Paradox Lost: The St. Petersburg Gamble Revisited

Barry Sopher

Rutgers University-New Brunswick

September 1, 2023

Abstract:

Daniel Bernoulli's famous account of a gamble whose expected value is a divergent sum, and that no one is willing to pay more than a modest sum to play, a finding referred to as the St. Petersburg Paradox, is a primary source for the idea of maximizing the expected (concave) utility of wealth, and hence of risk aversion, as a tool for analyzing decision making under uncertainty. In this paper I show that appropriate attention to the geometric distribution of the number of successes before a failure in the gamble (the number of tails flipped before a head is flipped) implies that the monetary winnings one should expect from a single play of the St. Petersburg gamble is $2, and not an arbitrarily large payoff. Interestingly, the argument and calculation that Bernoulli made to show that, using a (base 2) logarithmic function of wealth, the gamble should be seen to be worth 2 in monetary terms, is exactly the calculation needed (but for a completely different reason) to show that the expected number of successes in playing the gamble is 1, and thus that the "expected" value of the gamble is 2. That is, one should expect to have one success before having a failure, and thus one should expect to get the prize, $2, associated with one success. .

More generally, I show that the expected payoff Bernoulli identified for the monetary gamble is applicable only in the sense that the average payoff as the number of plays of the gamble increases also increases. Technically speaking, the limit defined by the expected value does not exist. It is not infinite. Rather, it is simply divergent, meaning, in our context, that for any finite number one might choose, there is some number of plays of the gamble that one can expect to yield an average payoff higher than that finite number.

I argue that the St. Petersburg Gamble is, thus, not really a useful foundation upon which to build a theory of risk bearing. This is not calling into question anything about the economic theory of risk bearing as it now exists, as that theory does not, in fact, rely upon Bernoulli's account being correct. But it does suggest that a deeper appreciation for the physical nature of dynamic situations such as the St. Petersburg Gamble should precede conjectures about what sort of preferences might account for the behavior one observes.

This is a revised version of a paper originally prepared for presentation at the Economic Science Association Annual North American Meeting in Tucson, AZ on October 30, 2021 and the International Meeting of the Economic Science Association in Cambridge, MA, on June 15, 2022. I acknowledge funding from the Laboratory for Economic Behavior, Institutions and Design at Rutgers for this research. I have benefited from discussions with Yunqing Zhu, who also programmed the simulations reported on here. I especially want to thank Xiye Yang, who suggested consideration of waiting time as appropriate to understanding the St. Petersburg gamble. I thank Jacqueline Gelman, Mark Ryan, James Hadley, Jiayi Lu and George Nahkla for their input at various stages of this project.

1. Preliminaries

In 1738, Daniel Bernoulli proposed the first account of what later came to be called risk aversion. The most useful part of his account, to modern eyes, concerns the advantages of diversifying risk, and the basis on which to determine whether to buy insurance. But the more famous part of his account concerns a certain game of chance. In the present paper I will be strictly concerned with this game of chance. In the game, a coin is flipped repeatedly until it comes up "Heads." If there are R Tails before a Head comes up, then the gambler wins a prize of 2^R. The game proposed has an "infinite" expected value (more on this below), but both intuition and a certain amount of empirical evidence (if largely hypothetical) shows that a typical person is unwilling to pay more than a modest sum to play the game. The prevailing idea at that time, due to Fermat and Pascal, was that the mathematical expectation of wealth should always determine value, thus the paradox. Bernoulli's proposed resolution of this paradox was that maximization of what we would now call expected utility of wealth, with a concave function of wealth, can rationalize the modest sums that people are typically, if only hypothetically, prepared to pay to play this game.

In this paper I demonstrate by Monte Carlo methods and by consideration of the sequential nature of the play of the gamble, that there is in fact no way to reliably achieve a payoff of very great magnitude from a single play of the proposed game of chance. By considering finite-period truncated versions of the game that Bernoulli considered, I show that one must play an ensemble of many instances of the game, the number of which is systematically related to the length, T, of the truncated game in question, in order to have access to the theoretical expected value of the gamble. Ultimately, allowing the length of the game, T, to approach infinity, one must play an ensemble of size R approaching infinitely many instances of the game in order to approach the infinite expected value Bernoulli identified. I am able to make precise what it means for an ensemble to be “large enough,” and “not large enough,” to obtain access to the expected value of a particular truncated game. As a general rule, the “expected value” of the gambles considered here, due to the non-Gaussian nature of the distributions involved, is not what one should expect to earn.

Concerning the sequential nature of the game, I note that, for the indefinitely repeated game, the distribution of the number of “successes” (tails) until one has a “failure” (heads) is geometric, so it is clear that what one should expect to earn from a single play of this game is quite finite, indeed: The mean number of successes is one, and thus one really should “expect” to earn $2 from playing the gamble. If one considers playing an ensemble of R plays of the gamble, by virtue of the fact that a sum of geometric random variables is Negative Binomial with parameters R and p=1/2 (the number of failures and the probability of success, respectively), then one can expect to have R successes before one has accumulated R failures. A straightforward accounting exercise in section 7 of the paper shows that, even with the indefinitely repeated version of the game, the number of plays, R, strictly limits what one can expect to earn from playing the game.

The limited experimental evidence that we have on the St. Petersburg Gamble, in fact, reflects a rather sensible view on the part of respondents as to what it is possible to win from playing the game, with no needed recourse to arguments about the utility of wealth or risk aversion. In fact, there is more evidence of apparent risk-loving behavior than of risk averse behavior in the experiments, to the extent that people are willing to pay more than $2 to play the game. I note that my findings on the St. Petersburg gamble do not in any way invalidate anything in the theory of risk bearing. Rather, they simply show that the St. Petersburg Gamble does not provide a useful basis for establishing a firm foundation for the notion of risk preferences, and that, in fact, it is not needed to make sense of people's view of the game. A theory of risk bearing may be, and has been, derived from consideration of simple gambles. Consideration of more complex stochastic processes, such as appear in the St. Petersburg Gamble, may pose other challenges for us, but disposing of risk preferences would be missing the point.

2. The Details

The St. Petersburg Gamble is the name I will use to refer to a game of chance, originally proposed by Nicholas Bernoulli to the mathematician Montmort in 1713, and later discussed by his cousin Daniel Bernoulli (1954) in 1738, which is a foundation myth, of sorts, for the theory of decision making under uncertainty. The game goes as follows (using dollars rather than ducats as the currency of account in the gamble): A offers an initial prize of $1 to B. B is allowed to flip a fair coin. If the coin comes up Heads, then B can take the $1, and the game is over. If the coin comes up Tails, then A doubles the prize to $2, and B is allowed to flip the coin again. If the coin comes up Heads, then B take the $2, and the game is over. If the coin come up Tails, then A doubles the prize to $4, and B is allowed to flip the coin once again. The game continues in this way, with the prize being doubled by A for each time the coin comes up Tails, only ending at the point when the coin finally comes up Heads. At that point, B takes the current prize, and the game is over. It is easy to see, as Bernoulli observed, that the expected value of this game is a divergent sum. That is, the probability of the game ending in the i-th round is $1/2^i$ and the prize to be taken if the game ends in the i-th round is 2^{i-1} so the expected value is (1/2)*1 + (1/4)*2 + (1/8)*4 + (1/16)*8 +…+ ½ + ½ + ½… , which is a divergent sum. I am careful to say the sum is divergent, rather than to say the sum is infinite, as much of the discussion (and confusion) about the St. Petersburg gamble seems to come from thinking of "infinity" as something like a number, which of course it is not.

Bernoulli argued in various ways that no sensible person should actually be willing to pay very much to play this gamble. First, in the general development of his theory, prior to consideration of this specific gamble, he asserted that utility losses must equal utility gains in a fair gamble, and arrives somewhat torturously at the conclusion that utility must be logarithmic (Bernoulli assumed this would be a base 2 logarithm). This is in marked contrast to the modern-day idea that one might, in principal, be able to elicit a utility function from an individual (say, by a reference lottery technique), and there is no presumption that such a function must be logarithmic, or even concave. In any event, a logarithmic utility function implies that prizes above a certain level (about what one would get with a string of 15 Tails), when combined with the probability of getting to that point, would add essentially nothing to the expected sum. On consideration of the specific gamble, he suggested that ever seeing more than 20 or 30 Tails in a row, say, is so unlikely that one can ignore any thing after, say, the 24^{th} term in the progression, giving a value of 12 or so. Finally, he argued that there is, essentially, declining marginal utility of wealth, apart from any considerations of risk as such. As each one of these explanations is a bit different, it is an odd performance, in many ways. Except in arguing that something other than just wealth be used in making calculations of the value of the gamble, and that the appropriate thing to do, one way or the other, is to essentially discount larger and larger possible values of wealth, the different explanations are not entirely consistent. That he arrived at several different "solutions," each implying that a "modest" value should be given to the gamble in question, shows that he was at pains to actually understand what was going on: he clearly felt that he knew what the answer should be (i.e., something “small”), and went through some intellectual acrobatics to try to show this.

What I will show in what follows is that, actually, the theoretical expected value of the gamble is not the appropriate focus, simply because the gambler is not in a position, in a single play of this game, to get anything like the theoretical expected value of the gamble. More to the point, a thousand gamblers, each playing the game once, could expect to have, as a median outcome, $1, and 97% of them could expect to have $16, at best, and generally less than that. The remaining 3% would get something more than $16, but it is very hard to be precise about what exactly they might get. Overall, we can say that the mean outcome for 1000 gamblers playing the game would be about $6 (for reasons that will become clear in due course), but the mean is a very unstable and rather unpredictable quantity for this game, for any finite sample. More to the point, the mean of 1000 gamblers, each playing the game once, is not a very good indicator of what an individual player could expect to get. As this observation is novel and probably somewhat surprising to most readers, I ask the reader’s patience as I spell out below what is going on.

