FURTHER EXPLORATIONS OF THE ST. PETERSBURG PARADOX WITH MAPLE

Dr. Joel C. Fowler
Mathematics Department, Southern Polytechnic State University
1100 South Marietta Parkway
Marietta, GA 30060-2896
jfowler@spsu.edu

Abstract: The Saint Petersburg Paradox is a well known example of how human behavior can diverge from mathematics. In the usual game a fair coin is tossed repeatedly until the first head occurs. The payoff is $1 if the game terminates in one flip, $2 if in two flips, $4 if in three, and so on, doubling the payoff for each additional flip required to obtain the first head. The expec
are willing to pa
explore, through
likelihood of makin
experiments led
to achieve a reas
the problem, pla
questions. The m
that the computa
games is needed
human sense that
take an unaccep
undergraduate special project.

For the standard St. Petersburg game, with a payoff of $\$2^{n-1}$ when the first head occurs on the n-th flip, the expected payoff is:

$$1\left(\frac{1}{2}\right)+2\left(\frac{1}{4}\right)+4\left(\frac{1}{8}\right)+8\left(\frac{1}{16}\right)+\ldots=\sum_{n=0}^{\infty} 2^{n-1}\left(\frac{1}{2^n}\right)=\sum_{n=0}^{\infty}\frac{1}{2}\ ,$$

which is a divergent series. Hence the expected value is considered to be infinite. The paradox is that few people are willing to pay more than a few dollars to play this game, even after the expected value computation is presented. The question is why the human psychology of the game is so far from the mathematics.

Many explanations that have been offered to explain the paradox. Most are based on expected value computations in which the payoffs are adjusted, in some fashion, to reflect what might be a human perception of the values of the payoffs. Here we take a different approach and consider just the risk of loss in playing the game many times. That is, rather than work with expected value, we focus on estimating the likelihood of a profit after N plays of the game, for a fixed cost to play per game, M.

The Maple packages we shall use are:
> with(RandomTools); with(plots); with(Statistics);

Using the built in Random Variable features of Maple we can simulate play easily.

> X := RandomVariable(Geometric(.5));
> Sample(2^X,10);
[1., 4., 8., 2., 4., 2., 1., 2., 1., 2.]

We can also graph the behavior of the average payoff over time for small values of N with:

> Games := 1000; Values := Sample(2^X, Games): RunningAverage := i -> add((Values[j]), j=1..i)/i: pointplot({seq([n,RunningAverage(n)], n=1..Games)});

