The Precautionary Principle

Yaneer Bar-Yam[1], Rupert Read[2], Nassim Nicholas Taleb[3]

Abstract
We present a non-naive, fragility-based, version of the precautionary principle, put it under formal statistical and probabilistic structure of "ruin" problems.

Keywords
Precautionary Principle

[1] *New England Complex Systems Institute*
[2] *School of Philosophy, University of East Anglia*
[3] *School of Engineering, New York University*

Contents

1. What we mean by a nonnaive PP

The Precautionary Principle (PP) is often presented in forms such that it can be used naively to justify any act of paranoia and interventionism (which in fact would create more risk); we are instead putting a formal statistical and risk-based structure around the idea, grounding it in rigorous probability theory and the properties of complex systems.

In other words, the PP is meant to help make decisions that ensure survival when statistical evidence is limited because it has not had time to show up, by focusing on the adverse effects of "absence of evidence".

2. Ruin not variations, systemic not local

The PP is not there to make life comfortable, rather to avoid a certain class of what is called in probability and insurance "ruin" problems.

For nature, the "ruin" is ecocide: an irreversible termination of life at some scale, which could be the planet.

3. PP is not risk management

There are three layers:

- Precautionary Principle
- Risk management (when downside risks are known or limited)
- Risk aversion (at the level of individual risk preferences)

It is important to avoid conflation of the PP and risk management. For instance, many worry about nuclear energy and its effects, and understandably so: but because of the known statistical structure of most of its problems and the absence of systemic consequences at small enough scales, at such scales the problem is better left to risk management than to the precautionary principle (whereas GMOs outside the laboratory are PP-salient even at very small scales: see 13.1, below).

If someone stands against a specific nuclear power plant in his area, this position would come largely from probabilistic risk management (or from ethics: a serious problem with all nuclear power is the temporal assymetry of its effects). Others who are as yet unsure can make decisions one nuclear incident at a time; but one should bear in mind that, as nuclear scales up, and the risk of incidents present or future grows, the PP becomes more and more relevant to it. (These points are explored further in 13.2, below.)

Also, it is an error is to think that risk management is not as serious or urgent as PP; it is, but it is not part of the unconditional nature of the precautionary principle.

4. Why Ruin is Serious Business

The risk of ruin is not sustainable, like a resource that gets depleted in the long term (even in the short term). By the ruin theorems, if you incur a tiny probability of ruin, as a "one-off" risk, survive it, then repeat the exposure, you will eventually go bust with probability 1.

The good news is that some classes of risk are deemed to be practically of probability zero: the planet took about close to zero risks of ecocide in trillions of variations over 3 billion years, otherwise we would not have been here.[1]

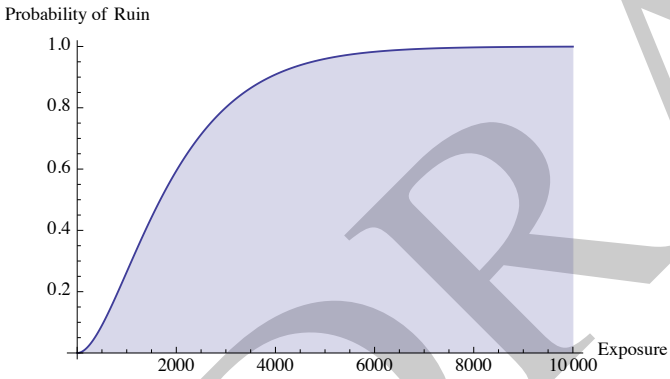

Figure 1. Why Ruin is not a Renewable Resource. No matter how small the probability, in time, something bound to hit the ruin barrier is about guaranteed to hit it.

5. Precautionary vs. Evidentiary

The idea of precautionary action is the avoidance of adverse consequences. This clashes with evidentiary action (from statistics) as one does not need evidence for precautionary action. Evidence may come too late. The non-naive PP bridges the gap between the two with the notion of fragility and concavity to a stressor (explained further down).

There are limitations of controlled experiments to reveal complexity of real world conditions and outcomes (see the "limits to phenomenology" paper, Bar-Yam, 2013). There are mathematical limitations to predictability in a complex system, "in the wild", and the effect of releasing new material in the ecology, which is why focusing on the difference between local (or isolated) and systemic threats is a central aspect of our warnings.

[1] The probability is practically close to zero, even adjusting for 1) survivorship bias, 2) risks taken by the system in its early stages, 3) such catastropic events such as the Permian mass extinction.

6. Is it True That Humans Overreact to Small Probabilities Hence the PP Would Feed Into Human Biases?

The argument is incorrect, derived by asking people questions about single probability events (Taleb Tetlock, 2014). True, researchers show that humans overreact to small probabilities, but those researchers failed to include the consequences which humans underestimate.

