the first throw is 1 instead of 2, the fair prize is $n/2$.

released, not 128; it is the problem of the last game. Obviously it is by no means certain, indeed it is very unlikely, that in the case of 127 coin tosses, these are distributed exactly with 64 on the first toss, 32 on the second…. And so on. But as the number of outcomes (and therefore of throws) increases, the law of large numbers tells us that the distributions will tend, in probability, to become ever closer to the progression in base 2, in the sense that the probability, that is the difference between the average revenue of the games and *n*, tends to zero (more precisely the difference is less than an ε small at will).

## *Appendix 2*

### *Buffon's experiment*

Georges-Louis Leclerc, Comte de Buffon (1777, English translation and commentary in Hey et al. 2010), the famous naturalist, wrote some essays on probability and dealt with the St. Petersburg paradox; he put forward all the hypotheses developed at that time: decreasing marginal utility[21], the non-existence of a bookmaker with infinite wealth, the impossibility of the game lasting indefinitely, the zeroing of odds greater than 1/10,000. But he did something that no other scholar had done: an experiment. He instructed a boy to flip a coin, observing how many times "heads" came out on the first toss, how many per second and so on, up to 2048 games.
Buffon's results:

| Number of launches | Empirical outcomes | Theoretical outcomes | Spread% |
|---|---|---|---|
| 1 | 1060 | 1024 | +3,52 |
| 2 | 494 | 512 | -3,52 |
| 3 | 232 | 256 | -9,32 |
| 4 | 137 | 128 | +7,03 |
| 5 | 56 | 64 | -12,5 |
| 6 | 29 | 32 | -9,38 |
| 7 | 25 | 16 | +56,25 |
| 8 | 8 | 8 | 0 |
| 9 | 6 | 4 | +100 |
| 10 | 0 | 2 | -100 |
| 11 | 0 | 1 | -100 |

[21] Buffon also proposed a function such as: y = x/(a+x), where a is the initial wealth and x the expected gain. The best known function is obviously the logarithmic Bernoulli function. Even before the publication of Bernoulli's work, Gabriel Cramer, in a letter to Bernoulli's uncle, had suggested the square root of x.

As one can see, the percentage differences tend to increase passing from the first launch to the following ones (except for the eighth launch where, by chance, there is coincidence).
It can be added that the games released in the first six launches (overall only seven games less than the theoretical ones) account for only 20% of the lower final payoff. It is the final tosses, and in particular the fact that Buffon did not went beyond the ninth toss, that determined the result of a 10.7% lower payoff than the theoretical one (9.82 against 11)[22].

*Vivian' test*

Let's now take a leap of more than a century (we are now in the computer age) and examine Vivian's simulation exercise (2013, pp. 357-358) with 268,435,456 games ($2^{28}$).

| Number of launches | Empirical outcomes | Theoretical outcomes | Spread% |
|---|---|---|---|
| 1 | 134.217.760 | 134.217.728 | 2,3841e-7 |
| 2 | 67.108.768 | 67.108.864 | -1,4305e-6 |
| 3 | 33.555.680 | 33.554.432 | 3,71933e-0,5 |
| 4 | 16.773.216 | 16.777.216 | -0,00023842 |
| 5 | 8.390.976 | 8.388.608 | 0,00028229 |
| 6 | 4.195.936 | 4.194.304 | 0,0003891 |
| 7 | 2.096.288 | 2.097.152 | -0,00041198 |
| 8 | 1.048.608 | 1.048.576 | 3,05176e-5 |
| 9 | 522.880 | 524.288 | -0,00268555 |
| 10 | 262.592 | 262.144 | 0,00170898 |
| 11 | 130.240 | 131.072 | -0,00634766 |
| 12 | 68.352 | 65.536 | 0,04296875 |
| 13 | 31.744 | 32.768 | -0,03125 |
| 14 | 16.192 | 16.384 | -0,01171875 |

[22] Buffon realized that the games were not 2048 (211) but 2047, like the theoretical outcomes. However, he did not attach importance to it (Paladini 2017).

