

# An interesting game

Analysis by bonanova, 28 Nov 2013

A game is proposed where \$1 is bet, and each pull of the handle returns randomly a payoff of .7 .8 .9 1.1 1.2 or 1.5.

In one method of play, winnings are taken after each pull, and a fresh \$1 is bet on the next pull. Analysis is straightforward, with an expected payoff equal to the arithmetic mean **AM** of the six individual payoffs, or  $31/30 = 1.03333 \dots$ . This is a slow and steady payoff with very small chance of going bust, even if winnings are returned in discrete amounts (pennies.)

A second method of play bets the winnings after each pull repeatedly. After  $p$  pulls the player quits and takes his winnings, only at the very end of play. The first method is additive, the second is multiplicative. One expects then the geometric mean **GM** of the six payoffs to play a role. **GM** here is the sixth root of .99792 - a number very close to unity: 0.9996530325. And that is a frequent result. What is interesting about this game is that it does not always happen that way. And that means you are almost certain of winning if you play the game repeatedly, say  $n$  times. Why? The answer to that question is the subject of what follows.

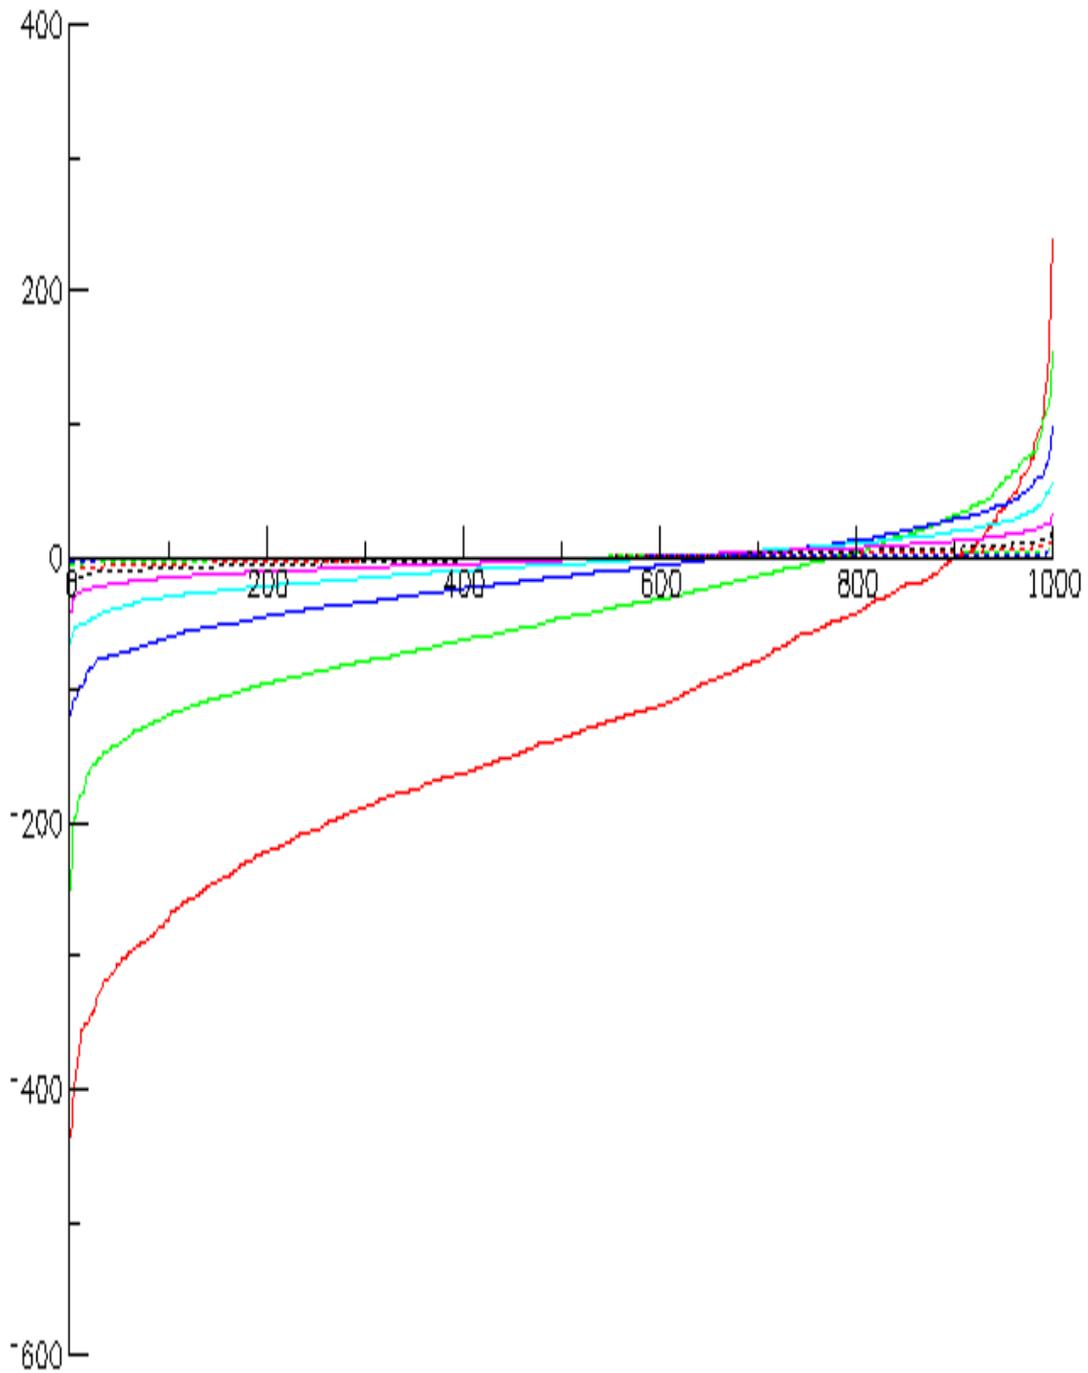
The product of  $p$  random selections of the six payoffs is almost sure to give you a large number at some point. On average, three times out of ten that product will be greater than unity. And the size of the greatest of those "winning" selections grows greater than exponentially with  $p$ . Once that number exceeds  $n$  you win, after splitting the winnings among the players.

