



Nassim Nicholas Taleb

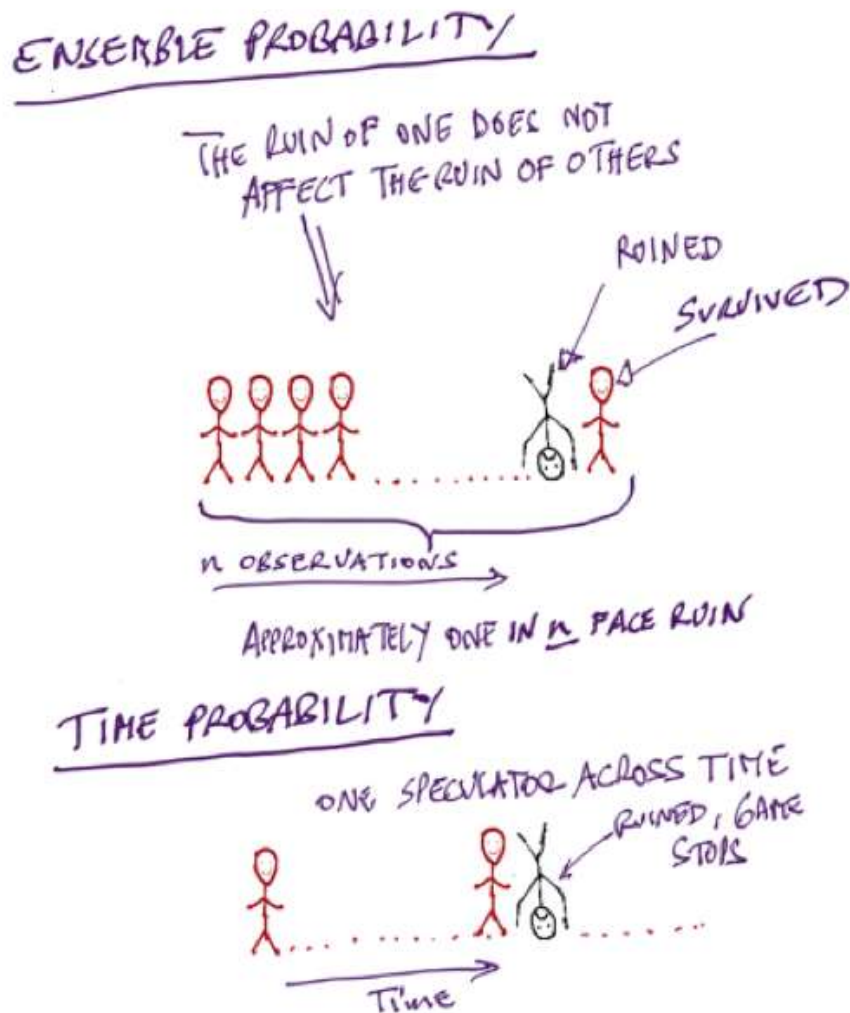
Follow

Aug 25, 2017 · 12 min read

The Logic of Risk Taking

A central chapter that crystallizes all my work. In forth. Skin in the Game

Time to explain ergodicity, ruin and (again) rationality. Recall from the previous chapter that to do science (and other nice things) requires survival t not the other way around?



The difference between 100 people going to a casino and one person going to a casino 100 times, i.e. between (path dependent) and conventionally understood probability. The mistake has persisted in economics and psychology since age immemorial.

Consider the following thought experiment.

First case, one hundred persons go to a Casino, to gamble a certain set amount each and have complimentary gin and tonic –as shown in the cartoon in Figure x. Some may lose, some may win, and we can infer at the end of the day what the “edge” is, that is, calculate the returns simply by counting the money left with the people who return. We can thus figure out if the casino is properly pricing the odds. Now assume that gambler number 28 goes bust. Will gambler number 29 be affected? No.

You can safely calculate, from your sample, that about 1% of the gamblers will go bust. And if you keep playing and playing, you will be expected have about the same ratio, 1% of gamblers over that time window.

Now compare to the second case in the thought experiment. One person, your cousin Theodorus Ibn Warqa, goes to the Casino a hundred days in a row, starting with a set amount. On day 28 cousin Theodorus Ibn Warqa is bust. Will there be day 29? No. He has hit an uncle point; there is *no game no more*.

