ERGODICITY AND INVESTMENTS

Part 2

by John Schuyler

Ergodicity

Ole Peters is a well-credentialed Ph.D. physicist. He is a Fellow at the London Mathematical Laboratory and the Principal Investigator of its economics program. Peters is also an External Professor at the famed Santa Fe Institute.

Ergodicity is a concept from statistical mechanics that also applies to economics. The implicit assumption is that the EV of an observable value is the mean of the observed outcomes. Peters shows this is not always the case.

He demonstrated a person starting with $100 and playing a lottery of 100 chained flips, with the outcome of the prior flip applied to the next flip. This experiment appears ergodic in that:

- All 2^{100} (=1.27 x 10^{30}) flip sequences are possible. Sufficient random sampling will eventually experience all possible sequences.
- The player has an incredible EV ending value of $13,150 for the $100 investment.

However,

- **The person almost surely loses most of his original investment.**
- This is because, in that EV calculation, there are very, very small probabilities of extremely large gains. A simulation, even with a fast computer, will almost never realize one of the extreme gains important to the *EV*.

Peters and others view non-ergodicity as a possible explanation of—or a contribution to—worldwide wealth inequality.

The Coin Flip Bets

Recall the bet or gamble from Part 1. Call it an investment, if you prefer. The coin flip has a 0.5 chance of Heads = Success. If successful, your bet value increases by 50%. Otherwise, it fails, and your value falls by 40%.

Assume that you can *chain* similar bets, with the next bet amount the adjusted stake value after the prior bet.

Probability Tree Examples

We assume an *EMV* decision policy for the thought experiments.

- PV discounting is unimportant because a bet outcome is realized immediately after a flip.
- Risk aversion isn't important until the stake or account amounts become significant, positive or negative. We ignore risk aversion for this experiment.

There is more about risk aversion later, with a more significant starting bet amount.

Single-bet

Consider a $100 bet on a single coin flip. Heads (the coin lands face up) gains you $50, and Tails (lands face down) loses you $40.

Here is the decision tree:

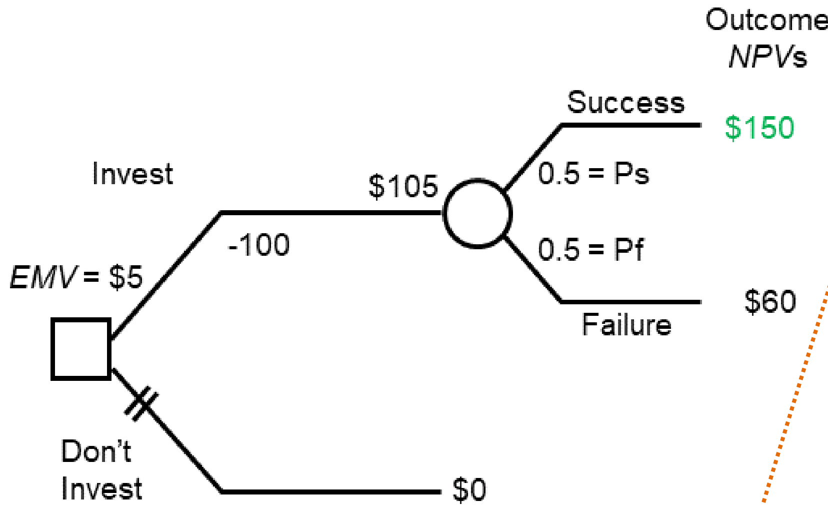

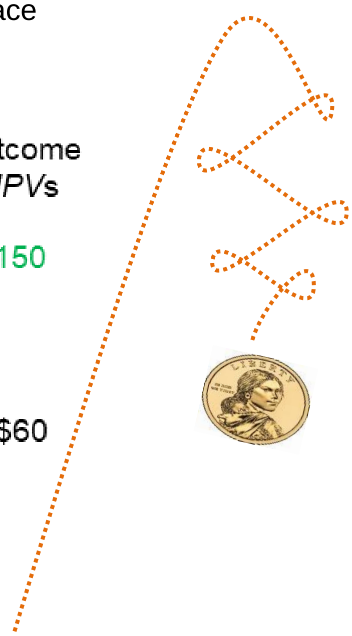

Would you accept this bet?

The *EMV* is $5. According to the *EMV* decision policy,[1] The gain of Success is greater than the loss of Failure. Invest (bet) is a *good decision* unless you are extraordinarily risk-averse and cannot afford to lose $40.

Following are 2-, 3-, and 100-flip sequence examples. In each:

- The initial bet is $100.
- With an even number of bets, the most likely flips outcome is the EV number of winning flips and the EV number of losing flips.
- As chained, the balance after a flip becomes the bet for the next flip.

---

[1] *EMV* is *expected monetary value*, the probability-weighed outcome value. The outcome value is usually a *net present value* (*NPV*), but this example ignores time value. The decision policy is to choose the alternative with the greatest *EMV*.

Two-flip Example

Consider an experiment of flipping a coin twice, where the second flip starts with the outcome value from the first. You start with $100.

