



THE BLACK SWAN
The Impact of the Highly Improbable

NASSIM NICHOLAS TALEB



RANDOM HOUSE TRADE PAPERBACKS
NEW YORK

PROLOGUE

ON THE PLUMAGE OF BIRDS

Before the discovery of Australia, people in the Old World were convinced that *all* swans were white, an unassailable belief as it seemed completely confirmed by empirical evidence. The sighting of the first black swan might have been an interesting surprise for a few ornithologists (and others extremely concerned with the coloring of birds), but that is not where the significance of the story lies. It illustrates a severe limitation to our learning from observations or experience and the fragility of our knowledge. One single observation can invalidate a general statement derived from millennia of confirmatory sightings of millions of white swans. All you need is one single (and, I am told, quite ugly) black bird.*

I push one step beyond this philosophical-logical question into an empirical reality, and one that has obsessed me since childhood.† What we

* The spread of camera cell phones has afforded me a large collection of pictures of black swans sent by traveling readers. Last Christmas I also got a case of Black Swan Wine (not my favorite), a videotape (I don't watch videos), and two books. I prefer the pictures.

† I used the logical metaphor of the black swan (not capitalized) for Black Swan Events (capitalized), but this problem should not be confused with the logical problem raised by many philosophers. This is not so much about exceptions as it is about the oversize role of extreme events in many domains in life. Furthermore, the logical problem is about the possibility of the exception (black swan); mine is about the *role* of the exceptional event (Black Swan) leading to the degradation of predictability and the need to be robust to negative Black Swans and exposed to positive ones.

call here a Black Swan (and capitalize it) is an event with the following three attributes.

① First, it is an *outlier*, as it lies outside the realm of regular expectations, because nothing in the past can convincingly point to its possibility. Second, it carries an extreme impact (unlike the bird). Third, in spite of its outlier status, human nature makes us concoct explanations for its occurrence *after* the fact, making it explainable and predictable.

I stop and summarize the triplet: rarity, extreme impact, and retrospective (though not prospective) predictability.* A small number of Black Swans explain almost everything in our world, from the success of ideas and religions, to the dynamics of historical events, to elements of our own personal lives. Ever since we left the Pleistocene, some ten millennia ago, the effect of these Black Swans has been increasing. It started accelerating during the industrial revolution, as the world started getting more complicated, while ordinary events, the ones we study and discuss and try to predict from reading the newspapers, have become increasingly inconsequential.

Just imagine how little your understanding of the world on the eve of the events of 1914 would have helped you guess what was to happen next. (Don't cheat by using the explanations drilled into your cranium by your dull high school teacher.) How about the rise of Hitler and the subsequent war? How about the precipitous demise of the Soviet bloc? How about the consequences of the rise of Islamic fundamentalism? How about the effect of the spread of the Internet? How about the market crash of 1987 (and the more unexpected recovery)? Fads, epidemics, fashion, ideas, the emergence of art genres and schools. All follow these Black Swan dynamics. Literally, just about everything of significance around you might qualify.

This combination of low predictability and large impact makes the Black Swan a great puzzle; but that is not yet the core concern of this book. Add to this phenomenon the fact that we tend to act as if it does not exist! I don't mean just you, your cousin Joey, and me, but almost all "social scientists" who, for over a century, have operated under the false belief that their tools could measure uncertainty. For the applications of the sciences of uncertainty to real-world problems has had ridiculous effects; I have been privileged to see it in finance and economics. Go ask your

* The highly expected *not happening* is also a Black Swan. Note that, by symmetry, the occurrence of a highly improbable event is the equivalent of the nonoccurrence of a highly probable one.

portfolio manager for his definition of "risk," and odds are that he will supply you with a *measure* that *excludes* the possibility of the Black Swan—hence one that has no better predictive value for assessing the total risks than astrology (we will see how they dress up the intellectual fraud with mathematics). This problem is endemic in social matters.

