CHAPTER ONE

1 The Swans Revisited

THE BLACK SWAN OF the popular imagination is one that swoops down from a clear blue sky, creating massive disorder in a very short amount of time. We expect it to be sudden and dramatic. The archetypical Black Swan is perhaps the 9/11 attack on the twin towers in New York in 2001. Virtually nobody could have been able to imagine such a thing. It was simply not on the mental map that something like that should even exist. Yet it happened, and in a single stroke, the world was a different place. The path we were on changed. The attack led to a whole new security apparatus, the war on terror, and the war in Iraq, to mention but a few of its consequences.

Actually, the 'out of the blue' aspect is not part of the original framework. Some Swans cited by Taleb take years if not centuries to play out. According to Taleb, Black Swans have just three attributes, none of which refers to suddenness. First, they are highly improbable. Second, they are highly consequential. Third, they make perfect sense after the fact.[1] When people talk about Black Swans it is usually the first two aspects they focus on, as if the term were essentially shorthand for low probability high impact risks. Simplifying in this way is

[1]The Black Swan criterion where the 9/11 attack might falter is the ex-post explanation part. Black Swans are supposed to make total sense after the fact, once our brain gets to work connecting the dots, which ends up giving it an air of inevitability. At least I struggle to connect some of those dots. Of course, we now know there are Islamic warriors engaged in a cosmic war that ends either in their destruction or in that of their enemies. However, that they would choose that means of meting out their punishment, and that it could be pulled off, remains unfathomable to this day.

wholly consistent with the reason that the Black Swan problem exists in the first place, reflecting as it does our tendency to reduce the number of dimensions of the phenomenon before us down to something more tractable and convenient.

Equating Black Swans with 'mere' low probability high impact risk, however, is to do the concept significant injustice. In reality, the Black Swan framework is valuable because it represents an altogether different way of approaching the world. Taleb asks us to reconsider some of our core assumptions about the very nature of the randomness we face as decision-makers and the inferences we make based on what we can observe. Furthermore, he brings our attention to the crucial role of expectations and attitudes in dealing with uncertainty. The problem, Taleb explains, is one of not being humble enough with respect to the limitations of our knowledge. If we believe the world consists of a certain kind of randomness and that we can have mastery over it, we may be in for some pretty bad surprises if those beliefs do not conform with reality. We can try to impose crisp and stylized ideas that appeal to our aesthetic sensibilities as much as we want, but the chaotic world we live in refuses to bend. This insistence on abstract beauty is what Taleb has in mind when he labels something as 'Platonic', after the famed Greek philosopher who saw loveliness in order and maintained that it could be superimposed on the messy reality we can observe with our senses (Taleb, 2007, p.19).

THE NATURE OF RANDOMNESS

Randomness refers to unpredictability. It applies whenever the outcome for some variable, such as the number of visitors to the Louvre on a given weekday, cannot be known with certainty beforehand. It is a function of our inability to know and predict the future. Try as we might, we never seem able to build those perfect forecasting algorithms that get it right all the time. In fact, as Taleb is at pains to point out, our overall track record in forecasting is awful (more on this later).

Why is there a general failure to predict what the future will bring? To answer this question, first consider that one very basic source of randomness is the physical world itself, which is constantly changing through processes that we do not fully comprehend. Science marches on, chipping away at the ignorance that produces apparent randomness. But despite the many laws of nature that have been uncovered, we never know where the next lightning will strike or how ocean currents will respond to changes in melting ice sheets. In the end, there are too many variables and too many complicated feedback

loops in these highly dynamic systems. On top of that there is human civilization itself. While once rudimentary and mostly local, over time society has become complex beyond imagination. Technical innovations have made possible advanced systems that increasingly connect people across different parts of the globe. It is *fundamentally unknowable* what outcomes these vast and interconnected systems of interacting people and technologies will produce. Human agency by itself ensures why the future keeps bringing so many surprises, as the 9/11 attack illustrates. It should be clear that we are up against a complexity that is beyond our ability to predict successfully.

The difficulties we face in predicting the future is related to the problem of induction, a classic problem in philosophy. While data can certainly teach us a great deal about the workings of the world, the philosopher and sceptic David Hume made us realize that we cannot arrive at *secure* knowledge on the basis of empirical observations. The problem of induction says that no matter how many observations you obtain, you cannot know for sure that the observed pattern is going to hold in the future. This inherent limitation is at the heart of the Black Swan concept. Any knowledge obtained through observation, Taleb says, is fragile. It is what the Black Swan metaphor itself is meant to convey. Recall that millions of observations on white swans had seemingly verified the notion that all swans are white, and it only took one observation of a black one to falsify it. Along the same lines, Peter Bernstein (1996) observed in his epic story about risk that: '. . . history repeats itself, *but only for the most part*'[2] (emphasis added). This sentence really sums it all up and explains why induction is treacherous ground for making assumptions about the future.

Once we capitulate to the fact that we cannot predict the future, the next best thing would be to be able to characterize randomness itself, i.e. describe it. In that way, we would have some idea about the scope for deviations from what we expect. A description of randomness would involve some degree of quantification of things like the range within which the values of a variable can be assumed to fall and how the outcomes are distributed within that range (frequencies). We might occasionally find such descriptions of random processes to be practically relevant insofar as they help us make informed decisions and our future wellbeing depends on the outcome of the variable in question. They are potentially helpful, for example, in coming up with a reasonable analysis of the trade-off between risk and return in different kinds of investment situations.