In the real world (outside of psychological experiments), what matters is probability $\times$ consequences. And consequences in the ecological domains can have thick tails. Overreacting to probabilities is not irrational at all when the effect is large and in the presence of ruin.

7. What We Mean by Thick Tails

The ruin problem is made simpler when we consider that only a certain class of probabilistic structures cause problems.

There are two types of probabilities: ones where the event is accompanied with well behaved mild effect (Gaussian, or Mediocristan), the other where small probability are associated with large and unpredictable consequences that have no characteristic scale (Thick Tails or Extremistan). In other words if you take a series of events, under the class of fat tails the sum will be entirely dominated by a single one: one day will have more variations than a long history.

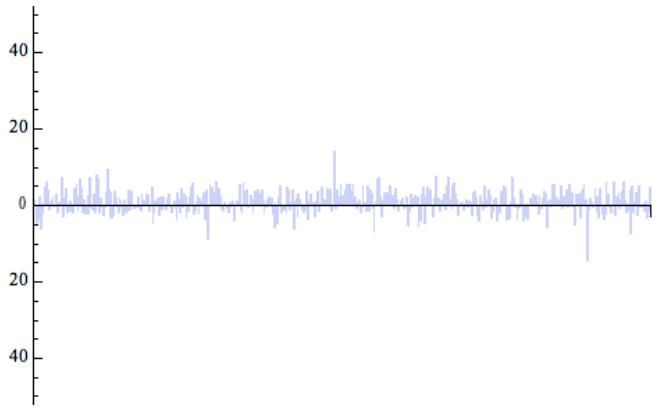

Figure 2. Thin Tails from Tinkering, Bottom Up, Broad Design. Mother Nature: no single variation represents a large share of the sum of the total variations. Even occasional mass extinctions are a blip in the total variations.

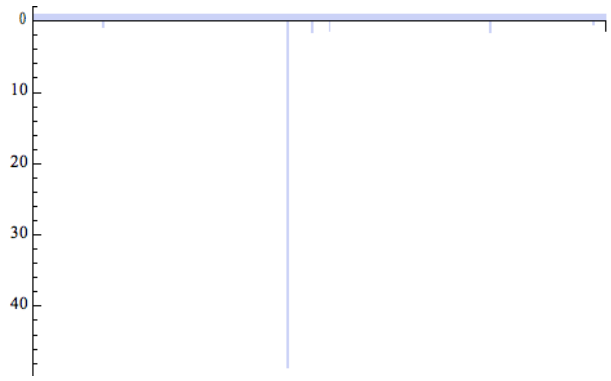

Figure 3. Fat Tails from Systemic Effects, Top-down, Concentrated Design Human made clustering of variations, where a single deviation will eventually dominate the sum. An example would be pandemics

8. Why interdependence brings fat tails

Fat tails result (among other things) from the interdependence of components, leading to aggregate variations becoming much more severe than individual ones. Interdependence disrupts the functioning of the central limit theorem, by which the aggregate is more stable than the sum of the parts. Whether components are independent or interdependent matters a lot to systemic disasters such as pandemics or generalized crises. The interdependence increases the probability of ruin, to the point of certainty.

9. The Scale —or Complexity— Argument

At the systemic scale, "in the large" nature has thin tails, even if tails are thick at the micro level. For proof see Taleb, 2014. Complex systems scientists talk about properties changing with scale, the large vs. the small.

By a statistical argument, had nature not produced thin-tailed variations in the aggregate or macro level, we would not be here today. A single one in the trillions, perhaps the trillions of trillions, of variations would have terminated life on the planet. Figures 2 and 2 3 show the difference between the two separate statistical properties.

10. Fragility as Nonlinear Response

Everything that survived is necessarily nonlinear to harm.

This explains why scale matters and why polluting small in an isolated way is not part of PP—except when it leads to systemic effects.

If I fall from a height of 10 meters I am injured more than 10 times than if I fell from a height of 1 meter, or more than 1000 times than if I fell from a height of 1 centimeter, hence I am fragile. Every additional meter, up to the point of my destruction, hurts me more than the previous one.

This nonlinear response is central for everything on planet earth, from objects to ideas to companies to technologies.

Another example. If I am hit with a big stone I will be harmed a lot more than if I were pelted serially with pebbles of the same weight. Everything that is fragile and still in existence (that is, unbroken), will be harmed more by a certain stressor of intensity X than by k times a stressor of intensity X/k, up to the point of breaking.

