Chapter 12. On the Difference Between a Large Stone and a Thousand Pebbles

How to punish with a stone —I landed early (once) — Why attics are always useful — On the great benefits of avoiding Heathrow unless you have a guitar —

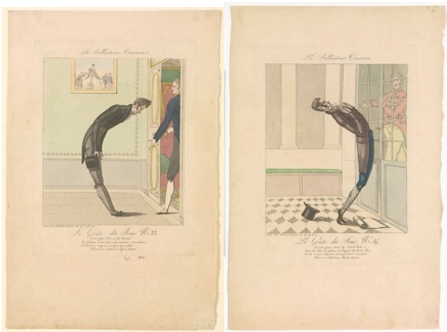

Figure 5- The solicitor knocking on doors in concave (left) and in convex (right) position. He illustrates the two forms of nonlinearity; if he were "linear" he would be upright, standing straight. This chapter will show —a refinement of Seneca's asymmetry — how one position (the convex) represents antifragility in all its forms, the other fragility (the concave) and how we can easily detect and even measure fragility by evaluating how humped "how convex" or "how concave" the courtier is standing.

I noticed looking at a coffee cup that it did not like volatility or variability or action. It just wanted calm and to be left alone in the tranquility of the home study-library. I also figured out that I had wasted a decade and a half not having noticed the point —that fragility was simply vulnerability to volatility.[1] This realization was a huge personal embarrassment for me since my specialty was the link between volatility and nonlinearity; I know, I know, a very strange specialty. So let us start with the result.

A Simple Rule to Detect the Fragile

A story, present in the rabbinical literature (*Midrash Tehillim*), probably originating from earlier Near Eastern lore, says the following. A king, angry at his son, swore that he would crush him with a large stone. After he calmed down he realized he was in trouble as a king who breaks his oath is unfit to rule. His sage advisor came up with a solution. Have the stone cut into very small pebbles, and have the mischievous son pelted with them.

The difference between a thousand pebbles and a large stone of equivalent weight is a potent illustration of how fragility stems from nonlinear effects. Nonlinear? Take for now that nonlinear means that the response is not straightforward and not a straight line, and if you double, say, the dose, you get a lot more or a lot less than double the effect —if I throw at someone's head a 10 pound stone, it will cause more than twice the harm of a 5 pounds one, more than five times the harm of a one pound, etc. It is simple: if you draw a line on a graph, with harm on the vertical axis, and size of stone on the horizontal axis, it will be curved, not a straight line. That is a refinement of asymmetry.

Now the very simple point, in fact, that allows for a detection of fragility:

> *For the fragile shocks bring higher harm as their intensity increases* (up to a certain level).

---

[1] The nontechnical reader can skip this chapter without any loss: the definition of antifragility from Seneca's asymmetry is amply sufficient for a literary read of the rest of the book. This is a more technical rephrasing of it with a few applications.

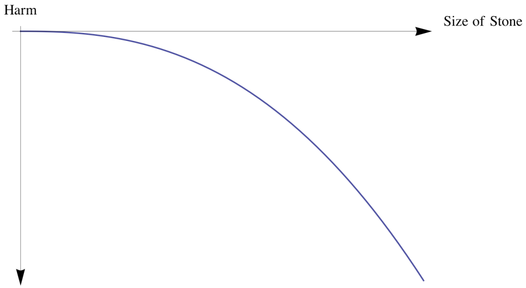

Figure 6- The King and His Son. The harm from the size of the stone as a function of the size of the stone (up to a point). Every additional weight of the stone harms more than the previous one. You see nonlinearity (the harm curves inwards, with a steeper and steeper vertical slope).

The example is shown in Figure 5. Let us generalize. Your car is fragile. If you drive it into the wall at fifty miles per hour, it would cause more damage than if you drove it into the same wall ten times at five mph. The harm at fifty miles per hour is more than ten times the harm at five mph.

Other examples. Drinking seven bottles of wine (Bordeaux) in one sitting, then water for the remaining six days is more harmful than drinking one bottle of wine a day for seven days (spread out in two glasses per meal). Every additional glass of wine harms you more than the preceding one, hence your system is fragile to alcoholic consumption. Letting a porcelain cup drop on the floor from a height of one foot (about thirty centimeters) is worse than twelve times the damage from a drop from a height of one inch (2 and a half centimeters).

Jumping from a height of thirty feet (ten meters) brings more than ten times the harm of jumping from a height of three feet (one meter) —actually thirty feet seems to be the cutoff point for death from freefall.

