

## Black Swans<sup>1</sup>

February 2, 2009

Richard S. Newrock

From the International Herald Tribune, November, 10, 2008:

- "We've got the right people in place, as well as good risk management and controls." – E. Stanley O'Neal, the CEO of Merrill Lynch in 2005.  
“[...] as 2006 drew to a close [...] Merrill's decision to invest heavily in the mortgage industry was paying off handsomely. [...] on December 30<sup>th</sup> of that year, it [...] doubled-down by paying 1.3 billion dollars for First Franklin, a lender specializing in risky mortgages.”
  
- From “Diploma to Nowhere,” data for the State of Ohio for 2003:
  - 37% of all first-time freshman in two and four year institutions took at least one remedial course.
  - 101.7 million dollars were the cost to Ohio institutions for remedial education.The data for the other 49 states is little different.
  
- On the 24<sup>th</sup> of November 1859, “On the Origin of Species” by Charles Darwin was published.

Credit default swaps ... Toxic mortgages ... Poor public school education ... Darwin. Oddly enough, these are all related; they all belong to a special class of events now known as Black Swans.

---

<sup>1</sup> This paper is based on the work of Nassim Nicholas Taleb, in particular, “The Black Swan.” (Random House, 2007). .

I began to learn about the highly improbable and Black Swans a few years ago – and decided to present a paper to this esteemed group. The paper was to be delivered a year ago December. Unfortunately, I had my own Black Swan event – a heart attack – and Albert gave me a year’s reprieve. But given the state of the economy, my timing is even better – except that I won’t seem as prescient.

This paper is about highly improbable events – and why they are improbable – or at least, seem so. I won’t say much about their impact and I will not be discussing how to predict or plan for them I will focus solely on why humans don’t account for them properly. And let me warn you from the outset: the topic is difficult and the paper lengthy. So those of you who have imbibed too much liquid refreshment may find it difficult to follow – I urge you to begin your nap now.

The use of the term “Black Swan” for an event of extremely low probability was, I believe, coined by Nassim Nicholas Taleb, a Wall Street “Quant” or Quantitative Analyst and Professor of Mathematics. This paper is a synopsis of some of his work. The term “Black Swan refers back several hundred years to the Age of Exploration, when naturalists held the unassailable belief that “all swans were white.” This was confirmed time and again by all the available evidence – no one ever observed anything but white swans. But when Australia was discovered, so were black swans. Of course, the birds’ plumage isn’t the point. The issue is the severe limitation humans have for learning by observation and experience, and, ultimately, how fragile our knowledge is. One single observation can destroy a rule based on some millennia of observation – a fact well-known to physicists.<sup>2</sup>

A Black Swan event has three characteristics:

- It is an outlier – nothing from the past really points to its possibility;
- It creates a major impact on something; and
- In spite of being an outlier, humans try to concoct an explanation for its occurrence. This is always done after the fact, as we attempt to make it explainable and predictable.

---

<sup>2</sup> Consider the evolution of the principle of Conservation of Energy, as an example.

Rarity, extreme impact and retrospective predictability.

Black Swan events are becoming ever more important as ordinary, predictable events become less and less consequential. Think of some of the life-changing events of the past several decades: the fall of the Soviet Bloc; the failure of savings banks; the current economic crisis; the rise of Islamic fundamentalism; the use of cell phones; and the spread of the internet. The combination of low predictability with large impact makes Black Swans interesting. But we – homo sapiens – act as if they do not exist. This is especially true in the social sciences and, more specifically, in the finance and gaming industries, where people generally act under the incorrect belief that they use tools that can measure uncertainty. Telephone your portfolio manager and ask him for a definition of “risk.” The odds are overwhelming that he will supply a definition that excludes Black Swans – and is therefore useless. Or ask the owner of a casino to define risk – and you get the same worthless answer. Historians are, of course, far worse. They are notorious for looking backwards, cherry-picking past events to identify harbingers of a Black Swan.

We have a deep question here – why are we blind to randomness? Why do we all – scientists and laymen alike – focus on the minutia and not on the significant? Or, as Taleb asks, “Why does reading the newspaper decrease your knowledge of the world?”

Each of our lives is the cumulative effect of a handful of large events. Think about your own life and count the significant events, the technological innovations, the new inventions since you were born. How many came to you on a set schedule? How often did meeting your mate, choosing your occupation, achieving your success, your children’s success, occur according to plan?

With Black Swans, what you don’t know is far more important than what you do know. Black Swans are caused or exacerbated by their being unexpected. Consider the events of 9-11. If they had been expected they wouldn’t have happened – we would have taken steps to prevent them. What you know becomes inconsequential if your enemy knows that you know it.