One of the characteristics of this game is that you never go bust. If your initial \$1 stake is infinitely divisible, you always have something to bet. For large  $p$  your stake can become infinitesimally small, but it never becomes zero. This suggests that a good way to analyze the game is to look at the logarithm of the outcome. In this domain, zero is break-even; positive is win; and negative is lose.

Let's simulate the game for  $n=1000$  (players) and  $p$  (pulls) varying from one hundred to one million in nine roughly equal logarithmic steps: 100, 300, 1000, 3000, 10,000, 30,000, 100,000, 300,000, and 1,000,000. Results should be stationary with respect to  $n$ , (so long as  $n >$  say 100) but vary greatly with  $p$ .

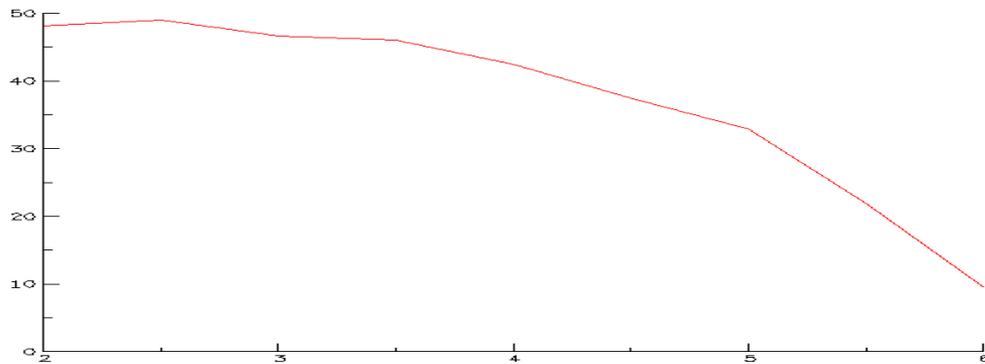
What we see playing out in these calculations is the fact that for each of the  $n$  players the worst outcome is  $(.7)^p$ , and that consequence only loses your initial \$1 bet. But the upside limit is  $(1.5)^p$ . Any reasonable fraction of that number makes you a big winner. For convenience, we order the  $n$  outcomes for each  $p$  from smallest to largest. For the raw (logarithmic) data, most of the action is seen to take place below the x-axis. That is, most of the nine thousand ( $9 \times n$ ) games return less than \$1 bet. In fact 62.8% of the games lose money.