Here is the probability tree:

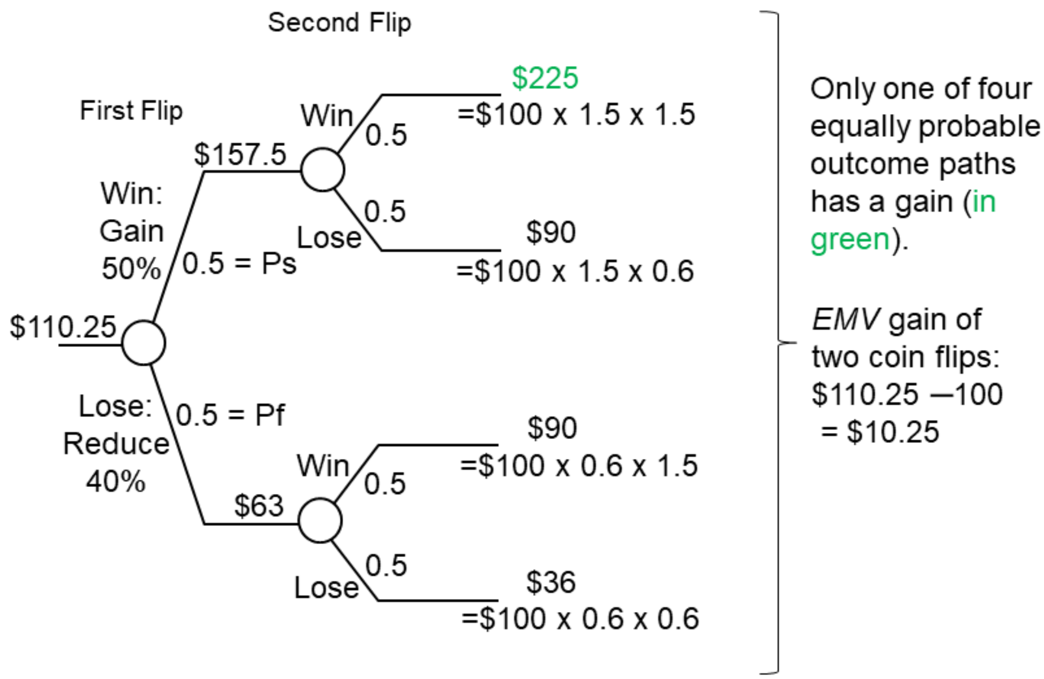

- The profit contribution of an *n*-flip path sequence is:

  Gain = $\$100 \times 1.5^{n} \times 0.6^{2-n} - \$100$

  where *n* is the number of winning flips.

- Rolling back (a.k.a. back-solving) the probability tree gives *EV* end balance =$110.25 at the root node. So, betting on a two-coin-flip strategy appears to be a good investment. On average, bets of this kind would provide you with an average gain of $10.25. However, also note that 3/4 of the time—the bottom three paths through the tree—you will lose value.

- This and the next experiments show with chained flips:
  - *EMV* is positive for every flip.
  - This counterintuitive behavior: The *base case* is the outcome value with the EV number of wins. One win and one loss. As illustrated, this results in a $10 loss.

I call the difference, *stochastic variance* (*SV*). This is a key component of *variance analysis*, detailing the difference between a) the *stochastic* (*probabilistic*) forecast *EMV* used for project or plan approval and b) the base case *NPV* estimate from a conventional deterministic model with EV inputs.

Three-Flip Example

Expanding the probability tree to three successive gambles produces a tree of $2^3 = 8$ paths. This gamble may appeal more than the 2-flip tree because there is a higher probability of winning (0.5, four of the eight outcomes) and a higher *EMV*.