3. The St. Petersburg Gamble as a Stochastic Process

The St. Petersburg Gamble (SPG) is generally viewed as a gamble with a probability mass function and is treated as if one is simply making a draw from a distribution, much as one might draw from an urn. The fact that it is a stochastic process that plays out over time, dynamically, is not much appreciated, as we shall see. One way to view the SPG is as a Markov process with absorbing states. An absorbing state is reached by flipping a Head at any point in the game. Once you flip a Head, the game is over. This can be contrasted with another possible game with binomial transition probabilities (let's call it the Repeated St. Petersburg Gamble). In the Repeated St. Petersburg Gamble, you start with a $1 possible prize. You flip a coin, and if it comes up Heads, you take the $1 prize, but if it comes up Tails, the prize is doubled to $2 and you flip the coin again. Having flipped Tails once and having the prize doubled to $2, if you flip a Head on the next try, you take the $2 prize. In either case, though, you have the option to "start over" and play the game again (with an additional payment to play, of course). As long as you continue to get Tails, the prize keeps being doubled, but as soon as you get Heads, after collecting the current prize, the prize goes back to being $1 and the process begins again. You never have to leave the game, but you have to pay once more to continue. You can leave the game at any time (you can "pay as you go," anteing up after each flip of Heads). Some recent work has pointed out that the St. Petersburg Gamble is, in fact, non-ergodic (Peters (2011), Amadou and Peters (2018), and Peters and Gel-Mann (2015)), which is indeed the root explanation for the regularities that we find. The Repeated St. Petersburg Gamble is, in fact, ergodic if played infinitely often, as there is a single recurrent class of states (it is possible to reach every state from every other state in finite time, and thus, over time, every state will be visited infinitely often). Playing the SPG R times in sequence is, in fact, equivalent to simply deciding to play the SPG R times "up front," provided, of course, the games are played independently. This shows, at a very essential level, what the non-ergodicity of the SPG leads to: The fact that "Head" is an absorbing state means that states involving a larger number of sequential "Tails" than has been reached when Head occurs cannot be reached in a single play of the game. Playing the game repeatedly makes it possible, eventually, to exceed the largest number of successes (tails) that one has so far achieved, though it may take a very long time.

4. Empirical Evidence from Experiments on the St. Petersburg Gamble

Hayden and Platt (2009), in one of only two incentivized experiments involving the St. Petersburg Gamble that I have been able to uncover, report that the median choice was $1.50, and the modal choices were $1 (about 30% of subjects) and $2 (about 20% of subjects) for an amount to pay to play the gamble. Roughly 15% were willing to pay less than $1, and roughly 25% were equally divided between being willing to pay between $2 and $8, or paying more than $8 (more detail than this is not reported in the paper, including if, and how much, subjects may have earned from playing the game). Subjects similarly were only willing to pay $1.50 or $2 to play truncated versions of the gamble, with the game lasting 5, 10 or 15 periods, at most. Cox, et al. (2009) found that when offered the chance to play truncated versions of the gamble (varying from 1 period to 10 periods), subjects were increasingly unlikely to be willing to pay as much as the expected value of a gamble as the number of periods increased, ranging from about 90% for the 1 period game, to 10% for the 10 period game. Note that the median outcome for a truncated game is $1, just as for the full SPG, so this evidence is consistent with subjects being generally unwilling be pay more than the median value of the game. In what follows, though, I will show more precisely what the value of the SPG is, and contrast it with what the value of the SPG has been argued to be in the past.

5. The Truncated St. Petersburg Gamble.

As I will make extensive use of the truncated version of the SPG, it is worth making a few observations. First, the traditional focus in writing on the St. Petersburg Paradox has been on the infinite expected value of the gamble. But the fact is, one is going to end up with something finite, so a good deal of the wheel-spinning that has been done, concerning unbounded utility and so forth, may be off the mark. (See Samuelson (1977) for an overview of this work.) Aumann's (1977) account of the need for utility to be unbounded is rather compelling: if there is an outcome y with infinite utility, then you would be willing to accept a gamble with a tiny probability of y and a complementary probability of death to any other outcome x with finite utility. But if we could show that the nature of the truncated game, in which the expected value is finite and calculable, is a sufficient to guide our thinking about the infinitely (or indefinitely) lived version of the gamble, as far as what a rational, or at least sensible, person would be willing to pay, then perhaps we would have cut the Gordian knot. This, in fact, is what I will do: specifically, we will show that for a single play of the SPG, the distribution of outcomes is essentially the same as for a finite length, truncated version of the game. Only by, essentially, being prepared to play (and pay for) additional plays of the SPG would a player be able to expect an outcome like the expected value of the game, and the number of additional plays that are required to achieve this expected value grows exponentially in T, the length of the game.

To be specific, the T-period truncated version of the SPG (T-SPG) is identical to the usual SPG, except that there is a maximum number of successes ("Tails") = T-1, that a player can have. With T-1 consecutive successes, the stakes double to become 2^{T-1}. Instead of flipping the coin one more time, the player simply takes the new prize of 2^{T-1}. In the normal SPG, there is a probability of $1/2^T$ of getting that prize, and the same probability of getting something more than 2^{T-1}. In the T-SPG, there is thus a probability of 2 x $(1/2^T) = 1/2^{T-1}$ of getting 2^{T-1}, which is also the probability of getting the next smallest prize, 2^{T-2}. Thus, the probability mass on the last two possible prizes is actually the same in the T-SPG. The expected value of the T-SPG is thus

(T+1)/2, since the last term in the summation that determines the expected value is a 1 rather than a ½, as it is for the first T-1 stages. Figure 1 illustrates this for the 5 period truncated SPG and the 10 period truncated SPG. The mass is identical for the first four outcomes—1, 2, 4 and 8. For the outcome of $16, the five-period game has twice the mass of the 10 period game. The 10 period game distributes the extra mass that is on 16 for the five period game to larger outcomes---half of it going to $32, a quarter to $64, etc. (the larger outcomes are not shown).

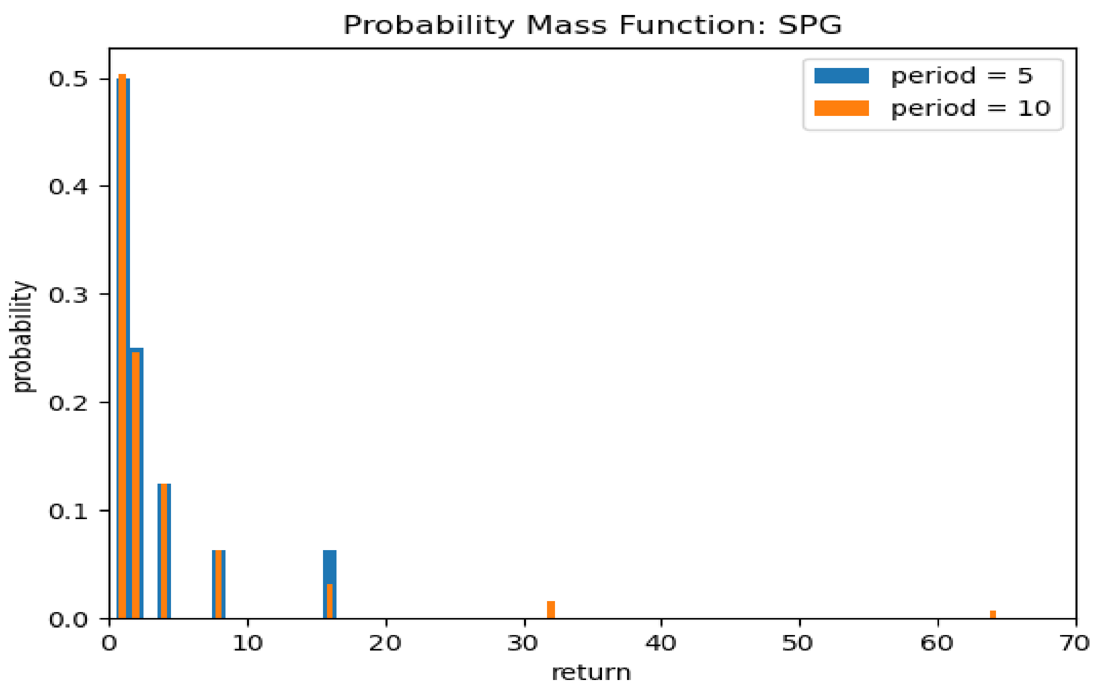

Figure 1: Probability Mass Functions for Truncated Gambles

6. Simulations of the SPG and T-SPG

For the T-SPG the expected value is (T+1)/2. Obviously as T grows, the expected value of the gamble grows. But what is the distribution of outcomes that can be achieved from playing the game once? It is not a distribution with mean (T+1)/2, except when T=1. When T=1, everyone gets 1, so the mean is 1. If T=2, the median is 1 and the mode is 1, but the mean is, in general, NOT 1.5. In general, for T>1, the median and the mode will be 1, but the mean will be something rather unstable—technically, the underlying Markov process at work is *nonergodic*, as already noted. The intuition for this is that as T grows, one needs to have larger and larger "samples" of plays of the game in order to experience the requisite (i.e. "expected") proportion of the higher possible outcomes. The number of states in the Markov process is actually growing as T grows. The more precise and technical reason why sampling yields a better outcome, on average, for the gambler, is that a sample of R instances of the SPG can be viewed as a sum of Geometric random variables (i.e., the distribution of the number of Tails before a Head is flipped). This sum has a Negative Binomial distribution with parameters R and p=1/2, so the mean, which is pR/(1-p) = R. That is, for example, the mean number of "successes" (Tails) before seeing 4 "failures" (Heads) will be 4 for R=4 replications of the SPG. For R=1, the mean number of successes before seeing a single failure is 1. So taking larger samples (playing more instances of the game) makes a quantitative difference in what you are likely to achieve, as an average outcome, and, of course, will also lead to a wider distribution of possible outcomes (higher variance). Later in the paper I will conduct statistical tests to show that as the game grows longer but the sample size stays the same, the sampling distribution of the average payoff of the T-SPG stays the same. Intuitively, there is more upside potential when you sample the game more often, so you are more likely to get the occasional good outcome. But for a given sample size, increasing the length of the game only increases the expected payoff to a point, after which there is no further improvement. Since we are actually interested in the expected *payoff*, and not the expected number of success before R failures, as such, some work needs to be done to account for the how the expected number of successes will fall. First, I will conduct some simulations to make concrete what needs to be explained, after which I will do this accounting exercise.