And there are huge payoffs in the unknown, where you have little to lose and much to gain. Consider venture capital investments or scientific discoveries – rare events result in new billion dollar companies or Nobel prizes. And contrary to most social scientists’ wisdom, very few new discoveries or technologies arise from design and planning. They are Black Swans.

But, in our very nature there are impediments: “we don’t learn that we don’t learn.” Most of us scorn the abstract. What did we learn from 9-11? Did we learn that some events are not predictable? No. Did we learn that there are built-in defects to conventional wisdom? No. We learned precise rules for avoiding terrorists on airplanes and tall buildings.

What are our brains for? They don’t appear to be intended for introspective thinking. Why not? Because your introspective ancestor would have been eaten by a lion while my non-thinking, but quicker-reacting forefather was running for cover. Thinking is time-consuming and energy-intensive. We humans were preceded by 100 million years of non-thinking, quick-reacting mammals. We’ve spent only tens of thousands of years as a thinking species – and when we do think more often than not it is about ways to get laid.

Our minds like to focus on well-defined forms. We tend to mistake the map for the territory. Once we form mental constructs, we favor them over messier, less tractable ideas. Maps, our mental models and constructions, are not always wrong – just, in certain applications, wrongly used. The problem is you do not know beforehand where the map will be wrong and where mistakes can lead to severe consequences. Think of a mental model as if it were a very helpful prescription drug – but one with random and extremely severe side effects.

We all have certain mind-sets. Where they intersect with reality, where the gap between what you know and what you don’t know exists, a gap that can be very wide, that is where Black Swans are “produced.” It is our drive to focus on what makes sense to us, to form mental maps, that gets us into trouble. We need to be imaginative to survive; we generally lack imagination and we often repress it in others.

David Brooks, in the NY Times, referred to this when he wrote:

“Perceiving a situation seems, at first glimpse, a remarkably simple operation. You just look and see what’s around. But the operation that seems the most simple is actually the most complex, it’s just that most of the action takes place below the level of awareness. Looking at and perceiving the world is an active process of meaning-making that shapes and biases the rest of the decision-making chain”

“Below the level of awareness” are key words here, as we’ll see.

Taleb refers to the “Triplet of Opacity.” History is opaque; you see the results but not the generator. When the mind comes into contact with history three things happen:

- There is an illusion of understanding. You think you know what is going on in a world that is essentially random;
- There is a retrospective distortion – we can only assess matters after the fact. History is much more organized in books than in reality. Scientific discovery occurs in leaps and jumps but the descriptions of science are lucid and mono-directional; and
- Factual information is overvalued and learned people make mental models far too soon.

History, science and society do not move smoothly from one nexus to the next; they don’t crawl, they jump.

- No one predicted Christianity would become the dominant religion of the western world. Roman scribes of the early era didn’t even mention it – apparently the ideas of a heretical Jew were not deemed to have staying power.
- Seven centuries later a collection of disreputable horsemen rode from the desert and spread Islamic law from India to Spain. It was absolutely unpredictable.
- The rise and impact of the Internet and the World Wide Web were completely unforeseen.

Using a mathematical analogy, societies are chaotic and fractile. They jump suddenly from fracture to fracture. But historians tend to think of history as predictable with little increments of progress.

To fully explore this I need to introduce a few concepts. We'll start with scalability, as applied to professions, and extrema, as applied to the world around us.

Some professions are not scalable – you get paid by the hour or the piece. Your earnings are limited by how much you actually work. The efforts of cobblers, plumbers, mechanics, physicians and lawyers cannot be scaled. A prostitute is paid by the trick and there are only so many tricks she can turn a day. And, she must be present to perform the service. Similarly, a chef-owner can do no better than filling all his tables once or twice a night. In such professions, no matter how highly you are paid, your revenue depends more on your *efforts* and less on the *quality* of your decisions. The work to be done is predictable; it may vary slightly but no one day is any more significant than any other. That is, your income isn't going to be Black Swan driven.

But some professions allow you to multiply your income significantly with little or no extra effort. If you are a movie or TV actor, a member of a rock band or a country singer, you only have to do the work once; sound engineers and projectionists do the rest of the work. This is true for writers, and in many cases, artists. The author expends the same effort writing a book and getting one person to read it as in getting several million to read it. The same is true for the finance industry. The same analysis and preparation go into buying 10 shares of stock as in buying 100,000 shares. So we separate professions into unscalable and scalable: Prostitutes and artists, doctors and movie stars. “Obviously” scalable professions are better. You can really earn a lot. But there is a huge problem. A scalable profession is only good if you are successful. Scalable professions are highly competitive and always produce huge inequalities; the results are essentially random, and large disparities exist between effort and reward.