Note that this is a simple expansion of the foundational asymmetry we saw in Chapter x, where we used Seneca's thinking as a pretext to talk about nonlinearity. Asymmetry is necessarily nonlinearity. More harm than benefits: simply an increase in intensity brings more harm than a corresponding decrease offers benefits.

Why Is Fragility Nonlinear?

Let me explain the central argument —why is fragility necessarily in the nonlinear and not in the linear? That was the intuition from the coffee cup I mentioned in the Prologue. Just as with the large stone hurting more than the equivalent weight in pebbles, if, for a human, jumping one millimeter (an impact of small force) caused an exact linear fraction of the damage of, say, jumping to the ground from thirty feet, then the person would be already dead from cumulative harm. Actually a simple computation shows that he would have expired within hours from touching objects or pacing in his living room, given the multitude of such stressors and their total effect. The fragility that comes from linearity is immediately visible, so we rule it out because the object would be already broken and the person already dead. This leaves us with the following: what is fragile is something that is both unbroken and subjected to nonlinear effects —and extreme, rare events since hits of large size (or high speed) are rarer than ones of small size (and slow speed).

Let me rephrase it, in connection with Black Swans and extreme events. There are a lot more ordinary events than extreme events. In the financial markets, there are at least ten thousand time more events of .1% than events of 10%. There are close to eight thousand micro-earthquakes daily on planet earth, that is, those below 2 on the Richter scale —about three million a year. These are totally harmless, and, with three million per year, you would need them to be so. But shocks of intensity 6 and higher on the scale make the newspapers. Take objects such as coffee cups get a lot of hits, a million more hits of (to take an arbitrary measure), say, one hundredth of a pound per square inch than hits of a hundred pounds per square inch. Accordingly, we are necessarily immune to the *cumulative* effect of small deviations, or shocks of very small magnitude, which implies that these affect us disproportionally less (that is, nonlinearly less) than larger ones.

Let me try again and re-express my previous rule.

For the fragile, the cumulative effect of small shocks is smaller than the single effect of an equivalent single large shock.

This leaves me with the principle that the fragile is what is hurt a lot more by extreme events than a succession of intermediate ones. Finito —and there is *no other* definition.

Now let us flip the argument and consider the antifragile. Antifragility too is grounded in nonlinearties, nonlinear responses.

For the antifragile, shocks bring more benefits (equivalently, less harm) as their intensity increases —up to a point.

A simple case —known heuristically by weightlifters. Recall the discussion of bodyguard episode in Chapter x, in which I only focused on the maximum I could do. Lifting one hundred pounds once brings more benefits than lifting fifty pounds twice, and certainly a lot more than lifting one pound a hundred times. Benefits here are in weightlifter terms; strengthening the body, muscle mass, and bar-fight looks rather than resistance and ability to run a Marathon. The second fifty pounds play a larger role, hence the nonlinear (that is, we will see, *convexity*) effect. Every additional pound brings more benefits, until one gets close to the limit, what weightlifters call "failure".[1]

For now, note the reach of this simple curve: if affects about just anything in sight, even medical error, government size, innovation —anything that touches uncertainty. And it helps put the "plumbing" behind the statements on size and concentration in book 2.

When to Smile and When to Frown

Nonlinearity comes into two kinds. Concave (curves inward) as in the case of the king and the stone or its opposite convex (curves outwards). And of course, mixed, with concave and convex sections.

Figures and show the following simplifications of nonlinearity: the convex and the concave resemble a smile and a frown, respectively.

[1] Actually there are different muscle fibers, each one responds to different sets of conditions with varied asymmetries of responses. The so-called "fast-twitch" fibers, the ones used to lift very heavy objects, are very antifragile, as they are convex to weight. And they die in the absence of intensity.

Figure 7 The different types of nonlinearities. The convex (left) and the concave (right). The convex curves outward, the concave inward.

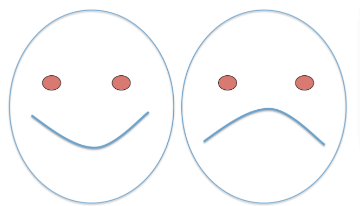

Figure 8 Smile ! A better way to understand convexity and concavity. What curves inwards looks like a smile —and what curves outwards makes a sad face. The convex (left) is antifragile, the concave (right) is fragile (has negative convexity effects).

I use the term "convexity effect" for both, in order to simplify the vocabulary, saying positive convexity effects and negative convexity effects.