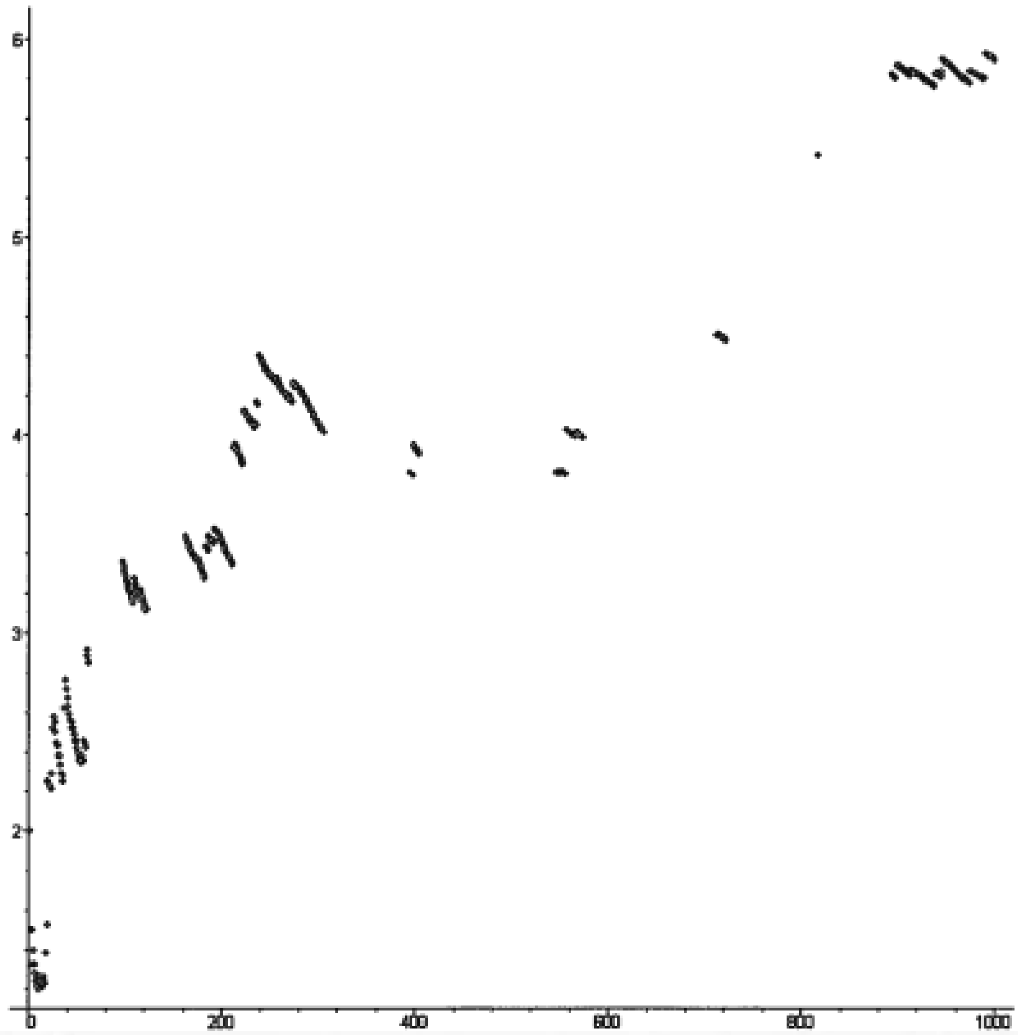

However, the speed of Maple limits this approach to values of N that are too small to reveal the nature of how the average payoffs grow as the number of games increases. We need a different approach. But the above graph is still revealing. Note that the behavior of the average over time is primarily periods of slow decline, punctuated with occasional, sometimes substantial, jumps upward. The periods of slow decline are intervals of repeated low payoffs, particularly 1's. In fact, roughly half of the N games played will result in a payoff of \$1. This leads to the idea of basing a simulation of N games on the number of games that result in a payoff of just \$1.

Note that the number of games, out of N, with a payoff of \$1 is a binomial random variable with N trials and a p-value of .5. And the remaining non-\$1 payoff games are simply another set of St. Petersburg games, but with doubled payoffs (i.e. \$2, \$4, \$8, ...). So if we let X(N) be the random variable for total payoff for N plays of the St. Petersburg game, and B(N) a binomial random variable for N trials and p = .5, we have:

$$X(N) = B(N) + 2\,X(N - B(N))$$

We may use this relationship recursively to produce a very efficient simulation that can be used for a very large numbers of games. Since the mean of B(N) is N/2, with a distribution that is approximately normal with a relatively small standard deviation, each iteration cuts the size of N roughly in half. In Maple the code for the total and average payoffs are:

```
> TotalPayoffSimul := proc(N::integer)
   description "simulate the total payoff for N plays";
     Payoff := 0; TempN := N;  Digits := ceil( evalf(log10(N)) )+2;
     for Counter from 0 while TempN > 0 do
      if TempN > 1000 then RV := RandomVariable( Normal(TempN/2,sqrt(TempN)/2) )
       else RV := RandomVariable( Binomial(TempN,1/2) ) end if;
      Value := round( Sample( RV, 1)[1] );
      Payoff := Payoff + Value*(2^Counter);
      TempN := TempN - Value;
   end do; Payoff;
  end proc;

> AvPayoffSimul := n -> evalf(TotalPayoffSimul(n)/n, 4);
```

Investigating the running time for this procedure, as well as typical payoffs for some very large values of N yields:

```
> st:= time(): [seq(AvPayoffSimul(10^n),n=15..25)]; time() - st;
  [30.45, 26.60, 30.04, 29.80, 35.15, 34.11, 37.59, 36.37, 38.86, 41.92, 55.55]
                          2.985
```

Thus we see, experimentally, that the procedure is quite efficient. It took less than 3 seconds to simulate the play of 10^{15}, 10^{16}, . . . , through 10^{25} games. Additionally, it appears that the growth rate of the average payoff is quite slow. An average payoff of around \$50 for 10^{25} games is extraordinarily low. To place the size of this number in perspective, 10^{25} games, played at 1,000,000 games per second, would take approximately 3×10^{11} years. This is well beyond current estimates of the age of the universe, which are around 1.4×10^{10} years.

