

Multiplicative Gains, Non-ergodic Utility, and the Just One More Paradox (with Supplemental Information)

John K. Myers (✉ john@scrier.org)

Space Coast Research Institute of Economic Reasoning

Research Article

Keywords: multiplicative gains, arithmetic gains, game theory, ergodicity, utility, expectation, Hero Or Heroin, American Roulette, Just One More paradox

Posted Date: February 22nd, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-237495/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Audience: General technical, economic, policy-making

Multiplicative Gains, Non-ergodic Utility, and the Just One More Paradox

John K. Myers *john (at) scrier.org*

Abstract. Interest in multiplicative vs. additive returns on bets has been revived by Peters¹, who proposes ergodicity and added noise are useful in understanding utility preferences. Peters requires a Monte Carlo simulation to demonstrate empirically a supposed paradox that arithmetic expectation is inappropriate for multiplicative gain distribution forecasting. Here I formalize the r operator notation, which significantly simplifies multiplicative problems, as an extension of the arithmetic group's Δ/d discrete and continuous operators into the multiplicative semigroup. I show how the annihilating (absorbing) element of the multiplicative semigroup at 0, not $\pm\infty$, may be used to conveniently represent nonlinear sequence occurrences, such as running out of money, without the need for special computer rules outside the mathematics. I use this to solve Peters' expected-value paradox elegantly, without ergodicities nor noise. But Peters misses the real paradox, called "Just One More": the outcome of an advantageous additive gamble is *identical* to the outcome of a similar disadvantageous multiplicative gamble, after one trial; hence, by induction, an agent will keep playing. I propose games "Hero or Heroin" and "American Roulette" to highlight this paradox. This may help in explaining addiction. The Supplement contains further visualizations and arguments against the need and applicability of ergodics for utility. The results contribute to the understanding of repeated multiplicative gambles with annihilating states, and their utility.

Summary. The results of betting a wealth bankroll, say \$100, on something as simple as a fair coin-flip, and then receiving, say, either additive +\$50 for heads / -\$40 for tails, or multiplicative +50% for heads / -40% for tails, can have subtle implications. In particular, repeated additive payoffs of this same bet result in larger and more probable wins, whereas repeated multiplicative payoffs, using the current bankroll, surprisingly result in larger and more probable losses; but much higher lottery payoffs for speculators. A major decision is whether to take a bet or not, based on the usefulness or *utility* of possible outcomes. Peters argues that *ergodicity*—roughly, path independence, so that any one random path through the system is representative of the whole—is needed to determine choices, and that psychology-based utility theory can't work. Without proper operators, computations on multiplicative versus additive gains may be difficult to understand. Here I show how to work with multiplicative gains without using any ergodics nor noise. First I formalize a useful notation for multiplicative gains on series by using groups. I explore annihilating states (representing insolvency, or death), which destroy ergodic assumptions. And I present a new paradox, highlighting the difference between incremental additive gains and multiplicative series, along with two new games illustrating preference results different from those assumed by the community. The results should help clarify multiplicative gains, such as are found in financial markets and military contests, and may be useful in understanding addictive behavior. In particular, it's anticipated that nuclear safety and economic cooperation calculations may benefit from these explorations.

The *r* operator

Everyone's familiar with the prefix operator Δx for discrete x , and dx for continuous x . This provides a handy and intuitive notation for additive increments. But what if you wish to talk about multiplicative increments? Does there exist a similar handy notation? And why would this be useful?

Many authors have used stand-alone variables, such as α_s^2 , o_i^3 , x_i^4 , or $(1+b)^5$, to derive Kelly's formula or the information/entropy of multiplicative outcomes. Here I contribute by formalizing multiplicative increment as an operator, and showing how it derives from a group extension. Interesting applications follow in succeeding paragraphs, so feel free to skip the derivation.

We dip briefly into groups. Real numbers or integer numbers, taken together with our familiar addition, I'll call the **additive group**. Real numbers or integer numbers, taken together with our familiar multiplication, I'll call the **multiplicative semigroup**. (It's not a "group", because including the 0 element causes unusual effects, as we'll soon see.)

So instead of Δ or d , let's use roman *r* for the **multiplicative increment operator** for discrete series, and italic *r* for continuous.

Thus, *r/r* in the multiplicative semigroup is extended from Δ/d in the additive group.

We recall that additive returns on bets are best expressed with the familiar delta notation. Given a series $\{x[0], x[1], x[2]...x[T]\}$ as a function of discrete time T , then the delta prefix operator comes from the inverse of the addition operator over the series,

$$\Delta x \stackrel{\text{def}}{=} (x[T+1] - x[T]) \quad (1)$$

and its continuous form, from continuous time t ,

$$dx \stackrel{\text{def}}{=} (x(t+\epsilon) - x(t)) \quad (2)$$

Then the *r* operator (for *ratio*) is similarly defined from the inverse of the multiplication operator,

$$r_x \stackrel{\text{def}}{=} (x[T+1] / x[T]) \quad (3)$$

or its continuous form, using italics *r*, from continuous time t ,

$$r_x \stackrel{\text{def}}{=} (x(t+\epsilon) / x(t)) \quad (4)$$

Although it would be nice to be able to write prefix r_x as rx , similar to dx or Δx , this would get confusing when we postmultiply by x . And writing $r(x)$, to emphasize that it's a unary operator, is usually too much, especially for concatenation. So the x glyph has to be lowered, subscripted, or made implicit.