Figure 2 shows the results of a simulation of the SPG. We start by replicating a graph shown in Hayden and Platt (2009).   The vertical axis shows the mean dollar payoff from playing the number of instances of the game shown on the horizontal axis. The graph shows the median outcome from a sample of 1,000 replications of samples of the given size. For example, sampling 512 instances of the game at a time, and saving the average outcome of the 512 instances for 1000 replications, the median average payoff is about $5.50. Note that $512=2^9$. As outlined in section 5, this is the reciprocal of the probability of getting a prize of 2^9 in the 10-period truncated SPG.   Note also that the expected value of the 10-period truncated SPG is $(10+1)/2=5.5$. Hayden and Platt (2009) noted that the median outcome increases as the size of the samples increase, but they had no further observations related to the significance of this graph, aside from arguing that decision makers might focus on the median outcome of a gamble rather than its expected value in deciding what the gamble is worth. Recall that this is the full SPG, which can go on indefinitely, as long as Tails continue to come up, and the expected value of this gamble is widely believed to be infinite. But here is a simple graph showing that (i) a measure of central tendency of the sampling distribution for the gamble, even with a large ensemble of instances of the game being played at once, is rather small, and (ii) this measure of central tendency appears to be closely related to the truncated versions of the game, via the relationship $R=2^{T-1}$. Specifically, this value of R roughly associates the median of the expected payoff value with the theoretical expected value of the T-period truncated SPG.

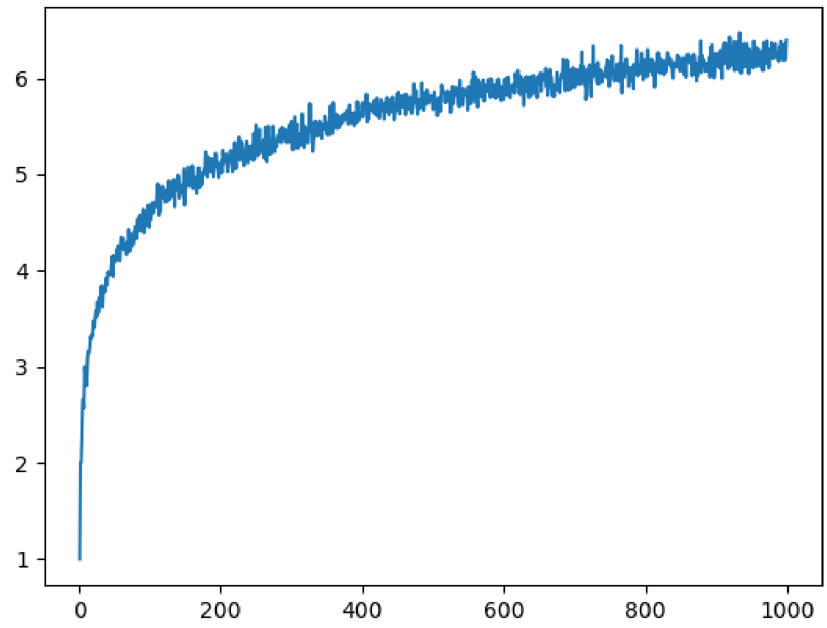

Figure 2: Median Average Payoff for Samples of Size R

Figure 3 explores further this relationship by considering the average payoff achieved when samples of a given fixed size are used in T-truncated SPGs that vary from 1 to 50 periods in length. The vertical axis shows the average return, and the horizontal axis shows the fixed length of games being considered. The upward-sloping red line is the theoretical expected value of the T-truncated games. The orange line is the median of the average payoff achieved if the sample is of size R=512 (the graph is generated with 1000 replications, as if 1000 different people each played the game 512 times). The length of a truncated game is shown on the horizontal axis. Note that the red and orange lines coincide for game sizes up to 10 periods long, after which they diverge, with the median outcome remaining roughly constant for game lengths of longer than 10 periods. Similar to the behavior of the median, other quantiles of the distribution of average outcomes are similarly limited by the size of the sample considered. The fifth percentile, shown in blue, is consistently around 4 for all game lengths equal to or greater than 10, while the 95th percentile, shown in green, is less stable, shooting up quickly at first for game lengths a little above 10, but eventually settling down to vary around 20 or so. As noted above, this is all derivable, in principle, from the fact that the underlying behavior of heads and tails is a negative binomial distribution, with parameters R=512 and p=1/2. The unstable behavior for game lengths significantly above 10 is interesting, and will be the focus of some formal hypothesis testing that I will report on later in the paper.

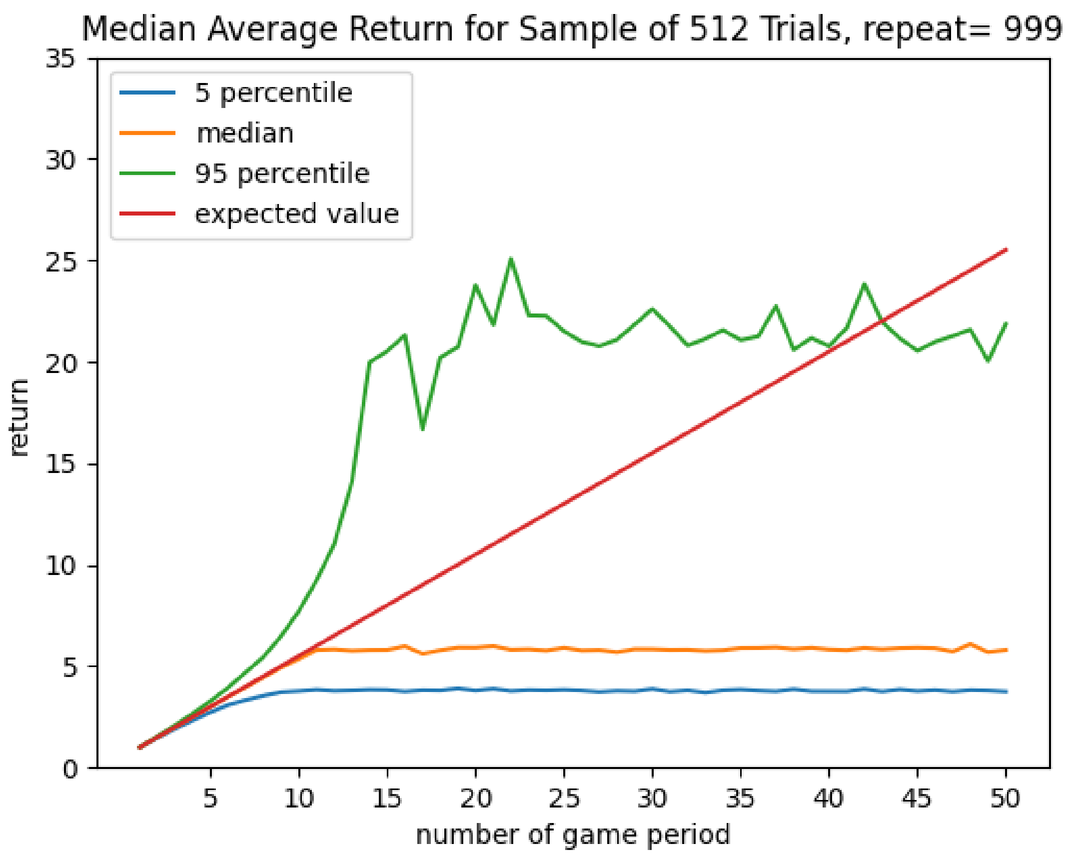

Figure 3: Median Average Return for a sample of a 512

For now, I will just note that the above analysis for R=512 can be repeated for any value of R. In particular, for any R= 2^{T-1}, the median line will coincide very closely with the theoretical expected value line up to a game of length T, and will roughly correspond to (T+1)/2 (which is the expected value of the T-truncated game) even for game lengths greater than T. The fifth and ninety-fifth percentile lines will be different, but rather regular. Of particular interest is the distribution of "average" payoffs when R=1, i.e., when the game is played

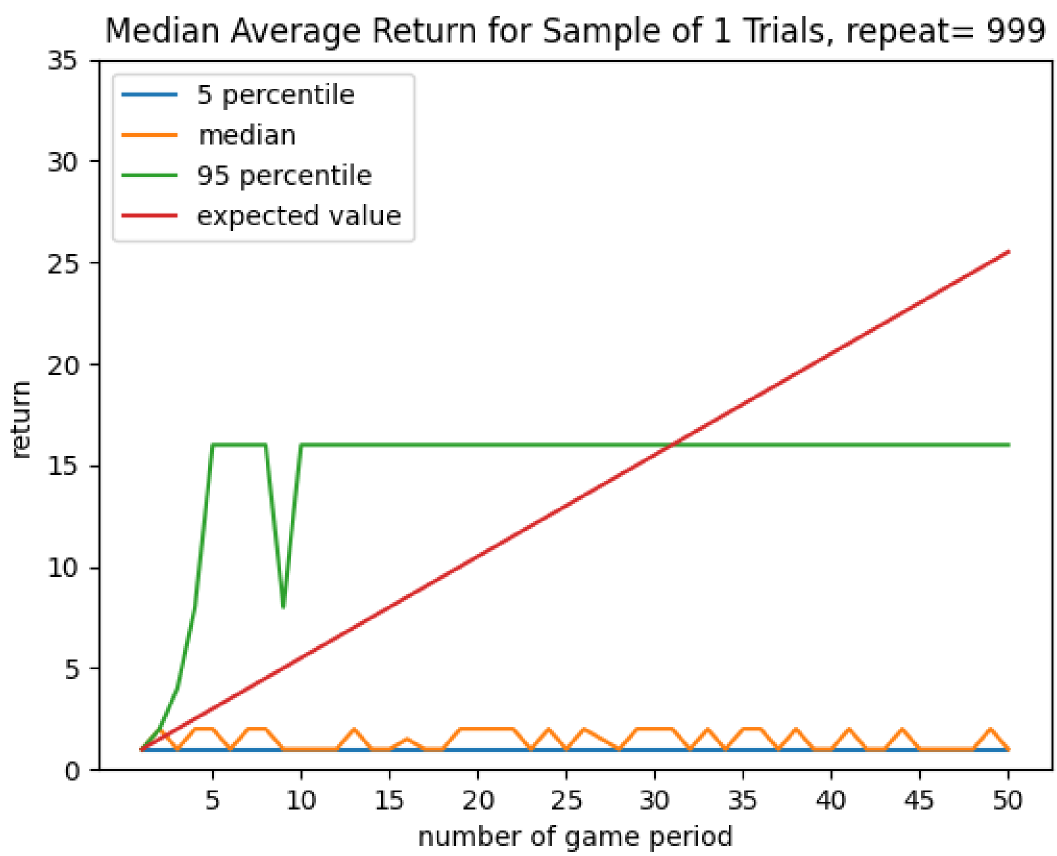

Figure 4: Median Average Return for a sample of a Single Trial

only once, as in the usual story of the St. Petersburg gamble. This case is illustrated in Figure 4. The median (when the sample is replicated 1000 times) is either 1 or 2, the fifth percentile is 1, and the 95th percentile is usually 16 for any game length longer than about 5. This is a direct reflection of the probability mass function for the truncated game, which has cumulative mass of .97 on outcomes less than or equal to 16—the number of replications does not need to be very big for this quantile to be very stable. It bears emphasizing that even for very long games, this 95th percentile does not shift. The clear implication is that the St. Petersburg gamble is simply not a very good bet, if you can only play it once. One need not be risk or loss averse to decline to pay more than a sum that is very modest indeed—like $1—to play this game, simply because something like $1or $2 is what you are most likely to get (so you just break even or earn a dollar) and something much bigger is exceedingly unlikely to materialize.