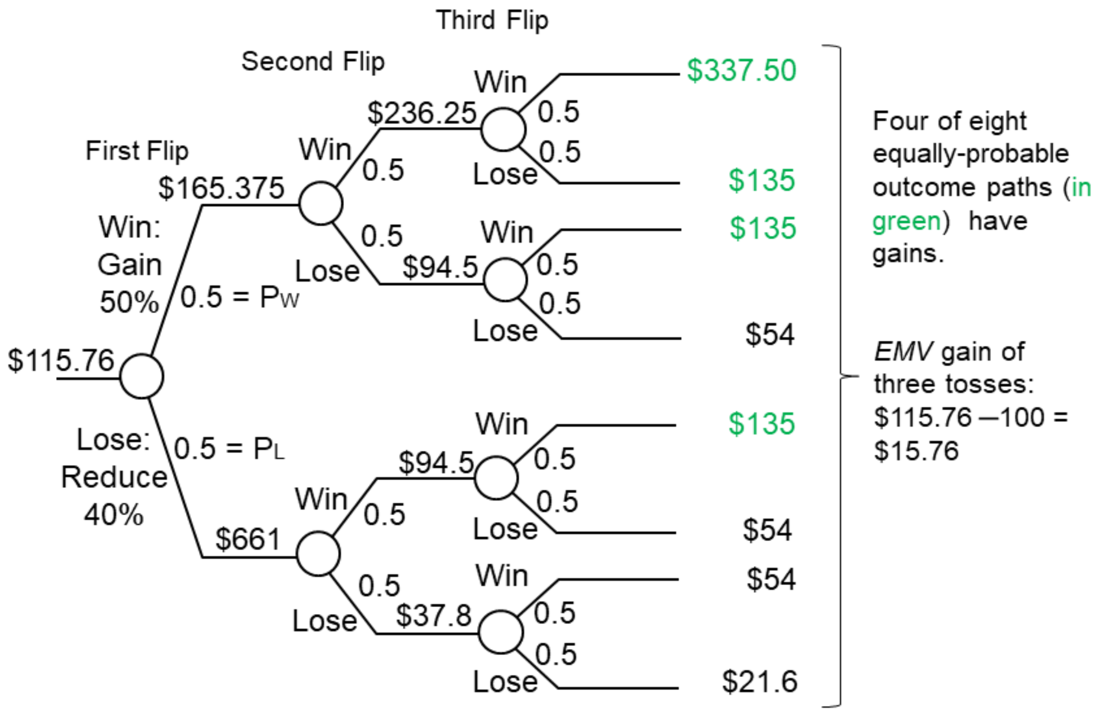

A 100- Flip Gamble

Let's now look at the more dramatic result of an experiment of flipping a fair coin 100 times with the same win-and-lose factors as before. This is more like the example by Ole Peters.

Again, suppose that you start with $100. The amount after every flip carries to the next flip. The number of heads can be any integer between 0 and 100. The most likely and the EV of wins (heads) is 50. The next page shows this outcome is worth $13,150.

Now, there are $2^{100} \cong 1.27\times10^{30}$ possible sequences in the 100 coin flips. Fully drawing and solving the probability tree is impossible.

The sequence of heads and tails does not affect the outcome, only the aggregate numbers of wins and losses.

A key idea in ergodicity economics is that a reasonable number of value trajectories (trials), say several thousand, can represent a system's properties. That is the foundation idea behind Monte Carlo simulation (MCS or simulation): the sample mean of many trials approaches the true mean (*EV*). This works for most models. **However, the sample mean does not work here.**

> This simple example helps explain. In 100 flips, the EV number of Heads or Tails is 50. Assume the investment is $100. If the flips—in any sequence—are 50 Heads and 50 Tails (the most likely outcome), the end value shows the investor losing almost all of her money:

Base case:

Deterministic case with EV number of wins (Heads) and losses (Tails):
End Value = $\$100 \times 1.5^{50} \times 0.6^{50} = \$100 \times 6.3762\times10^{8} \times 8.0828\times10^{-12} = \$0.51$

Yet, the probability-weighted outcome is:

Stochastic case:
EV End Value = $\$100 \times (0.5 \times 1.5 + 0.5 \times 0.6)^{100} = \$100 \times 1.05^{100} = \$13{,}150$

That is quite a difference!

Complex calculations often require MCS and many trials. This simple experiment is easy to model with MCS. However, the solution would take many human lifetimes to converge adequately.[2] Fortunately, there is an easy and exact solution method.

Binomial distributions represent coin flips or other two-outcome experiments with independence and a constant probability of Success on each try. Excel has a function for this. Here is an example calculating the probability of 50 successes in 100 tries when each try has a 0.5 chance of Success:

`=BINOM.DIST(50,100,0.5,0)` = 0.079589

50 Heads is the mode and most probable outcome.

The calculations are exact. Below is a table of potential outcomes (with some rows omitted for compactness). Computer programs don't easily produce subscripts and superscripts. A convention is to use "e"s for the scientific notation exponent. For example, the probability of all wins is 7.889×10^{-31} = 7.889e-31.

[2] Someday, a quantum computer may be suited to the task.

```
EMV = $13,050                                         amounts are $

                                       |...... for the Flips Portfolio ......|
Wins Losses  Probabilty   Cum Prob      End NPV   Contrib EV       Cum EV
  0    100   7.889e-31   7.889e-31   6.533e-21    5.154e-51    5.154e-51
  5     95   5.939e-23   6.262e-23   6.380e-19    3.789e-41    3.869e-41
 10     90   1.366e-17   1.532e-17   6.231e-17    8.508e-34    8.897e-34
 20     80   4.228e-10   5.580e-10   5.942e-13    2.512e-22    2.786e-22
 30     70   2.317e-05   3.925e-05   5.667e-09    1.313e-13    1.577e-13
 40     60     0.01084     0.02844   5.404e-05     5.86e-07    7.905e-07
 50     50     0.07959      0.5398      0.5154      0.04102       0.0665
 55     45     0.04847      0.8644       50.33        2.440        4.547
 56     44     0.03895      0.9033      125.82        4.901        9.448
 60     40     0.01084      0.9824        4915        53.30        120.9
 70     30   2.317e-05   0.9999839    4.687e+7       1086.1         5427
 71     29   9.790e-06   0.9999937    1.172e+8       1147.3         6575
 72     28   3.943e-06   0.9999977    2.938e+8       1155.2         7730
 73     27   1.512e-06   0.9999992    7.324e+8       1107.8         8838
 74     26   5.519e-07   0.9999997    1.831e+9       1010.5         9848
 80     20   4.228e-10   0.9999999   4.470e+11       189.01        12896
 90     10   1.366e-17           1   4.263e+15      0.05821        13150
 95      5   5.939e-23           1   4.163e+17    2.473e-05        13150
 97      3   1.276e-25           1   2.602e+18    3.319e-07        13150
 98      2   3.905e-27           1   6.505e+18    2.540e-08        13150
 99      1   7.889e-29           1   1.626e+19    1.283e-09        13150
100      0   7.889e-31           1   4.066e+19    3.207e-11        13150

                EMV = EV end value $13,150 - Investment $100 = $13,050
```

A chart of all possible outcomes is two pages below.