No matter how good he is or how alert your cousin Theodorus Ibn Warqa can be, you can safely calculate that he has a 100% probability of eventually going bust.

The probabilities of success from the collection of people does not apply to cousin Theodorus Ibn Warqa. Let us call the first set *ensemble* probability, and the second one *time* probability (since one is concerned with a collection of people and the other with a single person through time). Now, when you read material by finance professors, finance gurus or your local bank making investment recommendations *based on the long term returns of the market*, beware. Even if their forecast were true (it isn't), no person can get the returns of the market unless he has infinite pockets and no uncle points. They are conflating ensemble probability and time probability. If the investor has to eventually *reduce* his exposure because of losses, or because of retirement, or because he remarried his neighbor's wife, or because he changed his mind about life, his returns will be divorced from those of the market, period.

We saw with the earlier comment by Warren Buffett that, literally, anyone who survived in the risk taking business has a version of “in order to succeed, you must *first* survive.” My own version has been: “never cross a river if it is *on average* four feet deep.” I effectively organized all my life around the point that *sequence matters* and the

presence of ruin does not allow cost-benefit analyses; but it never hit me that the flaw in decision theory was so deep. Until came out of nowhere a paper by the physicist Ole Peters, working with the great Murray Gell-Mann. They presented a version of the difference between the ensemble and the time probabilities with a similar thought experiment as mine above, and showed that about everything in social science about probability is flawed. Deeply flawed. Very deeply flawed. For, in the quarter millennia since the formulation by the mathematician Jacob Bernoulli, and one that became standard, almost all people involved in decision theory made a severe mistake. Everyone? Not quite: every economist, but not everyone: the applied mathematicians Claude Shannon, Ed Thorp, and the physicist J.-L. Kelly of the Kelly Criterion got it right. They also got it in a very simple way. The father of insurance mathematics, the Swedish applied mathematician Harald Cramér also got the point. And, more than two decades ago, practitioners such as Mark Spitznagel and myself build our entire business careers around it. (I personally get it right in words and when I trade and decisions, and detect when ergodicity is violated, but I never explicitly got the overall mathematical structure –ergodicity is actually discussed in *Fooled by Randomness*). Spitznagel and I even started an entire business to help investors eliminate uncle points so they can get the returns of the market. While I retired to do some flaneuring, Mark continued at his Universa relentlessly (and successfully, while all others have failed). Mark and I have been frustrated by economists who, not getting ergodicity, keep saying that worrying about the tails is “irrational”.



Mark Spitznagel explaining (politely) that finance professors are probability challenged

Now there is a skin in the game problem in the blindness to the point. The idea I just presented is very very simple. But how come nobody for 250 years got it? *Skin in the game, skin in the game.*

It looks like you need a lot of intelligence to figure probabilistic things out when you don't have skin in the game. There are things one can only get if one has some risk on the line: what I said above is, in retrospect, obvious. But to figure it out for an overeducated nonpractitioner is hard. Unless one is a genius, that is have the clarity of mind to see through the mud, or have such a profound command of probability theory to see through the nonsense. Now, certifiably, Murray Gell-Mann is a genius (and, likely, Peters). Gell-Mann is a famed physicist, with Nobel, and discovered the subatomic particles he himself called *quarks*. Peters said that when he presented the idea to him, "he got it instantly". Claude Shannon, Ed Thorp, Kelly and Cramér are, no doubt, geniuses –I can vouch for this unmistakable clarity of mind combined with depth of thinking that juts out when in conversation with Thorp. These people could get it without skin in the game. But economists, psychologists and decision-theorists have no genius (unless one counts the polymath Herb Simon who did some psychology on the side) and odds are will never have one. Adding people without fundamental insights does not sum up to insight; looking for clarity in these fields is like looking for aesthetic in the attic of a highly disorganized electrician.



The clarity of the mind of psychologists and economists

Ergodicity

As we saw, a situation is deemed non ergodic here when observed past probabilities do not apply to future processes. There is a "stop" somewhere, an absorbing barrier that prevents people *with skin in the game* from emerging from it –and to which the system will invariably tend. Let us call these situations "ruin", as the entity cannot emerge

from the condition. The central problem is that if there is a possibility of ruin, cost benefit analyses are no longer possible.^[i]

Consider a more extreme example than the Casino experiment. Assume a collection of people play Russian Roulette a single time for a million dollars –this is the central story in *Fooled by Randomness*. About five out of six will make money. If someone used a standard cost-benefit analysis, he would have claimed that one has 83.33% chance of gains, for an “expected” average return per shot of \$833,333. But if you played Russian roulette more than once, you are deemed to end up in the cemetery. Your expected return is ... not computable.