[2]Bernstein, P. L., 1996. *Against the Gods: The remarkable story of risk*. John Wiley & Sons: New York.

When characterizing randomness, a useful first distinction is between uncertainty and known odds.[3] Uncertainty simply means that the odds are not known, indeed cannot be known. When randomness is of this sort, there is no way of knowing with certainty the range of outcomes and their respective probabilities. Known odds, in contrast, means that we have fixed the range of outcomes and the associated probabilities. The go-to example is the roll of a dice, in which the six possible outcomes have equal probabilities. Drawing balls with different colours out of an urn is another favourite textbook example of controlled randomness.

Uncertainty, it turns out, is what the world has to offer. In fact, known odds hardly exist outside man-made games. This is the case for exactly the same reasons that forecasting is generally unsuccessful: there are some hard limits to our theoretical knowledge of the world.[4] There is ample data, for sure, which partly makes up for it. But the world generates only one observable outcome at a time, out of an infinite number of possibilities, through mechanisms and interactions that are beyond our grasp. There is nothing to say that we should be able to objectively pinpoint the odds of real-world phenomena. Whenever a bookie, for example, offers you odds on the outcome of the next presidential election, it is a highly subjective estimate (tweaked in favour of the bookie).

Whenever data exists, it is of course possible to try to use it to come up with descriptions of the randomness in a stochastic process. Chances are that we can 'fit' the data to one of the many options available in our library of theoretical probability distributions. Once we have, we have seemingly succeeded in our quest to describe randomness, or to turn it into something *resembling* known odds. This is the frequentist approach to statistical inference, in which observed frequencies in the data provide the basis for probability approximations. Failure rates for a certain kind of manufacturing process, for example, can serve as a reasonably reliable indication of the probability of failure in the future.

It is important to see, however, that even when we are able to work with large quantities of data, we are still in the realm of uncertainty. The data frequencies

[3]Frank Knight (1921) first made this distinction and referred to known odds as 'risk'. This epithet is unfortunate and will not be adhered to in order to avoid unnecessary confusion. In the present book, risk is construed of as the value of a random variable on which our well-being depends (such as corporate performance) falling below some aspirational or critical level (such as the level needed for debt servicing). Risk is thus a function of uncertainty, but has nothing to do with whether odds are known a priori or not. Knight, F. H., 1921. *Risk, uncertainty, and profit*. Hart, Schaffner & Marx: New York, NY.

[4]I leave out a consideration of particles at the subatomic level, which, according to important theories in physics, are governed by pure randomness. This randomness, say the same theories, can be described in precise, mathematical terms (i.e. the odds can be known).

typically only approximate one of the theoretical distributions. What is more, the way we collect, structure, and analyse these data points determines how we end up characterizing the random process and therefore the probabilities we assign to different outcomes. To the untrained eye, they might seem like objective and neutral probabilities because they are data-driven and obtained by 'scientists'. However, there is always some degree of subjectivity involved in the parameterization. The model used to describe the process could end up looking different depending on who designs it. Hand a large dataset over to ten scientists and ask them what the probability of a certain outcome is, and you may well get ten different answers. Because of the problem of induction, as discussed, there is always the possibility that the dataset, i.e. history, is a completely misleading guide to the future. Whenever we approximate probabilities using data, we *assume* that the data points we use are representative for describing the future.

THE MOVING TAIL

At this point, we are ready to conclude that the basic nature of randomness is uncertainty. Known odds, probabilities in the purest sense of the word, are an interesting man-made exception to that rule. If we accept that uncertainty is what we are dealing with, a natural follow-up question is: What is uncertainty *like*? A distinction we will make in this regard is between 'benign' and 'wild' uncertainty.[5] Benign uncertainty means that we do not have perfect knowledge of the underlying process that generates the outcomes we observe, but the observations nonetheless behave *as if* they conform to some statistical process that we are able to recognize. Classic examples of this are the distribution of things like height and IQ in a population, which the normal distribution seems to approximate quite well.

While the normal distribution is often highlighted in discussions about 'well-behaved' stochastic processes, many other theoretical distributions appear to describe real-world phenomena with some accuracy. There is nothing, therefore, in the concept of benign uncertainty that rules out deviations from the normal distribution, such as fat tails or skews. It merely means that the data largely fits the assumptions of *some* theoretical distribution and appears to do so consistently over time. It is as if we have a grip on randomness.

[5]In Taleb's terminology, wild uncertainty is found in a place called Extremistan, whereas benign uncertainty harbours in Mediocristan (Taleb, 2007, p. 35).

Wild uncertainty, in contrast, means that there is scope for a more dramatic type of sea change. Now we are dealing with outcomes that represent a clear break with the past and a violation of our expectations as to what was even supposed to be possible. Imagine long stretches of calm and repetition punctured by some extreme event. In these cases, what happened did not resemble the past in the least. Key words to look out for when identifying wild uncertainty is 'unprecedented', 'unheard of', and 'inconceivable', because they (while overused) signal that we might be dealing with a new situation, something that sends us off on a new path.

The crucial aspect of wild uncertainty is precisely that the tails of the distributions are in flux. In other words, the historically observed minimum and maximum outcomes can be surpassed at any given time. I will refer to the idea of an ever-changing tail of a distribution as *The Moving Tail*. With wild uncertainty, an observation may come along that is outside the established range – *by a lot*. Such an event means that the tail of the distribution just assumed a very different shape. Put another way, there was a qualitative shift in the tail. Everything we thought we knew about the variable in question turned out to be not even in the ballpark.