Why does asymmetry map to convexity or concavity? Simply if for a given variation, you have more upside than downside and you draw the curve, it will be convex; the opposite for the concave. Figure x shows the asymmetry re-expressed in terms of nonlinearities. It also shows the Magical effect of mathematics that allowed us to treat steak tartare, entrepreneurship, and financial risk in the same breath: the convex graph turns into concave, simply when one puts a minus sign in front of it. For instance, Fat Tony had the exact opposite payoff than, say, a bank or financial institution in a certain transaction: he made a buck whenever they lost one, and vice versa. The profits and losses are mirror image of each other at the end of the day, except that one is the minus sign times the other.

Figure x also shows why the convex *likes volatility*. If you earn more than you lose from fluctuations, you want a lot of fluctuations.

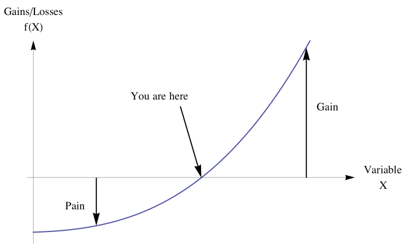

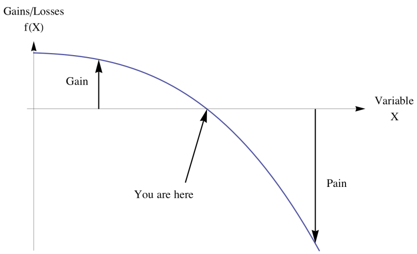

Figure 9 ***Pain More than Gain, or Gain More Than Pain****. Assume you start from the "you are here" spot. In the first case, should the variable x increase, i.e., move right on the horizontal axis, the gains (vertical axis) are larger than the losses encountered by moving left, i.e. an equivalent decrease in the variable x. The graph illustrates how positive asymmetry (first graph) turns into convex (inward curving) and negative asymmetry turn into concave (outward curving). To repeat, for a set deviation in a variable, in equivalent amounts in both directions, the convex gains more than it loses, and the reverse for the concave.*

Why is the Concave Hurt by Black Swan Events?

Now the idea that has inhabited me all my life — I never realized it could show that clearly when put in graphical form. Figure x illustrates the effect of harm and the

unexpected. The more concave an exposure, the more harm from the unexpected, and disproportionately so. So very large deviations have a disproportionately larger and larger effect.

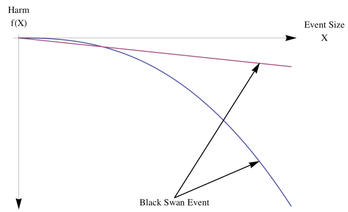

Figure 10- Two exposures, one linear, one nonlinear (with negative convexity). An unexpected event affects the nonlinear disproportionately more. The larger the event, the larger the difference

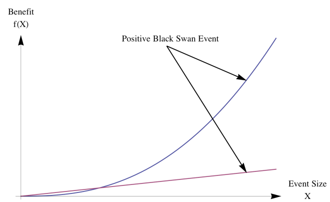

Figure 11- The mirror situation of figure x, with positive events disproportionately benefiting the convex over the linear.

Next, let us apply this very simple technique to the detection of fragility and position in the Triad.

TRAFFIC IN NEW YORK

Let us apply "convexity effects" to things around us Traffic is highly nonlinear. When I take the day flight from New York to London, and I leave my residence around 5 AM (yes, I know), it takes me around 26 minutes to reach the British Air terminal at JFK airport. At that time, New York is empty, eerily non-New York. When I leave my place at 6 AM for the later flight, there is almost no difference in travel time although traffic is a bit denser. One can add more and more cars on the highway, with no or minimal impact concerning time spent in traffic.

Then, a mystery, increase the number of cars by 10% and watch the travel time jumping by up by 50% (I am using approximate numbers). Look at the convexity effect at work: the average number of cars does not matter at all for traffic speed. If you have 90,000 cars for one hour, then 110,000 cars for another hour, traffic would be much slower than if you had 100,000 cars for two hours. Note that travel time is a negative, so I count it as a cost, like an expense, and a rise is a bad thing.

So travel cost is fragile to the *volatility* of the number of cars on the highway; it does not depend so much on their average number. Every additional car increases travel time more than the previous one.

This is a hint to a central problem of the world today, that of the misunderstanding of nonlinear response by those involved in creating “efficiencies” and “optimization” of systems. For instance, European airports and railroads are stretched, seeming overly efficient. They operate at close to maximal capacity, with minimal redundancies and idle capacity, hence acceptable costs; but a small additional congestion, say 5% more planes in the sky owing to a tiny backlog can set chaos in airports and cause scenes of unhappy travelers camping on floors; their only solace the sight of some bearded fellow playing French folk songs on his guitar.