Repetition of Exposures

Let us see why “statistical testing” and “scientific” statements are highly insufficient in the presence of ruin problems and repetition of exposures. If one claimed that there is “statistical evidence that the plane is safe”, with a 98% confidence level (statistics are meaningless without such confidence band), and acted on it, practically no experienced pilot would be alive today. In my war with the Monsanto machine, the advocates of genetically modified organisms (transgenics) kept countering me with benefit analyses (which were often bogus and doctored up), not tail risk analyses for *repeated* exposures.

Psychologists determine our “paranoia” or “risk aversion” by subjecting a person to a single experiment –then declare that humans are rationally challenged as there is an innate tendency to “overestimate” small probabilities. It is as if the person will never again take any personal tail risk! Recall that academics in social science are ... dynamically challenged. Nobody could see the grandmother-obvious inconsistency of such behavior with our ingrained daily life logic. Smoking a single cigarette is extremely benign, so a cost-benefit analysis would deem one irrational to give up so much pleasure for so little risk! But it is the act of smoking that kills, with a certain number of pack per year, tens of thousand of cigarettes –in other words, repeated serial exposure.

Beyond, in real life, every single bit of risk you take adds up to reduce your life expectancy. If you climb mountains *and* ride a motorcycle *and* hang around the mob *and* fly your own small plane *and* drink absinthe, your life expectancy is considerably reduced although not a single action will have a meaningful effect. This idea of repetition makes paranoia about some low probability events perfectly rational. But we

do not need to be overly paranoid about ourselves; we need to shift some of our worries about bigger things.

Note: The flaw in psychology papers is to believe that the subject doesn't take any other tail risks anywhere outside the experiment and will never take tail risks again. The idea of "loss aversion" have not been thought through properly –it is not measurable the way it has been measured (if at all measurable). Say you ask a subject how much he would pay to insure a 1% probability of losing \$100. You are trying to figure out how much he is "overpaying" for "risk aversion" or something even more stupid, "loss aversion". But you cannot possibly ignore all the other present and future financial risks he will be taking. You need to figure out other risks in the real world: if he has a car outside that can be scratched, if he has a financial portfolio that can lose money, if he has a bakery that may risk a fine, if he has a child in college who may cost unexpectedly more, if he can be laid off. All these risks add up and the attitude of the subject reflects them all. Ruin is indivisible and invariant to the source of randomness that may cause it.

I believe that risk aversion does not exist: what we observe is, simply a residual of ergodicity.

Who is "You"?

Let us return to the notion of "tribe" of Chapter x. The defects people get from studying modern thought is that they develop the illusion that each one of us is a single unit, without seeing the contradiction in their own behavior. In fact I've sampled ninety people in seminars and asked them: "what's the worst thing that happen to you?" Eighty-eight people answered "my death".

This can only be the worst case situation for a psychopath. For then, I asked those who deemed that the worst case is their own death: "Is your death *plus* that of your children, nephews, cousins, cat, dogs, parakeet and hamster (if you have any of the above) worse than just *your* death? Invariably, yes. "Is your death *plus* your children, nephews, cousins (...) *plus* all of humanity worse than just *your* death? Yes, of course. Then how can your death be the worst possible outcome?[1]

Thus we get the point that individual ruin is not as big a deal as the collective one. And of course ecocide, the irreversible destruction of the environment, is the big one to worry about.



Hierarchy of risks – Taking personal risks to save the collective are both “courage” and “prudence” since you are lowering risks for the collective

To use the ergodic framework: My death at Russian roulette is not ergodic for me but it is ergodic for the system. The precautionary principle, in the formulation I did with a few colleagues, is precisely about the highest layer.

About every time I discuss the precautionary principle, some overeducated pundit suggests that “we cross the street by taking risks”, so why worry so much about the system? This sophistry usually causes a bit of anger on my part. Aside from the fact that the risk of being killed as a pedestrian is one per 47,000 years, the point is that my death is never the worst case scenario *unless* it correlates to that of others.


I have a finite shelf life, humanity should have an infinite duration.

Or

I am renewable, not humanity or the ecosystem.