An illustration of wild uncertainty and of a tail in flux is provided by 'the Texas freeze', which refers to a series of severe blizzards that took place in February 2021, spanning a 10-day period. The blizzards and the accompanying low temperatures badly damaged physical structures, and among those afflicted were wellheads and generators related to the production and distribution of electricity. As the freeze set in, demand soared as people scrambled to get hold of whatever electricity they could to stay warm and keep their businesses going. In an attempt to bring more capacity to the market, the operator of the Texas power grid, Ercot, hiked the price of electricity to the legally mandated price ceiling of 9,000 $/MWh. The price had touched that ceiling on prior occasions – but only for a combined total of three hours. The extremeness of this event lay in the fact that Ercot kept it at this level for almost 90 consecutive hours.[6] A normal trading range leading up to this point had been somewhere between 20–40 $/MWh.

Any analysis of this market prior to February 2021 would have construed tail risk as being about short-lived spikes, which, when averaged out over several trading days, implied no serious market distress. The Texas freeze shifted the tail. It was a Black Swan. The consequences for market participants were

[6]www.griddy.com, accessed 10 July 2021.

massive,[7] and there was nothing in the historical experience that convincingly pointed to the possibility that the price could or would remain at its maximum for 90 hours. After the fact, it looked obvious that something like that could happen. Prolonged winter freezes in Texas are very rare, but with the climate getting more extreme by the day, why not?

The 'by a lot' is actually an important qualifier of wild uncertainty. To see why, consider that whenever we have a dataset, some of the observations will represent the tail of the distribution. They are large but rare deviations from some more normal state. Let us say that we have, in a given dataset, a handful of observations that can be said to constitute the tail. There will be, by construction, a minimum and a maximum value, which are the most extreme values that history has had to offer so far.

Unless we are talking about a truly truncated distribution, like income having zero as the lower limit, it is a potential mistake to think that the 'true' underlying data-generating process is somehow capped by the observed minimum and maximum values. If we feed all the observations we have into a statistical software, we can ask it to analyse which random process that most plausibly generated the patterns in the data. Now, if we immediately take the process identified by the programme and draw random values based on it in a simulation, it will come up with a distribution that contains outcomes that *go beyond* the lowest/highest observed values in the dataset without the probability of that dropping to virtually zero. This will always happen as long as the approach is to assume that there is some underlying random process generating the data and use real data to approximate it. It is as if the software doing the fitting 'gets it' that if we have observed certain extreme values, even more extreme observations cannot be ruled out. If we have observed a drop in the S&P 500 of minus 58% over a certain period of time, who would say that a drop of minus 60% is outside the realm of possibilities? The simulated extremes will lie somewhere to the left (right) of the minimum (maximum) observed in the data. The tail we model in this way will encompass the observed tail *and then some*.

The upshot of this discussion is that experiencing an outlier that is only *somewhat* more extreme than the hitherto observed minimum/maximum should fall within the realm of benign uncertainty. We should not be surprised or taken aback by it. There is an implied probability of that, meaningfully

[7]Some retail customers had opted out of the standard fixed-rate utility plans and instead bought their electricity from businesses that passed wholesale prices on to them directly. Some of these customers racked up electricity bills in the range of $8,000–10,000 in a matter of days (Winter storm fallout sends Texas power firm Griddy into Bankruptcy. *Financial Times*, 15 March 2021).

separate from zero, being handed to us by the fitted distribution. We have to add 'by a lot' for it to count as wild uncertainty, because then the tail has shifted dramatically and in a way that was by no means implied by the historical track record. It is an outlier so extreme that it has a probability of effectively zero, even when the underlying random process we use to form a view of the future has been fitted to all the tail events in the historical track record.

Under conditions of wild uncertainty, it is clear that the concept of probability starts looking increasingly subjective and unverifiable. Indeed, Taleb calls probability 'the mother of all abstract concepts' (Taleb, 2007, p. 133) and maintains that we cannot calculate probabilities of shocks (Taleb, 2012, p. 8).[8] It is important to see, though, that his scorn is reserved mostly for those who insist on using the symmetric normal distributions and its close relatives. The properties of the normal are seductive because we can derive, with relative ease, all sorts of interesting results, but it is, Taleb maintains, positively dangerous as a guide to decision-making in a world of wild uncertainty. Why? Primarily because of how it rules out extreme outliers and blinds us to them. A key feature of the normal distribution is that its tails quickly get thinner the further removed from the mean you move, which implies that their likelihood of happening gets lower and lower. In fact, as we move away from the mean, the assigned probabilities drop very fast – much too fast, in Taleb's view (Taleb, 2007, p. 234). The stock market crash in October 1987, for example, saw a return of minus 20.5%. The odds of a drop of at least that magnitude would have been roughly *one in a trillion* according to the normal. In other words, anyone going by that distribution would have considered it, for practical purposes, an impossible event.