The central idea of this book concerns our blindness with respect to randomness, particularly the large deviations: Why do we, scientists or nonscientists, hotshots or regular Joes, tend to see the pennies instead of the dollars? Why do we keep focusing on the minutiae, not the possible significant large events, in spite of the obvious evidence of their huge influence? And, if you follow my argument, why does reading the newspaper actually decrease your knowledge of the world?

It is easy to see that life is the cumulative effect of a handful of significant shocks. It is not so hard to identify the role of Black Swans, from your armchair (or bar stool). Go through the following exercise. Look into your own existence. Count the significant events, the technological changes, and the inventions that have taken place in our environment since you were born and compare them to what was expected before their advent. How many of them came on a schedule? Look into your own personal life, to your choice of profession, say, or meeting your mate, your exile from your country of origin, the betrayals you faced, your sudden enrichment or impoverishment. How often did these things occur according to plan?

What You Do Not Know

Black Swan logic makes what you don't know far more relevant than what you do know.* Consider that many Black Swans can be caused and exacerbated by *their being unexpected.*

Think of the terrorist attack of September 11, 2001: had the risk been reasonably *conceivable* on September 10, it would not have happened. If such a possibility were deemed worthy of attention, fighter planes would have circled the sky above the twin towers, airplanes would have had

* The Black Swan is the result of collective and individual epistemic limitations (or distortions), mostly confidence in knowledge; it is not an objective phenomenon. The most severe mistake made in the interpretation of my Black Swan is to try to define an "objective Black Swan" that would be invariant in the eyes of all observers. The events of September 11, 2001, were a Black Swan for the victims, but certainly not to the perpetrators. The Postscript provides an additional discussion of the point.

locked bulletproof doors, and the attack would not have taken place, period. Something else might have taken place. What? I don't know.

Isn't it strange to see an event happening precisely because it was not supposed to happen? What kind of defense do we have against that? Whatever you come to know (that New York is an easy terrorist target, for instance) may become inconsequential if your enemy knows that you know it. It may be odd that, in such a strategic game, what you know can be truly inconsequential.*

This extends to all businesses. Think about the "secret recipe" to making a killing in the restaurant business. If it were known and obvious, then someone next door would have already come up with the idea and it would have become generic. The next killing in the restaurant industry needs to be an idea that is not easily conceived of by the current population of restaurateurs. It has to be at some distance from expectations. The more unexpected the success of such a venture, the smaller the number of competitors, and the more successful the entrepreneur who implements the idea. The same applies to the shoe and the book businesses—or any kind of entrepreneurship. The same applies to scientific theories—nobody has interest in listening to trivialities. The payoff of a human venture is, in general, inversely proportional to what it is expected to be.

Consider the Indian Ocean tsunami of December 2004. Had it been expected, it would not have caused the damage it did—the areas affected would have been less populated, an early warning system would have been put in place. What you know cannot really hurt you.

Experts and "Empty Suits"

The inability to predict outliers implies the inability to predict the course of history, given the share of these events in the dynamics of events.

But we act as though we are able to predict historical events, or, even worse, as if we are able to change the course of history. We produce thirty-year projections of social security deficits and oil prices without realizing that we cannot even predict these for next summer—our cumulative prediction errors for political and economic events are so monstrous that every time I look at the empirical record I have to pinch myself to verify

* The Idea of Robustness: Why do we formulate theories leading to projections and forecasts without focusing on the robustness of these theories and the consequences of the errors? It is much easier to deal with the Black Swan problem if we focus on robustness to errors rather than improving predictions.

that I am not dreaming. What is surprising is not the magnitude of our forecast errors, but our absence of awareness of it. This is all the more worrisome when we engage in deadly conflicts: wars are fundamentally unpredictable (and we do not know it). Owing to this misunderstanding of the causal chains between policy and actions, we can easily trigger Black Swans thanks to aggressive ignorance—like a child playing with a chemistry kit.

Our inability to predict in environments subjected to the Black Swan, coupled with a general lack of the awareness of this state of affairs, means that certain professionals, while believing they are experts, are in fact not. Based on their empirical record, they do not know more about their subject matter than the general population, but they are much better at narrating—or, worse, at smoking you with complicated mathematical models. They are also more likely to wear a tie.