Here is a plot of the entire simulation.



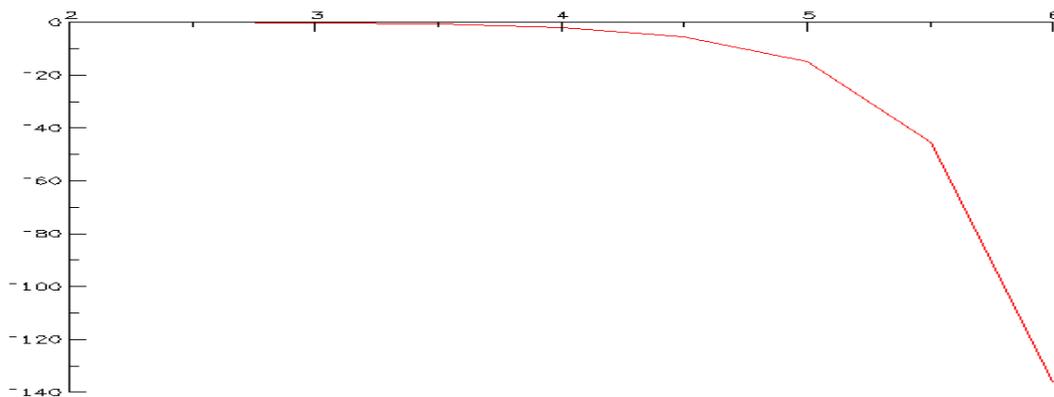
Logarithmic payoffs:  $n=1000$  players; nine values of  $p$ =number of pulls.  
 $p=100, 300, 1K$  (all dotted),  $3K, 10K, 30K, 100K, 300K$  and  $1M$  (solid) pulls.  
 Negative values denote losers; Positive=winners; Zero=break-even.

For all 9000 games, 63% lose money. Let's break that down by value of  $p$ . For an individual pull, of course, it's 50%. There are equal numbers (3) of payoffs above and below unity. That percentage falls below 50, however, after only the second pull. And for one million pulls, it drops to just 10%. But remember we only need one good payoff: Here's the complete story.



Percent of players that win vs.  $\log(2-6)$  of  $p$  = number of pulls (100-1,000,000). 100 - 48%; 1000 - 46%; 10,000 - 42%; 100,000 - 33%; 1,000,000 - 9.6%. With more pulls, fewer players win. The curve continues asymptotically to zero without ever reaching zero. However, the 1000 players always (collectively) win.

Another measure of success is the median payoff for each value of  $p$ . If we draw a vertical line at  $n=500$  on the first graph, where as many players did worse as did better, we get these values (which still seem pessimistic):