The first priority, therefore, is to avoid the normal distribution like the plague. In its place, if we still feel compelled to work with probabilities, Taleb offers the idea of fractals. Fractals refer to a geometrical pattern in which the structure and shape of an object remains similar across different scales. The practical implication is that the likelihood of extreme events decreases at a much slower rate. If one subscribes to this view, the probability of finding an exceptionally large oil field is not materially lower than a large or medium-sized one because the geological processes that generate them are scale-independent. This relation between frequency and size is associated with the so-called power law distributions, which we will relate to socio-economic processes in Chapter 7. According to Taleb, the idea of fractals should be our

[8]Taleb, N. N., 2012. *Antifragile: Things that gain from disorder*. Random House: New York.

default, the baseline mental model for how probabilities change as we move further out on the tail (Taleb, 2007, p. 262).

In many cases, we lack data that we can explore for mapping out the tail of a random process. In this kind of setting, uncertainty tends to be wildly out of the gate. Technological innovation fits right into this picture, because it brings novelty and injects it into the existing, already volatile, world order. New dynamics are set in motion, triggering unintended consequences and side effects that ripple through the system in an unpredictable fashion. Because we keep innovating, we also keep changing the rules of the game, forever adding to the complexity. Two Black Swans that have sprung from the onward march of technology are the emergence of the internet and the more recent invasion of social media and mobile phones into our lives. There was no existing dataset that we could have studied prior to them that might have suggested that such transformations of our reality were about to happen. Or, more importantly, that they were even possibilities at all. To appreciate how technologies that we are completely immersed in today and take for granted are actually Black Swans, cases of wild uncertainty, consider the words of Professor Adam Alter of New York University:

> 'Just go back twenty years [to 2000] . . . imagine you could speak to people and say, hey, you are going to go to the restaurant and everyone's going to be sitting isolated and looking at a small device, and then they're going to go back home and spend four hours looking at that device, and then you're going to wake up in the morning and look at that device . . . and people are going to be willing to have body parts broken to preserve the integrity of that device . . . people would say that is crazy'[9]

Alter's thought experiment of going back 20 years in time and imagining talking to people about something highly consequential that later happened is a useful one for deciding whether something is to be considered a Black Swan. If you imagine their reaction to what you describe would be that it is ridiculous or inconceivable, chances are that you have found one.

THE ROLE OF EXPECTATIONS

To continue our story, it becomes clear that any characterizations of random processes will be increasingly subjective as we move away from data-driven

[9]Episode #1564 of the Joe Rogan Experience. Alter is referring to the claim made by some, mostly the young, that they would rather have a finger broken than their phone smashed.

approaches. We leave the world of inference from data and enter the realm of the imagination. Our faculties for reasoning and logic can partly make up for a lack of data – we can figure certain stuff out. When the imagination fails us, we have those truest of Black Swans, the inconceivable ones, the 'unknown unknowns'. We have already mentioned the 9/11 attack as being in this category. In a similar way, the collapse of the Soviet Union was utterly unthinkable to the Western intelligentsia and political establishment at the time. George Kennan, an American diplomat and historian, commented as follows, based on a review of the history of international affairs in the modern era:

> '[It is] hard to think of any event more strange and startling, and at first glance inexplicable, than the sudden and total disintegration and disappearance . . . of the great power known successively as the Russian Empire and then the Soviet Union.'[10]

That is, nobody expected the Soviet Union to crumble at this point in time. One of the most crucial aspects of Black Swans is that they are always measured against expectations and prior knowledge. This is an underappreciated point. As noted, most people use the term loosely, largely equating it with high-impact outcomes that were somehow shocking to us. With the considerable difference, perhaps, that calling it a Black Swan provides an air of complete unpredictability and that, therefore, one is not to be blamed for what just happened. Getting tail risk wrong may be an indictable offense. But Black Swans? They seem to absolve everyone of any responsibility for what went down, because nobody could have seen it coming.

The habit mentioned earlier of equating Black Swans with 'mere' tail risk misses out on what is perhaps its most important dimension, namely the expectations we had going into the situation. Because of the role of expectations, what is a Swan to you may not be one to me. It is, in Taleb's preferred terminology, a 'sucker's problem'. Naïve and ignorant individuals are more prone to experience Black Swans simply because they fail to form realistic expectations. To illustrate this idea, Taleb uses the example of a turkey somewhere in the US as Thanksgiving approaches. Having walked about generously fed for its entire life, the turkey is unsuspecting of the calamity that is about to befall it. The butcher, however, is obviously not unsuspecting, and he is therefore not in for a Black Swan – exactly the same event, but wildly diverging expectations.

---

[10]For this quote, and other useful perspectives, read the article *"Everything You Think You Know About the Collapse of the Soviet Union Is Wrong"* on foreignpolicy.com. Everything you think you know about the collapse of the Soviet Union is wrong, *Foreign Policy*.

The relativity of Black Swans has wide implications. Whenever a high impact event occurs, this may or may not be shocking. The more interesting discussion to be had is about who was attuned to this possibility and who was caught out? A Black Swan always requires a vantage point. To the suckers, it appears as if the tail just moved, but not necessarily to someone who sees the world a bit differently. Whenever we hear the term Black Swan mentioned, therefore, what should immediately spring to mind is the follow-up question 'Well, a Black Swan to whom?'

The Covid-19 pandemic is a case in point. Was this a Black Swan? It certainly meets the criterion of being a highly consequential event. Interestingly, Taleb himself has gone on record saying that C-19 was *not* a Black Swan. His argument is that there is a history of pandemics, based on which popular films were made well before C-19. A basic analysis of the connectivity of the modern world (i.e. means of travel) would also have pointed to the obvious plausibility of a global pandemic. Respected institutions issued reports warning of global pandemics already in the early 2000s. Bill Gates gave a thoughtful talk on the subject in 2015, also issuing words of warning for those willing to listen.[11] These considerations may well have sensitized students of history who also had the wisdom to internalize the possibility that this could happen in their own lifetime.