Black Swans being unpredictable, we need to adjust to their existence (rather than naïvely try to predict them). There are so many things we can do if we focus on antiknowledge, or what we do not know. Among many other benefits, you can set yourself up to collect serendipitous Black Swans (of the positive kind) by maximizing your exposure to them. Indeed, in some domains—such as scientific discovery and venture capital investments—there is a disproportionate payoff from the unknown, since you typically have little to lose and plenty to gain from a rare event. We will see that, contrary to social-science wisdom, almost no discovery, no technologies of note, came from design and planning—they were just Black Swans. The strategy for the discoverers and entrepreneurs is to rely less on top-down planning and focus on maximum tinkering and recognizing opportunities when they present themselves. So I disagree with the followers of Marx and those of Adam Smith: the reason free markets work is because they allow people to be lucky, thanks to aggressive trial and error, not by giving rewards or "incentives" for skill. The strategy is, then, to tinker as much as possible and try to collect as many Black Swan opportunities as you can.

Learning to Learn

Another related human impediment comes from excessive focus on what we do know: we tend to learn the precise, not the general.

What did people learn from the 9/11 episode? Did they learn that some events, owing to their dynamics, stand largely outside the realm of the pre-

dictable? No. Did they learn the built-in defect of conventional wisdom? No. What did they figure out? They learned precise rules for avoiding Islamic prototerrorists and tall buildings. Many keep reminding me that it is important for us to be practical and take tangible steps rather than to “theorize” about knowledge. The story of the Maginot Line shows how we are conditioned to be specific. The French, after the Great War, built a wall along the previous German invasion route to prevent reinvasion—Hitler just (almost) effortlessly went around it. The French had been excellent students of history; they just learned with too much precision. They were too practical and exceedingly focused for their own safety.

We do not spontaneously learn that we don't learn that we don't learn. The problem lies in the structure of our minds: we don't learn rules, just facts, and only facts. Metarules (such as the rule that we have a tendency to not learn rules) we don't seem to be good at getting. We scorn the abstract; we scorn it with passion.

Why? It is necessary here, as it is my agenda in the rest of this book, both to stand conventional wisdom on its head and to show how inapplicable it is to our modern, complex, and increasingly *recursive* environment.*

But there is a deeper question: What are our minds made for? It looks as if we have the wrong user's manual. Our minds do not seem made to think and introspect; if they were, things would be easier for us today, but then we would not be here today and I would not have been here to talk about it—my counterfactual, introspective, and hard-thinking ancestor would have been eaten by a lion while his nonthinking but faster-reacting cousin would have run for cover. Consider that thinking is time-consuming and generally a great waste of energy, that our predecessors spent more than a hundred million years as nonthinking mammals and that in the blip in our history during which we have used our brain we have used it on subjects too peripheral to matter. Evidence shows that we do much

* *Recursive* here means that the world in which we live has an increasing number of feedback loops, causing events to be the cause of more events (say, people buy a book *because* other people bought it), thus generating snowballs and arbitrary and unpredictable planet-wide winner-take-all effects. We live in an environment where information flows too rapidly, accelerating such epidemics. Likewise, events can happen *because* they are not supposed to happen. (Our intuitions are made for an environment with simpler causes and effects and slowly moving information.) This type of randomness did not prevail during the Pleistocene, as socioeconomic life was far simpler then.

less thinking than we believe we do—except, of course, when we think about it.