Notes:

- Abbreviations in the table:

CumProb	cumulative probability through the row's number of flips
Contrib EV	portion of *EV* end value contributed by this row's outcome
Cum EV	cumulative *EV* end value through the row's number of flips

- Fun fact: Starting with $100, the end value **with 100 successes** is $\$4.07\times10^{22}$ (that's 41 sextillion). This is almost **100 million times all the money in the world!** (total global wealth = $454T = $\$4.54\times10^{14}$).[3]
- MATLAB® computing platform (a product of The Mathworks®, Inc.) performed the calculations (64-bit double-precision) and charts.
- The $13,050 *EMV* is 2,610 times the $5 EV gain on the initial $100 bet. By this metric, the 100-flip lottery seems an incredible investment.

  However …
- Referring to the 55 wins row: With 100 flips, this lottery requires 56 or more wins (heads) to profit from the $100 initial bet amount. There is an 86.4% chance of 55 or fewer wins and losing value.

[3] The best and latest source that I found for total world wealth: *Global Wealth Report 2023*, Credit Suisse https://www.credit-suisse.com/about-us/en/reports-research/global-wealth-report.html

- Tip-offs suggesting non-ergodic behavior are these extreme statistics:

`EV =`	**`$13,150`**	mean
`SD =`	**`$60,328,925`**	standard deviation
`CV =`	**`SD/EV = 463,358`**	coefficient of variation
`Skewness =`	**`119,388,967`**	asymmetry measure (long right tail)
`Kurtosis =`	**`2.04e+18`**	peakedness measure (sharp peak)

These values are so extreme that they may be meaningless in their usual context.

This chart represents key values in the large table of possible outcomes.

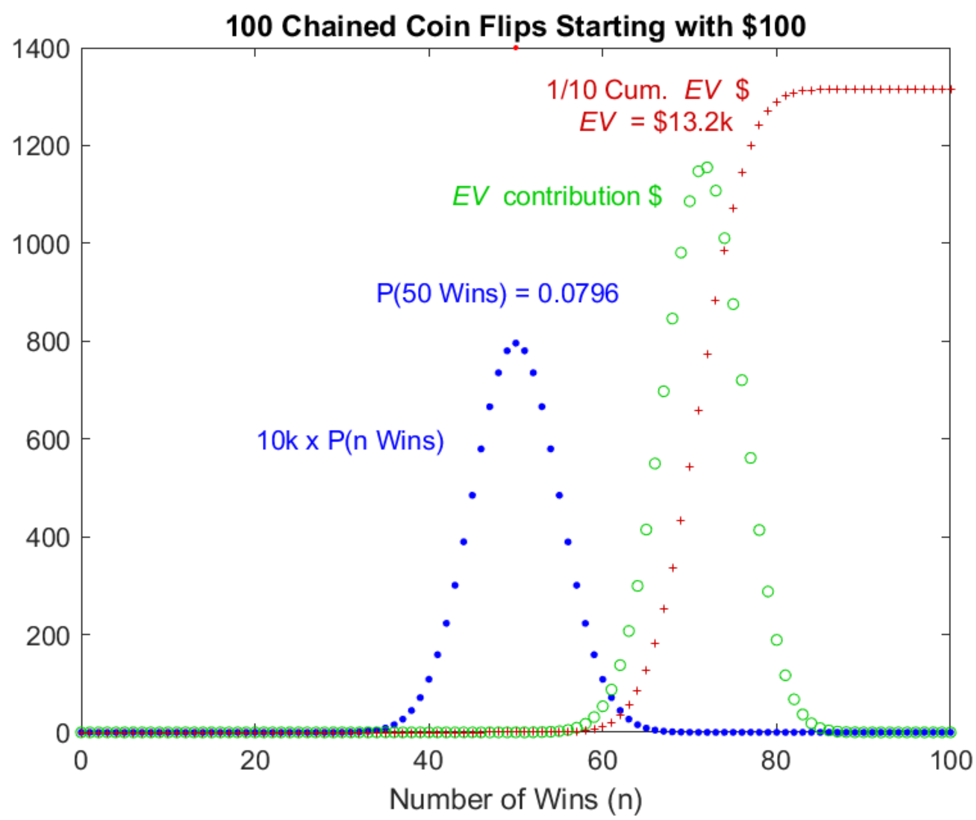

Coin experiment with 100 flips. Outcome values contributing to *EV* (middle, green distribution) do not become significant until about 59 wins (upper-tail of the blue Number of Wins (left, blue) distribution. Almost all *EMV* contributions are from trials with flips 60 to 80. Most trials will lose most of the bet amount, and a rare few (≥80 wins) will have huge payoffs.