Another way to see what is going on is to consider histograms of the average payoffs for given sample size (for example, R=16 or R=512), for truncated games as the game length increases.

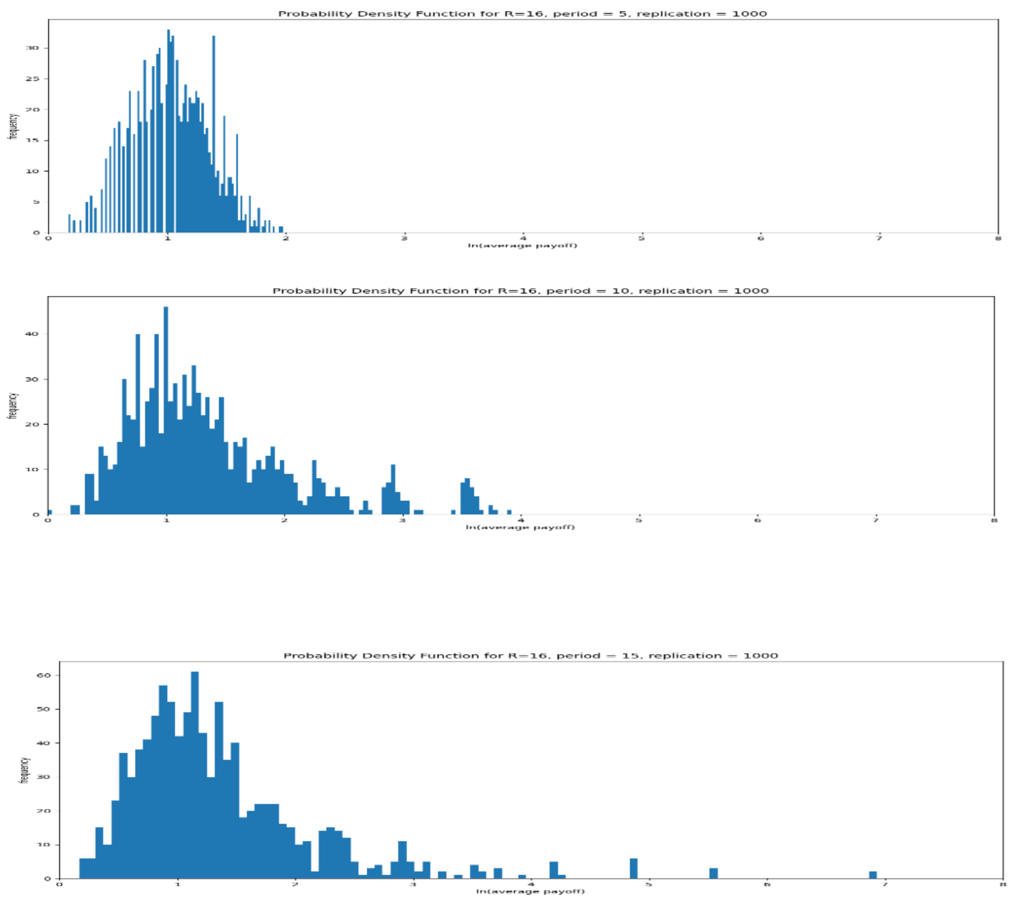

Figure 5: Distribution of Average Payoff for a Given Sample Size (R=16) as Game Lengthens (Horizontal axis is ln(average payoff), to allow display of the full distribution.)

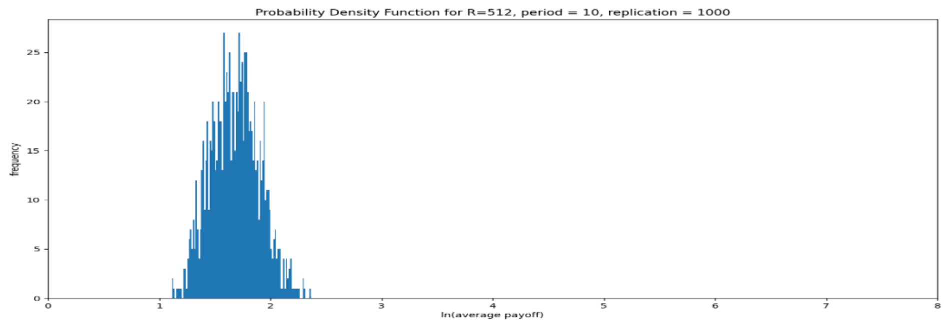

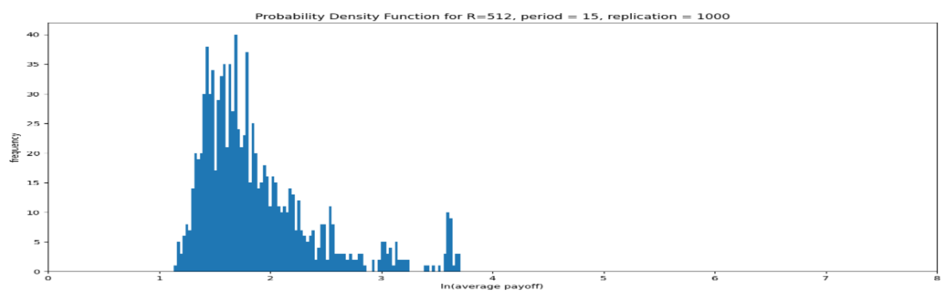

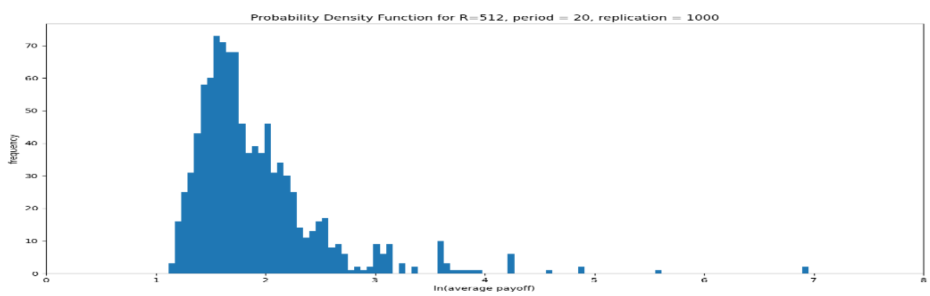

Figure 6: Distribution of Average Payoff for a Given Sample Size (R=512) as Game Lengthens (Horizontal axis is ln(average payoff), to allow display of the full distribution.)

Visually, the distributions shown in Figure 5 and 6 have a similar “core” of outcomes, with a lumpy right tail as the game length increases. In Table 1 I report the results of formal statistical tests between the empirical cumulative distribution functions associated with given sample sizes of 16 (appropriate for the T=5 truncated game) and 512 (appropriate for the T=10 period game), with comparisons between the “baseline” game length and various longer game lengths.   We use the Kolmogorov-Smirnov Test, which is equal to the largest pairwise difference in the cumulative functions over the entire support of the empirical cumulative distributions functions.

Table 1: Number of Rejections for the Kolmogorov-Smirnov Test for Cumulative Distributions of Average Earnings for Truncated Games of Different Length, for a Fixed Sample Size

Replications fixed at R=16, # of Kolmogorv-Smirnov test rejections at the 5% level over 100 repeats, each involving 1000 replications of the sample of 16. Periods refers to the number of periods in a truncated SPG of that length.

5 v 10 periods	10 v 15 periods	15 v 20 periods	20 v 100 periods	10 v 1000 periods	20 v 1000 periods
100	2	3	4	2	3

Replications fixed at R=512, # of Kolmogorv-Smirnov test rejectionvs at the 5% level over 100 repeats, each involving 1000 replications of the sample of 512. Periods refers to the number of periods in a truncated SPG of that length.

10 v 15 periods	15 v 20 periods	20 vs 25 periods	25 v 100 periods	15 v 1000 periods	25 v 1000 periods
100	4	4	3	5	3

Overall, a standard pattern emerges from these tests: Any truncated game of length T + 5 for a sample size of R= 2^{T-1} will have a distribution of average payoffs significantly different from that of the T period game. However, there is no detectable difference between the distribution of average payoffs for any two games, each of different length, but at least T + 5 in length. For example, the distribution of payoffs (when the sample size is 512 for all games) is different between the 10 period and, pairwise, the 15, 20, 25, 50, 100 and 1000 period games. But there is no difference between the distribution of average payoffs between any two adjacent game lengths when both are larger than 10 (e.g., between 15 and 20, between 20 and 25, between 25 and 50, etc.) In fact, this will generally work for any two game lengths, as long as both are at least T + 5 periods long—for example, for games of length 15 and 16, or for games of length 17 and 23. A similar exercise using T=5 as the baseline game and R=16 as the sample size yields the same result. Any two games of length at least 10 have mean earnings distributions that are not different. Repeated simulations show that these statistical results are robust: comparisons between pairs of adjacent game lengths of length T + 5 or greater are significantly different in fewer than 5%, on average, of cases when the exercise is repeated 100 times. Whether the exact margin of 5 that we have identified for T=5 and T=10 period games would hold up for larger values of T (with associated sample size of R= 2^{T-1} is not immediately obvious, but further (time consuming) simulations would demonstrate this. We don't actually need to do more simulations, though, as this can be established in a constructive fashion, using urn-based framework, which I will develop in section 7.