These definitions can be rewritten, using subscripted stand-alone variables, as generative formulae

$$x[T+1] = \Delta x[T] + x[T] = \Delta_T + x[T] \quad (5a)$$

$$x(t+\epsilon) = dx + x(t) \quad (5b)$$

or

$$x[T+1] = r_x[T] \cdot x[T] = r_T \cdot x[T] = r_T x_T \quad (6a)$$

$$x(t+\epsilon) = r_x \cdot x(t) = r_t \cdot x(t) = r_t x \quad (6b)$$

which, with left-hand sides equal, are showing that

$$r_T = (1 + \Delta_T/x[T]) \quad (7)$$

or, in the continuous form,

$$r_t = (1 + dx/x(t)). \quad (8)$$

Thus, any Δx may be converted into its equivalent rx , assuming a non-zero starting point x ; and any finite rx can be converted into its equivalent Δx . Note that a negative Δx that's less than x will still result in a positive rx ; if you lose half your money, you've just multiplied times a positive $+0.5$.

(But what about a zero starting point x ? Multiplicative growth assumes once you reach x as true 0, you always stay at 0. You've been kicked out of the game. So successive growth rates after $x = 0$ don't matter; they may be classed as undefined, indeterminate, or don't care.)

We recall that rate of arithmetic growth, in discrete or continuous measures, is the familiar

$$x' \stackrel{\text{def}}{=} \Delta x / \Delta t \quad \text{or} \quad x' \stackrel{\text{def}}{=} dx / dt. \quad (9)$$

Then the rate of geometric growth, in discrete or continuous measures, will be the corresponding

$$x' \stackrel{\text{def}}{=} r_x^{1/\Delta t} = r^{1/\Delta t} \quad \text{or} \quad x' \stackrel{\text{def}}{=} r_x^{1/dt} = r^{1/dt}. \quad (10)$$

being glossed as a superscript close-single-quote. The additive straight division by dt turns into an exponential division by dt .

This has the potential for forming a new calculus based on the multiplicative semigroup instead of the additive group.

Thus, in business calculations, the familiar expression for compounded-interest future value as a function of present value, interest rate i , and period n

$$FV = PV (1 + i)^n \quad (11)$$

can be more clearly written

$$FV = PV (r)^n. \quad (12)$$

with $r = (1 + i)$.

Note that, since stand-alone r and r are simply factors, they both commute and associate. They also concatenate. So,

$$x[n] = r_0 r_1 \dots r_{n-1} x[0] = r_{n-1} \dots r_1 r_0 x[0], \quad \text{and} \quad (13)$$

$$x'_{0 \rightarrow n} = (r_0 r_1 \dots r_{n-1})^{1/n} = (r_{n-1} \dots r_1 r_0)^{1/n} \quad (14)$$

Note of course that $x[0] \rightarrow x[n]$ requires n increments, but has $(n+1)$ x terms, because of fenceposts.

Annihilating state $r = 0$ can model permanently losing

In the examples in the Supplements, I have to use a special-exception computer rule to handle cases when the player has run out of money in additive cases and can't play the game any more. But what if there were a way to represent this automatically with r in the multiplicative case?

Defunct people can't play the game. Since any delta can be converted into an r , some may use this as an excuse to say that multiplicative gains can therefore be bound by the rules of additive gains. But here's the exception: The theory introduces an absorbing state at a bankroll of 0, representing insolvency or death; this is the multiplicative semigroup's *annihilating element* (absorbing element).

Since the *additive* group's annihilating element occurs at \pm infinity, its effects have been masked, and have not yet been included in decision theory. The multiplicative annihilating element at 0 introduces a realistic nonlinearity into the progressions, vaguely similar to the way a ReLU function⁶ introduces realistic nonlinearities into neural nets; see catastrophe theory⁷ for better nonlinearities.

Peters' expected-value paradox solved by inspection using r , without ergodicities

Peters proposes a paradox with a fair coin toss either winning 50% of one's wealth on heads (H), or losing 40% of one's wealth on tails (T). He represents this as (his eq.(2)):

$$\Delta x = \begin{cases} \Delta x_H = +0.5x, & p_H = 1/2 \\ \Delta x_T = -0.4x, & p_T = 1/2 \end{cases} \quad (15)$$

states that the expected value is $\langle \Delta x \rangle = +0.05x$, and therefore people will *want* to take this bet. Although Peters notices the difference between additive and multiplicative payoffs, he has to use a Monte Carlo simulation to show why this is actually a poor bet; nonlinear transformations of wealth, differing in additive vs. multiplicative cases to show why some people might not take this; *ergodic observables* to come up with a new explanation; and then *introduction of noise* to justify the ergodics.