When did scalable professions first appear? When did things occur that introduced significant disparities in income and wealth? An example is Edison's invention of sound recording. One hundred and fifty years ago, if you wanted to hear music, you listened to local people. Any decent pianist could give concerts and charge enough to earn a living. And he or she needed to

be present at every performance. The work could not be stored and replayed. Then came recording and reproduction. Now, I download a music file from Amazon to my iPod, and listen to Vladimir Horowitz whenever and wherever I want. Indeed, I'd rather spend \$15 for a CD by Horowitz than \$3 for one by an equally talented but unknown CCM student. The invention of recording made the music profession scalable.

But scalability didn't start with recording. It probably started with evolution. Evolution is scalable – the DNA that wins, for whatever reason, reproduces. Other DNA vanishes. But that doesn't concern us as we are focused here on society. As far as societies are concerned, the big transition was probably the invention of the alphabet, which enables us to store information and reproduce it. And then came the printing press. Now we didn't need bards and troubadours. We reproduce stories and ideas with extremely high fidelity, without limit and with no expenditure of additional energy by the author. The author doesn't even have to be alive for the subsequent performances. Indeed death is often a great career move for authors<sup>3</sup>. Probably applies to artists as well.

In the old days, everyone had an audience. A story teller had a market, just like a cobbler or a butcher – and no one from far away could displace him and take away his livelihood. Now a few, often lucky people, who for some reason get early attention, quickly supplant everyone else on the bookshelves. One can think of other inventions that created scalable professions – movies, stock tickers, lithography, among others.

This scalable-non-scalable concept also helps us distinguish between two types of randomness. Consider a gedanken – or thought – experiment. Gather together 1,000 randomly selected people. Cincinnatians, New Yorkers, Parisians, vegetarians, vegans, Darwinists, terrorists, Dervishes, Arabs, anthropologists, artists, and lawyers. Throw in a few Creationists, if you want. Weigh them and determine the average weight. Now add heaviest man you know – assume he weighs 500 pounds. He will represent only ½ of 1 % of the average weight of 1,000 people. Percentage-wise, even the heaviest person on the planet wouldn't represent much more. Go to

---

<sup>3</sup> Taleb's joke not mine!

10,000 randomly selected people and the heaviest person on the planet's contribution to the average weight is negligible. This obviously works for height, also. Or the food you consume in a year – two days feasting – Christmas and Thanksgiving – doesn't much affect your yearly caloric intake. Or shoe sizes. The point is, for these sorts of things, if the sample is sufficiently large, no individual event can significantly affect the average. Furthermore, If I told you I knew someone who weighed 50 tons or was 500 feet tall you'd not believe me, based on your own experiences.

Now, continuing the gedanken experiment, let's total the net worth of those 1,000 people. Now add Warren Buffet, who has a net worth of 40 or so billion dollars (or did). He alone represents 99.9% of the total capital of our randomly selected group. In fact, if we total the net worth of the 1000 people, plus that of all those in this room, the total is likely to be less than the rounding error on Buffet's net worth.

This works with book sales as well. Take a random selection of 1000 authors and add J. K. Rowling, who has sold several hundred million books. The others collectively may have a couple of hundred thousand readers. You can do this with academic citations, company size, incomes, CD's sold, etcetera. In all these situations one single observation can disproportionately impact the aggregate.

So social matters, which are man-made, are somehow different from physical matters. One is an average world; the other is an extreme one. Taleb calls them Mediocristan and Extremistan, playing on the definition of "stan." Extremistan produces Black Swans. If you live in Mediocristan, you can be comfortable with what you've measured, and what you have learned from the measurements. Provided you *know* that you are in the Mediocristan, it is not possible to have one single event that dominates a phenomenon; there are no Black Swans in Mediocristan. But if you live in Extremistan, you have big problems. You cannot derive useful averages from a random sample, since so much depends on single observations. One unit can affect the total in a very disproportionate way and you must always be suspicious of any conclusions that you derive from data. We have two kinds of randomness.

I note parenthetically that luck plays a huge role in Extremistan. But you need to be a bit careful, especially in things that require a technical skill – in piano playing or surgery, skill matters. Real inequities occur in Extremistan when somebody perceived as being only marginally better than everyone else, gets nearly all of the pie.

We can apply scalability to the stans. Mediocristan includes weight, height, food intake, baker's and prostitutes incomes, mortality rates, auto accidents, casino gambling, and so on. Extremistan includes total wealth, books sales per author, CD sales, celebrity name recognition, references in Google, deaths in war or terrorist incidents, and financial markets, among others. Medioscristan is where we have the mundane, the obvious, and the predictable. Extremistan is where we have the accidental, the singular, the unseen, and the unpredicted.