However, casual observation suggests that most of us have not reached such a state of immaculate wisdom. Many do not read books and have never heard of connectivity. All of our egocentric biases make going into denial about pandemics the easiest thing in the world, something which Albert Camus, the French philosopher, understood well:

> 'Everybody knows that pestilences have a way of recurring in the world; yet somehow we find it hard to believe in ones that crash down on our heads from a blue sky.'[12]

For most of us, watching a film about something makes it seem even more unreal. It puts it in the same category as Bruce Willis drilling holes in meteorites about to smash into Earth – pure entertainment.[13] If so, consuming films

[11]The next outbreak? We're not ready, Bill Gates – YouTube.
[12]Camus, A., 1947. *The Plague*. Librairie Gallimard.
[13]I should be careful here. The Earth is under threat from rocks in space. While earthlings' primary tactic, if detected in time, is likely to be to try to nudge such an object on to another course using force. Perhaps, however, the response will also involve some attempt at sending people up there to deal with the problem on site. Let us leave all options on the table.

may turn us into even bigger suckers because they warp our expectations. Since we are now discounting zombies heavily as just movie entertainment, woe on us the day they show up on the doorstep, because we have not prepared one bit for that eventuality!

At any rate, even if we had read about pandemics and realized that the probability of another one is clearly not zero, there is something about the magnitude of the consequences. Just like the 'business-as-usual' risk, Black Swans have two dimensions: possibility and consequence. Granting the possibility of something is a binary situation: a recognition that such a thing could happen (as opposed to saying there is no way it ever could). Even if we are willing to entertain the possibility of something, we could still be suckers with respect to the consequences once the event is unleashed. This is why C-19, for most people, was a genuine Black Swan. Pandemics, sure, I think I heard about that in school. But who could have imagined entire countries shutting down? Spending months on end in lockdowns? Tourism coming to a near stand-still? Mad dashes – knife fights even – for toilet rolls? It would take a pretty serious student of history, and one with a very fertile imagination at that, to envision the severity of the consequences along so many paths. Therefore, for a substantial majority of the planet's inhabitants, C-19 was a Black Swan. Along the same lines, it would be questionable not to label the French Revolution with its decidedly wild consequences a Black Swan just because there had been revolutions prior to that.

When it comes to the possibility of something happening, history by now offers a pretty impressive palette of different types of events. We might think of these as 'known unknowns', as it is clearly within our reach to form an understanding of them. What even the most creative and superbly educated mind cannot envision, however, is the magnitude of the consequences of those events as they would play out in the present day. The dynamics are truly impossible to imagine, and hence the consequences. History only repeats itself for the most part. With respect to the consequences, most of us are suckers, especially when it comes to our own lives and times. Whenever something impactful but entirely conceivable hits our vicinity, we are stunned no matter what.

WHAT MAKES US SUCKERS?

The previous section referred briefly to the 'not in our lifetime' perspective. Let us linger on this point for a bit, as it is one of the keys to understanding Black Swans and why we are essentially born suckers. Most of us will freely admit that humanity is in for one disaster or other. Sooner or later,

that asteroid will knock us out of our pants, for sure, but it is always later, somewhere out in the distant future. It is not going to happen in my lifetime. Why? Because I am somehow special. Stuff only happens to other people, whereas I am destined to lead a glorious and comfortable existence. Based on such egocentric beliefs, we might coolly concede that in the larger scheme of things, something is sure to happen, yet almost completely discount the possibility as far as our lifetime and corner of the world is concerned. Not that we always say so publicly or even think in those terms outright, it is more of a tacit assumption.

This 'because I'm special' protective mechanism goes a long way in setting the stage for Black Swans. However, it is only one of the many ways in which our outlook is warped, which brings us to the long catalogue of biases that have been identified and described by scholars. A bias can be said to be a predisposition to make a mistake in a decision-making situation, because it leads us away from the decision that would be taken by a rational and well-informed person who diligently weighs pros and cons. What biases tend to have in common is that they make us more of a sucker than we need to be. They are a staple of business books nowadays (those on risk management in particular) and may bore the educated reader. Since they are so fundamental to the concept of Black Swans, we must briefly review them nonetheless. What follows is a non-exhaustive list of certain well-documented biases that in various ways contribute to the Black Swan phenomenon. As is commonly pointed out, these biases have been mostly to our advantage over the long evolutionary haul, but are often liabilities in the unnatural and complex environment we find ourselves in today.

- **The narrative fallacy**

  In explaining why we are so poorly equipped to deal with randomness, Taleb focuses on what he refers to as 'the narrative fallacy', which he defines as 'our need to fit a story or pattern to a series of connected or disconnected facts' (Taleb, 2007, p. 309). We invent causes for things that we observe in order to satisfy our need for coherent explanations. It turns out that we do not suffer dissonance gladly, so our brain will supply any number of rationalizations to connect the dots. By reducing the number of dimensions of the problem at hand and creating a neat narrative, things become more orderly. Everything starts to hang together and make sense, and that is how the dissonance is resolved. Since we are lazy as well, we often converge on the rationalization that satisfies our craving with the least amount of resistance. However, when we force causal

interpretations on our reality, and invent stories that satisfy our need for explanations, we make ourselves blind to powerful mechanisms that lie outside these simple narratives.