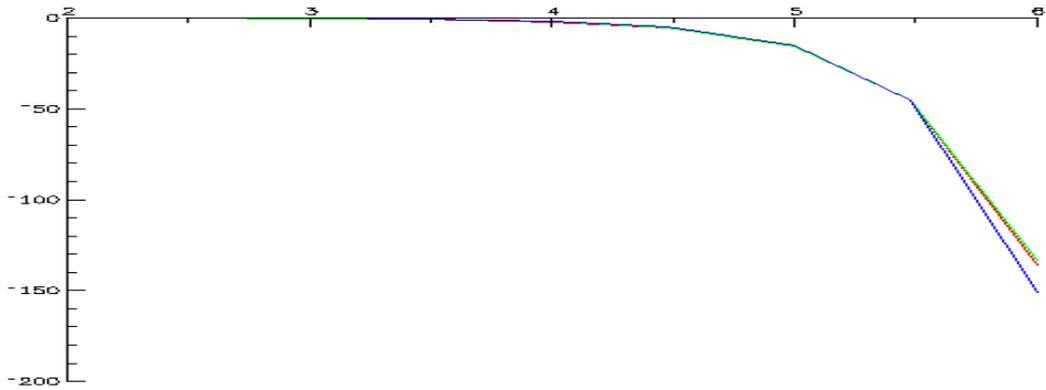


Median (logarithmic) payoff vs.  $\log(p)$ .  $p$  is 100 - 1,000,000. ( $\log p$  is 2 - 6.) Median payoff stays relatively close to unity for several orders of magnitude then falls drastically. For  $p=1000$  median payoff is  $\sim 0.56$  ( $\log = -0.24$ ); for  $p=1,000,000$  median payoff is  $\sim 10^{-136}$  ( $\log \sim -136$ ).

We can take the average (of logarithmic data) as well. It tracks the median very closely, telling us that logarithmic payoffs occur symmetrically above and below the median value. The logarithmic data have no skew.

We can also calculate the logarithm of the  $p^{\text{th}}$  power of the geometric mean of the six payoffs. We find that this also tracks the average and median values. Thus, we see how **GM** governs the results: it's logarithm dictates the median/average of the logarithmic winnings.

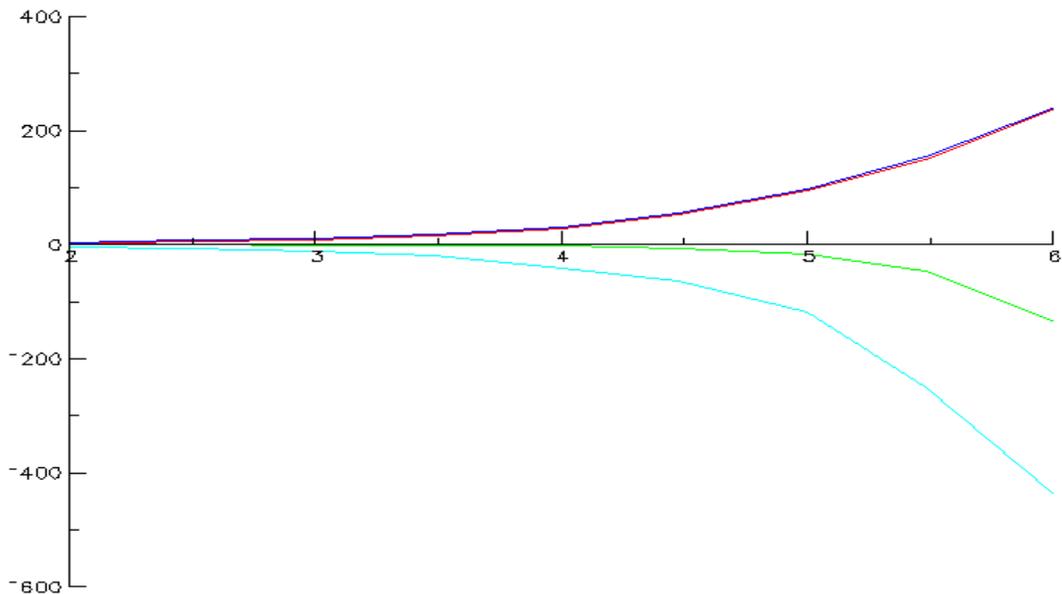
Here is a plot of all three values.



Average and Median of logarithmic results for  $p = 100 - 1,000,000$ . The curve in blue is  $\log(GM^p)$  shown for comparison.

What also can be seen from the raw data is that with increasing  $p$ , although fewer of the  $n$  players come out ahead individually, the winnings of the successful players increase enormously. So if the  $n$  players agree to share the winnings (like a group that buys a lotto ticket) they should play for the largest possible value of  $p$ . To make this intuitive, consider that if only one player wins, only once, and wins at least  $\$n$ , the group wins. That is, the each person will win at least  $(\text{Max payoff})/n$ .