| | | | |
|---|---|---|---|
| 15 | 8.320 | 8.192 | 0,015625 |
| 16 | 4.288 | 4.096 | 0,046875 |
| 17 | 2.080 | 2.048 | 0,015625 |
| 18 | 864 | 1.024 | -0,15625 |
| 19 | 256 | 512 | -0,5 |
| 20 | 160 | 265 | -0,375 |
| 21 | 64 | 128 | -0,5 |
| 22 | 32 | 64 | -0,5 |
| 23 | 64 | 32 | 1 |
| 24 | 64 | 16 | 3 |
| 25 | 0 | 8 | -1 |
| 26 | 32 | 4 | 7 |
| 27 | 0 | 2 | -1 |
| 28 | 0 | 1 | -1 |
| 29 | 0 | 1 | -1 |

The games released up to the 11th launch differ from the theoretical ones with percentages much lower than 1%, up to the 17th. The games released at the 18th launch are 864 empirical against 1024 theoretical ones, with a difference of 15.6%; the number of games released in subsequent launches will have larger and larger spreads. Until the 21st throw, the sum of the games released in the exercise differs from the theoretical one by only 64 games[23].

But, despite being less, the remaining empirical games are concentrated on the 24th throw (64 against 16 theoretical), and above all on the 26th throw (32 against 4 theoretical), so the average empirical payoff (34.04 against 28 theoretical) is 21.6% higher. Therefore the 192 empirical games[24], that come up between the 22nd and the 26th throw, are those that determine the result[25].

Thus in the case of Buffon's simulation the player, by agreeing to pay 11-9,82 monetary units per game, would have lost 2,417 monetary units; the dealer, in the case of Vivian's simulation, by agreeing to paying 6,04

[23] I.e. 268.435.264 empirical instead of 268.435.328 theoretical outcomes.

[24] That is a 7,15256e$^{-7}$ percentage.

[25] In Vivian's hypothesis, the last game comes out not on the first launch but on the 29th, thus resulting in a payoff of 30, and in any case creating a loss of over a billion monetary units.

(34,04–28) units per game[26], would have lost 1,626,718,863 currency units. Only by establishing that the monetary unit of the game is one millionth part of a euro (or dollar), the dealer could avoid a catastrophic loss.
In conclusion, 98% of the games in the case of Buffon ($2^{11}$) and 99.99999% in the case of Vivian (228), are distributed not very far from the theoretical distribution; most of the overall final difference is in fact determined by the games released respectively from the seventh launch onwards in the case of Buffon and from the 23th onwards in the case of Vivian. The results of these two simulations coincide with the 30,000 MC simulations illustrated in the text. The conclusions are therefore: I) as the number of games increases, most of them do not significantly differ from the theoretical distribution; but II) the overall average price depends crucially on the few games of the highest throws, often it is a single game. This fact reflects the law of large numbers.

*Appendix 3*

Two mathematicians (Klyve - Lauren 2011), who know nothing about the St. Petersburg paradox or Feller's limit theorem, having read, in a text on the history of mathematics, of Buffon's experiment, decide to repeat it, finding a payoff average equal to 18.32, much higher than Buffon's 9.82. Getting a suspicion, they tried the following number of games:

| Number of launches | Empirical outcomes | Theoretical outcomes | Spread% |
|---|---|---|---|
| 1000 | 7.80 | 64 | 17,5 |
| 2048 | 18.70 | 11 | 70 |
| 5000 | 15.49 | 12,29 | 26 |
| 10000 | 14.93 | 13,29 | 12,3 |
| 50000 | 15.60 | 15,61 | 0 |
| 100000 | 21.15 | 16,61 | 27,3 |
| 500000 | 19.46 | 18,93 | 2,8 |
| 1000000 | 25.69 | 19,93 | 28,9 |

They thus discover that there is a tendency for the average payoff to increase with the number of games (from 7.80 in the case of a thousand to 25.69 in the case of one million), and that this was a well-known fact, at

[26] 11 and 28 are the *n* of the two games, which according to Feller constitute the “fair” price of the game. The figures are halved if, as was usual in Buffon's time, we start from one monetary unit at the first throw instead of two.

least since Feller's time[27]. But in addition to the fact that the average payoff grows as the number of games increases, it is also noted that the differences with respect to the theoretical payoff are very large. For example, the payoff obtained in the case of 100,000 games (21.15) is higher (by 27.3%) than the theoretical one, and also higher than that of 500,000 games (19.46), which is instead rather close to the theoretical one (18.93).