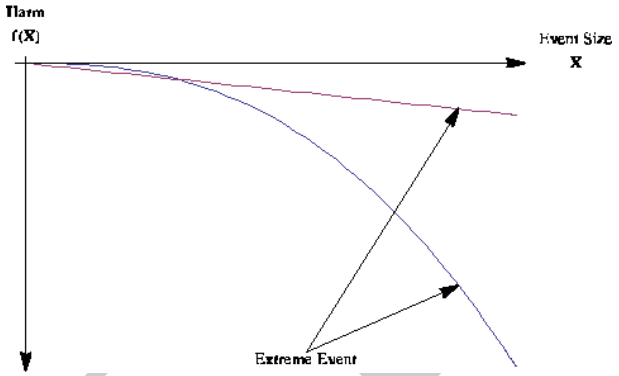

Figure 4. The nonlinear response compared to the linear.

11. Why is fragility a general rule?

This has something to do with the statistical structure of stressors, with small deviations much, much more frequent than large ones. Look at the coffee cup on the table: there are millions of recorded earthquakes every year. Simply, if the coffee cup were linearly sensitive to earthquakes, it would not have existed at all as it would have been broken in the early stages of the graph.

12. Fragility, Dose response and the $1/n$ rule

Because of nonlinearities, it is better to spread pollutants, or more generally our effect on the planet, across the broadest number of uncorrelated sources of harm, rather than concentrate.

13. Why are GMOS to be put under PP but not the nuclear?

13.1 GMOS

Genetically Modified Organisms, GMOs fall squarely under PP not because of the harm to the consumer because of their systemic risk on the system.

Top-down modifications to the system (through GMOs) are categorically and statistically different from bottom up ones (regular farming, progressive tinkering with crops, etc.) There is no comparison between the tinkering of selective breeding and the top-down engineering of arbitrarily taking a gene from an organism and putting it into another. Saying that such a product is natural misses the statistical process by which things become "natural".

What people miss is that the modification of crops impacts everyone and exports the error from the local to the global. I

do not wish to pay—or have my descendants pay—for errors by executives of Monsanto. We should exert the precautionary principle there—our non-naive version—simply because we would only discover errors after considerable and irreversible environmental damage.

13.2 Nuclear

In large quantities we should worry about an unseen risk from nuclear energy and certainly invoke the PP. In small quantities it may be OK and a matter of risk management—how small we should determine, making sure threats never cease to be local. Keep in mind that small mistakes with the storage of the nuclear are compounded by the length of time they stay around. The same with fossil fuels. The same with other sources of pollution.

The general idea is that we should limit pollution to small, very small doses, and increases the sources rather than multiply the doses even if the "scientists" promoting the sources deem any of them "safe".

14. How about the risk of famine without GMOs?

Invoking the risk of "famine" as an alternative to GMOs is a deceitful strategy, no different from urging people to play Russian roulette in order to get out of poverty.

And calling the GMO approach "scientific" betrays a very poor—indeed warped—understanding of probabilistic payoffs and risk management.

15. How About Preventive Strikes?

Preventive action needs to be limited to correcting situations *via negativa* and bring them back in line with a statistical structure that avoids ruin.

16. Fallacious arguments against PP

Next is a continously updated list of the arguments against PP that we find flawed but sounddbite- like.

16.1 Crossing the road (the paralysis argument)

Many have countered with "nothing is ever totally safe", "I take risks crossing the road every day, so listening to you I should stay home in a state of paralysis,". First answer (in *The Black Swan*) we don't cross the street blindfolded. A deeper answer is that death from road accidents are effectively thin-tailed: I do not incur the risk of generalized human extinction by crossing the street —a human life is bounded and its unavoidable termination is part of the logic of the system.

In fact the very idea of the PP is to avoid such frivolous focus.

The paralysis argument is also used to present our idea as not compatible with progress. Untrue: tinkering, bottom up progress where mistakes are bounded is *how* true progress has taken place in history.

16.2 The Loch Ness fallacy

Many have countered that we have no evidence that the Loch Ness monster doesn't exist, and to take the argument of evidence of absence being different from absence of evidence, we should act as if the Loch Ness existed. The argument is a corruption of the absence of evidence problem (paranoia is not risk management) and certainly not part of the PP, rather part of risk management.

16.3 The fallacy of misusing the naturalistic fallacy

Some people invoke "the naturalistic fallacy", which is limited to the moral domain. We are using nature for the statistical significance, not to make claims of how things "should be organized" or deriving the "ought". It is simply that scientifically, a large n cannot be ignored.

In other words we are not saying nature is the smartest possible, we are saying that time is smarter than GMO engineers. Plain statistical significance.

Cass Sunstein —a critic of the precautionary principle — claims (in a series of papers that are totally oblivious to, or uninformed of, the notion of fat tails) a "false belief that nature is benign" on the part of agents.[2] The entire method of analysis misses both the statistical significance of nature and the fact that it is not required to believe in the perfection of nature, or in the "benign" attributes, rather, in its statistical power.