- **Confirmation bias**

  This is one of the leading causes of Swan-blindness discussed in The Black Swan, where Taleb refers to confirmation as 'a dangerous error' (Taleb, 2007, p. 51). It has to do with the general tendency to adopt a theory or idea and then start to look for evidence that corroborates it. When we suffer from this bias, all the incoming data seems, as if by magic, to confirm that the belief we hold is correct; that the theory we are so fond of is indeed true. Whatever instances contradict the theory are brushed aside or ignored, or re-interpreted (tweaked) in a way that supports our pre-existing beliefs. Out the window goes Karl Popper's idea of falsification, the true marker of science and open inquiry. Using falsification as a criterion, a theory is *discarded* if evidence contradicting it becomes undeniable. In the specific context of managing risks, the confirmation bias is a problem because we will be too prone to interpret incoming observations of stability to suggest that the future will be similarly benign.

- **The optimistic bias**

  Research has shown that humans tend to view the world as more benign than it really is. Consequently, in a decision-making situation, people tend to produce plans and forecasts that are unrealistically close to a best-case scenario.[14] The evidence shows that this is a bias with major consequences for risk taking. In the words of Professor Daniel Kahneman (2011): 'The evidence suggests that an optimistic bias plays a role – sometimes the dominant role – whenever individuals or institutions voluntarily take on significant risks. More often than not, risk takers underestimate the odds they face, and do not invest sufficient effort to find out what they are.'[15] Pondering on extreme and possibly calamitous outcomes will clearly not be a priority for an individual with an optimistic bent. Taking a consistently rosy view distorts expectations and therefore invites the Black Swan.

[14]This bias co-exists with a negativity bias that has us putting undue focus on negative aspects and overreacting to them. For example, if you receive 100 great reviews for something you did, but only one negative, you are quite likely to spend the rest of the week ruminating over that negative review.

[15]Kahneman, D., 2011. *Thinking fast and slow*. Farrar, Straus and Giroux: New York.

- **The myopia bias**

  Myopia, in the literature on the psychology of judgement, refers to the tendency to focus more on short-term consequences than long-term implications. Because of our desire for instant gratification, we tend to place much less weight on future gains and losses relative to those in the near-term. Professors Meyer and Kunreuther call this the most 'crippling' of all biases, resulting in gross underpreparedness for disasters that could have been mitigated with relatively simple measures.[16] This was the case, for example, with the tsunami in the Indian Ocean in 2004. Only a few years prior, in Thailand, relatively inexpensive mitigation measures had been discussed – and dismissed. The reason? There were many reasons, but among other things, there was a worry that it might cause unnecessary alarm among tourists. Such miniscule short-term benefits got the upper hand in preparing for events with colossal consequences.

- **The overconfidence bias**

  Humans are prone to overrate their own abilities and the level of control they have over a situation. The typical way of exemplifying this tendency is to point to the fact that nearly everyone considers himself an above-average driver. Taleb prefers the more humorous example of how most French people rate themselves well above the rest in terms of the art of love-making (Taleb, 2007, p. 153). As for the effect of overconfidence on decision-making, it is profound – and not in a favourable way. Professor Scott Plous (1993) argues that a large number of catastrophic events, such as the Chernobyl nuclear accident and the Space Shuttle Challenger explosion, can be traced to overconfidence. He offers the following summary: 'No problem [. . .] in decision-making is more prevalent and more potentially catastrophic than overconfidence.'[17] Overconfidence has been used to explain a wide range of observed phenomena, such as entrepreneurial market entry and trading in financial markets, despite available data suggesting high failure rates.

Considering the above, one is inclined to agree with Taleb when he remarks that '. . . it is as if we have the wrong user's manual' (Taleb, 2007, prologue xxii) for navigating successfully in a world of wild uncertainty. We crave simple but coherent narratives. We value elegant theories and become committed to them.

---

[16]Meyer, R. and H. Kunreuther, 2017. *The Ostricht Paradox: Why we underprepare for disaster*. Wharton School Press.

[17]Plous, S., 1993. *The psychology of judgment and decision making*. McGraw-Hill: New York.

We think we are special and that the world around us is benign. We are equipped with a mind that was created for an existence with much fewer variables and more direct cause-and-effect mechanisms. Reflecting deeply about interconnected systems was not key to survival in our evolutionary past. In a somewhat shocking passage, Taleb says that 'our minds do not seem made to think and introspect' because, historically speaking, it has been 'a great waste of energy' (ibid.).

In fact, information, which potentially helps us rise above sucker-status, is costly to acquire and process. Imagine that I bring up the possibility of nuclear terror affecting a major US city. Such a scenario involves hundreds of thousands of dead and an upheaval of life as we knew it, before even considering what the countermeasures might be. Any firm with operations in the US is likely to be greatly affected by this calamity. Now what is your gut reaction to this proposed topic of conversation? In all likelihood, your kneejerk reaction is to immediately try to shut it down. The sheer unpleasantness of the topic makes us not want to go there, even for a brief moment of time. It is too much to take in, and frankly too boring, so, to save us the mental energy, we are perfectly willing to resort to the handy tactic of denial.