We can see applications of the point across economic domains: central banks can print money; they print and print with no effect (and claim the “safety” of such measure) then, "unexpectedly", the printing causes a jump in inflation. Many economic results are completely cancelled by convexity effects —and the happy news is that we know why. Alas, the tools (and culture) of policy makers are based on the overly linear, ignoring these hidden effects. They call it “approximation”. When you hear of “second order” effect, it means convexity is causing the failure of approximation to represent the real story.

I have put a (very hypothetical) graph of the response of traffic to cars on the road in Figure x. Note for now the curved shape of the graph. It curves inward.

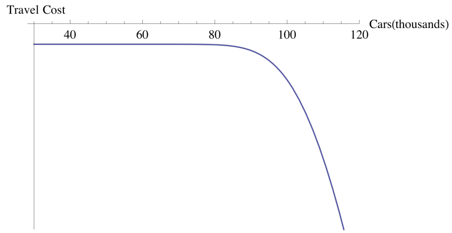

Figure 12- The graph (vertical) shows how the author’s travel time (and travel costs) to JFK depend, beyond a certain point, nonlinearly on the number of cars. We show travel costs as curving inward —concave, not a good thing.

Someone Call New York City Officials

An apt illustration of how convexity effects affect an overoptimized system, along with misforecasting large deviations is this simple story of a underestimation made by New York City officials of the effect of a line closure on traffic jam. This error is remarkably general: a small modification with compounded results in a system that is extremely stretched, hence fragile.

One Saturday evening in November 2011, I drove to New York City to meet the philosopher Paul Boghossian for dinner in the village —typically a forty minute trip. Ironically, I was meeting him to talk about my book, this book, and more particularly, the ideas on redundancy in systems. I have been advocating the injection of redundancy in people’s lives and had been boasting to him and others that, since my new year resolution of 2007, I have never been late to anything, not even by a minute (well, almost). Recall in Chapter x my advocacy of redundancies as aggressive stance. Well, such personal discipline forces me to build buffers, and, as I carry a notebook, it allowed me to write an entire book of aphorisms. And long visits to bookstores. Or I can sit in a café and read hate mail. And of course, no stress as I have no fear of being late. But the greatest benefit of such discipline is that it prevents me from cramming appointments (typically, appointments are neither useful nor pleasant). Actually by another rule of personal discipline I do not make appointments (outside of lectures) except the very same morning as a date on the calendar makes me feel like a prisoner, but that’s another story.

At some point as I hit Midtown, traffic stopped around 6 P.M. Completely. By 8 P.M. I had hardly moved a few blocks. So even my “redundancy buffer” failed to let me keep the so-far unbroken resolution. Then after relearning to operate the noisy cacophonic thing called the radio, I started figuring out what had happened: New York City had authorized a film company to use the 59th Street Bridge, blocking part of it, assuming that it would be no problem on a Saturday. And the small traffic problem turned into mayhem, owing to the multiplicative effects. What they felt would be at the worst a few minutes delays was multiplied by two orders of magnitude; minutes became hours. Simply, the authorities running New York City did not understand nonlinearities.

This is the central problem of modern complexity: these type of errors compound, multiply, swell, with an effect that only goes in one direction —the wrong direction.

WHERE MORE IS DIFFERENT

Another intuitive way to look at convexity effects: consider the scaling property. If you double the exposure to something, do you more than double the harm it will cause? If so, then this is a situation of fragility. Otherwise, you are antifragile.

The point has been aptly expressed by P.W. Anderson with the title his paper: “more is different”[xciv]. And what scientists involved in complexity call “emerging properties” is the nonlinear result of adding units, as the sum becomes increasingly different from the parts. Just look how different the large stone is from the pebbles: the latter have the same weight and the same general shape, but that’s about it. Likewise, we saw in Chapter x that a city is not a large village; a corporation is not a larger small business. We also saw how the randomness changes in nature from Mediocristan to Extremistan, how a state was not a large village, and many alterations that come from size —and speed. All these show nonlinearity in action.

A “Balanced Meal”

Another example of missing the hidden dimension, that is, variability. We are currently told by the Soviet-Harvard U.S. health authorities to eat set quantities of nutrients (total calories, protein, vitamins, etc.) every day, in some recommended amounts, say the optimal amounts of each. Every food item has a “percentage daily allowance”. Aside from the total lack of empirical rigor in the way these recommendations are currently derived (more on that in the medical chapters),

there is another sloppiness in the edict: an insistence in the discourse on the *regularity*. Those recommending the nutritional policies fail to understand that “steadily” getting your calories and nutrients throughout the day, with “balanced” composition and metronomic regularity does not necessarily carry the same effect as having them unevenly or randomly distributed, say by having a lot of proteins one day, fasting completely another, feasting the third, etc.