A NEW KIND OF INGRATITUDE

It is quite saddening to think of those people who have been mistreated by history. There were the *poètes maudits*, like Edgar Allan Poe or Arthur Rimbaud, scorned by society and later worshipped and force-fed to school-children. (There are even schools named after high school dropouts.) Alas, this recognition came a little too late for the poet to get a serotonin kick out of it, or to prop up his romantic life on earth. But there are even more mistreated heroes—the very sad category of those who we do not know were heroes, who saved our lives, who helped us avoid disasters. They left no traces and did not even know that they were making a contribution. We remember the martyrs who died for a cause that we knew about, never those no less effective in their contribution but whose cause we were never aware of—precisely because they were successful. Our ingratitude toward the *poètes maudits* fades completely in front of this other type of thanklessness. This is a far more vicious kind of ingratitude: the feeling of uselessness on the part of the silent hero. I will illustrate with the following thought experiment.

Assume that a legislator with courage, influence, intellect, vision, and perseverance manages to enact a law that goes into universal effect and employment on September 10, 2001; it imposes the continuously locked bulletproof doors in every cockpit (at high costs to the struggling airlines)—just in case terrorists decide to use planes to attack the World Trade Center in New York City. I know this is lunacy, but it is just a thought experiment (I am aware that there may be no such thing as a legislator with intellect, courage, vision, and perseverance; this is the point of the thought experiment). The legislation is not a popular measure among the airline personnel, as it complicates their lives. But it would certainly have prevented 9/11.

The person who imposed locks on cockpit doors gets no statues in public squares, not so much as a quick mention of his contribution in his obituary. “Joe Smith, who helped avoid the disaster of 9/11, died of complications of liver disease.” Seeing how superfluous his measure was, and how it squandered resources, the public, with great help from airline pilots, might well boot him out of office. *Vox clamantis in deserto*. He will

If we had stopped COVID w/ annoying measures

retire depressed, with a great sense of failure. He will die with the impression of having done nothing useful. I wish I could go attend his funeral, but, reader, I can't find him. And yet, recognition can be quite a pump. Believe me, even those who genuinely claim that they do not believe in recognition, and that they separate labor from the fruits of labor, actually get a serotonin kick from it. See how the silent hero is rewarded: even his own hormonal system will conspire to offer no reward.

Now consider again the events of 9/11. In their aftermath, who got the recognition? Those you saw in the media, on television performing heroic acts, and those whom you saw trying to give you the impression that they were performing heroic acts. The latter category includes someone like the New York Stock Exchange chairman Richard Grasso, who "saved the stock exchange" and received a huge bonus for his contribution (the equivalent of several *thousand* average salaries). All he had to do was be there to ring the opening bell on television—the television that, we will see, is the carrier of unfairness and a major cause of Black Swan blindness.

Who gets rewarded, the central banker who avoids a recession or the one who comes to "correct" his predecessors' faults and happens to be there during some economic recovery? Who is more valuable, the politician who avoids a war or the one who starts a new one (and is lucky enough to win)?

It is the same logic reversal we saw earlier with the value of what we don't know; everybody knows that you need more prevention than treatment, but few reward acts of prevention. We glorify those who left their names in history books at the expense of those contributors about whom our books are silent. We humans are not just a superficial race (this may be curable to some extent); we are a very unfair one.

LIFE IS VERY UNUSUAL

This is a book about uncertainty; to this author, the rare event *equals* uncertainty. This may seem like a strong statement—that we need to principally study the rare and extreme events in order to figure out common ones—but I will make myself clear as follows. There are two possible ways to approach phenomena. The first is to rule out the extraordinary and focus on the "normal." The examiner leaves aside "outliers" and studies ordinary cases. The second approach is to consider that in order to understand a phenomenon, one needs first to consider the extremes—

particularly if, like the Black Swan, they carry an extraordinary cumulative effect.

I don't particularly care about the usual. If you want to get an idea of a friend's temperament, ethics, and personal elegance, you need to look at him under the tests of severe circumstances, not under the regular rosy glow of daily life. Can you assess the danger a criminal poses by examining only what he does on an *ordinary* day? Can we understand health without considering wild diseases and epidemics? Indeed the normal is often irrelevant.

Almost everything in social life is produced by rare but consequential shocks and jumps; all the while almost everything studied about social life focuses on the "normal," particularly with "bell curve" methods of inference that tell you close to nothing. Why? Because the bell curve ignores large deviations, cannot handle them, yet makes us confident that we have tamed uncertainty. Its nickname in this book is GIF, Great Intellectual Fraud.