16.4 The "Butterfly in India" fallacy

The statement "if I move my finger to scratch my nose, by the butterfly-in-India effect, owing to nonlinearities, I may terminate life on earth" is known to be flawed, but there has been no explanation of why it is flawed. Our thesis, can rebut it with the argument that in the aggregate nature has seen trillions of these fingers moving.

16.5 The potato fallacy

Many species were abruptly introduced into the Old World starting in the 16th Century. These did not cause environmental consequences. Some use that fact in defense of GMOs. The argument is fallacious at two levels:

First, by the fragility argument, potatoes, tomatoes, and similar "New World" goods were developed locally through progressive bottom-up tinkering in a complex system with its interractions, and had they had an impact on the environment, it would have caused adverse interractions.

Second, a counterexample is not evidence in the risk domain , as we see next with the Russian roulette fallacy.

16.6 The Russian roulette fallacy (the counterexamples in the risk domain)

The potato example, assuming it wasn't flawed and potatoes had been generated by some scientists, would still not be sufficient. Nobody says "look, the other day there was no

[2] Sunstein, Cass R., Beyond the Precautionary Principle (January 2003). U Chicago Law Economics, Ol in Working Paper No. 149; U of Chicago, Public Law Working Paper No. 38.

war, so we don't need an army", as we know better in real-life domains. Nobody argues that a giant Russian roulette with many barrels is "safe" and great money making opportunity because it didn't blow someone's brain up the other day.

More generally one needs a large sample for claims of *absence of risk* in the presence of a small probability of ruin, while a single, "n=1" example would be sufficient to counter the claims of safety —this is the Black Swan argument.

16.7 The "Are you a biologist?" fallacy

To understand the gambler's ruin problem with the miscalibration of roulette betting, we ask a probabilist, not a carpenter. No amount of expertise in carpentry can replace probabilisitic rigor in understanding the properties of long sequences of small probability bets.

There has been an "expert problem", a very poor record historically in understanding the risks of innovations in biological products, from misestimated risks of biofuels to transfat to nicotine, to a variety of products. Consider the recent major drug recalls such as Thalidomide, Fen-Phen, Tylenol, Vioxx —all of these show chronic Black Swan blindness on the part of the specialist. Yet these risks were local not systemic: with the systemic the recall happens too late, which is why we need this strong version of the PP.

16.8 The pathologization fallacy

Often narrow models reveal biases that, in fact, turn out to be rational positions, except that it is the modeler who is using an incomplete representation. Often the modelers are not familiar of the dynamics of complex systems or use Gaussian statistical methods that do not take into account of fat-tails and make inferences that would not be acceptable under different classes of probability distributions. Many biases such as the ones used by Cass Sunstein about the overestimation of the probabilities of rare events in fact correspond to the testers using a bad probability model that is thin-tailed. See *Silent Risk*, Taleb (2014) for a deeper discussion.

It became popular to claim irrationality for GMO and other skepticism on the part of the general public —not realizing that there is in fact an "expert problem" and such skepticism is healthy and even necessary for survival. For instance, in *The Rational Animal* [3], the author pathologize people for not accepting GMOs although "the World Health Organization has never found evidence of ill effects" a standard confusion of evidence of absence and absence of evidence. Such a pathologizing is similar to behavioral researchers labeling hyperbolic discounting as "irrational" when in fact it is largely the researcher who has a very narrow model and richer models make the "irrationality" go away).

These researchers fail to understand that humans may have precautionary principles against systemic risks, and can be skeptical of the untested for deeply rational reasons.

[3] Kenrick, D. T. (2013). *The rational animal: How evolution made us smarter than we think*. Basic Books.

References, Further Reading, & Technical Backup

Bar-Yam, Y., 2013, The Limits of Phenomenology: From Behaviorism to Drug Testing and Engineering Design, arXiv 1308.3094

Rauch, E. M. and Y. Bar-Yam, 2006, Long-range interactions and evolutionary stability in a predator-prey system, *Physical Review* E 73, 020903

Taleb, N.N., 2014, *Silent Risk: Lectures on Fat Tails, (Anti) Fragility, and Asymmetric Exposures*, SSRN

Taleb, N.N. and Tetlock, P.E., 2014, On the Difference between Binary Prediction and True Exposure with Implications for Forecasting Tournaments and Decision Making Research http://dx.doi.org/10.2139/ssrn.2284964

Conflicts of Interest

One of the authors (Taleb) reports having received monetary compensation for lecturing on risk management and Black Swan risks by the Institute of Nuclear Power Operations, INPO, the main association in the United States, in 2011, in the wake of the Fukushima accident.