Because the procedure is so efficient, we may gather quite bit of data for particular large values of N and examine the experimental distribution of average payoffs. For $N = 10^{20}$, we gather data for 500 experimental simulations via the command:

```
> N := 10^20 ; Trials := 500 ;  C := [ seq(AvPayoffSimul(N), n=1..Trials ] ;
```

which generates the vector C in around 12 minutes. The histogram for C is:

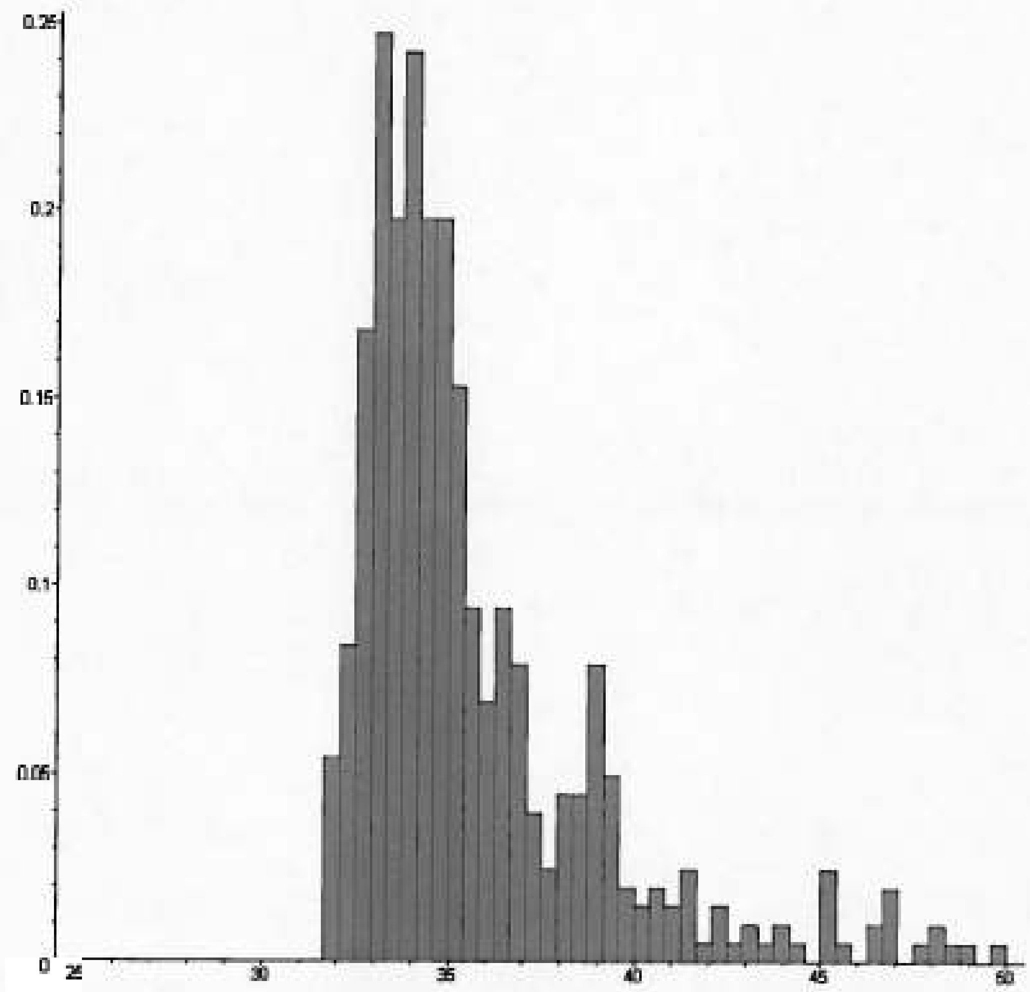

Note that samples of this type can be used to estimate the actual probabilities of exceeding a particular average payoff using confidence interval methods. With this sample of size 500 and a confidence level of 99%, the error in such an estimate is approximately .05 . For example, from the data in vector C, used to generate the above histogram, we find that an estimate of the probability of a payoff of at least \$40 for 10^{20} plays is only .118 .