Even worse, as I have shown in *Antifragile*, the fragility of the components is required to ensure the solidity of the system. If humans were immortals, they would go extinct from an accident, or from a gradual buildup of misfitness. But shorter shelf life for humans allows genetic changes to accompany the variability in the environment.

Courage And Precaution Aren't Opposite



Aristotle

How can courage and prudence be both classical virtues? Virtue, as presented in Aristotle's *Nichomachean Ethics* includes: *sophrosyne* (σωφροσύνη), prudence, a form of sound judgment he called more broadly *phronesis*. Aren't these inconsistent with courage?

In our framework, they are not at all. They are actually, as Fat Tony would say, the same ting. How?

I can exercise courage to save a collection of kids from drowning, and it would also correspond to some form of prudence. I am sacrificing a lower layer in Figure x for the sake of a higher one.

Courage, according to the Greek ideal that Aristotle inherited—say the Homeric and the ones conveyed through Solon, Pericles, and Thucydides, is never a selfish action:

Courage is when you sacrifice your own wellbeing for the sake of the survival of a layer higher than yours.

As we can see it fits into our table of preserving the sustainability of the system.

A foolish gambler is not committing an act of courage, especially if he is risking other people's funds or has a family to feed. And other forms of sterile courage aren't really courage.[2]

Rationality, again

The last chapter presented rationality in terms of actual decisions, not what is called "beliefs" as these may be adapted to prevent us in the most convincing way to avoid things that threaten systemic survival. If superstitions is what it takes, not only there is absolutely no violation of the axioms of rationality there, but it would be technically irrational to stand in its way.

Let us return to Warren Buffett. He did not make his billions by cost benefit analysis, rather, simply by establishing a high filter, then picking opportunities that pass such threshold. "The difference between successful people and really successful people is that really successful people say *no* to almost everything." He wrote. Likewise our wiring might be adapted to "say no" to tail risk. For there are zillion ways to make money without taking tail risk. There are zillion ways to solve problems (say feed the world) without complicated technologies that entail fragility and an unknown possibility of tail risks.

Indeed, it doesn't cost us much to refuse some new shoddy technologies. It doesn't cost me much to go with my "refined paranoia", even if wrong. For all it takes is for my paranoia to be right once, and it would have saved my life.

Love Some Risks

Antifragile revolves around the idea that people confuse risk of ruin with variations – a simplification that violates a deeper, more rigorous logic of things. It makes the case for risk loving, systematic "convex" tinkering, taking a lot of risks that don't have tail risks but offer tail profits. Volatile things are not necessarily risky, and the reverse. Jumping from a bench would be good for you and your bones, while falling from the twenty-second floor will never be so. Small injuries will be beneficial, never larger ones. Fearmongering about some class of events is fearmongering; about others it is not. Risk and ruin are different *things*.

Technical Notes

[1] Actually, I usually joke my death *plus* someone I don't like such as the psychologist Steven Pinker surviving is worse than just my death.

[2] To show the inanity of social science, they have to muster up the sensationalism of "mirror neurons"

[i] The following question arises. Ergodicity is not statistically identifiable, not observable, and there is no test for time series that gives ergodicity, similar to Dickey-Fuller for stationarity (or Phillips-Perron for integration order). More crucially: if your result is obtained from the observation of a times series, how can you make claims about the ensemble probability measure?

The answer is similar to arbitrage, which has no statistical test but, crucially, has a probability measure determined *ex ante* (the "no free lunch" argument). Further, consider the argument of a "self-financing" strategy, via, say, dynamic hedging. At the limit we assume that the law of large numbers will compress the returns and that no loss and no absorbing barrier will ever be reached. It satisfies our criterion of ergodicity but does not have a statistically obtained measure. Further, almost all the literature on intertemporal investments/consumption requires absence of ruin.

We are not asserting that a given security or random process is ergodic, but that, given that its ensemble probability (obtained by cross-sectional methods, assumed via subjective probabilities, or, simply, determined by arbitrage arguments), a risk-taking strategy should conform to such properties. So ergodicity concerns the function of the random variable or process, not the process itself. And the function should not allow ruin.

In other words, assuming the SP500 has a certain expected return "alpha", an ergodic strategy would generate a strategy, say Kelly Criterion, to capture the assumed alpha. If it doesn't, because of absorbing barrier or something else, it is not ergodic.