Interestingly, the median of average earnings achieved in an T period truncated SPG when R= 2^{T-1} is not quite equal to the theoretical expected value for the game of that length, and the median increases for games slightly longer than T, but settle down to a stable number for games of length T+2 or greater. For example, the median for a T=5 period game is 2.88 (vs. the theoretical expected value of 3). For T=6, the median is 3.27, and for T>6 the median is 3.28. Similarly, for T=10 the median is 5.37 (vs. the theoretical expected value of 5.5). For T=11 the median is 5.73, and for N>11 the median is 5.8. So the medians of the sampling distribution are the same for any two games of length > T+2, even though it is only for games of at least length T+5 (as far as we can now tell) that the distributions are the same. I would be surprised if this sort of relationship failed to hold for much larger values of T (and samples of $R= 2^{T-1}$).

One piece of insight as what underlies the regularity just noted can be gleaned from a closer look at the right end of the distribution of average payoffs. Figures 7a to 7c shows the

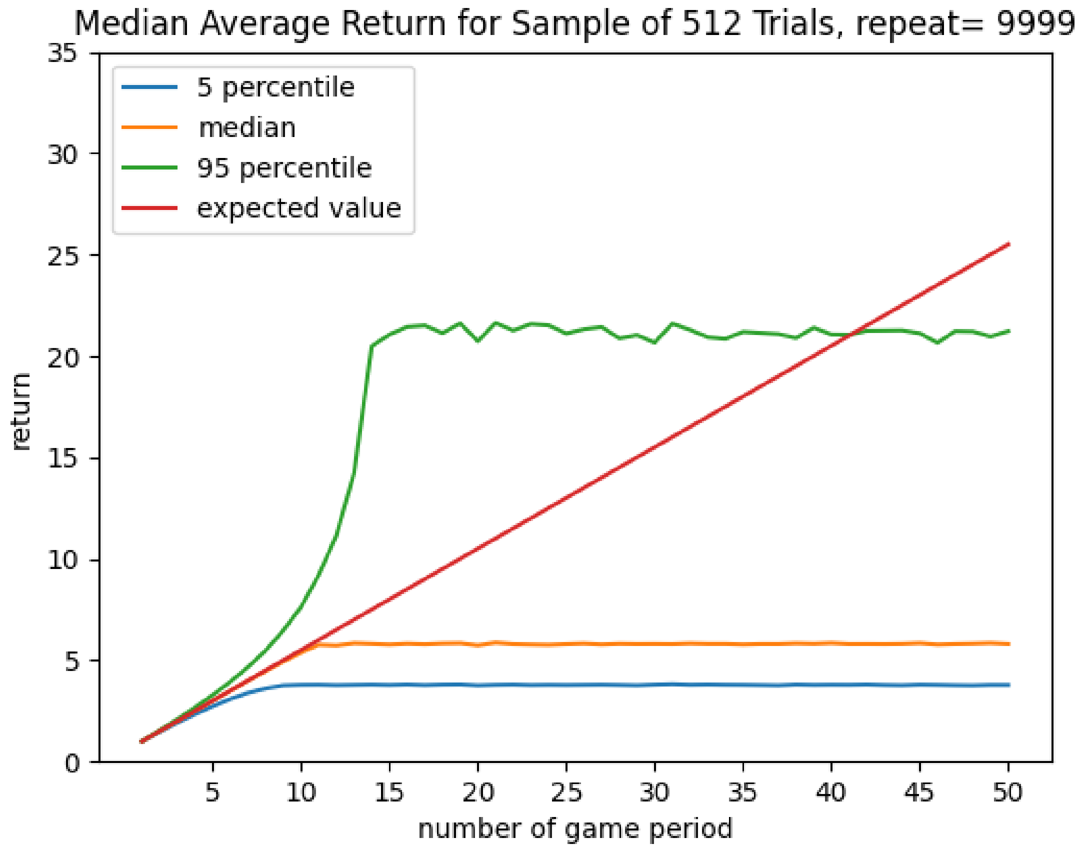

Figure 7a: Median Average Return for Sample of 512 Trials, repeated 10,000 Times

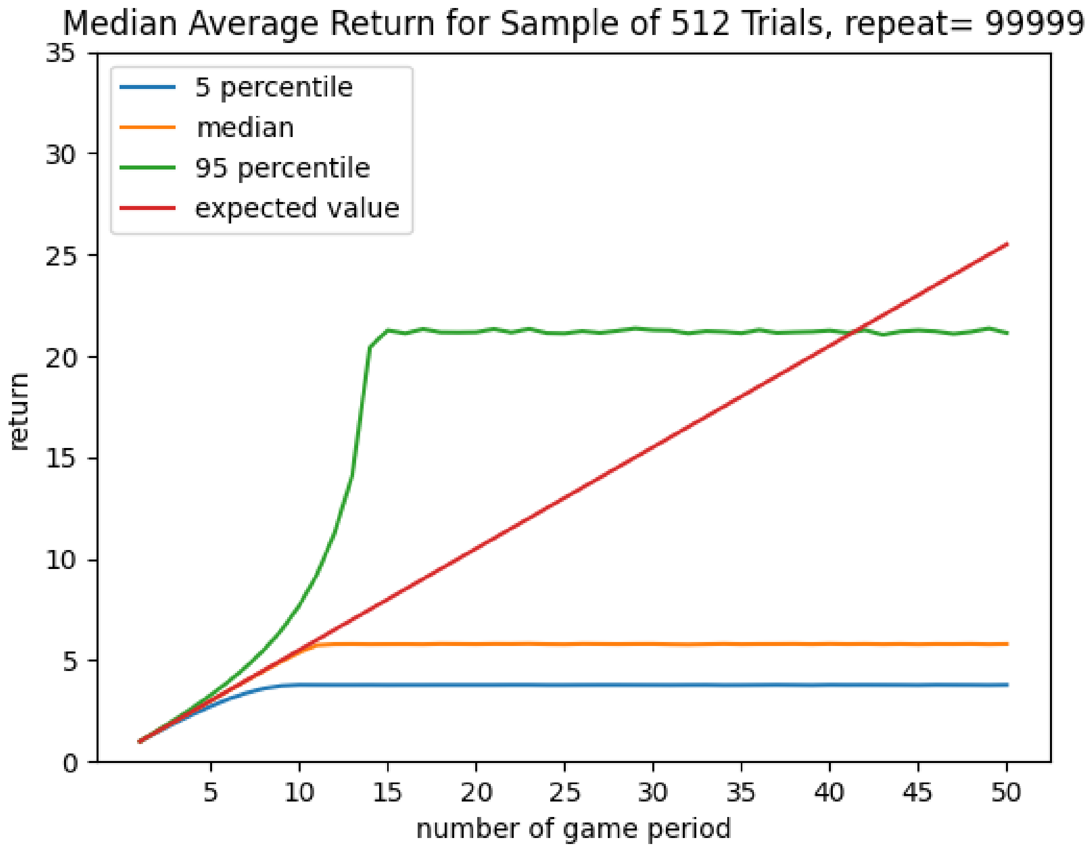

Figure 7b: Median Average Return for Sample of 512 Trials, repeated 100,000 Times

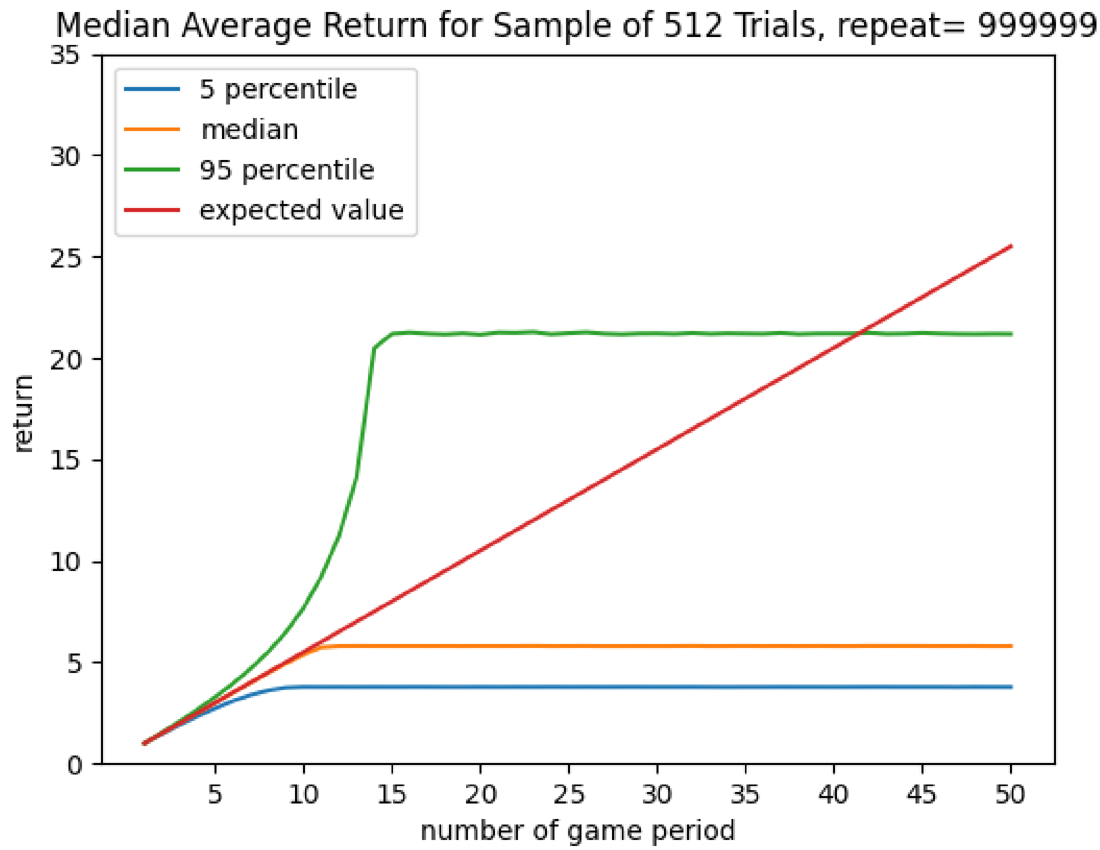

Figure 7c: Median Average Return for Sample of 512 Trials, repeated 1,000,000 Times

of the 5th, 50th and 95th percentiles of average payoffs over 10,000, 100,000 and 1,000,000 separate replications (Figure 3 showed the same statistics for 1,000 replications). The resulting smoothing that we see in Figure 6 compared to Figure 3 is exactly as one would expect, and the 95th percentile line, in particular, is reassuring in showing that very extreme outcomes are, statistically, well contained as the length of the game played is increased. I note one specific thing, that the 95th percentile line "levels off" right at 15 periods, and varies very little after that, consonant with the result that the distribution of the average return from playing games of length 16 and greater is not different from that of the 15 period game.