Going further, we can generate average payoff data for a variety of values of N and plot the growth trend. The command:

```
> f := k -> modp(k,50)+21;  Payoff := N -> AvPayoffSimul(2^f(N)); Trials := 500; X :=
Vector(Trials, f) ; Y := Vector(Trials, Payoff);  [seq([ 2^X[n], Y[n] ]  , n=1..Trials)];
```

loads the length 500 vector X with the values 21, 22, 23, . . . , 70, each assumed 10 times, while the length 500 vector Y is loaded with corresponding simulated average payoffs for playing $N = 2^{21}, 2^{22}, 2^{23}, \ldots, 2^{70}$ games, again with each value assumed 10 times. Thus a plot of Y against X is a plot of experimental data for playing N games that relates the average payoff for N games to $\log_2(N)$. This command took approximately 7 minutes to run. A plot of the output is:

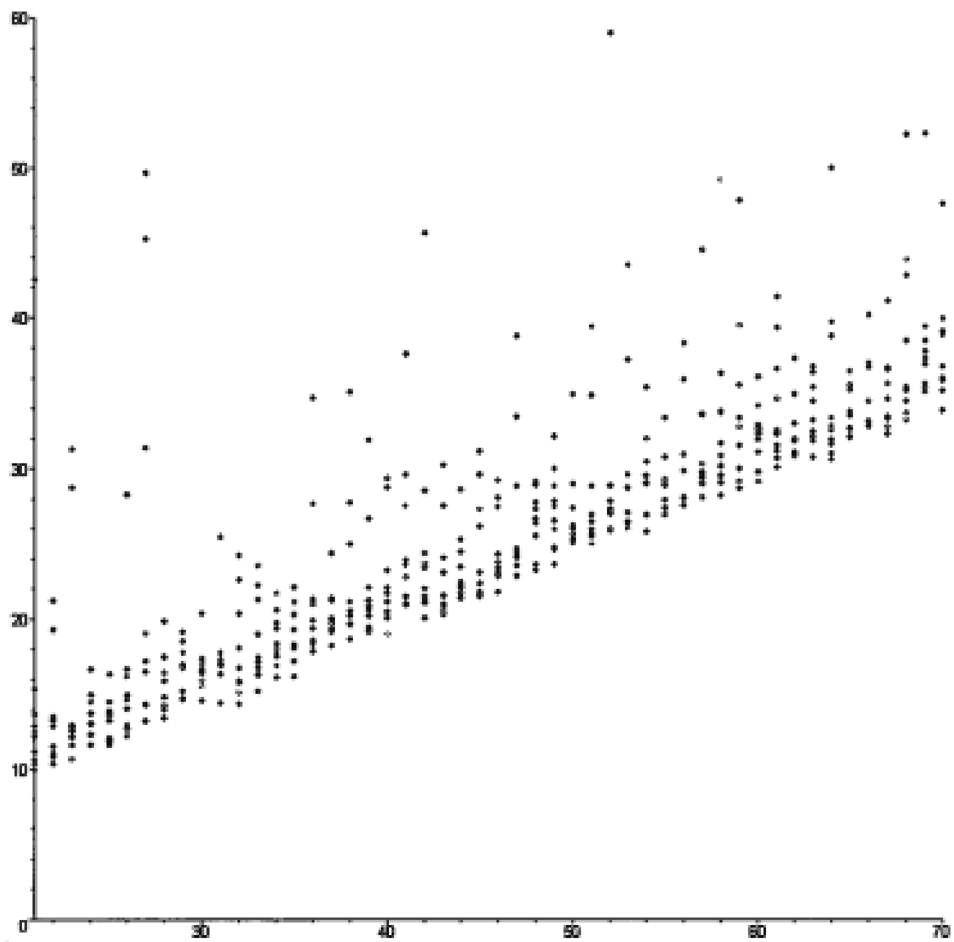

Trimming the top 3% of outliers, since we are interested in likely payoffs, and performing a linear regression yields the experimental relationship:

$$\text{AvePayoff}(N) \sim 2.7163 + .5056 \log_2(N) .$$

Solving for N and rounding yields the general rule of thumb:

$$N \approx 4^{\text{AvePayoff} - 5}$$

That the number of games to make a particular average payoff likely is a base 4 exponential, in that average payoff, is not unexpected. Returning to the recursive relationship on which we have based the simulations we have:

$$X(N) = B(N) + 2\,X(N - B(N)) \;,$$

where X(N) is the total payoff from playing N games and B(N) is a binomial random variable with N trials and p = .5 . Proceeding informally, if we replace B(N) by its mean, N / 2 , then this becomes:

$$X(N) \approx N/2 \;+\; 2\,X(N/2) \;.$$

Dividing by N then yields:

$$\frac{X(N)}{N} \approx 1/2 \;+\; \frac{X(N/2)}{N/2} \;.$$

Noting that the first and last terms are simply average payoffs for N and N / 2 games we then have:

$$\text{AvePayoff}(N) \approx 1/2 \;+\; \text{AvePayoff}(N/2)$$

So, in very loose terms, the likely average payoff gains 1/2 as the number of games doubles. Thus, to achieve an average payoff of some number V would require 2V doublings of the number of games, i.e. $2^{2V} = 4^V$.

While the replacement of B(N) by its mean, N / 2, is overly casual, mathematically, it is worth noting that B(N) may be expected to lie relatively close to its mean. In fact, B(N) should be no farther than $2\sqrt{N}$ from N / 2, with probability .99996, since this is 4 standard deviations from the mean for in an approximately normal distribution.

To further check this exponential relationship, we can produce data using our simulation proceudre for $N = 4^{95}$ plays of the St. Petersburg game, just as we did earlier for 10^{20} games. This should correspond, via our exponential relationship, to an average likely payoff of around \$100.

The command:

```
> N := 4^95 ; Trials := 500 ; E := [ seq( AvPayoffSimul(N) , n=1..Trials ) ];
```

produces the data for 500 simulations in vector E. This command required 68 minutes to run. A histogram for the data in vector E is:

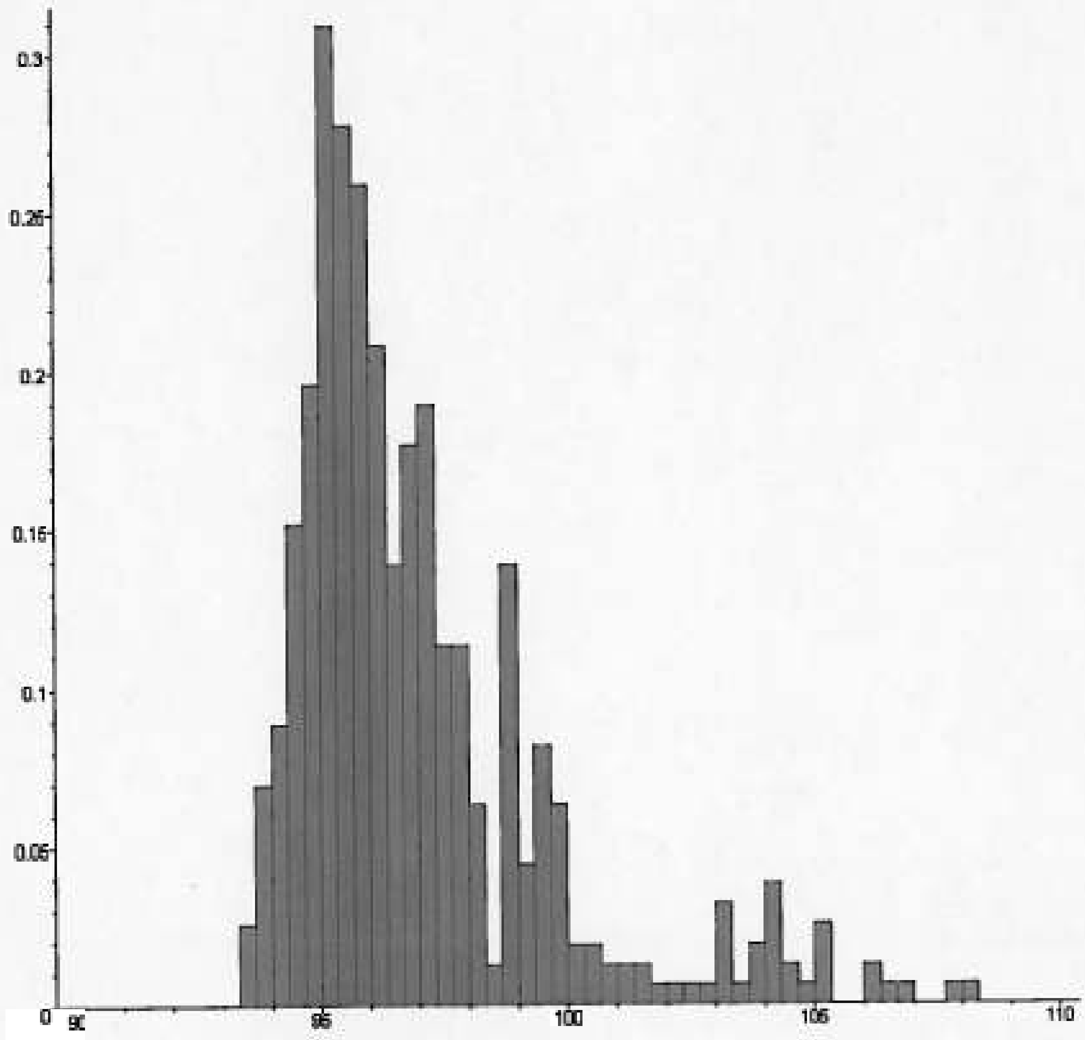

Using this data we may estimate (again with 99% confidence that the error is no more than .05) the likelihood of achieving an average payoff of at least \$95 at .836; and the likelihood of an average payoff of at least \$100 at .138 . In other words, if one were to play $N = 4^{95}$ St. Petersburg games, paying \$95 per play, the chances of losing money overall would be only around 16%. If, however, the cost were raised by just \$5 to \$100 per play, then the likelihood of losing money overall jumps to 86%. That the bulk of the distribution lies within such a tight range is rather remarkable. In any case, this example supports our experimental relationship that the number of games needed to make a particular average payoff likely is a base 4 exponential.

This also sheds light on the psychology of the St. Petersburg paradox. Purely from an expected value standpoint, one should happily pay $100 per play of the game (and, in fact, any finite amount of money). But the reaction of most people to paying $100 per play (since, in the long run, they'll come out ahead) is concern over all of the losses they'll experience while waiting for the big payoffs to come in and wipe out those losses. That is, they won't pay $100 per play (or anywhere near that) because they feel they're too likely to lose money. In fact, their intuition is correct. From the above, at $100 per play, one would need to play much more than $N = 4^{95}$ games to have a reasonable chance of not losing money. But this is a cosmically large number of games. Played at a rate of 1,000,000 games per second, 4^{95} games would require around 5×10^{43} years to play. This is well beyond some estimates of how long matter will exist in our universe, which are approximately 10^{30} years. To put it another way, paying $100 per play would be a tremendous mistake, since there is virtually no possibility of ever playing as many games as would be needed to have a reasonable chance of at least breaking even.

On a smaller scale, the typical amounts people are actually willing to pay per play are usually no more than $8. To determine what this payoff translates into, as far as the number of games needed for likely profit is concerned, we can again use the simulation procedure. We find experimentally that, to have a roughly even chance of not losing money, when paying $8 per game, the number of games needed is on the order of 10,000. It's reasonable to assume that virtually no one would have the patience to flip a coin an average of 20,000 times in order to complete that many games. So by that measure, the human estimate of only a few dollars, as a reasonable amount to pay per play, is actually correct. That is, the amount a person is willing to pay is intuitively linked to their sense of how long it would probably take to earn a profit, and that intuition appears to be somewhat accurate. An interesting question to pursue would be to see how an individual's answer might change, if the play of the game is moved from a tabletop to a high speed computer simulation in they could be assured of playing hundreds of thousands of games in seconds.

In conclusion, from the standpoint of risk of loss, there actually is no St. Petersburg paradox. Human intuition turns out to be relatively close to the mathematics. In fact, one might say that, from this perspective, the only paradox relating to the St. Petersburg game is why so much emphasis is placed on expected value as the best means of measuring how valuable the game payoffs are.