The fact of having obtained, in the case of 2,048 games, a payoff equal to 18.70, therefore double than that of Buffon, pushes the two authors to carry out a MC simulation with a million interactions. They find a distribution not only strongly asymmetrical, progressively decreasing, but with numerous teeth and also empty intervals.

The result obtained amazes them; the surprise emerges from their words: "We were taken aback at the resulting bizarre distribution. When we showed the graph around, several colleagues expressed surprise. One questioned our random number generator! What sense can we make of this picture? We immediately see that the most likely values are indeed those near $\log_2(2048)$, as expected. However, the rest of the graph is surprising. Its comb-like, fractaline quality demands explanation". Indeed, if they had been aware of Rodriguez's (2005) work[28], at least the strong distribution asymmetry should not have surprised them. Furthermore, repeating the 2,048 games a million times involves simulating over two billion games; now the fact that the average payoff increases as the number of games increases implies that, if the overall average payoff were not to deviate too much from the theoretical one of two billion, that is $2^{30.9}$, it would be close to 30.9. Since the modal value of the distribution is around 11, it is evident that the payoff distribution must be distributed in a highly asymmetrical way, reaching some 330 payoffs, that is thirty times greater than the modal value.

Anyway, the fractal structure (even if imperfectly so) of the distribution constitutes the novelty that emerges from the MC simulation of the 2,048 games. Klyve and Lauren's explanation focuses on the teeth of the distribution. They note, correctly, that the average payoff of a Buffon-style experiment is determined, above all, by the highest value reached by a given game. The implicit hypothesis is that the games are distributed, even if not perfectly, according to the theoretical distribution. For example, looking at a tooth at the average payoff of 42, they notice that it can be the result of a game that comes out on the sixteenth throw. Now $2^{16}/2048=32$. If the other 2047 games come out at various throws according to the theoretical distribution (even if not perfectly), so to have an average payoff around 10, the tooth at 42 finds a logical explanation. But we can

[27] But Buffon had already had the intuition of this fact, even if he hadn't developed this intuition.

[28] However Rodriguez was on the wrong track, trying to find a rule concerning the growth of the mean payoff of the *N* games.

have the same result with two games that come out at the fifteenth launch, or three games that come out at the fourteenth launch; the probability of these events is the same. This determines a thickening of games with payoff 42.
Repeating with $2^{17}$ and $2^{18}$ we identify two other teeth of the distribution. It is interesting to note that, repeating the MC simulation of Buffon's games for a million times, Olivero (2016) finds a similar distribution, although not identical, to that of the two authors, with teeth located at the same payoff levels.
The other peculiar aspect is that of the presence of empty spaces, that is, ranges of values in which no payoff has been released; the same phenomenon is found in Olivero's simulation, even if the coincidence is approximate. Empty ranges are found between 180 and 200, between 220 and 260, between 305 and 325. Now that the phenomenon occurs in the case of a very limited number of games it is easily explained. For example, with two games it is not possible to obtain an overall payoff of 14, or 22 or 26. Even with three games, it is not possible to have a payoff of 46 (32+14), or 54 (32+22) or 58 (32+26). But here we have empty intervals of whole tens or more. Furthermore, MC simulations performed by Olivero (2016) show that even with games of size $2^{13}$, $2^{14}$, $2^{15}$, with some small differences, the gaps occur in corresponding intervals.
The explanation is due to the fact that there are distributions which, although possible, are however highly improbable, because they would imply a too marked deviation from the theoretical distribution[29]; if the games are $2^n$, it is very unlikely that those that come out in the first launches differ significantly from $2^n/2$ and $2^n/4$, as we have now seen by examining the two distributions of Buffon and Vivian's games. Only a small percentage of games (the last ones) deviate more from the theoretical distribution.