7. Connecting the Sample Size Results with the Expected Payoff Results

The expected number of successes before R failures is R, as noted above. I now provide a concrete and constructive demonstration of what this fact implies about the nature of what the expected payoff is for R plays of the SPG. For any finite period T-SPG we can think of a play of the game as being a draw from an urn containing 2^{T-1} balls, each numbered with the size of the prize that one will win if one draws that ball from the urn. For example, if T=10, there are 512 balls. Half of the balls, 256, will be numbered with $1, 128 with $2, 64 with $4, 32 with $8, 16 with $16, 8 with $32, 4 with $64, 2 with $128, 1 with $256, and another 1 with $512. This last ball is due to the fact that we put all of the remaining probability weight on the last prize, rather than continuing to halve it up *ad infinitum*. Now, a single draw from this urn could yield any one of the prizes available, but it is most likely to yield $1. On the other hand, 512 independent draws from this urn (with replacement, of course) is just the right number to yield, on average over many many trials, the average payoff of $5.50, which is the theoretical expected value for this urn. If the 512 draws are done by a single individual, then that individual can expect to win $5.50, as an average payoff. If the 512 draws are done by 512 different individuals, the average winnings of the group of 512 can be expected to be $5.50, but the "average individual" will not earn that much. But the key point here is that a single individual can expected to earn (T+1)/2 in the T-period SPG by playing it 2^{T-1} times.

Reasoning in a different direction, suppose that a single individual still draws 512 times, but this time from the urn that represents the T=11 period truncated game. This urn would have 1024 balls. Half of the balls, 512, will be numbered with $1, 256 with $2, 128 with $4, 64 with $8, 32 with $16, 16 with $32, 8 with $64, 4 with $128, 2 with $256, 1 with $512, and another 1 with $1024. The key idea here is that as there are only 512 draws to be done, the *expected* number of $1 balls drawn will be 256, the *expected* number of $2 balls drawn will be 128, and so forth—that is, the expected number of balls with each possible dollar amount drawn will be the same as the actual number of balls in the T=10 urn, until we get down to the last two, the $512 ball and the $1024 ball, each of which have an expected number of draws of ½. The intuition here is that there are not enough draws being done to achieve the theoretical expected value of the T=11 truncated SPG. You can expect to get mostly the same value out of this gamble as you get from the T=10 period gamble, except that sometimes you will get the $512 ball, while other times you will get the $1024 ball on one of your 512 draws from the urn (averaging over many such ensemble bets of 512, of course).

If we consider urns representing the T=12 period gamble, the T=13 period gamble, and so on, the problem continues: the expected number of draws of balls labelled $2048 (for T=12), $4096 (for T=13), and so on, are smaller and smaller, and only occasionally will be drawn in an ensemble of 512 draws from the urn, and this will be insufficient for the player to achieve the expected value that is (theoretically) possible. Figure 8 illustrates an example of average earnings achieved by a player who plays each game of length T, with T ranging from 1 to 50 periods. The purple line represents the theoretical expected value for each game, while the red line shows the actual average earnings achieved. The dashed line shows the expected value for the T=10 period game. The average earnings tend to be near the expected value for the 10 period game, though there are deviations, with deviations above being naturally larger than those below $5.50, due to the exponential growth of possible payoffs as T increases. But expected earnings of the 10 period game are a center of gravity, of sorts, for the average earnings achieved.

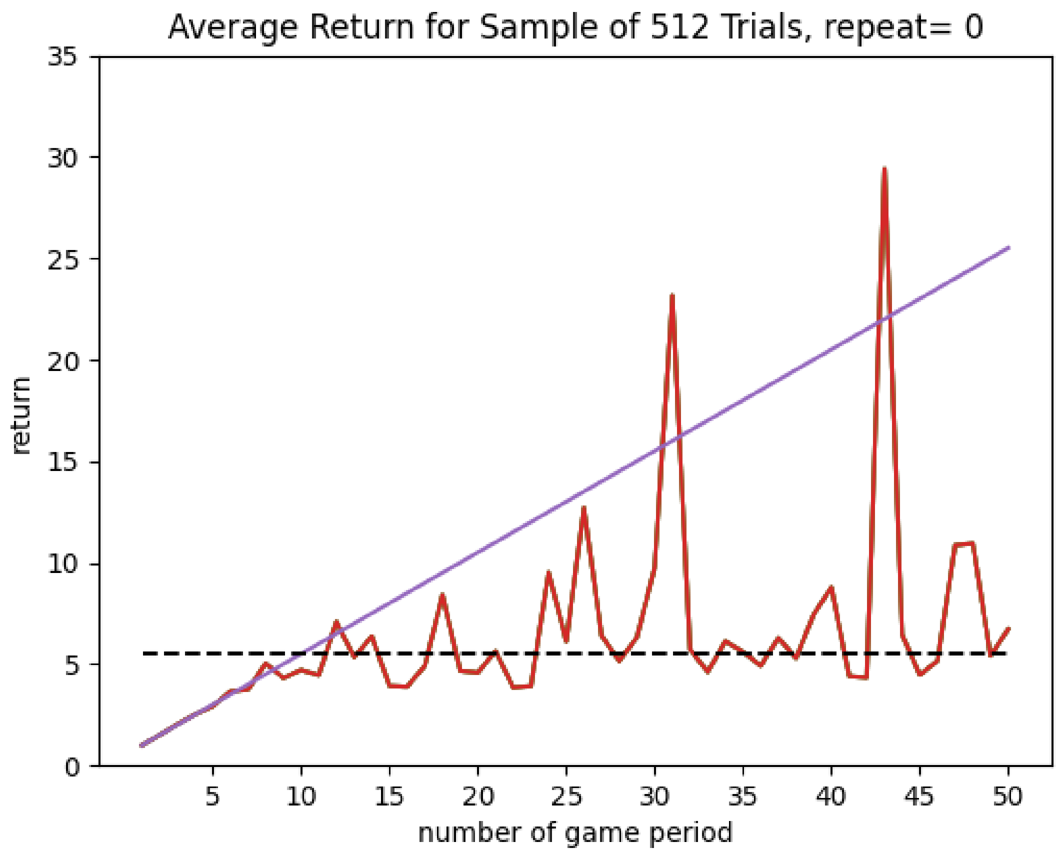

Figure 8: Average Earnings for a Sample of 512 Trials for Each Gamble of length T=1 to T=50.

Using the idea of urns containing numbered balls to represent different T-SPG games, we can get a very concrete understanding of why the result that, with a sample of size $R=2^{T-1}$, the distribution of average payoffs for any two games of length at least T+5 are not different, is a very general result. As discussed earlier, when allocating the expected R "successes" to draws from an urn, the number of successes is exactly "enough" to be able to achieve the theoretical expected value of (T+1)/2 when there are $R=2^{T-1}$ balls in the urn, but when there are more balls (i.e., when the urn represents a longer T-SPG game), the number of draws being made from the urn will be insufficient to achieve the higher expected value that, theoretically, is associated with the longer game. Table 2 illustrates the situation when R=16, and Table 3 illustrates the situation when R=512. The tables show the number of balls labelled with each possible dollar prize that can be won in drawing from urns representing games of length 5 through 11 periods (in Table 2) and from urns representing games of length 10 through 16 periods (in Table 3). The far right-hand column of each table shows the expected number of balls of each denomination that will be drawn when R=16 (in Table 2) or when R=512 (in Table 3) is the number of draws. Of particular interest are expected draws less than 1, as that means that in a given ensemble of draws, a player is unlikely to consistently draw the higher prizes that are made available as the game is lengthened. As the length of the game increases, higher and higher prizes are possible, but there is still, essentially, only one draw of those prizes in a given ensemble of plays that we would expect to see.

A key observation now is to note that the last seven rows of Table 2 and Table 3 are exactly the same, except that the rows in Table 3 are associated with higher payoffs. The point here is that the differences between the distribution (think here of the cumulative distribution of average payoffs) as the game is lengthened are going to be the same, in an expected sense, regardless of the "scale" at which we are operating, where the scale is the size of the ensemble of bets, R. So it is not surprising that the Kolmogorov-Smirnov tests discussed earlier worked out so similarly for the two cases considered, R=16 and R=512. The difference between the distribution of average payoffs for a given level of R in two different length gambles comes down to what is happening in the right tail of the distributions. This is because, in an expected sense, the bulk of the distributions are going to be the same, since the expected number of draws of the lower dollar prizes than any two cases have in common are exactly the same. It is the new, larger prizes that become available when the game is lengthened that are going to make the difference. For any two adjacent urn representations (i.e., columns) in Table 2 or Table 3, the difference between them, in expected terms, is that the urn on the right always has one more (larger) prize, and likelihood of that prize and the next smaller prize being drawn are equal, while in the left hand column all of the likelihood is on the smaller of the two prizes. But as the games get longer, the total likelihood on these prizes is getting smaller, so that eventually, even if a larger prize is available, it is not detectable by the K-S test.