But the definition is ill-structured. Although it looks like an arithmetic payoff, it is actually a geometric payoff, and should be written as such. The correct formulation should be

$$r_x = \begin{cases} r_{x_H} (= +0.5x) = 1.5, & p_H = 1/2 \\ r_{x_T} (= -0.4x) = 0.6, & p_T = 1/2 \end{cases} \quad (16)$$

for the multiplicative case, or replacing current x with x_0 in the original formula for the additive case.

Then, in the case of two tosses of a fair coin, it's easy to see that, for additive gains,

$$x[i+2] \Leftarrow \Delta x_H + \Delta x_T + x[i] = +0.5 x_0 + -0.4 x_0 + x[i] = (+0.1 x_0) + x[i], \quad (17)$$

on average, for any i , by the law of large numbers; however, for two tosses with multiplicative gains,

$$x[i+2] \Leftarrow r_{x_H} r_{x_T} x[i] = (1.5) (0.6) x[i] = (0.9) x[i] \quad (18)$$

again on average, for any i , by the law of large numbers, with a fair coin. This comes from the classical definition of probability that equal outcomes come up equal times, in the long run. These will give the median results; full distributions are explored in the Supplement.

Thus, by inspection, we can see that repeated coin flips of the arithmetic payoff on average result in a win of $(+0.1 x_0)$ every two flips. A professional gambler will find this advantageous and will want to take this bet, repeating as often as possible for more and more winnings. Whereas, by inspection, we can see that repeated coin flips of the geometric payoff on average (median case) result in a lose, being a reduction to $(0.9) x[i]$ every two flips. A professional gambler will find this disadvantageous, and will not want to take this bet. Unless their personality is a speculator looking for a lottery payoff.

No Monte Carlo simulation is needed to determine this. Also, no ergodics nor noise are required to determine whether it is wise to take the gamble or not. The r operator helps this to fall out by inspection.

The Real Paradox: “Just One More”

We've seen how the +50/-40 additive bet represents a net win more than half of the time in the long run, and in general many people will play it as often as possible, in order to bank on that win. Whereas, the +50%/-40% multiplicative bet represents a net loss more than half of the time in the long run, and in general most people who understand mathematics, and who are not speculators wishing for a lottery ticket, will avoid it.

However, the first move of the game, in both cases, is *identical*.

One might think then that the recommendation would be to play the game only once, and then stop at the end of the first turn.

But then how is this resulting situation different, in the multiplicative case, from the additive case? At the start of the second turn you're still given a choice between winning 50% of your current bankroll, or losing 40%. You might just as well have come into the game fresh, at the start of the second term, with a bankroll of \$60 or \$150.

So the additive winning player will say, “Just one more turn”, in the case of the multiplicative game.

But, **by induction**, if you are always coming into the game with fresh eyes, the sequence of turns will never stop. The multiplicative game will end up looking like a sequence of additive games. And there is no reason for a person used to winning the additive game to ever pause in the multiplicative version. By induction, they should be identical. But they aren't.

I call this the “Just One More” paradox. Here I present two example games to illustrate this:

Hero or Heroin

In this game, a person has a choice between taking a heroin stimulus injection, which will make them feel like a million dollars for an hour, or being a hero and going “cold turkey”, refusing to play. The heroin has a 1 in 13 chance of killing them each time. What should they do?

American Roulette

Although European roulette wheels have 18 red slots, 18 black slots, and one green 0, American roulette wheels have an extra, special “00” slot. It seems this should have some extraordinary meaning.

Doctor et al⁸ ask rhetorically, “Would a person ever prefer a process that, after three rounds, diminishes wealth from US\$10,000 to 0.5 cents, over one that yields a 99.9% chance of US\$10,000,000 and otherwise US\$0?” Although this is fine for a game show, in which, if the player loses all money on the table, they can still go home, it turns a blind eye to the nature of losing *all* wealth. They still have family, friends, a home, a bank account, and possible jobs to turn to. So they have *not* wagered *all* their wealth; arguably just the top fraction. They still have wherewithal to fall back on.

So, to make it more interesting, let's say that a gambling house has set up an extreme wager. There's a special roulette wheel that has 1,000 slots. 999 of these will return \$10,000,000 on your \$10,000 bet. With the last 00 slot, the house has hired a special military agent, “Double-O Nothing”^a. You contractually agree that, if 00 comes up, he may kill all of your family and friends, bulldoze your home, wipe out your bank account, and kill anyone who might try to hire you. But remember, you will personally make ten million dollars if 00 doesn't come up.

Unfortunately, when given millions, you habitually pay for large parties each night, which take your

a Hat tip to Johnny Hart, creator of *The Wizard of Id*.

bankroll back to the \$10,000 you started with. You can play every day. Do you take the gamble? How many days do you play?

Many folks would in fact take this gamble, especially with the moral risk of although having their family and friends wiped out, they would survive. However, the details of repeated trials are unfortunate. Extinction is modeled well as a Poisson process; basically, you have an average of a thousand days of riotous living before life as you know it will end.

But perhaps *just one more* day might not hurt?

Addiction and Just One More

The Just One More paradox may be useful in understanding addictive behavior. Although the long-term effects of repeated plays may be deadly, the short-term effects appear beneficial. This can apply to both people and government bodies.