This is a denial of hormesis, the slight stressor of episodic deprivation. For a long time, nobody even bothered to try to figure out whether variability in distribution —the second order effect —mattered just as much as long term composition. Now research is starting to catch up to such very, very simple point. It turned out that the effect of variability in food sources and the nonlinearity in the physiological response is central to biological systems. Consuming no protein at all on Monday, and catching up on Wednesday seemingly causes a different —better— physiological response, possibly because the deprivation, as a stressor, activates some pathways that facilitate the subsequent absorption of the nutrients (or something similar). And, until a few recent (and disconnected) empirical studies, this convexity effect has been totally missed by science —though not by religions, ancestral heuristics, and traditions. And if scientists get some convexity effects (as we said about domain dependence, doctors, just like weightlifters, understand here and there nonlinearities in dose-response), the notion of convexity effects itself, though, appears completely missed from the language and methods.

Walk, Don’t Run

Another illustration, this time a situation that benefits from variation—positive convexity effects. Take two brothers, Castor and Polydeuces, who need to travel a mile. Castor walks the mile at a leisurely pace and arrives at destination in twenty minutes. Polydeuces spends fourteen minutes playing with his handheld device getting updates on the gossip, then runs the same mile in six minutes, to arrive at the same time as Castor.

So both persons have covered the exact same distance, in exactly the same time —same average. Castor who walked all the way presumably will not get the same health benefits and gains in strength as Polydeuces who sprinted. Health benefits are *convex* to speed (up to a point, of course).

The very idea of exercise is to gain from antifragility to workout stressors —as we saw, lifting weight, exerting exercise are just exploitations of convexity effects.

Small May Be Ugly, It Is Certainly Less fragile

We often hear the expression "small is beautiful". It is potent and appealing; many ideas have been offered in its support —mostly all of them anecdotal, romantic, or existential. Next, let us present it without our approach of *fragility equals concavity equals dislike of randomness* and see how we can measure such effect.

How to Be Squeezed

A squeeze is then people have no choice but to do something, and do it right away, regardless of the costs.

Your other half is to defend a doctoral thesis in the History of German Dance and you need to fly to Marburg to be present at such an important moment, meet the parents, and get formally engaged. You live New York and managed to buy an economy ticket to Frankfurt for $400 and you are excited about how cheap it is. But you need to go through London. Upon getting to New York's Kennedy airport, you are apprised by the airline agent that the flights to London are cancelled, sorry, delays due to backlog due to weather problems, that type of thing. Something about Heathrow's fragility. You can get a last minute flight to Frankfurt, but, now you need to pay $4000, close to ten times the price, and hurry as there are very few left. You fume, shout, curse, blame yourself, your upbringing and parents who taught you to save, then shell out the $4000. That's a squeeze.

We can apply the idea of convexity effects to these forced actions —in fact the exact same graph as the one about travel time — on the idea of size. When one is large, one becomes vulnerable to errors going in one direction, horrendous squeezes. Because of the squeeze effects are nonlinearly costlier.

To see how size becomes a handicap, consider the reasons one should not own an elephant as a pet, regardless of what emotional attachment you may have with the animal. Say you can afford an elephant as part of your post-promotion household budget and have one delivered to your backyard. Should there be a water shortage —hence a squeeze since you have no choice but to shell out the money for water — you would have to pay a higher and higher price for each additional gallon of water. That's fragility, right there, a negative convexity effect coming from getting too big. The unexpected cost, in percentage of the total, would be monstrous. Owning, say, a cat or a dog would not bring about such high unexpected additional costs over the regular at times of squeeze — the overruns taken as a percentage of the total costs would be very low.

In spite of what is studied in Business School concerning "economies of scale", size hurts you at times of stress; it is not a good idea to be large during difficult times. Some economists have been wondering why mergers of corporations do not appear to play out. The combined unit is now much larger, hence more powerful and according to the theories of economies of scale, it should be more "efficient". But the numbers show no gain from such increase in size, at best —that was already true in 1978, as Richard Roll then voiced the "hubris hypothesis" finding it irrational for companies to engage in mergers given the poor historical record. Recent data, more than three decades later, still confirm both the poor record of mergers and the same hubris as managers seem to ignore the bad economic aspect of the transaction. There appear to be something with size that is harmful to corporations.