As problems, extreme and abstract possibilities, remote from everyday practicalities, are not inspiring enough to energize us. They are out of sight and therefore out of mind. We are unable to maintain a focus on them for long enough. Our thoughts will gravitate towards something more tangible, some action that yields a more gratifying sense of accomplishment here and now. It often takes a herculean effort to process remote possibilities and we are rarely in the mood for it. They are therefore not necessarily 'unknown unknowns', rather they can be thought of as 'unknown knowables'. Unknown knowables is meant to convey that it is within our reach to form an understanding of the possibility and most of its consequence, but we fail to do so because of our laziness or disinterest. That makes it, for practical purposes, a Black Swan, at par with the unknown unknowns. At least to some, that is, because others might be prepared to take up the challenge.

THE RELATIVITY OF BLACK SWANS

Earlier in this chapter, we noted that the popular view of Black Swans is that they strike quickly and unexpectedly. Except that there is nothing in the Black Swan framework that says it has to be sudden or even happen within a reasonably short time-period, like a few months. In fact, many of the examples discussed in Taleb's book are episodes that may seem like distinct and well-delineated events in a history

book, but were prolonged affairs with a long lead-up. World Wars I and II are both in this category. The rise of Christianity is mentioned as another Black Swan event. A dominant Christianity would no doubt have appeared like an absurd proposition to someone living around the time of the birth of Jesus. Its consequences were certainly immense, so it meets this criterion too. It also took *centuries* to gain a foothold and start making its impact felt. The rise of the internet and social media were mentioned earlier as examples of technology-driven Black Swans. They too emerged gradually over many years, infiltrating our lives one small step after the other. Therefore, from the viewpoint of a decision-maker in the real world (which is the perspective that Taleb urges us to take) they were not instantaneous.

The fact that monumental changes can take a long time in gestation adds to the relativity of Black Swans. Those that are less wedded to specific ideas and more open to rewriting the story they tell themselves come around quicker to change. This introduces a strategic dimension to Black Swans, massive agents of change as they are. The observation to make is that when others refuse or are unable to see a changing reality, the value of being a non-conformist increases. The sucker status of those that you interact with competitively is a variable of interest, a theme we will come back to many times in this book.

Apart from the biases that shape our thinking, the relativity of Black Swans is also a matter of information and knowledge in the more traditional sense. The more we invest in high-quality information and capabilities for processing it, the wider our frame of reference will be, and the fewer Black Swans we will experience. Take Donald Trump's ascendancy to president as an example, which was a Black Swan to those who kept discounting him heavily. As his candidacy was first announced, many took it as a joke, and went on to seriously underestimate him throughout the race. The general assumption was that America would somehow come to its senses and see him for what he was, which, in their view, included being wholly unfit for office. However, those who had had their ears on the tracks and picked up on the dark undercurrents of American society knew better. In their eyes, the arrival of Trump was the culmination of a process that had been long in coming. They had realized that a substantial number of people had come to resent the system and thought it was rigged against them, feeling that nobody stood up for them except Trump. Out of spite, they were prepared to vote for someone who promised to shake the system up. This undercurrent unleashed the mindboggling Trump presidency, which culminated in the Capitol Hill insurrection. What a Black Swan, to all those who had not noticed the sentiment that had developed among broad layers of the US population.

For students of Swanology, the #1222 episode of the Joe Rogan Experience podcast is a gem. It illuminates several of the mechanisms involved in turning

some (most) of us into suckers, and highlights the enormous differences in expectations that can arise *even when we are looking at the same set of facts*. The topic of conversation is about the mother of all Black Swans, the wiping away of the existing world order by way of an asteroid impact. According to Rogan's guests, Graham Hancock and Randall Carlson, there is compelling evidence to suggest that there have been repeated cosmic impacts during the 180,000 years or so that anatomically modern humans have existed. Besides transforming the geography of the Earth, these events may have erased civilizations existing at the time so thoroughly that no archaeological evidence of them can be found today. Hancock and Carlson, labelled 'catastrophists' by some, argue that human civilization did not begin to stir around 12,000 years ago, which is the conventional view, but that it was *rebooted* following a near-complete destruction caused by an asteroid that smashed into Earth a number of centuries before. The asteroid impact they refer to was followed by extreme shifts in the global climate. One consequence was a near-instantaneous melting of the ice sheets and the flooding that resulted completely overwhelmed humans living at the time. (Fascinatingly, if true, this could explain the enduring myths handed down to us from deep history regarding an epic flooding.)

Mainstream science, however, adheres to a view referred to as 'gradualism', which holds that we can extrapolate backward in time from processes we are able to observe today. In this view, everything we see in the landscapes today are the result of gradual processes that have been going on for eons. As a result of the vested interest in this narrative by academics who have built a career on it, they have put up quite a resistance to the ideas put forth by Hancock and Carlson. This resistance takes us into Black Swan territory with respect to the civilization-destroying potential of cosmic impact. Below are some excerpts from their conversation. We do not have to take any sides with respect to the facts being discussed to enjoy the insights it brings about the Black Swan formation process.

GH: Whenever you propose a cataclysm and present evidence for it . . . you can be sure that you will be descended upon by a crowd of furious critics.

JR: As a species, we have amnesia.

JR: Why would they try to ignore something like that?

GH: When new information emerges that contradicts established theories . . . when you get very committed to a model or idea . . . you start to connect your personality to it, and any attack on it becomes an existential attack on you.

GH: Again and again, what we see is new facts being dismissed because they don't fit into the existing theory . . . this is a problem in the whole history of science. I've come to view archeology and history as more ideology, really, than science.