Applications

Modeling multiplicative games, along with prescriptive and descriptive utility preference determinations for different personalities, has many real-world uses. Among these are finance; medicine and cancer treatments; international trade; the exponential benefit of cooperation over competition; and nuclear safety.

Conclusions

Ergodicity is not needed to solve utility preference determination. It's not even applicable in real-world cases involving annihilating states, such as insolvency or death. Instead, using proper operators gives straightforward solutions. The r notation is formalized from groups; it yields an annihilating state at zero, instead of at infinity, which provides useful non-linear representation power. The difference between advantageous additive gains and disadvantageous multiplicative gains is subtle; it gives rise to the Just One More paradox, illustrated by the Hero Or Heroin and American Roulette games, which may help in understanding addictions. The results may be useful in fiscal policy, cooperative mutually-beneficial trade, and nuclear safety.

Statement of Conflicts of Interest

The author declares no competing interests.

Article References

- 1 Peters, O. The ergodicity problem in economics. *Nature Physics*, vol 15, Dec 2019, pp. 1216-1221
- 2 Kelly, J. A New Interpretation of Information Rate. *Bell System Technical Journal*, **35**, pp. 917-926. (1956)
- 3 Press, W., et al. *Numerical Recipes: The Art of Scientific Computing, Third Edition*. (Cambridge Univ. Press, New York. 2007.) pp. 755-758.
- 4 Cover, T. Universal Portfolios. *Mathematical Finance*, **1(1)**, pp. 1-29. (1991)
- 5 Wikipedia, Kelly criterion. At https://en.m.wikipedia.org/wiki/Kelly_criterion
- 6 Goodfellow, I., Bengio, Y., Courville, A. *Deep Learning*. (The MIT Press, Cambridge. 2016)
- 7 Wikipedia, Catastrophe Theory. At https://en.m.wikipedia.org/wiki/Catastrophe_theory
- 8 Doctor, J., Wakker, P., Wang, T. Economists' views on the ergodicity problem. *Nature Physics* **16**, 1168 (2020)

Supplementary Information on: Multiplicative Gains, Non-ergodic Utility, and the Just One More Paradox

John K. Myers *john (at) scrier.org*

Abstract. Peters⁹ proposes ergodicity and added noise are useful in understanding preferences to additive and multiplicative bets. He poses two apparent paradoxes—that expected utility theory requires that people can interact with copies of themselves in parallel universes, to determine utility; and that additive expected values don't model multiplicative gain sequences well. I review probability interpretations and use visualizations of outcome distributions to show possible worlds need exist only in the mind of the single decider at the present moment; utility is a function of a distribution, not of an average; and utility depends on the temperament of the decider, not ergodicity. When realistic annihilating states are introduced, such as the assumption that insolvency ends betting, I demonstrate the resulting distributions can easily be skewed and non-ergodic anyway. The results show classical subjective probability does not seem to need ergodics, but that games and utility functions over repeated multiplicative gambles are subtle and can benefit from visualizations.

Summary. The results of betting a wealth bankroll, say \$100, on something as simple as a fair coinflip, and then receiving, say, either additive +\$50 for heads / -\$40 for tails, or multiplicative +50% for heads / -40% for tails, can have subtle implications. A major decision is whether to take a bet or not, based on the usefulness or *utility* of possible outcomes. Peters argues *ergodicity* is needed to determine choices, and that psychology-based utility theory can't work because “expected utility theory implicitly assumes that individuals can interact with copies of themselves, effectively in parallel universes”. So it seems the usage of possible worlds in probabilistic decision-making, and in particular computations on multiplicative versus additive gains, can benefit from exploration. Here I show how to work with multiplicative gains without using any ergodics nor noise, and how probability interpretations demonstrate subjective probability usage over present possible worlds inside the mind of one deciding agent, without any copies of people in parallel universes in the future. I explore annihilating states (representing insolvency, or death), which destroy ergodic assumptions. I suggest life time as a grounding measure for utility comparisons. The results should help clarify multiplicative gains, such as are found in financial markets and military contests, and may be useful in understanding addictive behavior. In particular, it's anticipated that nuclear safety and economic cooperation calculations may benefit from these explorations.

Background

Multiplicative gains appear to be the same as additive gains, but the subtle difference produces profound differences in outcomes. Peters¹ proposes applying physical ergodic concepts to the problem of gambling-outcome utility preference determination. But the subtle differences are not made obvious. So Peters requires a random 1,000-toss Monte Carlo simulation to demonstrate empirically that the multiplicative bet is disadvantageous and does not follow the arithmetic expected value. He suggests adding Wiener noise to the process to get Brownian motion on the growth, in order to obtain supposedly ergodic observables, and suggests this is important in understanding utility preferences. He criticizes psychological utility determination arguments as “hard to constrain, and often circular”. He wishes to construct “ergodic observables for non-ergodic (growth) processes”. Ergodicity can also be understood as *path independence*—that a final result does not depend on its path through time, but only

upon its final position. But when a process *is* path-dependent, this breaks down, as we'll demonstrate. Finally, Peters proposes a probability paradox: How can different copies of myself, in multiple possible worlds in the future, coordinate to pool their value?