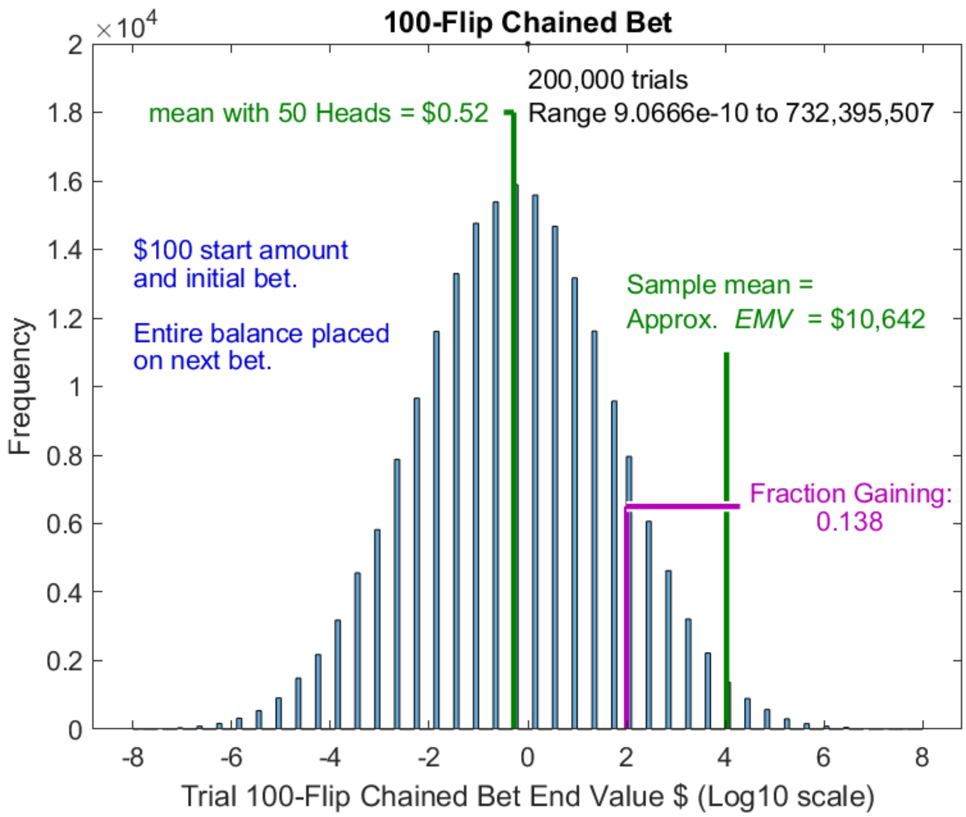

Includepicture “c:\\mfiles\\ergodicity\\ErgoMCSchained.png” *mergeformat

Here is a chart of the lottery distribution produced with MCS.

- Though the produced chart is a frequency distribution, the discrete outcomes are evident. Therefore, it is actually a *probability mass function* (PMF).
- The most likely outcome is 50 Heads and 50 Tails, and this outcome value is only $0.52.
- Betting $100 on a 100 flips sequence would produce an end value of $0.52 for the average person.
- Extreme outcomes are far off-scale. The ending value with 100 Tails is $6.533e-21, and 100 Heads is $4.066e19.
- Even with 200,000 trials, the $10,642 sample mean substantially differs from the true $13,150 mean.

Value Trajectories

Ergodic means encompassing all possibilities. A non-ergodic view represents outcomes that we can reasonably observe in real life.

Next, I’ve replicated a rather convincing chart from Peters (2019):