As with the idea of having elephant as pets, squeezes are much, much more expensive (relative to size) for large corporations. The gains from size are visible but the risks are hidden, and some concealed risks seems to bring frailties into the companies.

Large animals, like elephants, boa constrictors, mammoths, and others animals of size tend to become rapidly extinct. Aside from the squeeze, large animals are more fragile than small ones. An explanation by Jared Diamond "Why cats have nine lives" is as follows[xcv]. Large animals are naturally more fragile. If you throw a cat or a mouse from an elevation of several times their size, they manage to survive. Elephants, by comparison, break limbs very easily.

Kerviel and MicroKerviel

Let us look at a case study from vulgar finance, a field in which participants are very good at making mistakes. On January 21, 2008, the Parisian bank Societé Générale rushed to sell in the market close to seventy billion dollars of stocks, a very large amount for any single "fire sale". Markets were not very active (called "thin") as it was Martin Luther King day in the United States and markets worldwide dropped precipitously, close to ten percent, costing the company close to six billion dollars in losses just from their fire sale. The entire point of the squeeze is that they couldn't wait, and they had no option but to turn a sale into a fire sale. For they had, over the weekend, uncovered a fraud. Jerome Kerviel, a rogue back office employee, was playing with humongous sums in the market and hiding these exposure from the main computer system. They had no choice but to sell, immediately, sell these stocks they didn't know they owned.

Now to see the effect of fragility from size, look at Figure x showing losses as a function of quantity sold. A fire sale of 70 billion worth of stocks leads to a loss of 6 billion. But a fire sale of a tenth of the size, 7 billion would have no loss at all, as markets would absorb the quantities without panic, maybe without even noticing. So this tells us that if, instead of having one very large bank, with Monsieur Kerviel as a rogue trader, we had ten smaller banks, each with a proportional Monsieur Mini-Kerviel, and each had his rogue trading independently and at random times, the total losses for the ten banks would be close to nothing.

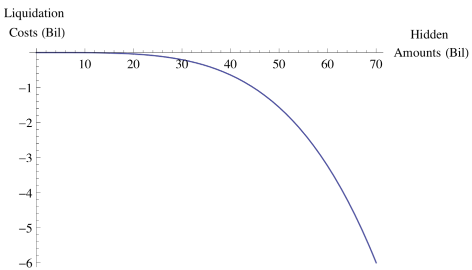

Figure 13 Small may be beautiful; it is certainly less fragile. The graph shows transaction costs as a function of the size of the error: they increase nonlinearly, and we can see the mega-fragility.

About a few weeks before the Kerviel episode, a French business school hired me to present to the board of executives of the Societé Générale meeting in Prague my ideas of Black Swan risks. [xcvi]In the eyes of the bankers, I was like a Jesuit preacher visiting Mecca in the middle of the annual Hajj —their "quants" and risk people hated me with passion and I regretted not having insisted on speaking in Arabic given that they had simultaneous translation. My talk was about how pseudo risk techniques à la Triffon —methods commonly used as I said to measure and predict events, that have never worked before and how we needed to focus on fragility and barbells. During the talk I was heckled relentlessly by Kerviel's boss and his colleague, the head of risk management. After my talk everyone ignored me, as if I were a Martian, with a "who brought this guy here" awkward situation (I had been selected by the school, not the bank). The only person who was nice to me

was the chairman as he mistook me for someone else and had no clue about what I was discussing.

So the reader can imagine my state of mind when, shortly after my return to New York, the Kerviel trading scandal broke. It was also teasing that I had to keep my mouth shut (which I did except for a few slips) for legal reasons.

Clearly, the subsequent *post mortem* analyses were plainly mistaken, attributing the problem to *bad* controls by the *bad* capitalistic system, and lack of vigilance on the part of the bank. It was not. Nor was it "greed" as we commonly blame. The problem is, primarily, size and the fragility that comes from size.

Always keep in mind the difference between a stone and its weight in pebbles. The Kerviel story is illustrative, but we can generalize and look at evidence across domains.

In project management, Bent Flyvbjerg has shown firm evidence that an increase in the size of projects maps to poor outcomes and higher and higher costs of delays as a proportion of the total budget. But there is a nuance: it is the size per segment of the project that matters, not the entire project—some projects can be divided into pieces, not others. Bridge and tunnel projects, involve monolithic planning, as these cannot be broken-up in small portions; their percentage costs overruns increase markedly with size. For roads, built by small segments, there is no serious size effect—as the project managers incur only small errors, and can adapt to them. Small segments go one small error at the time, with no serious role for squeezes.