Let's clarify these points, and then explore them with visualizations using only 10 tosses.

Kinds of Probability

A review of probability interpretations¹⁰ seems useful. “Probability” continues to be confusing, because the word is conflated to mean at least five different things. *Propensity* probability, in a Newtonian universe,^a is an attribute of a physical thing-in-itself right now having a propensity to yield different possible outcomes. It is unknowable directly; and can only be **modeled, in the mind of a decision-maker**, by logical assumptions, or by measurements/tests.^b *Classical probability* was a first attempt at logical assumptions to model and assign probability measures based on the principle of indifference, that cards or dice should get equal probabilities because one card logically comes up as often as another. More recently, this approach is backed by Jaynes's *maximum entropy*, which states if you don't have any information as to why one outcome should be preferred over another, then assign equal measures. *Frequentist probability (statistics)* attempts to model propensity descriptively by *counting* occurrences in the *past*; this assigns a measure over a set of past outcome observations in the exterior world. Of course, it runs into trouble when the set is small, empty, or non-representative of rare outcomes. *Subjective probability* believes that probability is all in the mind: it assigns a measure over a set of *interior* possible worlds. This is not necessarily consistent, non-fuzzy, nor even known. *Credence* is used to name the subjective probability assigned to a mentally-modeled event. This is often conflated with *confidence*,^c a second-order probability term modeling the width and intensity of the credence distribution, often approximated by *interval probabilities*, that measures *how sure* the decider is that the credence is accurate. You can believe that a coin is 50/50 fair, and still not have confidence that the casino has not introduced a weighted coin, until you've seen the coin in action many times. This leads to *pignistic probability*, or “betting”, which combines credence with confidence to come up with a probability scalar when the decision-maker is forced to decide on a bet and actually act or fold. Note this is often different from credence, due to real-world caution, concerning the slippage between believed subjective models and actual propensity. *Bayesian probability* is a prescriptive mathematics describing how a subjective modeling decider *should* act to assemble and update believed probabilities in a consistent manner. However, it requires priors, which usually ground out in either maximum-entropy or expert credence approaches. Finally, *Kolmogorov's probability calculus* rests in the realm of pure mathematics, which can be used by Bayesians and others to help solve concrete problems consistently. Notice that the domains of these—the present thing-in-itself, past observations in the world, models of potential future possible worlds in the mind, pure mathematics—are quite different, and so these are describing different things by trying to use the same word.

Utility Needs Parallel Worlds, But They're Not Where You Might Suppose

With this as background, *where* and *when* are the possible worlds assigned to wealth outcomes from betting 10 coin flips in the future? These are **believed models** of future projections that are **in the mind of the decision-maker, right now**, concerning the best possible models that the expert statistician, maximum-entropy mathematician, Bayesian, or subject-matter engineer advisor

-
- a This ignores possibilities in which we live in a “Gone With The Wind” universe—a dramatic creative work, a movie, in which spacetime is fixed—Feynman diagrams seem to support this; or a “Grand Theft Auto” universe run by a computer simulation; or a “Dungeons & Dragons” universe where observations are only placed when an observer looks for them.
- b Which are subject to Heisenberg uncertainty.
- c Different from the statistical usage of “confidence”.

components can come up with. Even then, they could be incomplete or wrong. They are used by the decision-maker in a pignistic way, along with confidence, to determine at each step how much to bet, or whether to take the money, go home, and choose not to play any more. Then the real world plays out along **one** particular path. In a Newtonian interpretation, there are *no* parallel universes in the *real* world. There is only *what actually happens*. Thus, there is no need to talk of individuals interacting with parallel copies of themselves.

Where, then, do these possible worlds come in? It's a two-step process. First the individual generates the distribution of predicted possible worlds in his or her mind, based on best information, as will be shown in the example figures. Only then does the deciding individual use utility preferences, over the *distribution*, *not* over the mean, to determine whether to go with the bet or not. In particular, worst-case analysis plays heavily, as shown in maximin theory. And these preferences are subjective.

There is much more that can be said about the epistemology of possible-world reasoning inside a single agent's mind, and indeed a 5-valued logic of {Believed Actual, Believed Possible, Believed Hypothetical, Believed Inconsistent, Unknown} can be determined over this. See ref¹¹ for further details.

Utility depends on the temperament of the decider

Options trading classically divides between the *hedgers* and *speculators*. *Hedgers*, such as farmers, want to make a known good profit on their crops and are willing to give up the uncertain chance of making a *great* profit in order to lock in good prices, fold, and go home. *Speculators* enjoy risk, like to buy lottery tickets, and are willing to *pay* for the chance to make a *great* profit, at the risk of losing cash in hand. Note that these deciders are both placing *differing* preferences on the *identical* probabilistic outcome distribution. The same distribution will attract or repel different types of personalities. So simple ergodic transforms on a single distribution are inadequate to predict subjective utility preferences—psychology is required.

Utility: Life Time and Opportunity Cost

People only have so much time left in their lives. This is a limited resource that can be used to ground derivations. Independently wealthy people have their money working for them, and earning income at a particular rate; whereas salaried people have to invest life time into earning money, but also know how fast they earn. Putting a bankroll at risk is then equivalent to betting the number of hours left in one's life needed to replace that amount of money if it's lost. Obviously this will be quite different for different people. However, this provides a quantitative approach towards defining utility, in terms of *life-hours*.