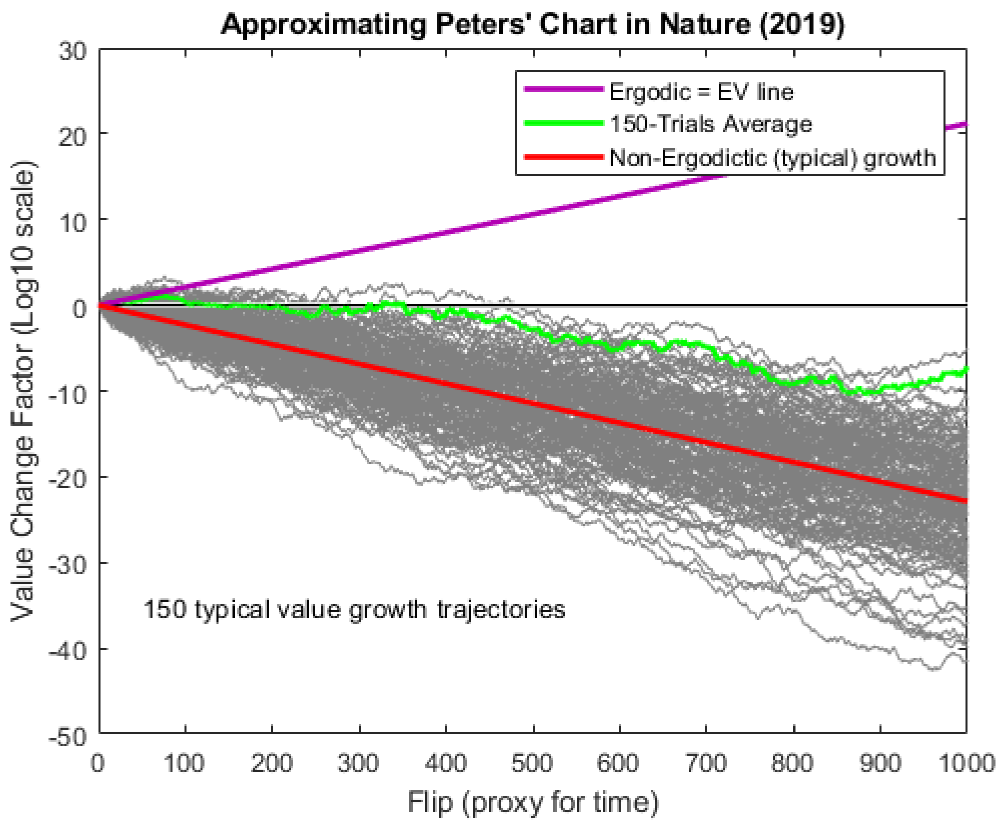

Typical results are far from the ergodic (expected value) projection. The typical trial trajectory will be somewhere in the gray cloud. The green trace is the average of one run's 150 trials (note the log scale): several better trajectories dominate the EV trajectory. The cloud appears centered around the red, non-ergodic line. Yet, the EV of all possible trajectories (ergodic) is the purple (topmost) *EV* line.

Notes:

- In this experiment, the initial bet amount is only $1, which is the horizontal reference line on the chart at $\$10^0$.
- The bet is a succession of 1000 chained flips (The large number of flips is for dramatic effect).
- The probability of losing value, at 557 or fewer wins, is 0.99864 (calculated precisely using the binomial distribution).
- The upward-sloping, purple, **ergodic** *EV* line is the unbiased forecast. This presumes, at each flip, the value grows by a factor of $0.5 \times 1.5 + 0.5 \times 0.6 = 1.05$. Or, after *n* flips, the *EV* is a point on that line.

  For example, the EV after 1000 flips is $1.05^{1000} = 1.546 \times 10^{21}$, the purple line value at the right edge of the chart.

  The problem is that the high-value trajectories that make the purple line the mean are extremely rare.

If the simulation runs for an astronomical number of trials (such as 100 / the probability of the best outcome = 100 / 0.5^{1000} = 1.07×10^{303}), we can expect the average value growth (green trajectory) to be close to the EV end-value line (purple).[4]

- However, for any reasonable number of trials, virtually all value trajectories will be somewhere in a cloud, approximately centered about the red, decreasing, **non-ergodic line**.
- Most of the *EV* is from cases where the number of wins is in the 65–78 range. Those outcomes occur, on average, in only 0.176% of 100 coin flips.
  - In the presented 10,000-trial by 150-coin flips run, the average (green) trajectory is near the top of the gray band. The top two or three trajectories dominate the average (note the log *y*-axis scale),
  - Each flip is a marginal bet. Each has an EV payoff of $.05 for every $1 invested. The trajectories cloud and average move quickly upward with a better probability of success or payoff ratio amounts.

Recap

Blaise Pascal introduced the expected value concept in 1654. Do these charts and values "**upend three centuries of economic thought**"? For the risk-neutral decision maker without a capital constraint, any positive *EMV* investment is worthwhile. Ergodicity explains a situation where the *EMV* decision policy can fail.

They do demonstrate the need for modeling with care:

- Objectively assess risks and possible outcomes.
- Size your bets based on your risk tolerance. Use your personal utility function if you have one.
- Beware of chains of modest investments. Instead, choose fewer, larger, and more attractive investments.

What Would John Do?

The probability tree for the 3-flip experiment seems appealing. It has a 50% chance of profit. And the EV gain is over 15% of the start amount:

$$\left(\frac{1.5+0.6}{2}\right)^3 = 1.1576$$

Ergodicity condemns strategies with many chained events, so there should be a constrained number of flips.

After some thought, I settled on this strategy:

1. Start the bet (investment) with, say, 3% of net worth (*NW*). The 3% reflects a degree risk tolerance, though this experiment isn't based on optimizing expected utility (*EU*). Both the *EMV* and *EU* decision rules would have the person continuing the bet without end.

---

[4] If a trial solution takes a microsecond of computer time, a run to get a reasonable EV approximation would require about 2.4×10^{279} times the 14 billion years age of the universe.

Flip the coin.

If heads, adjust *NW* by a factor of 0.97 + .03x1.5 = 1.015 (50% gain on the bet)
If tails, adjust *NW* by a factor of 0.97 + .03x0.6 = 0.988 (40% loss on the bet)

2. Take another flip with the revised *NW*, again betting 3% of updated *NW*.

   If heads, adjust *NW* by a factor of 0.97 + .03x1.5 = 1.015 (same as above)
   If tails, adjust *NW* by a factor of 0.97 + .03x0.6 = 0.988 (same as above

3. Stop betting after two successive Tail coin flips, else return to Step 2.

Here are 20 example trial flip sequences, out of 10,000 trials, each starting with $15,000.