Black Swan events are, by definition, surprises. In understanding them, the main issue is how to go from what we logically know from the data we gather – a series of specific occurrences – to general conclusions. Does the data from a series of past events allow us to learn about other properties of a system? Bertrand Russell offers an interesting view, often called Hume's problem by philosophy professors. Consider a domesticated turkey (Russell actually used a chicken, but we are all Americans here). Every single day, every single feeding, firm's up the bird's belief that friendly members of the human race are looking out for its well-being. But, suddenly, on the Wednesday before Thanksgiving, the turkey incurs a rather permanent revision of its belief structure – a Black Swan event. That's the Black Swan problem in a nutshell – is it possible to predict the future given knowledge of the past.

This issue is at the very heart of empirical knowledge. For the turkey, the feeling of safety reaches its maximum just when the risk is highest. Something that has worked in the past turns out to be, at best, irrelevant, and, at worst, fatal.

“But in all my experience, I have never been in any accident, of any sort worth speaking about. I have seen but one vessel in distress in all my years at sea. I never saw a wreck, I

have never been wrecked nor was I ever in any predicament that threatened to end in disaster of any sort.”

E. J. Smith, Captain, RMS Titanic.

It is important to realize that Black Swans are relative to one’s knowledge. To the turkey, the event before Thanksgiving is a Black Swan. It is not to the butcher, since it isn’t unexpected. In a gambler’s or swindler’s parlance, Black Swan events are sucker problems – they occur relative to your expectations. You can eliminate Black Swans with knowledge. Perhaps by science, but science creates Black Swans by giving people unwarranted confidence about being in control.

Note that Black Swans do not have to be instantaneous. Two of the events I mentioned at the beginning were not – the publication of Darwin’s book, for example. Its impact has been growing slowly for years. As has the fight for women’s rights beginning with the ERA and the subsequent migration of competent women out of school classrooms. Sometimes a new technology creeps up on us slowly – cell phones are a good example.

The central difficulty lies in learning from the past. It is extremely convenient for us to believe we live in Mediocristan – it allows us to rule out Black Swan events. And, we do tend to believe we live there. Black Swans either do not exist or are inconsequential in this land.

There are a number of consequences that arise from our blindness to Black Swans:

- We focus on segments of the seen and generalize to the unseen – the Error of Confirmation.
- We fool ourselves with stories that cater to our thirst for patterns – the Narrative Fallacy.
- We act as if Black Swans don’t exist – our brains are not programmed for Black Swans.
- What we see is not all there is – history hides Black Swans from us and gives us mistaken ideas about the odds of their occurrence.
- We tunnel – we tend to focus on a few well-defined sources of uncertainty.

Confirmation is ingrained in us and it is dangerous and misleading. Subtle differences in statements create vast differences in meaning. Examining the first 1,000 days of our turkey's life leads you to conclude there is "no evidence for a future Black Swan event." This is not the same as saying "there is evidence for no Black Swan events." The difference is vast but the two statements are so close in your mind that you confuse them. And, it can be demonstrated that you are more likely to remember the latter statement a few days later; humans make this mistake all the time. "Almost all terrorists are Muslims" is confused with "almost all Muslims are terrorists." Even if 99% of all terrorists are Muslims, they would represent less than one ten thousandth of a percent of all Muslims. Minorities in the US suffer from this as well – consider DWB or "driving while black." John Stuart Mill once complained that "I never meant to say that all Conservatives are generally stupid, but that stupid people are generally Conservative." This is a chronic problem; our inferential processes are not made for complicated wording, where the meaning changes radically with a few simple word exchanges. In a primitive environment "most killers are wild animals" and "most wild animals are killers" differ little as far as the consequences are concerned. But our intuition has not evolved for situations where such differences make a huge difference.

Additionally, we are unable to rapidly transfer knowledge from one arena to another. Everyone in this room probably took an exam – perhaps the SAT or ACT exams – which had something like this on it: "All zoogles are boogles. You see a boogle. Is it a zoogle." The answer of course is "not necessarily" and you answered it correctly. But you don't transfer that knowledge between arenas. How many of you get on an elevator with a black person and immediately feel a momentary chill of fear? It is the same question, but you don't rapidly transfer your knowledge from the theoretical to the practical. It is a trait of human nature. Knowledge, even when exact, does not lead to appropriate action because we tend to forget what we know or, more importantly, how to process it properly when we aren't paying attention. We are domain specific. How many people do you know who will spend five minutes driving around the parking lot looking for a spot close to the gym, and then head straight for the treadmill?

In the medical literature there is the acronym NED - “no evidence of disease.” There is no acronym END - “evidence of no disease.” But physicians routinely twist this around. Your physician telling you that there is “no evidence of cancer” is not the same as him saying there is “evidence of no cancer.” I have interacted with only one physician in Cincinnati who gets this right – Bruce Bracken.