Utility depends on the distribution

Assume a wise hedging farmer, who wants to feed his family. Assume a binary choice between planting rice, which has a certain chance of doubling his investment; vs. planting bicontinuous Menger sponges, which have a 0.001 lottery chance of earning millions or billions as next-generation batteries for EV electric vehicles, but have a 0.999 chance of losing 90% of the farm. It does not matter how arbitrarily large the lottery payoff is; the wise farmer is not going to bet the farm on a risky endeavor. So averaging expected utilities, even with severe logarithmic distortions, are not the whole answer to pignistic utility determination for hedgers. One needs to look at the entire overall distribution, including the mode, variance, upper and lower limits, etc., as inputs to psychological utility preference determination.

Ergodicity, Free Will, and Insolvency

Ergodicity only applies if the player is constrained to always finish all runs, and does not have the option of re-evaluating their position and choosing in the middle to fold and go home. This is highly unrealistic for most real-life games. In such a case, the bankroll amount of the player is frozen and is constant for the rest of the time series. This is a problem with ergodic-inspired approaches.

Ergodicity also only applies if there are no annihilating (absorbing) states, such as insolvency. In most games, you must have money to play, and if you run out of money, you are removed from the game. Even if a bank is willing to loan you money, at some point in most reasonable games there comes a time at which the bank draws the line and is unwilling to advance you more money. This then freezes the results at that point.

Note also that subtracting out annihilated states skews the distribution: you come up with a survivor bias, which shifts the mean. Also final outcomes are *path-dependent* (and therefore non-ergodic). If you go very high, then low, you can achieve successful results, different from if you go low, get annihilated, then would have gone very high. We'll see this in the examples.

Experiments

In order to compare additive vs. multiplicative games, instead of talking about esoteric growth functions, it is most useful to visualize the exhaustive results of all possible outcomes from 10 trials. This brings out features and realizations that may have been hidden.

All experiments are simulated inside the computer. In each experiment, a (virtual) player starts with a bankroll of \$100. A fair coin is tossed 10 times. Depending on whether the coin comes up heads or tails, the player receives a 50 amount or loses a 40 amount. The game is played for 10 tosses. All possible combinations are graphed, giving 2^{10} or 1024 different possible worlds. However, because outcomes are symmetric and discrete, in the unabsorbed cases there are only $(n+1) = 11$ different distinct results, both in the arithmetic and multiplicative cases. 1024 outcomes was chosen because it's close to 1000, giving an intuitive feel for percentages from outcome tallies (presented in the Tables). 10 trials is large enough to see tail effects, yet small enough to understand easily.

Peters requires 1,000 trials, not 10, for his Monte Carlo empirical demonstration. However, because he only simulates with 150 runs, which is not exhaustive like ours, he misses the hugely positive long tail effects important to speculators on the multiplicative case.

Experimental cases explore: a baseline additive game with unlimited loans and no annihilating states; the same game, with naive truncation *at the end* thus resetting all negative final outcomes to zero; the same game, but this time with a dynamic annihilating state *in the middle* that sticks any intermediate negative outcome at zero; and a standard multiplicative game, with the same percentage payoffs, but computed as a percentage of the *current* not initial bankroll.

Experiment 1. In Fig. 1, a fair coin is tossed 10 times, with arithmetic payoffs of +\$50 or -\$40 per toss, depending on heads or tails. As always, the player starts with a bankroll of \$100. Importantly, *the result percentage is fixed on the initial bankroll of \$100*, and the payoffs do not vary from +\$50/-\$40 even when the bankroll grows.

For the experiment to be realized, it's necessary to assume *unlimited bank loans*. If the player's bankroll goes negative, the player is *still allowed to keep playing*. This results in 17% of the final cases ending negative. Note that, with bank loans allowed, a player's random-walk path may wander negative in the middle of the 10 trials run, then correct back into positive territory.

The outcome is the familiar bell curve. There are $10+1 = 11$ distinct outcome columns. The extreme ends at +600 and -300 represent only 1 outcome out of 1024 cases apiece, and so are mostly too short to see. Each graphical box in a vertical column represents 10 outcomes, for easy counting. See Fig. 1.

This is the only experiment to go negative.

Means, medians, and modes are reported at the bottom of Tables 1 and 2. The additive experiment has a winning mean, median, and mode of \$150, all the same, \$50 **more** than the \$100 the player starts with. With 10 trials, a player ends up winning more than the starting bankroll 62.3% of the time. The expected final amount of the win, and the percentage of winning outcomes, only increases with the number of trials. So, a professional gambler would *take* this bet, and would attempt to play it as rapidly and as often as possible, in order to rack up both a large bankroll and large chances of winning.

Humans are intuitively used to additive gains. This is the game that most people intuitively feel they are playing.

Experiment 2. Fig. 2 shows the same additive experiment as Fig. 1, with the modification that negative results at the end are statically clipped to 0. That is, if the player ended up losing at the casino, at the end of the day all debts are forgiven to 0.