GH: There is an ideological view of how civilization developed, that there is this long, slow, gradual, politically correct rise . . . and here we are, the apex and pinnacle of this story, and gosh, we are so proud of ourselves and our achievements.

GH: They are in a state of denial and just don't want to recognize it.

JR: It is so sad. You count on these people to distribute the information but their ego gets involved in things . . . if you have an absolutely established narrative that you teach and you are unwilling to look at any possible deviation from that, you are saying, almost from an authority position, 'we know what happened and we know where we are going'.

RC: We're kind of in this mode now where there's a very large and growing political agenda around the idea that humans are the sole cause of global change . . . now we come along and say, no, there's actually been forces unleashed on this planet that utterly dwarfs anything we've done yet. What does that do to that paradigm?

GH: . . . south of Minnesota you [had] a heavily vegetated area covered with primal forest and that is what goes on fire and the reason it goes on fire is because when these impacts come in they generate huge amounts of heat . . . and it [sets] the world on fire.

JR: Oh that gives me goosebumps . . . a single afternoon all over the world and everything changes forever and it's [ruined] for a thousand years.

GH: We all need to know about this . . . this is our background, this is where we come from, the present order of the world has descended from that moment.

The conversation points to several biases at work that lull us into a state of ignorant bliss highly conducive to Black Swans. There is scientific dogma, which is a kind of confirmation bias in which incoming evidence is fitted to the received models of explanation. Any dissenters presenting alternative interpretations of the evidence are (according to Rogan's guests) ruthlessly clamped down on and stripped of their sources of funding. There is also the issue of political agendas, where the ruling ideology will only accept facts

that support the preferred narrative (currently, man as solely responsible for global warming). Finally, there is the sheer existential discomfort that results from having to ponder the fact that asteroid impacts are a regularly occurring phenomenon in deep history. Many of us just do not like being told that the Earth is essentially a sitting duck in space just waiting for the next swarm of asteroids to come our way. Others, in contrast, remain open to consider what the facts are trying to tell us no matter what conclusions they lead us to. A wide gap in expectations has opened up, and this matters, not least in the kind of competitive interaction we will take an interest in later.

MEET THE PREPPERS

Few would be more open to alternative interpretations of the facts than the so-called 'preppers'. These are the folks who take the issue of cataclysms very seriously, and who prepare for them with passion. In fact, you could accuse them of having an overactive imagination, of buying into conspiracy theories left and right. There is indeed a fringe element in this movement, with a clear anti-government stance and a penchant for organizing into paramilitary units, the so-called 'survivalists'. But Professor Bradley Garrett, who has carried out in-depth investigations of the phenomenon, argues that prepping is actually going mainstream, and has developed into a prospering multibillion dollar industry.[18] Interestingly, and perhaps worryingly, the elites have taken a keen interest in preparing for doomsday scenarios. There is now an active market in bunkers offering luxury and comfort, enabling you to sit out the end times in some style. Apparently, well-to-do tech-entrepreneurs are very active in the market because they foresee a breakdown of the social contract as technology continues to remove the need for millions of workers.

If preppers can be accused of having overactive imaginations, perhaps that makes up for the deficiency that the rest of us consistently display in this area. In fact, preppers could be rated the finest Black Swan spotters out there. There are levels to this, and they have explored remote and disastrous possibilities to a degree that the average person is not even close to. Given their orientation, they are not constrained by the need to conform or any stigma that may come from expressing unorthodox views. Instead, they revel in it. Consequently, they have taken the process of turning 'unknown unknowns' into 'known unknowns' as far as you reasonably could. If they have turned every

[18]Garrett, B., 2020. *Bunker: Building for the End Times*. Scribner: New York.

stone in search of what could trigger an upheaval, why not see what they have come up with? In Table 1.1 I have listed, for future reference, all the extreme events brought up by the preppers interviewed in Garrett's book.

A recurring theme in the prepper world-view is that many of the events in Table 1.1 will play out in an interlinked sequence. In the parlance, there will be 'ripple effects' where various parts of the system fall apart like domino bricks. I have organized the list loosely from meta-events that trigger a shock to the system, down to some of the societal consequences that could follow. There is no particular ordering implied here. Things further down the list could of course happen for reasons unrelated to those higher up, and the top entry

TABLE 1.1 The prepper's list

Nuclear war
Nuclear terror
Pandemics
Synthetic biotech (engineered pandemic by rogue nation/scientist)
Asteroid impact
Electro-magnetic pulse from sun (will 'fry' all electronics)
Man-made electro-magnetic impulse (ditto)
Mega volcano eruptions
Run-away technology (a bit vaguely)
Artificial intelligence turning hostile
Hurricanes and floodings and wildfires
Sea-levels rising (by 2040 Florida will be sea-floor according to one prepper)
Desertification (by 2040 Europe will be Saharan, same source)
Blackout of electric grids
Collapsing eco-systems
Crop failures
Cessation of global trade networks
Hyperinflation
Collapse of paper currency
Financial collapse
Government collapse
World without rule of law

could take us directly to the last, and so on (and government collapse could easily be the precipitating factor for nuclear havoc!). The permutations are endless. The point to be made is simply that misfortune rarely comes alone. When it hits the fan, as preppers like to say, things could quickly be going wrong on multiple fronts. It has long been known that risks are not independent of each other, and that we have to take these tendencies to co-vary this into account when designing risk management strategies. As per the preppers, the same appears to be true in the tail of things.