Table 2: Accounting for the Expected Number of Successes before a Given Number of Failures: R=16

Length F=R/B Balls in urn	T=5 F=1 B=16	T=6 F=1/2 B=32	T=7 F=1/4 B=64	T=8 F=1/8 B=128	T=9 F=1/16 B=256	T=10 F=1/32 B=512	T=11 F=1/64 B=1024	Expected # drawn if R=16	
$	# of Balls	# of Balls	# of Balls	# of Balls	# of Balls	# of Balls	# of Balls	= F*(# of Balls)	
1	8	16	32	64	128	256	512	8	
2	4	8	16	32	64	128	256	4	
4	2	4	8	16	32	64	128	2	
8	1	2	4	8	16	32	64	1	
16	1	1 (1/2)	2 (1/2)	4 (1/2)	8 (1/2)	16 (1/2)	32 (1/2)	1	
32	0	1 (1/2)	1 (1/4)	2 (1/4)	4 (1/4)	8 (1/4)	16 (1/4)		
64	0	0	1 (1/4)	1 (1/8)	2 (1/8)	4 (1/8)	8 (1/8)		
128	0	0	0	1 (1/8)	1 (1/16)	2 (1/16)	4 (1/16)		
256	0	0	0	0	1 (1/16)	1 (1/32)	2 (1/32)		
512	0	0	0	0	0	1 (1/32)	1 (1/64)		
1024	0	0	0	0	0	0	1 (1/64)		

Table 3: Accounting for the Expected Number of Successes before a Given Number of Failures: R=512

Length F=R/B Balls in urn	T=10 F=1 B=512	T=11 F=1/2 B=(1024)	T=12 F=1/4 B=2048	T=13 F=1/8 B=4096	T=14 F=1/16 B=8192	T=15 F=1/32 B=16384	T=16 F=1/64 B=32772	Expected # drawn if R=512	
$	# of Balls	# of Balls	# of Balls	# of Balls	# of Balls	# of Balls	# of Balls	= F*(# of Balls)	
1	256	512	1024	2048	4096	8192	16384	256	
2	128	256	512	1024	2048	4096	8192	128	
4	64	128	256	512	1024	2048	4096	64	
8	32	64	128	256	512	1024	2048	32	
16	16	32	64	128	256	512	1024	16	
32	8	16	32	64	128	256	512	8	
64	4	8	16	32	64	128	256	4	
128	2	4	8	16	32	64	128	2	
256	1	2	4	8	16	32	64	1	
512	1	1 (1/2)	2 (1/2)	4 (1/2)	8 (1/2)	16 (1/2)	32 (1/2)		
1024	0	1 (1/2)	1 (1/4)	2 (1/4)	4 (1/4)	8 (1/4)	16 (1/4)		
2048	0	0	1 (1/4)	1 (1/8)	2 (1/8)	4 (1/8)	8 (1/8)		
4096	0	0	0	1 (1/8)	1 (1/16)	2 (1/16)	4 (1/16)	1	
8192	0	0	0	0	1 (1/16)	1 (1/32)	2 (1/32)		
16384	0	0	0	0	0	1 (1/32)	1 (1/64)		
32772	0	0	0	0	0	0	1 (1/64)		

To make this more concrete, a look at the cumulative distribution functions associated with R=16 and T=5, 10, and 15, and separately with R=512 and T=10, 15, and 20, are revealing. Figures 9 and 10 contain these distributions, which are based on the same sort of 1000 replications of the games that were used in conducting the Kolmogorov-Smirnov tests reported on earlier. As for the histograms in Figures 5 and 6, the horizontal scale is ln(average earnings) to allow easy display of all outcomes.   Except for a shift in location, it is not possible to distinguish between these two sets of cumulative functions.

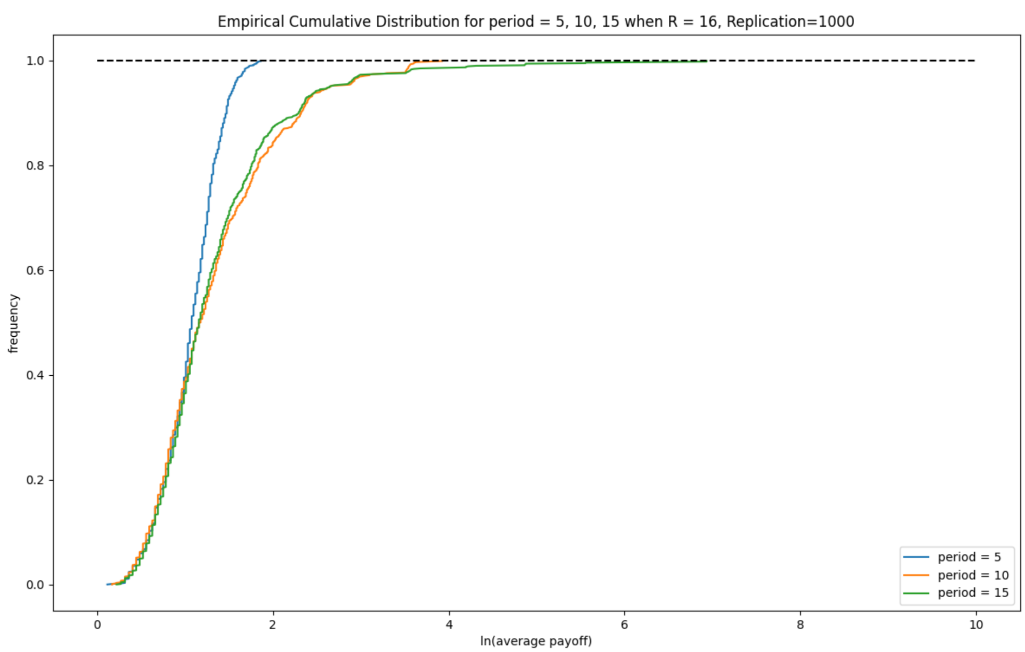

Figure 9: Cumulative Distributions for R=16, and T=5, 10 and 15

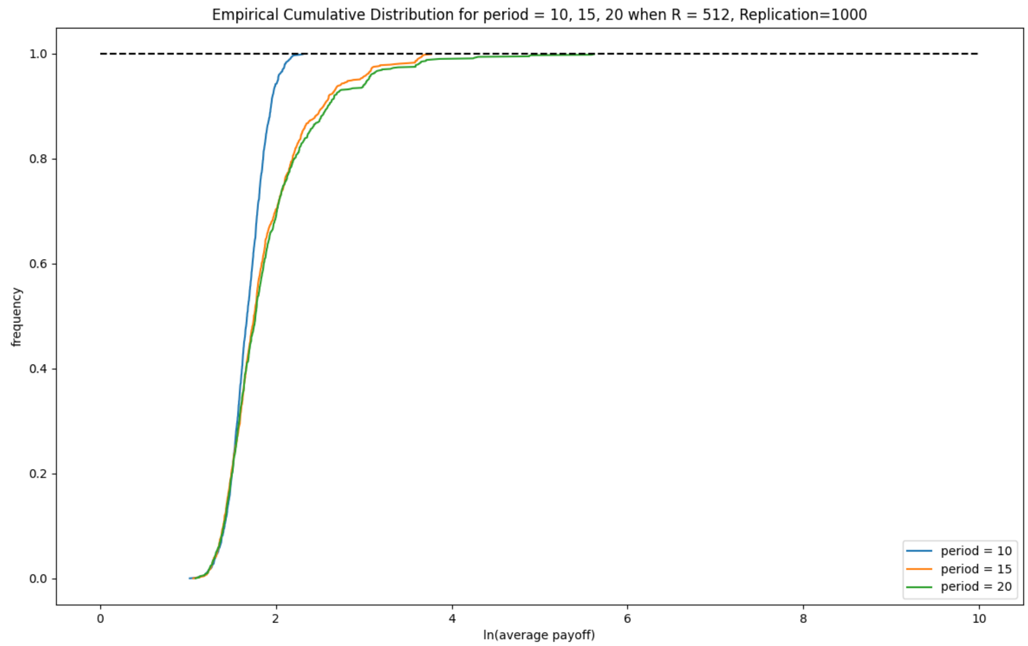

Figure 10: Cumulative Distributions for R=512 and T=10, 15, and 20

We can summarize our results in a condensed form by considering the effect of more people playing the SPG, and of each person playing the SPG more often. We focus here on the indefinitely repeated SPG, reasoning that, as just shown, the expected value of the T-SPG is approximately what a player of the regular SPG could expect to earn if $R=2^{(T-1)}$. Table 4 provides a schematic for thinking about these two dimensions. We will not refer to T in this table, but rather think of both N and R as indexes ranging from 0 upward in order to generate different sized ensembles of players (for a given play of the game) and different sized ensembles of plays of the game (for a given individual), via the quantities $2^{(N-1)}$ and $2^{(R-1)}$. A given box in the table shows the earnings that can be expected for (i) a given player who has played the number of replications of the SPG shown in the left hand column, in the upper left of the box, and (ii) the average earnings over the number of players playing the SPG, shown in the bottom row of the table, in the lower right of the box. For example, average earnings of a player playing the SPG 2^7 =128 times is 4.5 (highlighted in yellow), and the average over 2^3 =8 players playing the SPG once is 2.5 (highlighted in green).

The stochastic process does not "care" whether an ensemble of plays of the game are associated with different players, or with a single player, but, obviously, it matters to the individuals. As noted early in the paper, 1000 individuals (approximately 2^{10}, or N=11 in the table) playing the SPG once will earn, on average, about (N+1)/2 = 6, but the typical player will earn much less. On the other hand, a given individual playing the game 1000 times (approximately 2^{10}, or R=11 in the table) can expect to earn about (R+1)/2=6. The vertical dimension of this table is of great interest to an individual player, the horizontal dimension not at all, unless one can team up with other players in order to economize on entry fees and share in a higher average payoff. In general, the larger R is, the larger the average payoff to an individual player. Alternatively, the larger the set of other players, N, that one can combine with and split proceeds equally, the larger the average payoff each individual can achieve.