What is involved here is falsifying – being able to state with certainty that something is wrong. It is not easy to do. One of the problems is that when we make an hypothesis and try to verify it, we tend to ask questions for which a “yes” answer will support our hypothesis – we do not try to falsify it. Or put another way, once your mind is inhabited with a certain world-view, you tend to consider only instances that prove you right. Arabs and Israelis, looking and listening to the *same* news program will see very different stories in the same sequence of events. The same is true of Democrats and Republicans. Proponents of Intelligent Design are incapable of falsifying.

We inherited from our ancestors instincts and other thought processes that were adequate for survival on the mid-African veldt. But these instincts are not well-suited for survival in a much more complex environment. We live in Extremistan and our world is dominated by extremely rare events – a Black Swan occurs after thousands upon thousands of white ones. And as our society becomes more complex, and more connected, the impact of Black Swans increases tremendously and the sources of Black Swans multiply. In a primitive environment all we worried about was a new type of wild animal, weather changes, and new enemies. These events were sufficiently repeatable that we built up innate fears of them. Our instinct then allows us to make quick inferences to avoid those fearful things. But as a result, we are focused on a small number of sources of uncertainty. That helped us years ago and is ingrained in us. And that is our predicament.

We like stories. We like to summarize. We like to simplify, to reduce the dimensionality of a problem or issue. This leads us to the Narrative Fallacy. We are very vulnerable to our predilection for compact stories over raw truths. It distorts our mental representation of the world and is acutely incorrect when rare events are considered. The Narrative Fallacy refers to our

limited ability to examine sequences of events or facts *without weaving an explanation for them*, or forging a logical link that includes a relationship arrow. Explanations bind facts together, making them easier to remember and make sense. But when it makes us believe that we understand something, problems arise.

Not theorizing is difficult – it takes considerable mental effort to view facts and remember them without judging and creating an explanation. Making explanations is built-in – experiments have demonstrated that it is part of our biology, the default option as it were, so fighting it takes real effort.

But there is a deeper reason why we are inclined to narrate – the effect on information storage and retrieval. First, Information is costly to obtain. Second, it is costly to store. The more ordered, the less random, the more patterned a series of words or symbols is the easier it is to store. Try to describe the world around you with its details and you will attempt to weave a narrative thread into what you are saying. Third and finally, information is costly to manipulate and retrieve.

Humans have a hunger for rules that reduce the dimension of information so it can be squeezed into our heads – although the capacity of the brain appears to exceed  $10^{11}$  storage sites, our working memory is apparently much smaller. The more random information is, the greater the dimensionality, the harder it is to store. But if you find – or create – a pattern, you simply memorize the pattern. You would be hard-pressed to remember every word of in a dictionary of 100,000. You would have no problem memorizing a 100,000 word book composed of “Richard Gass spends far too much time running in circles,” repeated 10,000 times. Or a book of numbers generated by  $2^n+1$  where n varies from 0 to 99,999. You would just remember the summary or the formula. The more random information is, the harder it is to memorize and use it. The more you summarize, the more you order, the more you remove randomness. A pattern is clearly more compact than raw information. Unfortunately, the same conditions that drive us to simplify also pushes us to think the world is less random that is really is. And Black Swans are what we leave out of simplification.

Art and science are products of our need to simplify – to reduce dimensions and impose order. Stories and myths, paintings and photographs, spare us the complexity of the world and impart an order to it. The same desire for order pertains to science as well. Unlike art, the stated purpose of science is to get to the truth. Not to give you a feeling of organization or make you feel better – knowledge used as therapy.

We also impose causality onto our data – creating another reduction in the dimensionality of events. Narrativity has a chronologic dimension as well so both narrativity and causality make time flow in a single direction.

But memory and the arrow of time get mixed. We easily recall events in our past that fit a narrative and neglect those that do not appear to play a causal role in that narrative. We also know what happened subsequently to those events and it is impossible to ignore posterior information when solving a problem. Conventional wisdom holds that memory is like a serial recording device. It is not. It is dynamic, not static; new versions of the same thing are continually recorded, due to the power of posterior information. Memory is really a “self-serving dynamic revision machine”; old memories are constantly overwritten by new ones. You remember the last time you remembered an event and, without realizing it, change the story at every subsequent remembrance. Memories are revised along causal lines; we continually re-narrate past events in the light of what appears to make logical sense, long after the events have occurred. Memory is not fixed, constant and connected. What makes sense according to information obtained subsequently will be remembered more vividly. We invent some of our memories – a sore point in law.

How does narrativity affect our understanding of Black Swans? It messes up our projections of odds. Consider the following two statements given to so-called forecasting professionals. They were asked to imagine these two scenarios and estimate the odds of occurrence in a year:

A) A massive flood somewhere in America in which more than 1,000 people die.

B) An earthquake in California, causes massive flooding in which more than 1,000 people die.