Although unrealistic, this is apparently the method used in the Copenhagen experiment¹² discussed by Peters. Also note that the Copenhagen experiment excluded accountants, engineers, computer scientists, and anyone who might be good at numbers (p.15).

This skews the distribution. Naive static clipping of the results actually *increases* the mean to \$161, because the left side is moved forward. However, chances of winning remain the same at 62.3%.

Experiment 3. Fig. 3 shows the same additive experiment as Fig. 1, with the modification that negative results in the middle are dynamically frozen at 0. That is, without loans, if the player ran out of money, the play stopped and they were not allowed to play any more. With the same 1024 exhaustive possible outcomes, this has a substantially larger peak at 0 than Fig. 2, as it eats into columns that would have resulted in the middle and towards the right. In fact, the mode shifts from \$150 to \$0. The mean also shifts downwards, from \$150 to \$146.48. But there's still a 56.8% chance of winning.

Note that the computer program requires a special exception rule for this arithmetic experiment. Whereas, if a multiplicative experiment needs to create an annihilating state, it can simply set the r payoff reward to 0 inside the mathematics itself, without have to go outside to set exception rules.

Experiment 4. Fig. 4 shows a standard multiplicative gain experiment. A fair coin is tossed 10 times, with arithmetic payoffs of +50% or -40% of the *current* bankroll per toss, depending on heads or tails, again starting at \$100. There are again $10 + 1 = 11$ distinct results columns, with count tallies interestingly the same as in experiment 1. See Table 2. (Note that the single-outcome columns of \$0.60 and \$5,766.50 are too short to show up on the graph.) However, the left half of the bell curve is substantially compressed close to zero; whereas the right half of the curve is substantially stretched. It means that this series does not follow the central limit theorem.

Although the mean final outcome is \$162.89, this is skewed by the extremely long tail on the right; as both the median and mode outcomes are \$59.05, less than the starting bankroll of \$100. Indeed, a player can expect to win and take home more than their starting bankroll in only 37.7% of the cases. In 62.3% of the cases, a player repeating 10 times is going to lose.

Experiment 4 does not explore the multiplicative annihilating case when one of the outcome's $r = 0$, which would end up with more and more results on the 0 axis as repeated trials progress. And, with a geometric growth rate of $\sqrt[10]{0.9} \approx 0.949$, less than 1.0, the median of the trials will continue to

diminish, with repetitions, but the player will never go negative. If the growth rate were greater than 1.0, on the other hand, representing the mutual profitability of a win-win transaction, the median of the trials would grow, with repetitions, without limit.

Because the median bankroll continues to shrink over time, as the number of trials increases, most professional gamblers/investors will *stay away* from this bet.

Some high-risk speculators may be attracted to this, however, as it provides a good “lottery ticket”-- with a 1/1024 chance, betting \$100 can end up with \$5,766.50.

Discussion

There are a few details that might appear surprising upon first glance. First, the outcomes of Exp. 1 are divided into 11 columns, instead of being a continuous bell curve. This is a result of the play being only 10 discrete Bernoulli trials, and each trial having the same payoff conditions; so the trials commute, in the additive case, resulting in only 11 distinct final outcomes. Second, the outcomes of Exp. 4 are *also* divided into 11 columns. Of course, multiplicative cases commute also. Third, the scores (tallies) for each column, read left-to-right, are the *same* between Exp. 1 and Exp. 4. This is also a result of running Bernoulli trials, without any annihilating states. So the two curves are isomorphic.

However, we recall that building the probability distribution is only the first step before applying the personal utility function over the distribution. And money is money. So, whether the person is a hedger, a professional gambler, or a speculator, that person will be providing their *same* personal utility function over widely varying distribution results from the different experiments. Utility preferences depend upon the person and their personality, not upon the ergodics of the distribution.

Conclusions

Ergodicity is not needed to solve utility preference determination. It's not even applicable in real-world cases involving annihilating states, such as insolvency or death. And discussions between clones in future possible universes are not needed for probability distribution determination. Instead, probabilities are subjective, depend on the expertise of the decider, and are formed in the present from multiple possible future worlds imagined inside the mind of a single decider. Visualizing distributions comprehensively using even only 10 trials serves to make consequences of decisions transparent. Multiplicative gains should not be mistaken for arithmetic gains, and the Just One More paradox shows this has important implications for addiction, cooperative commerce, and nuclear safety, among others.

Supplement References

- 9 Peters, O. The ergodicity problem in economics. *Nature Physics*, vol 15, Dec 2019, pp. 1216-1221
- 10 Wikipedia, Probability interpretations. At https://en.m.wikipedia.org/wiki/Probability_interpretations
- 11 Myers, J. B-SURE: A Believed Situation and Uncertain-Action Representation Environment, *COLING-92, vol. III* (1992). pp. 961-965.
- 12 Meder, D., et al. Ergodicity-breaking reveals time optimal decision making in humans. Preprint at <https://arxiv.org/abs/1906.04652.pdf> .