Table 4: Average Payoffs as the Number of Players and Replications Per Player Grows, Full SPG

2^{R-1}	(R+1)/2 1.0	(R+1)/2 1.5	(R+1)/2 2.0	(R+1)/2 2.5	(R+1)/2 3.0	(R+1)/2 3.5	(R+1)/2 4.0	(R+1)/2 4.5	(R+1)/2 5.0	(R+1)/2 5.5	(R+1)/2 (N+1)/2
2^9	5.5 1.0	5.5 1.5	5.5 2.0	5.5 2.5	5.5 3.0	5.5 3.5	5.5 4.0	5.5 4.5	5.5 5.0	5.5 5.5	5.5 (N+1)/2
2^8	5.0 1.0	5.0 1.5	5.0 2.0	5.0 2.5	5.0 3.0	5.0 3.5	5.0 4.0	5.0 4.5	5.0 5.0	5.0 5.5	5.0 (N+1)/2
2^7	4.5 1.0	4.5 1.5	4.5 2.0	4.5 2.5	4.5 3.0	4.5 3.5	4.5 4.0	4.5 4.5	4.5 5.0	4.5 5.5	4.5 (N+1)/2
2^6	4.0 1.0	4.0 1.5	4.0 2.0	4.0 2.5	4.0 3.0	4.0 3.5	4.0 4.0	4.0 4.5	4.0 5.0	4.0 5.5	4.0 (N+1)/2
2^5	3.5 1.0	3.5 1.5	3.5 2.0	3.5 2.5	3.5 3.0	3.5 3.5	3.5 4.0	3.5 4.5	3.5 5.0	3.5 5.5	3.5 (N+1)/2
2^4	3.0 1.0	3.0 1.5	3.0 2.0	3.0 2.5	3.0 3.0	3.0 3.5	3.0 4.0	3.0 4.5	3.0 5.0	3.0 5.5	3.0 (N+1)/2
2^3	2.5 1.0	2.5 1.5	2.5 2.0	2.5 2.5	2.5 3.0	2.5 3.5	2.5 4.0	2.5 4.5	2.5 5.0	2.5 5.5	2.5 (N+1)/2
2^2	2.0 1.0	2.0 1.5	2.0 2.0	2.0 2.5	2.0 3.0	2.0 3.5	2.0 4.0	2.0 4.5	2.0 5.0	2.0 5.5	2.0 (N+1)/2
2^1	1.5 1.0	1.5 1.5	1.5 2.0	1.5 2.5	1.5 3.0	1.5 3.5	1.5 4.0	1.5 4.5	1.5 5.0	1.5 5.5	1.5 (N+1)/2
2^0	1.0 1.0	1.0 1.5	1.0 2.0	1.0 2.5	1.0 3.0	1.0 3.5	1.0 4.0	1.0 4.5	1.0 5.0	1.0 5.5	1.0 (N+1)/2
Replications ↑ **Players →**	2^0	2^1	2^2	2^3	2^4	2^5	2^6	2^7	2^8	2^9	2^{N-1}

8. Implications for Decision Making

We can formalize the preceding exploration of the nature of the sampling distribution for replications of the St. Petersburg Gamble for a rational decision maker as follows: For any T^* one might choose (ie, a truncated game of length T*), $R=2^{T^*-1}$ is the number of replications, or plays, of the T^*-long game needed to ensure, on average, that one will be able to achieve an average payoff of (nearly) $(T^*+1)/2$. For any $T > T^* + 2$ the median payoff of the T-long game will be (slightly greater than) $(T^*+1)/2$. More generally, then, and ignoring now the slight shortfall of median compared to the theoretical expected value noted above, for any finite R, your expected payoff is well approximated by the mean outcome of the game with T periods in length, with T=[ln(R)/ln(2)] +1. So, evidently, if offered a game with a sample size of R, the monetary value of the game is Value(R)={[ln(R)/ln(2)] + 2}/2, the expected value of the T= [ln(R)/ln(2)] +1 period game. Only with R approaching infinity will your average payoff approach infinity. For example, R=1 implies T=1, and R=2 implies T=2. So Value(1)=1, while Value(2)=1.5. Note that Value(346) = {[ln(346)/ln(2)] + 2}/2 = 5.2, for example, and since we can't have other than whole-number length games, slight adjustments to our formulae are needed. It can be confirmed that for sample of size 256, Value (256)=5 (i.e., for $R=2^9$) and, in general, when R is an integer power of 2, the answer will be an integer. So T= int{[ln(R)/ln(2)] +1}is the corrected definition, or Value(R)=int({[ln(R)/ln(2)] + 2}/2). One most efficiently gambles by choosing integer powers of 2 as the sample size.

The clear implication of the above is that the same reasoning will apply to the full, indefinitely repeated SPG, since having additional periods beyond T^* available is of no particular advantage. The very tiny probabilities of getting the very huge prizes associated with the event that the game reaches a stage $T>T^*$ are not just perceptually small (as Bernoulli argued) but actually small, and smaller than one thinks, in the sense that there are just not enough chances to get those small probability/large outcome events that are theoretically possible on a consistent basis. For example, a sample of R=512 allows for a prize that occurs with probability 1/512 on occasion (about one time, for the 10 period game), but is unlikely to allow for an event that occurs with probability 1/524,288, as in the 20 period game. That is the intuition: there simply are not enough trials to be able to see the full range of outcomes that are theoretically possible, and in particular, it is the biggest possible prizes that are least likely to be seen.

Note that I am not saying that risk aversion could not play a role in decision making here. But it is important to know what the concrete consequences of one's actions are likely to be. Up to now, in discussions of the St. Petersburg Gamble, the implicit or explicit assumption has been that one is dealing with a gamble where the expected value is readily "available." The preceding, I think, makes it clear that this is far from the case. Thinking in terms of the urns containing 2^{T-1} balls: as T goes to infinity, we need an urn containing an infinite number of balls for the "infinite prize" to be available for drawing, which of course is impossible. But even for large finite values of T, the number of plays of the game required to achieve the theoretical expect value of the T period game is huge (T=100 requires a sample of size 6.34×10^{29}, for example). It does appear that, if there were a bookmaker prepared to accept a large number of bets from an individual, then such an individual, with sufficiently "deep pockets," could assure him or herself with a high degree of certainty of a given level of earnings. But the pockets of the gambler would need to be very deep indeed. For example, to assure oneself of getting $5.50 on average, one would want to have a sample of R=512. To assure oneself of getting $10.50, on average, one would want to have a sample of R=524,288. How such gambles would be priced is an important question. The former could be expected to cost in the neighborhood of $2,500, the latter in the neighborhood of $5 million, if the market were competitive. Of course, all of this brings up the question of what one's attitude to risk might be. With so many replications, risk seems to be less of a concern: even though a sample of size $R=2^{T-1}$ does not guarantee exactly (T+1)/2, the simulations show pretty clearly that oversampling (playing more than R times) actually allows one to make the variance arbitrarily small. The nature of the pricing of the gamble and the role of competition suggests that there are zero economic profits to be earned here. One could expect those offering the gamble to try to obfuscate in the usual ways, tempting potential gamblers with the possible very large prizes that are available. One would also expect those offering the gamble to not be interested in allowing gamblers to choose large ensembles of bets simultaneously. By so doing, if they could manage to get something like the theoretical expected value of some truncated gamble as a selling price, but with an associated ensemble of bets that is below what we now know is needed for the gambler to achieve that average payoff, then there would be profits to be earned. For the gamblers' part, one could expect to see cooperative efforts to buy shares of an appropriately large ensemble of bets as a way to overcome the deep pockets problem, as suggested in section 7 in the discussion of Table 4.

Can anything be said about risk preferences in the context of the Saint Petersburg Gamble?

9. Conclusions

Getting back more directly to the issue that motivated this paper to begin with, what are we to make of the St. Petersburg Gamble, as originally proposed? I would propose the following: (i) We have been wrong to suppose that the St. Petersburg Gamble has an infinitely large expected value. "Arbitrarily large" might be a better way to describe what is possible, but what is possible is constrained by the size of the ensemble of bets one is able to purchase. (ii) The observation that the typical person is going to be willing to pay, at most, a very modest amount in order to play this gamble tells us essentially nothing about risk aversion. The scant empirical evidence that we have suggests that something like $1, or maybe $2, is about the most that a typical person would be willing to pay, and this actually looks like an income-maximizing strategy. If anything, there is more evidence for risk seeking behavior than there is for risk average behavior in the data. The fact is that very large payoffs are possible in this game, and some, for example, those willing to pay up to and beyond $8, as in the Hayden and Platt (2009) study, are actually paying much more than they should, if they hope to earn a profit. Likely, they are not aiming at a profit so much as hoping to get lucky and win a big prize. (iii) More attention paid to the dynamics of wealth accumulation in economics, quite apart from issues of risk aversion, as such, would appear to be in order. The basic framework for the study of risk bearing that we have is fine for a large set of questions: insurance, diversification of risk, and other questions that are essentially static in nature, with a simple ex ante vs. ex post, mediated by a probability density or mass function. But dynamic questions, of which the St. Petersburg Gamble is a prime example, have been surprisingly understudied—surprising, that is, in light of the extent to which what was thought to be true about this gamble, so familiar to us, was so wrong. Evidently we, as a profession, have been insufficiently curious to probe the issues, which actually are many. Just to give one example, the issue of how to best save for retirement, and the difficulties so many people have in doing this, is treated mainly as a behavioral problem of people being insufficiently patient and too present-oriented to do what is needed. But it may well be that it is more a question of people having a sufficient appreciation and understanding for the nature and power of compound interest, and this may apply equally well to public policy debates on the subject. We have been in the habit of saying that those unwilling to pay more than a modest amount to play this gamble are risk averse, possibly excessively so, only to belatedly realize that those people were doing something rather sensible.

10. References

Aumann, Robert (1977) The St. Petersburg Paradox: A discussion of some recent comments," Journal of Economic theory, 14, 443-445.

Bernoulli, Daniel (1954) "Exposition of a New Theory on the Measurement of Risk." *Econometrica* **22(1):** 23-36. https://www.jstor.org/stable/1909829

Cox, James, Vjollca Sadriaj, and Bodo Vogt (2009) "On the empirical relevance of the St. Petersburg Paradox." *Economics Bulletin* **29**(1): 214-220.

Hayden, Benjamin Y. and Michael L. Platt (2009) "The mean, the median and the St. Petersburg paradox, *Judgment and Decision Making*, **4(**4).

Peters, Ole (2011) "The time resolution of the St Petersburg Paradox." *Phil. Trans. R. Soc. A,* **369**(1956): 4913–4931.

Peters, Ole, and Adamou Alex. "The Time Interpretation of Expected Utility Theory." Papers 1801.03680, arXiv.org (January 12, 2018).

Peters, Ole and Murray Gell-Mann (2015) "Evaluating Gambles using Dynamics."

Samuelson, Paul (1977). "St. Petersburg Paradoxes: Defanged, Dissected, and Historically Described, *Journal of Economic Literature*, Vol. 15, No. 1, pp. 24-55. https://www.jstor.org/stable/2722712