Respondents said the first event was *less* likely to occur than the second. The earthquake in California is a readily imaginable cause for flooding, even though it lessens the universe of possible flood generators and US locations.

Which of these is more likely:

“Joey seemed happily married. He killed his wife.

“Joey seemed happily married. He killed his wife to get her inheritance.”

People will choose the second because a cause is given – even though the first is broader and can accommodate more causes – he was drunk, he went mad, she had an affair with the postman, he mistook her for his broker, etcetera. This tendency to pick the narrower statement over the broad, leads to pathologies in our decision-making. For example, it has been demonstrated that we will pay more for terrorist insurance than plain insurance – even though plain insurance includes terrorist insurance. The Black Swans we imagine, discuss and worry about, do not resemble those likely to occur – narration makes us worry about the wrong improbable events.

Some Black Swans are overblown in our minds Why? There are two types of Black Swans. Narrated Black Swans; these are present in the public eye and are discussed on TV by talking heads and their like on radio. And those no one talks about – you would be ashamed to discuss them in public because of their implausibility. Imagine if, in 1990, a Senator introduced a bill that required airlines to spend considerable sums to install bullet-proof cockpit doors to thwart terrorists? It’s human nature that the incidence of occurrence of narrated Black Swans is overestimated, while the second is severely underestimated. Lottery ticket buyers way overestimate their chances of winning because they visualize the huge payoff. They treat odds of 1 in a thousand the same 1 in ten million. Research shows that we *overreact* to low-probability outcomes when the possibility of an event is discussed with us and we become aware of it. We tend to estimate the role of rare events from Mediocristan correctly, but fail when it comes to variables in Extremistan. We are not skilled at gauging the effect of the improbable.

How this works in our minds is easy to see. Researchers have been able to map our thought processes into two modes, dubbed System 1 and System 2. System 1 is experiential. It is effortless, fast, automatic and parallel-processed. It is also opaque to us in that we do not know we are using it and it is prone to errors. It is emotional because it is quick, it produces short-cuts that allow us to function rapidly and effectively, and to survive. System 2 is cognitive. It is what we normally call thinking. It is what we do the classroom, in this room; it takes effort, It is reasoned, slow, logical and you are aware you are doing it. It makes fewer mistakes than System 1 because you correct the process as it goes along.

Our mistakes come from using System 1, when we believe we are using System 2. The main problem of System 1 is our lack of awareness of using it, since we react without thinking and introspection. Our confusing “no evidence of black swans” with “evidence of no black swans” shows System 1 in action. You have to make an effort to use System 2, to override your initial reactions. Nature makes you use System 1 to immediately get you out of harm’s way – to get you out of path of an oncoming tiger before you try to determine if it is an illusion of a tiger. System 1 protects us from danger far better than System 2. Indeed studies show that System 1 often reacts to the presence of danger before we are even aware of the danger.

Our misunderstanding of Black Swans can be attributed to our using System 1, imposing on ourselves the wrong event map. On a day-to-day basis we are not introspective enough to realize that our understanding of events is less than is warranted from a dispassionate observation of our experiences. We tend to forget Black Swans soon after they occur – they are too abstract for us. Instead we concentrate on the precise events that occurred. In Mediocristan narratives work – the past is likely to predict the future. In Extremistan, you do not have repetition and you need to be suspicious of the past and avoid easy and obvious narratives.

The Confirmation Bias and the Narrative Fallacy are internal mechanisms that get in the way of understanding improbable events. There are external mechanisms as well and one is particularly important: silent evidence.

Two thousand years ago Cicero presented the following story: Diagoras, a non-believer in the gods, was shown painted tablets bearing the portraits of some worshipers who prayed and survived a shipwreck. The implication was that praying protects you from drowning. Diagoras asked, “where are the pictures of those who prayed and then drowned?” The drowned worshipers would have some difficulty presenting their portraits. This is the sort of thing that makes people believe in miracles – the bias of silent evidence. Francis Bacon also mentions it, stating “And such is the way of all superstition, whether in astrology, dreams, omens, divine judgements, or the like.”

The problem? We don’t remember this. Silent evidence pervades everything connected to history, history in the sense of a succession of events seen through the lens of posterity. This bias pervades everything we do – the illusion of skill in many professions, factors in the success of ideas and religions, the nature vs. nurture debate, the logic of history, and our perception of the nature of Black Swans.

It is all too easy to avoid looking at the cemetery while concocting historical theories. But it is a problem with more than history – it is a problem with the way we gather evidence in every domain. There is a bias – a difference between what we see and what is there. Drowned worshipers do not write histories. Losers in history, people or ideas, do not write either. Fortunately, since it is a bias, it is possible to account for it by taking into account both the dead and the living.