How to Exit a Movie Theatre

Another example of the costs of a squeeze. Imagine how people exit a movie theatre. Someone shouts "fire", and you have a dozen persons squashed to death. So we have a fragility of the theater to size, stemming from the fact that every additional person exiting brings more and more trauma (such disproportional harm is a negative convexity effect). A thousand people exiting (or trying to exit) in one minute is not the same as the same number exiting in half an hour. Someone unfamiliar with the business who naively *optimizes* the size of the place (as Heathrow airport, for example) might miss the idea that smooth functioning at regular times is different from the rough functioning at times of stress.

It so happens that modern economic optimized life causes us to build larger and larger theaters, but with the exact same door. They no longer make this mistake too often while building cinemas, movie theaters and stadiums, but we tend to make the mistake in other domains, such as, for instance, natural resources and food supplies. Just consider that the price of wheat more than tripled in the years 2004-2007 in response to a small increase in demand, around one percent.[1]

PROJECTS AND PREDICTION

Why Planes Don't Arrive Early

Let us start as usual with a transportation problem, and generalize to other areas. Travelers (typically) do not like uncertainty —especially when they are on a set schedule. Why? There is a one way effect.

I've taken the very same London-New York flight most of my life. The flight takes about 7 hours, the equivalent of a short book plus a brief polite chat with a neighbor and a meal with Port wine, stilton cheese and crackers. I recall a few instances in which I arrived early, about twenty minutes, no more. But there have been instances in which I got there more than two and three hours late, and, in at least one instance, it has taken me more than two days to reach my destination.

Because travel time cannot be really negative, uncertainty tends to cause delays, making arrival time increase, almost never decrease. Or it makes arrival time just decrease by minutes, and increase by hours, an obvious asymmetry. Anything unexpected, any shock, any volatility is much more likely to extend the total flying time.

This also explains the irreversibility of time, in a way, if you consider the passage of time as an increase in disorder.

Let us now apply to concept to projects. Just as when you add uncertainty to a flight, the planes tend to land later, not earlier (and these laws of physics are so universal that they even work in Russia), when you add uncertainty to projects, they tend to cost more and take longer to complete. This applies to many, in fact, almost all projects.

The interpretation I had in the past was that a psychological bias, the underestimation of the random structure of the world, was the cause behind such

[1] The other problem is that of the misunderstanding of the nonlinearity of natural resources, or anything particularly scarce and vital. Economists have the so-called law of scarcity, by which things increase in value according to their demand —but they ignore the consequences of nonlinearities. My former thesis director Hélyette Geman and I are currently studying a "law of convexity" that makes commodities, particularly vital ones, even more dear than previously thought.

underestimation —projects take longer than planned because the estimates are too optimistic. We have evidence of such bias, called overconfidence. Decision scientists and business psychologists have theorized something called the "planning fallacy", in which they try to explain the fact that projects take longer rarely shorter using psychological factors.

But the puzzle was that such underestimation did not seem to exist in the past century or so and we were dealing with the very same humans, endowed with the same biases. Many large-scale projects one century and a half ago were completed on time; many of the tall buildings and monuments we see today are not just more elegant that modernistic structure, but they were completed within, and often ahead of schedule. These include not just the Empire State Building (still standing in New York), but such items such as the London Crystal Palace erected during the Great Exhibition of 1851, the hallmark of Victorian reign, based on the inventive ideas of a gardener. The Palace, which housed the exhibition, went from its organization to the grand opening in just nine months. The building took the form of a massive glass house, 1848 feet long by 454 feet wide; it was constructed from cast iron-frame components and glass made almost exclusively in Birmingham and Smethwick.

The obvious is usually missed here: the Crystal Palace project did not use computers, and the parts were built not far away from the source, with a small number of businesses involved in the food chain. Further, there were no business schools at the time to teach something called "project management" and increase overconfidence. There were no consulting firms. The agency problem (which we defined as the divergence between the interest of the agent and his client) was weak. In other words, it was a much more linear economy —less complex— than today. And we have more nonlinearities —asymmetries, convexities —in today's world.

Black Swan effects are necessarily increasing, as a result of complexity, interdependence between parts, globalization, and the beastly thing called "efficiency" that make people now sail too close to the wind. Add to that consultants and business schools. One problem somewhere can halt the entire project —so the projects tend to get as weak as their weakest link in the chain (an acute negative convexity effect). The world is getting less and less predictable, and we rely more and more on technologies that have errors and interactions that are harder to estimate, let alone predict.