Successes (Heads) are +1s and Failure (Tails) are -1s.
Each row is a trial ending with two Tails (-1, -1):

-1 +1 +1 -1 +1 +1 -1 -1	= $9,842 after 8 flips
+1 -1 -1	= $8,100 after 3 flips
+1 +1 +1 +1 -1 -1	= $27,338 after 6 flips
+1 -1 +1 +1 -1 -1	= $10,935 after 6 flips
+1 +1 +1 +1 +1 -1 +1 -1 -1	= $36,906 after 9 flips
+1 +1 +1 -1 -1	= $18,225 after 5 flips
+1 -1 -1	= $8,100 after 3 flips
-1 -1	= $5,400 after 2 flips
-1 +1 +1 -1 -1	= $7,290 after 5 flips
+1 +1 -1 -1	= $12,150 after 4 flips
-1 +1 -1 +1 +1 +1 +1 +1 -1 +1 -1 -1	= $19,929 after 12 flips
+1 -1 -1	= $8,100 after 3 flips
-1 +1 +1 +1 +1 +1 +1 -1 -1	= $36,906 after 9 flips
-1 -1	= $5,400 after 2 flips
+1 +1 +1 -1 +1 -1 -1	= $16,403 after 7 flips
+1 +1 -1 -1	= $12,150 after 4 flips
+1 +1 +1 +1 -1 +1 -1 +1 -1 -1	= $22,143 after +10 flips
+1 +1 -1 +1 +1 -1 -1	= $16,403 after 7 flips
+1 +1 +1 +1 +1 -1 +1 +1 +1 -1 +1 +1 +1 -1 -1	= $168,151 after 15 flips
+1 -1 +1 -1 -1	= $7,290 after 5 flips

The mean ending value was $29,062, almost twice the initial $15,000 stake. Yet most trials lost value. Only 25.5% of trials had a profit. This was a typical run, though different starting seeds and number of trials produced substantially different results.

The following two charts show the distributions and key statistics for the *highly-skewed* outcome distribution.

The dotted green, dashed blue, and solid red lines indicate the three most popular central measures: mode, median, and mean.

In most trials, losses consume almost all the initial bet funding. The few large outcomes among the 25.5% of profit cases make the strategy interesting and worthwhile.

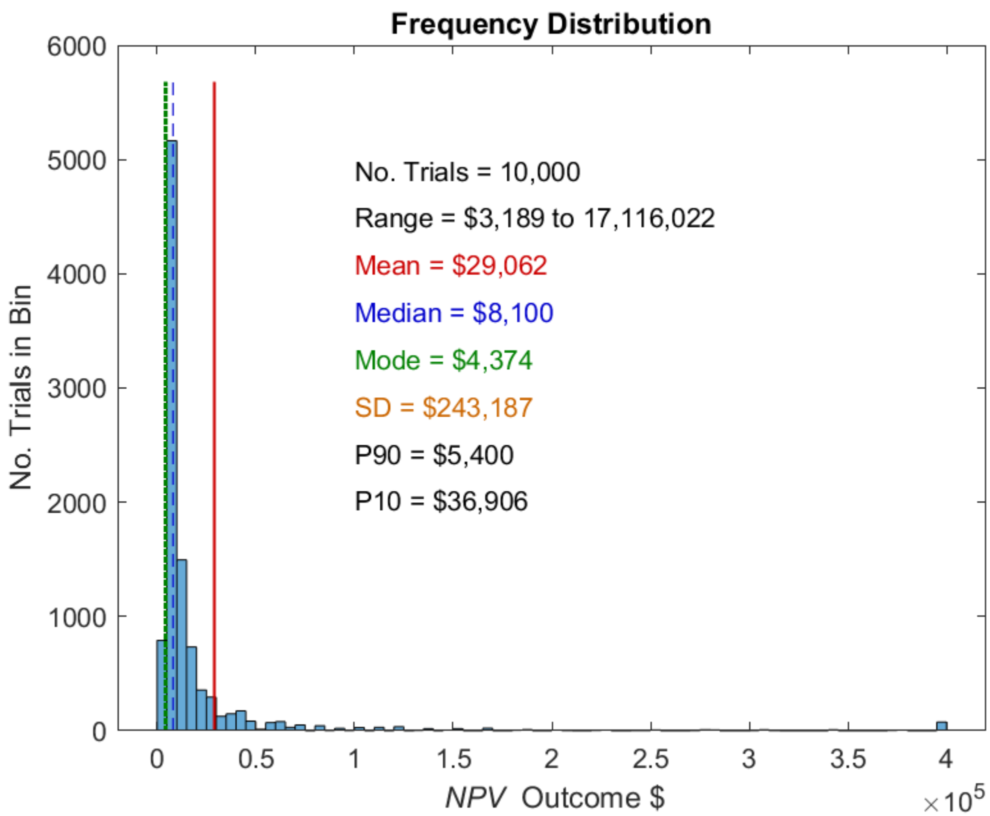

Includepicture “c:\\mfiles\\ergodicity\\ErgoJohnsBetFD.png” *mergeformat

Note: The little bump bar at $400k represents capping—for the graphic only—the few outcomes greater than $400k.