This is a concept that is worthy of some examples. Consider the Phoenicians, who, we were told, produced no literature, a strange thing since we believe that they invented the alphabet. On the basis of this lack of literature, historians concluded that they were more interested in commerce than the arts. The alphabet served the lower purpose of commercial record keeping rather than the more noble one of literature. But now we’ve learned that the Phoenicians indeed wrote a lot – on perishable pieces of papyrus that biodegraded before they could be copied to longer lasting parchment. The neglect of the silent evidence led to the wrong conclusion.

The neglect of silent evidence is especially important in areas with winner-take-all attributes. We need to look at more than the success stories. Consider the numbers of people who call themselves writers, but are temporarily asking “do you want fries with that.” But we don’t often see MD’s flipping burgers. So, while you can investigate surgeons by using the sample you see performing operations, you cannot investigate writers by looking at who is publishing. And this holds in any profession with superstar effects.

Once you start thinking of silent evidence you see it everywhere. Taleb presents an interesting observation about a second form of the problem of silent evidence. I will repeat it here, albeit with a Cincinnati bias. There was an article in the NY Times discussing the mounting threat of the Russian Mafia and how it was displacing the Italian Mafia. The reporter stated that the Russians were more brutal as a result of being “hardened by the Gulag.” Hardened by the Gulag? Does that make any sense? Let’s consider another gedanken experiment – while we are downtown let’s gather a random selection of rats. Sick ones, healthy ones, large ones, small ones, fat ones, thin ones. You’ll likely find them in Over-the-Rhine and in the kitchens of fine restaurants. Gather thousands of them so we form a very heterogeneous cohort. Now bring them to Dr. Gass’s laboratory and place them in a hermetically sealed chamber. Subject the rats to increasing doses of something dangerous – radiation, poison gas, Cincinnati chili, whatever. At every dosage level, the strongest rats left survive and the weakest ones die. Progressively you create, on average, a stronger and stronger collection of rats. Except for one salient fact: every rat that survives, including the strongest, will be weaker after treatment than before. Now turn the surviving rats loose in Fountain Square, call the Enquirer, and inform the reporter on the rodent beat of this significant change in the rat population. He will write a long and learned article on the social dynamics of rats and state somewhere that “those rats are now the bullies of the rat population. They run the show. Strengthened by their experience in the laboratory of Dr. Gass, they dominate.”

This bias hides best when its impact is greatest. The dead rats are invisible and the more lethal the treatment, the more dead rats, and the more fooled you are about a strengthening effect. The

more injurious the treatment, the larger the difference between the survivors and rest, and the more fooled we are about a strengthening effect.

Does crime pay? Newspapers report on criminals that get caught. There is no section in the Enquirer detailing the stories of the people who have committed crimes but have not been caught. This is also true for bribery, prostitution rings, drug trafficking, poisoning spouses, and joining the Literary Club. So, not only do we not have the full picture, but the one we have is biased. It is dominated by the properties of dumb criminals, the ones who get caught. The silent evidence might show that crime does pay.

What about beginners luck? There is a belief among gamblers that beginners are almost always lucky. Should each of us become a gambler for a short while and cash in? Well, no; the same illusion has occurred. A person who starts gambling will be either lucky or unlucky. The unlucky ones stop gambling and pursue other interests – fishing, scrabble, bird watching, writing papers, or some such. They drop out of the sample. The lucky ones continue to gamble and remember how lucky they were as beginners. The drop-outs are no longer part of the gambling community and don't speak about their experiences.

How does this silent evidence affect our decision making? Consider Hurricane Katrina. In its aftermath, after TV images of death and destruction, politicians promised aid to rebuild. But they didn't promise to use their own money – they're using the government's; that is, ours. But that money has to be taken away from someone else. Perhaps it comes from diabetes or cancer research. Dead patients don't vote and don't appear on TV while dying – they don't connect to our emotional systems. More of them die every day than died in Katrina, and they are the ones who need our help. The point is that we see obvious consequences but not the invisible, less obvious ones. A French humanist<sup>4</sup> stated it well: we see what governments do and therefore sing their praises – but we do not see the alternatives lost. The confirmation problem occurs – governments are great at telling you what they did but not what they did not do.

---

<sup>4</sup> Bastiat

Taleb applies this to 9-11. About 2500 people were killed by bin Laden's group. Their families benefitted, as fitting, by assistance from numerous charities. But, during the last three months of 2001, an additional 1,000 people died as silent victims of the terrorists – on the highways because they were afraid to fly. There is evidence of this and it occurs because the road is much more lethal than the sky. These families, victims of 9-11, got no charitable support.

People are killed every day by doctors who ignore silent evidence. A drug saves lives, with the risk of killing one or two people. There is a net benefit to society, but a doctor has no incentive to use the drug. The lawyers of the people killed will go the doctor like pit bulls, while the lives saved will not be counted. A life saved becomes a statistic and is invisible. A person hurt is an anecdote; anecdotes are obvious.