Fig. 1: Arithmetic returns on $+50/-40$, starting from 100, with unlimited loans. Ending dollars vs. total number of possible-world outcomes after 10 equiprobable coin-flip trials, showing all 1024 possible combinations. Each complete box cell represents 10 outcomes. Note characteristic Gaussian distribution. Majority of outcomes end up more than the start at 100.




Fig. 2: Arithmetic returns on $+50/-40$, starting from 100, with naive truncation afterwards at 0




Fig. 3: Arithmetic returns on $+50/-40$, starting from 100, with path sticking at 0 when it goes negative (no loans, and can't play further when run out of money). Results are path-dependent and non-ergodic. Note differences in lower bands from Fig. 2.

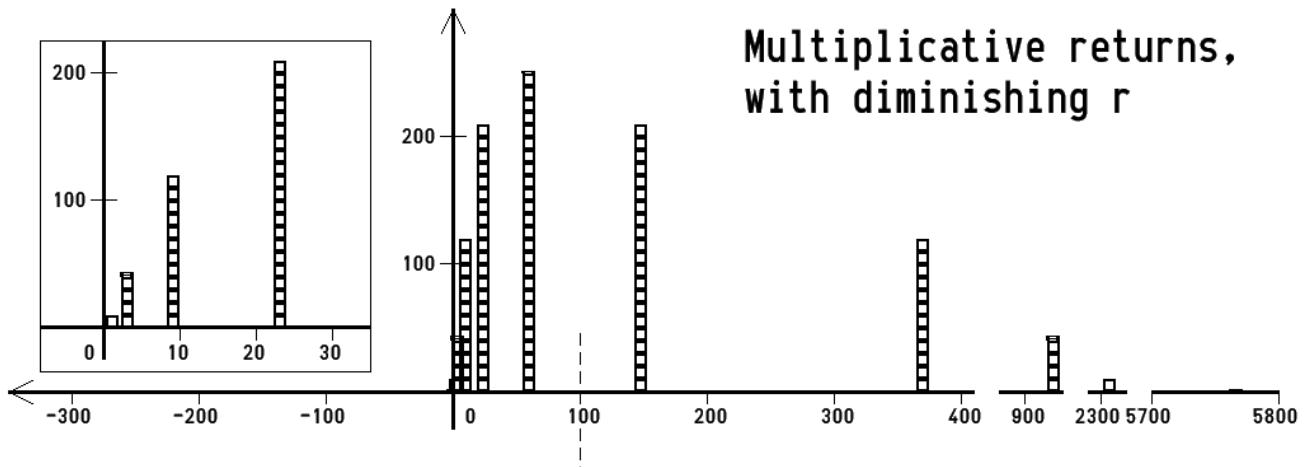


Fig. 4: Multiplicative returns of +50%/-40%, starting from 100, again 10 trials with 1024 outcomes. Left side of graph is severely compressed, and is replicated with zoomed abscissa in the inset; right side is severely expanded, requiring elided gaps in the axis display. Each box is 10 outcomes. Eleven distinct column counts are same as in Fig. 1. Majority of outcomes end up less than the start at 100.

Tables

Outcome	Arithmetic, with loans	Arithmetic, naive trunc.	Arithmetic, annihilating 0
-\$300	1		
-\$210	10		
-\$120	45		
-\$30	120		
\$0		176	324
\$60	210	210	118
\$150	252	252	207
\$240	210	210	200
\$330	120	120	119
\$420	45	45	45
\$510	10	10	10
\$600	1	1	1
Mean	\$150	\$161	\$146.48
Median	\$150	\$150	\$150
Mode	\$150	\$150	\$0
p(>100)	0.623	0.623	0.568

Table 1: Arithmetic payoff experiments, from Figs. 1,2,3.

Outcome	Multiplicative
\$0.60	1
\$1.51	10
\$3.78	45
\$9.45	120
\$23.62	210
\$59.05	252
\$147.62	210
\$369.06	120
\$922.64	45
\$2,306.60	10
\$5,766.50	1
Mean	\$162.89
Median	\$59.05
Mode	\$59.05
p(>100)	0.377

Table 2: Multiplicative payoff experiment, from Fig. 4

Figures




Figure 1

Arithmetic returns on +50/-40, starting from 100, with unlimited loans. Ending dollars vs. total number of possible-world outcomes after 10 equiprobable coin-flip trials, showing all 1024 possible combinations. Each complete box cell represents 10 outcomes. Note characteristic Gaussian distribution. Majority of outcomes end up more than the start at 100.




Figure 2

Arithmetic returns on +50/-40, starting from 100, with naive truncation afterwards at 0




Figure 3

Arithmetic returns on $+50\%/ -40\%$, starting from 100, with path sticking at 0 when it goes negative (no loans, and can't play further when run out of money). Results are path-dependent and non-ergodic. Note differences in lower bands from Fig. 2.




Figure 4

Multiplicative returns of $+50\%/ -40\%$, starting from 100, again 10 trials with 1024 outcomes. Left side of graph is severely compressed, and is replicated with zoomed abscissa in the inset; right side is severely expanded, requiring elided gaps in the axis display. Each box is 10 outcomes. Eleven distinct column counts are same as in Fig. 1. Majority of outcomes end up less than the start at 100.