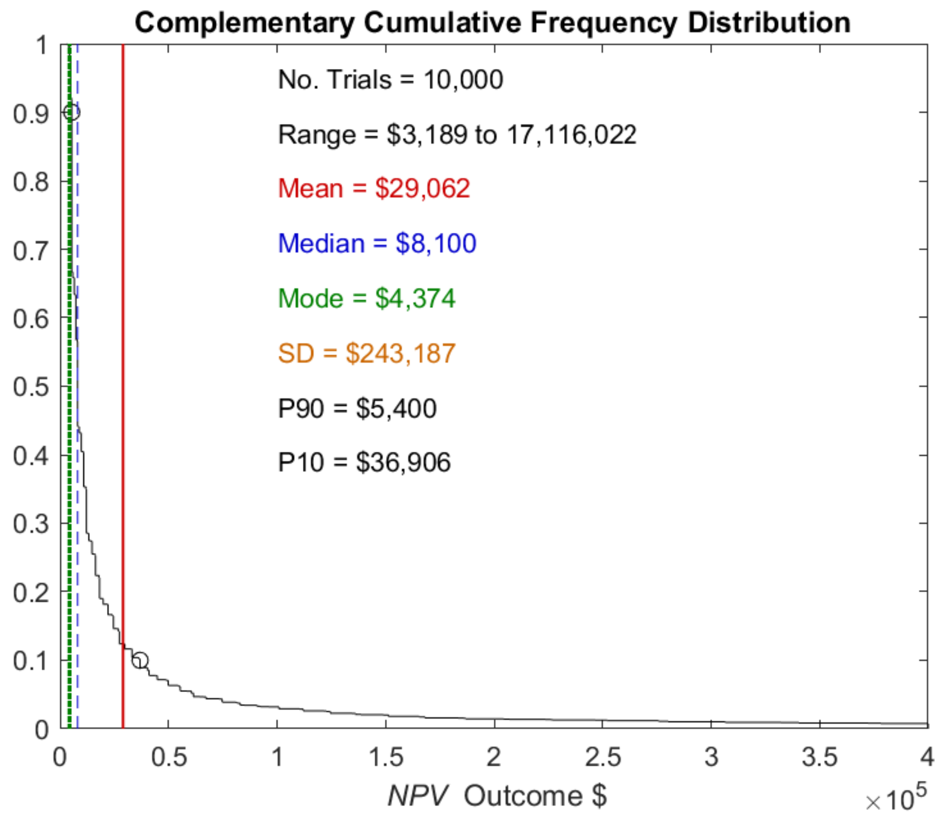

Includepicture "c:\\mfiles\\ergodicity\\ErgoJohnsBetCCFD.png" *mergeformat

John's bet strategy may be too risky for some readers. Despite that:

- EV Profit of $29,062 – 15,000 = $14,062

The solution is highly positively skewed (longer tail on the right):

- There is only a 25.5% chance of a profit.
- The median loss (*P50*) is $15,000 – 8,100 = $6,900.
- The most likely (mode) loss is $15,000 – 4,374 = $10,626.

For comparison, we are comparing the lottery to the alternative of holding the $15,000 in cash.

- The *P10* gain = $36,906 – 15,000 = $21,906.
- The high value realized in 10k trials was $211,015.

  The 100-Heads theoretical maximum is $15,000 x 1.5^{100} = $6.1 x 10^{21} with a probability of 7.9 x 10^{-31}

What can we learn about ergodicity and this simulation?

1. A common ranking metric for portfolio planning is the *profitability index*:

$$PI = \frac{EMV}{Investment} = \frac{.05}{1}$$

   Even positive *EMV* decisions can lead to ruin, especially if the investments have low *profitability indices* (*PI*s) and many reinvestments:

   a. The portfolio distribution is highly skewed.
   b. A loss is most likely.
   c. Portfolio success with many flips will be rare.

2. You will improve your portfolio and reduce the effect of ergodicity by:
   a. Finding investment opportunities with better *PI* s.
   b. Look for projects with higher probabilities of Success (Ps's)
   c. Avoid assembling a portfolio of projects with high correlations.

Links to References

https://www.bloomberg.com/news/articles/2020-12-11/everything-we-ve-learned-about-modern-economic-theory-is-wrong *Bloomberg News* (December 11, 2020)

https://doi.org/10.1038/s41567-019-0732-0 The source article: Peters, Ole, 2019, *Nature Physics*, **15** , 1216-1221.

https://taylorpearson.me/ergodicity/ includes a thought experiment by Nassim Taleb

https://en.wikipedia.org/wiki/Ergodicity

https://en.wikipedia.org/wiki/Ergodicity_economics

https://jasoncollins.blog/ergodicity-economics-a-primer/

Total money in the world (2020 report, as of end-2019)
https://www.credit-suisse.com/about-us/en/reports-research/global-wealth-report.html

■

Disclaimer: This note is not investment advice. Though the concept is general, it may not apply to your circumstance.

There is no representation, warranty, or guarantee of accuracy, adequacy, applicability, completeness, reliability, or suitability of the information to a particular situation.