## *Bibliography*

[29] For example, in the case of Buffon's 2,048 games, a payoff equal to 316 can be obtained with the following distribution, from the first roll to the seventeenth:

898, 301, 251, 199, 120, 85, 55, 35, 30, 25, 18, 11, 8, 6, 4, 1, 1

As one could see, the games that come out at the first launches are considerably far from the theoretical distribution (the first launch differs by 12.3% from the theoretical one and the second by 41.2%, when in Buffon's case the difference was only 3,5% in both launches, and in those of the simulations carried out by the writer, the differences were around 1% -2%). Continuing to carry out MC simulations, sooner or later a distribution of this type may occur, but it is indeed a very rare event.

**Bernoulli Daniel** (1738), Specimen theoriae novae de mensura sortis, *Comentarii Academiae Scientarium Imperialis Petropolitanae V*, 175-192; English translation 1954 by L. Sommer with footnotes by Karl Menger, Exposition of a new theory on the measurement of risk, *Econometrica* 22(1), pp. 23-36.

**Buffon Georges-Louis Leclerc** (1777), *Essai d'arithmétique morale*, Supplément a l'histoire naturelle IV, 1777.

**Chow Soo Hong – Epstein Larry** (1988), The Law of Large Numbers and the Attractiveness of compound Gambles, *Journal of Risk and Uncertainty*, I, pp. 125-132.

**Feller William** (1945), Note on the Law of Large Numbers and 'Fair' Games." *Annals of Mathematical Statistics.* 16, 301–304.

**Feller William** (1950), *An Introduction to Probability Theory and Its Applications*, vol I, John Wiley, New York.

**Foley Eden** (2010), A New Look at the St. Petersburg Paradox*, Department of Mathematics and Statistics, Northern Kentucky University, AMS.*

**Hey John-Neugebauer Tibor-Pasca Carmen** (2010), introduzione alla traduzione in inglese degli *Essai d'arithmétique morale* di Buffon, LSF Research Working Series n. 10-06.

**Klive, Dominic - Lauren Anne** (2011). An empirical approach to the St Petersburg paradox. *Mathematical Association of America*, 42(4):260-264.

**Marshall, Alfred.** *Principles of Economics* London: Macmillan and Co., Ltd., 1890.

**von Neumann John – Morgenstern Oscar** (1944), *Theory of Games and Economic Behavior*, Princeton University Press.

**Olivero Alexander** (2016), *A Computational and Theoretical Exploration of the St. Petersburg Paradox*, Undergraduate Honors Thesis Collection, Paper 317, Butler University.

**Paladini Ruggero** (2017), Il paradosso di S. Pietroburgo, una rassegna, *Public Finance Research Paper n. 2,* Istituto di Economia e Finanza, Sezione del DIGEF, Sapienza Università di Roma.

**Pianca Paolo** (2007) *The St. Petersburg Paradox: Historical Exposition, an Application to Growth Stocks and Some Simulation Approaches,* Quaderno di Didattica n. 24, Department of Applied Mathematics, University of Venice.

**Rodriguez Ricardo** (2006), Finite sequences of St. Petersburg games: inferences from a simulation study, *Journal of Statistical Computation and Simulation* vol. 76 n. 10 pp. 925-933.

**Samuelson Paul** (1963), Risk and Uncertainty: A Fallacy of Large Numbers, *Scientia* 6th series, April-May, pp.1-6.

**Samuelson Paul** (1977), St. Petersburg Paradoxes: Defanged, Dissected and Historically Described." *Journal of Economic Literature.* 15(1), pp. 24-55, 1977.

**Vivian Robert** (2013), *Ending the mith of the St. Petersburg paradox,* MPRA Paper n. 50515, University of Witwatersrand.

http://www.digef.uniroma1.it/pubblicazioni

E-PFRP N. 44

2020