Evolutionary fitness is something that is often aggrandized. The more someone is unfamiliar with Black Swans, the more he or she believes in the optimal working of evolution. But they ignore the silent evidence. Evolution is a series of flukes or mutations, some good but most bad. We only see the good. But in the short term it is not obvious which traits are good for you. The concept that we are here, that evolution did a great job, is bogus in the light of silent evidence. And in a Black Swan environment, a rare event (an asteroid strike perhaps) can shake up a species after a long run of fitness.

Finally, since I am a physicist and I cannot help but mention the anthropological-cosmological argument. But I will leave that to the other physicists in the room to explain.

A successful person – in any field – will always try to convince you that his or her achievements could not be accidental or due to luck. A gambler who picks the right number at roulette seven times will explain that the odds against such a streak are so long that something special must have happened, some divine intervention has occurred (Intelligent Design proponents fall into this trap all the time). Or perhaps he believes that he has special skills at picking numbers or some such ESP. But if you take into account the number of gamblers and the number of roulette plays, such a streak is bound to happen. This leads what may be called the reference point

argument – do not compute odds from the vantage point of a winner, but from that of all those who started in the cohort. From the reference point of the winner, a long string of wins is too extraordinary to be explained by luck, skill is needed. I would note that Joe Dimaggio's and Pete Rose's hitting streaks fall into this category as well.

This greatly weakens the notion of "because" that we scientists often employ and that historians invariably misuse. We need to learn to accept that there may be no "because," that "because" is "fuzzy." This makes us a bit nervous as we humans like the notion of causality – we are explanation-seeking animals and like to believe that everything has an identifiable cause. Yet there may not be a visible "because"; frequently there is nothing, not even a spectrum of possible explanations. Silent evidence masks this.

The other external problem of importance is education – and the way we do it. We are not taught how to think correctly about these kinds of probability. What we are taught is almost always based on games – especially craps, but often poker, roulette – and the ubiquitous Gaussian bell curve. But, if we assume that random chance has anything to do with mathematics at all, in the real world it is likely not Gaussian - but Mandelbrotian.

Consider the casino - a nice, self-contained environment where Gaussian mathematics and gaming probabilities should be all that one needs to make a profit. A casino's risk management is generally geared toward two things – setting up rules and policies for betting in the games, and stopping cheaters. Now, casinos are very diversified – they have many tables for each game, and many games. There are limits on the amounts one can bet on anyone play, and the rules of the casino are designed to always favor the house. So they really don't have to worry about the occasional lucky hit or series of lucky hits by a gambler. Absent cheating, the aggregate income of a casino is very stable. They do have to worry a bit about "whales," those gamblers who bet millions, are flown in by the casino, and get special treatment and rules. But they monitor them closely.

To keep their income stream stable, casinos will spend millions on stopping cheaters and they use some very sophisticated electronic surveillance and computer systems to do it. For those of you who saw the movie "Ocean's 13," the systems imagined there are not far from the truth. But in spite of the sophisticated anti-cheating systems, in spite of the control of the whales, in spite of the diversity of the games, the four largest losses one Las Vegas casino took, or narrowly avoided, fell completely outside their models – they were Black Swans:

- First, it lost 100 million dollars when an irreplaceable performer in their main show was mauled by a tiger.
- Second, a disgruntled contractor, offended by the settlement he was offered for work he had done, attempted to place dynamite around the support pillars in the casino's basement. He failed.
- Third, for some unknown reason, a casino employee failed to send several years worth of gamblers' profit forms to the IRS. They incurred a massive fine.
- Fourth, the casino owner's daughter was kidnaped and he violated Federal Law by dipping into the casino's coffers for the ransom. The casino incurred a huge fine.

The dollar value placed on these four Black Swans swamps the risks given by the standard model by about 1000 to 1. The casino spent hundred's of millions on gambling theory and high tech surveillance, while nearly all of its risks came from outside their models. And we still teach our students about uncertainty and probability from gambling examples.

The stories we make up, our mental bias's, our unwillingness to consider external evidence, and our poor teaching – all these result in our inability to comprehend and act on highly improbable events.

Yogi Berra has a saying:

“It's tough to make predictions, especially about the future.”

Berra knows something about randomness since he was a practitioner of uncertainty, regularly facing random outcomes. Berra is also said to have said

“The future ain’t what it used to be.”

He’s correct. The gains in our ability to predict and model the world around us are being dwarfed by increases in complexity of the world, and the increased probability of Black Swans. But as a further example of why knowledge can be less than useful, I leave you with this: Berra never said the first statement, the physicist Niels Bohr did. And many people have said the second, including people predating Berra.