And the information economy is the culprit. Bent Flyvbjerg, the one of bridges and road earlier in the chapter, showed another result. The problem of cost overruns and delays is much more acute in the presence of Information

Technologies (IT), as computer projects cause a large share of these costs overruns and it is better to focus on these principally.[xcvii] But even outside of these IT-heavy projects, we tend to have very severe delays.

But the logic is simple: again negative convexity effects are the main culprit, a direct and visible cause. There is an asymmetry in the way errors hit you— the same as with travel.

No psychologist who discussed the "planning fallacy" has realized that, at the core, it is not essentially a psychological problem, not an issue with human errors, but part of the nonlinear structure of the projects. Just as time cannot be negative, a three month project cannot be completed in zero or negative time. So, on a time line going left to right, errors add to the right end, not the left end of it. If uncertainty were linear we would observe some projects completed extremely early (just as we would arrive sometimes early, sometimes late). But this is not the case.

Wars, Deficits, and Bonds

The Great War was estimated to last only a few month; by the time it was over it got France and Britain heavily in debt; they incurred at least ten times what they thought their financial costs would be, aside from all the horrors, sufferings, and destructions. The same of course for the second war added to the U.K. debt, causing it to become heavily indebted, mostly to the United States.

In the United States the prime example remains the Iraq war, expected by George W. Bush and his friends to cost thirty to sixty billions, and so far taking into account all the indirect costs, may have swelled to more than two trillion —indirect costs multiply, causing chains, explosive chains of interactions, all going in the same direction of more costs, not less. Complexity plus asymmetry (plus such types as George W. Bush), once again, leading to explosive errors.

The larger the military, the disproportionally larger the cost overruns.

But wars —with more than twentyfold errors — are only illustrative of the way governments underestimate explosive nonlinearities (convexity effects) and why they should not be trusted with finances or any large scale decisions. Indeed, governments do not need wars at all to get us in trouble with deficits: the underestimation of the costs of their projects is chronic for the very same reason ninety-eight percent of modern projects have overruns. They just end up spending more than they tell us. This has led me to install a governmental golden rule: no borrowing allowed, forced fiscal balance.

WHERE THE "EFFICIENT" IS NOT EFFICIENT

We can easily see the costs of fragility swelling in front of us, visible to the naked eye. Global disaster costs are today more than three times what they were in the 1980s, adjusting for inflation. The effect, noted a while ago by the visionary researcher on extreme events Daniel Zajdenweber, seems to be accelerating[xcviii]. The economy can get more and more "efficient", but fragility is causing costs of errors to be higher.

And naive cost-benefit analyses, can be a bit harmful, an effect that of course swells with size. For instance, the French focused on nuclear energy as it seemed "clean" and cheap. And "optimal" on a computer screen. Then, after the wake-up call of the Fukushima disaster of 2011, they realized that they needed additional safety features and scrambled to add them, at any cost. In a way this is similar to the squeeze I mentioned earlier: they are forced to invest, regardless of price. This additional safety came at a heavy cost. Such additional expense was not part of the initial cost-benefit that entered the initial decision and looked good on a computer screen. So when deciding on one source of fuel against another, or similar comparisons, we do not realize that model error may hit one side more than the other.

Pollution and Harm to the Planet

From this we can generate a simple ecological policy. We know that fossil fuels are harmful in a nonlinear way. The harm is necessarily concave (if a little bit of it is devoid of harm, a lot can cause climatic disturbances). While on epistemological grounds, because of opacity, I do not need to believe in anthropogenic climate change (caused by humans) in order to be ecologically conservative, we can put these convexity effects to use in producing a risk-management rule for pollution. Simply, just as with size, split your sources of pollution among many other natural sources. The harm from polluting with ten different sources is smaller than the equivalent pollution from a single source.

Well, let's look at nature's mechanisms at regulating the concentration. We modern humans do go to the supermarket to purchase the same items, say tuna, and because of sticky modern habits, rigidity of factories, extreme consumption of tuna can hurt other animals, mess with the ecosystem, and lead species to extinction. But ancestral humans did it differently[xcix]. Jennifer Dunne, a complexity researcher who studies hunter gatherers, examined evidence about the behavior of the Aleuts, a North American native tribe, for which we have ample data, covering five millennia. They exhibit a remarkable lack of concentration in their predatorial behavior, with a behavior of prey-switching. Whenever they started to become low on a resource, they switched to another one, as if to preserve the ecosystem. So they understood convexity effects —or, rather, their habits did.

Conclusion

To conclude this chapter, fragility in any domain, from a coffee cup to an organism, to a political system, to the size of a firm, to delays in airports, resides in the nonlinear.

Let us take this fragility into another application: model error.