PLATO AND THE NERD

At the start of the Jewish revolt in the first century of our era, much of the Jews' anger was caused by the Romans' insistence on putting a statue of Caligula in their temple in Jerusalem in exchange for placing a statue of the Jewish god Yahweh in Roman temples. The Romans did not realize that what the Jews (and the subsequent Levantine monotheists) meant by *god* was abstract, all embracing, and had nothing to do with the anthropomorphic, too human representation that Romans had in mind when they said *deus*. Critically, the Jewish god did not lend himself to symbolic representation. Likewise, what many people commoditize and label as "unknown," "improbable," or "uncertain" is not the same thing to me; it is not a concrete and precise category of knowledge, a *nerdified* field, but its opposite; it is the lack (and limitations) of knowledge. It is the exact contrary of knowledge; one should learn to avoid using terms made for knowledge to describe its opposite.

What I call *Platonicity*, after the ideas (and personality) of the philosopher Plato, is our tendency to mistake the map for the territory, to focus on pure and well-defined "forms," whether objects, like triangles, or social notions, like utopias (societies built according to some blueprint of what "makes sense"), even nationalities. When these ideas and crisp constructs

inhabit our minds, we privilege them over other less elegant objects, those with messier and less tractable structures (an idea that I will elaborate progressively throughout this book).

Platonicity is what makes us think that we understand more than we actually do. But this does not happen everywhere. I am not saying that Platonic forms don't exist. Models and constructions, these intellectual maps of reality, are not always wrong; they are wrong only in some specific applications. The difficulty is that a) you do not know beforehand (only after the fact) *where* the map will be wrong, and b) the mistakes can lead to severe consequences. These models are like potentially helpful medicines that carry random but very severe side effects.

The *Platonic fold* is the explosive boundary where the Platonic mind-set enters in contact with messy reality, where the gap between what you know and what you think you know becomes dangerously wide. It is here that the Black Swan is produced.

TOO DULL TO WRITE ABOUT

It was said that the artistic filmmaker Luchino Visconti made sure that when actors pointed at a closed box meant to contain jewels, there were real jewels inside. It could be an effective way to make actors live their part. I think that Visconti's gesture may also come out of a plain sense of aesthetics and a desire for authenticity—somehow it may not feel right to fool the viewer.

This is an essay expressing a primary idea; it is neither the recycling nor repackaging of other people's thoughts. An essay is an impulsive meditation, not science reporting. I apologize if I skip a few obvious topics in this book out of the conviction that what is too dull for me to write about might be too dull for the reader to read. (Also, to avoid dullness may help to filter out the nonessential.)

Talk is cheap. Someone who took too many philosophy classes in college (or perhaps not enough) might object that the sighting of a Black Swan does not invalidate the theory that *all swans are white* since such a black bird is not technically a swan since whiteness to him may be the essential property of a swan. Indeed those who read too much Wittgenstein (and writings about comments about Wittgenstein) may be under the impression that language problems are important. They may certainly be important to attain prominence in philosophy departments, but they are

something we, practitioners and decision makers in the real world, *leave for the weekend*. As I explain in the chapter called "The Uncertainty of the Phony," for all of their intellectual appeal, these niceties have no serious implications Monday to Friday as opposed to more substantial (but neglected) matters. People in the classroom, not having faced many true situations of decision making under uncertainty, do not realize what is important and what is not—even those who are scholars of uncertainty (or *particularly* those who are scholars of uncertainty). What I call the practice of uncertainty can be piracy, commodity speculation, professional gambling, working in some branches of the Mafia, or just plain serial entrepreneurship. Thus I rail against "sterile skepticism," the kind we can do nothing about, and against the exceedingly theoretical language problems that have made much of modern philosophy largely irrelevant to what is derisively called the "general public." (In the past, for better or worse, those rare philosophers and thinkers who were not self-standing depended on a patron's support. Today academics in abstract disciplines depend on one another's opinion, without external checks, with the severe occasional pathological result of turning their pursuits into insular prowess-showing contests. Whatever the shortcomings of the old system, at least it enforced *some* standard of relevance.)