So let's see what the payoffs look like as  $p$  increases. The next graph shows the max, min and average (logarithmic) payoff together with the log of the average (actual) payoff:



Logarithmic payoffs vs.  $\log(p)$  for  $n=1000$ . Outer curves are Max and Min; green is Median. Red curve is  $\log(\text{avg of actual payoffs})$ . The red curve lies very close to the maximum value. This is because the average payoff can't be less than  $\text{max payoff}/1000$ . So  $\log(\text{avg})$  is  $> \log(\text{max})-3$ .

We've seen that  $p$  affects the results enormously.

What role does  $n$  play? The curves are stationary with respect to  $n$ . This means that you only need enough tries (or players) to get at least one winning game. That number is not even close to 1000.

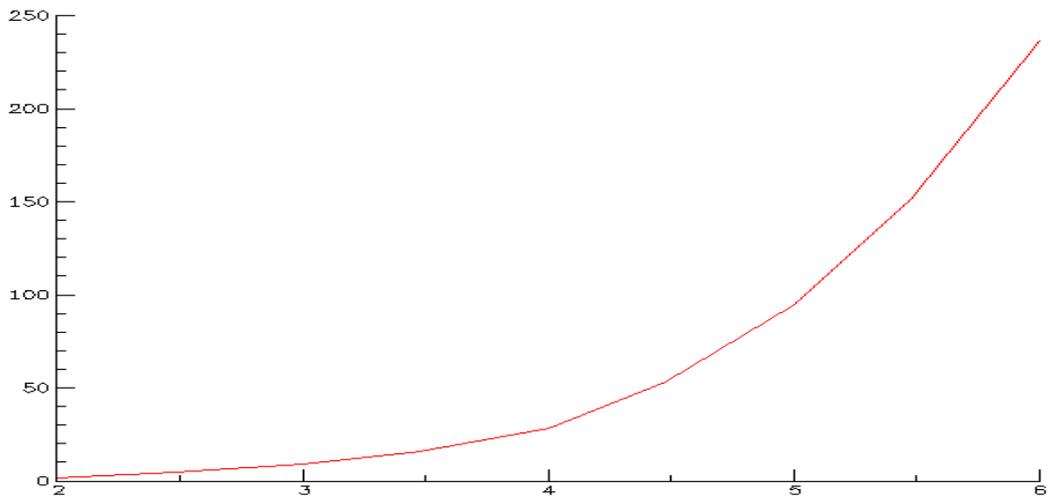
Inspection of the (log) payoff graph shows for each  $p$  a (log) payoff interval lying between the max and min curves. Perhaps 1/3 of that interval is above the line. Players would not populate that regions randomly (in the first graph, they would populate the x-axis uniformly): in fact 90% of them will lose for  $\log p = 6$ . But for a modest number, say  $n=20$ , you'd have a decent chance of getting one or two into the win column, and the best result is likely to be huge. Suppose its log is 50, out of a range of roughly -400 to +200. That's good enough for the following advice to be of use:

Get nineteen friends, find the game, and pull the handle a million times. Now!

At 3 seconds per pull, working eight hours a day, 5 days a week, in just 21 weeks you'd get say  $\$10^{50}$ . That's  $\$10^{44}$  per pull. Split it 10 ways, it's still  $\$10^{43}$  per pull. You could take a 21-week leave from your job for that.

Or take a more modest  $p=1000$  pulls, where  $n=10$  gets you into the big money at least once. With only nine friends, in less than one hour, you could collect \$2 billion. Divide that  $n=10$  ways and it still a cool \$200 million.

Here's the complete payoff picture.



Log of average payoff vs. log (2-6) of  $p$  = number of pulls (100-1,000,000). Average payoffs for 100, 1000, 10,000, 100,000 and 1 million pulls are 27.1,  $1.87 \times 10^9$ ,  $2.36 \times 10^{28}$ ,  $3.21 \times 10^{94}$  and  $2.67 \times 10^{236}$ , respectively.