The philosopher Edna Ullmann-Margalit detected an inconsistency in this book and asked me to justify the use of the precise metaphor of a Black Swan to describe the unknown, the abstract, and imprecise uncertain—white ravens, pink elephants, or evaporating denizens of a remote planet orbiting Tau Ceti. Indeed, she caught me red handed. There is a contradiction; this book is a story, and I prefer to use stories and vignettes to illustrate our gullibility about stories and our preference for the dangerous compression of narratives.*

You need a story to displace a story. Metaphors and stories are far more potent (alas) than ideas; they are also easier to remember and more fun to read. If I have to go after what I call the narrative disciplines, my best tool is a narrative.

Ideas come and go, stories stay.

* The metaphor of the black swan is not at all a modern one—contrary to its usual attribution to Popper, Mill, Hume, and others. I selected it because it corresponds to the ancient idea of a "rare bird." The Latin poet Juvenal refers to a "bird as rare as the black swan"—*rara avis in terris nigroque simillima cygno*.

THE BOTTOM LINE

The beast in this book is not just the bell curve and the self-deceiving statistician, nor the Platonified scholar who needs theories to fool himself with. It is the drive to “focus” on what makes sense to us. Living on our planet, today, requires a lot more imagination than we are made to have. We lack imagination and repress it in others.

Note that I am not relying in this book on the beastly method of collecting selective “corroborating evidence.” For reasons I explain in Chapter 5, I call this overload of examples naïve empiricism—successions of anecdotes selected to fit a story do not constitute evidence. Anyone looking for confirmation will find enough of it to deceive himself—and no doubt his peers.* The Black Swan idea is based on the structure of randomness in empirical reality.

T To summarize: in this (personal) essay, I stick my neck out and make a claim, against many of our habits of thought, that our world is dominated by the extreme, the unknown, and the very improbable (improbable according to our current knowledge)—and all the while we spend our time engaged in small talk, focusing on the known, and the repeated. This implies the need to use the extreme event as a starting point and not treat it as an exception to be pushed under the rug. I also make the bolder (and more annoying) claim that in spite of our progress and the growth in knowledge, or perhaps *because* of such progress and growth, the future will be increasingly less predictable, while both human nature and social “science” seem to conspire to hide the idea from us.

Chapters Map

The sequence of this book follows a simple logic; it flows from what can be labeled purely literary (in subject and treatment) to what can be deemed entirely scientific (in subject, though not in treatment). Psychology will be mostly present in Part One and in the early part of Part Two; business and natural science will be dealt with mostly in the second half of Part Two and in Part Three. Part One, “Umberto Eco’s AntilibRARY,” is mostly

* It is also naïve empiricism to provide, in support of some argument, series of eloquent confirmatory quotes by dead authorities. By searching, you can always find someone who made a well-sounding statement that confirms your point of view—and, on every topic, it is possible to find another dead thinker who said the exact opposite. Almost all of my non-Yogi Berra quotes are from people I disagree with.

about how we perceive historical and current events and what distortions are present in such perception. Part Two, “We Just Can’t Predict,” is about our errors in dealing with the future and the unadvertised limitations of some “sciences”—and what to do about these limitations. Part Three, “Those Gray Swans of Extremistan,” goes deeper into the topic of extreme events, explains how the bell curve (that great intellectual fraud) is generated, and reviews the ideas in the natural and social sciences loosely lumped under the label “complexity.” Part Four, “The End,” will be very short.

I derived an unexpected amount of enjoyment writing this book—in fact, it just wrote itself—and I hope that the reader will experience the same. I confess that I got hooked on this withdrawal into pure ideas after the constraints of an active and transactional life. After this book is published, my aim is to spend time away from the clutter of public activities in order to think about my philosophical-scientific idea in total tranquillity.