

Essentials

COPYRIGHTED MATERIAL

■ Chapter 1 ■

Playing God

■ *It's tough to model human action* ■ *Finance is not as religious as physics* ■ *Black Swans make things harder* ■ *The markets are not Normal and the past is a faulty guide* ■ *Should we care that theorists persist?* ■

Economists (particularly those involved in financial research) are often accused of suffering from an acute case of “physics envy.” If only the economic landscape could be as mathematically tractable as the physical landscape. If only terrifyingly precise theoretical predictions held in economics as well as they do in physics. If only we could also be deemed scientists.

Economics, of course, is not physics. For one very simple, yet inevitably powerful reason: In one case the laws are immutably God-made and thus permanently exact (all one has to do is go find them and, with luck, express their structure down on paper); in the other, the rules are dictated not by God, but by His creatures, us humble humans. And if there is something that we know about ourselves is that, when it comes to economic activity (which of course includes the financial markets), we tend to be reliably unreliable. Our behavior is not set in stone, preprogrammed, preordained. It is not law-abiding, but rather entirely anarchic, ever changing. While the physical terrain is characterized by its divine lawfulness, the human-determined economic domain is shaped by pagan lawlessness.

Few have explained this dichotomy better than Emanuel Derman, a former top Goldman Sachs executive and now a professor at Columbia University, and someone who as a leading “quant” has spent a big chunk of his professional life trying to determine whether the markets are mathematically tamable. Derman, who has a PhD in physics and is a globally revered expert, once offered the following beautifully stated clarifications: *“It’s not that physics is better, but rather that finance is harder. In physics you are playing against God, and He doesn’t change His laws very often. In finance, you are playing against God’s creatures, agents who value assets based on their ephemeral opinions.”*

That is, while accurate modeling and forecasting may be possible (and, naturally, desirable) in the physical world, they are likely to be impossible (and possibly entirely undesirable) in the financial world. The eventual level of asset prices will depend on the actions of millions of individual investors, constantly buying and selling. Can anybody honestly claim to be able to register such behavior with a few equations? Who knows why and when people would revert to dumping an asset, or to accumulate it? Can any type of math capture those wild spirits?

Where will the yield curve be tomorrow? That will depend on bond prices, which in turn depend on the actions of people buying and selling bonds. Where will stock prices be next week? That will depend solely on human action, too. Where will the dollar be next month? Supply and demand. Can we really aspire to predict those actions? Seems far-fetched, and Derman agrees: *“No mathematical model can capture the intricacies of human psychology. Watching people put too much faith in the power of formalism and mathematics, I saw that if you listen to the models’ siren song for too long, you may end up on the rocks or in the whirlpool.”*

Physicists can search for truth because in the physical world truth really exists. Once one of nature’s explicitly mechanical laws is discovered by a clever scientist, it can be relied upon not to change. Ever. But there are no immutable laws when it comes to the values of financial assets. No permanent rules set at the time of genesis. No divine inevitability. In finance, there is no truth. A new reality is created every minute through the unpredictable actions of utility-seeking humans.

Let Derman deliver the final nail in the coffin: *“As a physicist, when you propose a model of Nature, you are pretending you can guess the structure created by God. Perhaps it is possible because God doesn’t pretend. But as a quant, when you propose a new model of value, you are pretending you can guess the structure created by other people. As you say that to yourself, if you are honest,*

your heart sinks. You are just a poor pretender and you know immediately there is no chance at all that you are truly right. When you take on other people, you are pretending you can comprehend other pretenders, a much more difficult task.”

As a financial modeler you are trying to guess what other people are going to do. But their eventual actions will depend on what they think *you* are going to do. So you have to correctly guess what other people are going to guess regarding your own future actions. Plausible? Like the notion of gravitation suddenly ceasing to work.

And yet financial economists and financial mathematicians have, for at least the past 50 years (and with particular ardor for the last three decades), devoted their considerable talents and energy to theoreticizing the markets into systems of equations, statistical symbols, and Greek letters. They have embarked on a quest to formalize and axiomatize finance that is Taliban-like in its dogmatism and resoluteness. God may not have a say in the markets (a direct one, at least), but that’s no reason for us to feel unattended and uncared-for. Finance theorists have shown an indefatigable resolution to step in His shoes and fill the vacuum, all too willingly enacting laws and principles that are held by those ivory towers worldwide (and throughout a non-insignificant number of nonacademic posts, including many trading floors, treasury departments, regulatory agencies, and newspapers) with the same submissiveness as that shown by churchgoing parishioners.

Call me an unrepentant atheist, but I find myself among those who seriously doubt the validity of the mathematically charged financial prophecies. I believe that humans are so unpredictable when it comes to their dealings in equities, currencies, bonds, or mortgages (by this I don’t mean random, I mean unforeseeable, undetectable; stating that the markets are random, as some well-established theories do, would imply us knowing how humans behave in the marketplace) that not even a real Prophet could untangle such conundrums. It is not that the theorists are not brilliant or that the tools are wrong per se. I just don’t think that financial markets can be quantitatively understood, synthesized, and predicted. Any more than one can quantitatively understand, synthesize, and predict, say, the future sexual activities of a group of diverse and unrelated strangers (in fact, this may be far easier than in the case of the markets, where new, potentially influential information constantly shows up; where the actions of some people affect the actions of the rest; and

where somewhat predictable physiological necessities and personal health levels do not shape the outcome.)

As much as academics may want, as Derman puts it, to subjugate the market with axioms and theorems, the market can (and will) do anything it likes. While atoms and planets have no choice but to follow their divinely preordained paths, economic agents enjoy much greater freedom and have a tendency to stubbornly and rebelliously refuse to bow down to the authority of the mathematical sheriff. Any financial economist who attempts to scribble the market's equivalent of the Ten Commandments from the isolated confines of a university office is presupposing that the outside world will obediently oblige (at all times), thus transforming theory into reality, and theory into law. Some claim that this is indeed not wholly implausible (those who defend the *performativity* of theory, whereby the existence of a model molds reality towards compliance with the former's tenets). But, in principle, it sounds quite presumptuous to count on the unquestioned compliance of financial players, if only for the obvious reason that a large group of them may have never been aware of the theory's very existence, let alone understand it and agree with it. Apples and particles, in contrast, didn't need to wait for Newton and Einstein to publish their conjectures before they could fall from trees and move randomly inside a gas.

Emanuel Derman once wrote that "*There is an almost religious quality to the pursuit of physics that stems from its transcendent qualities. . . . It's hard not to have a sense of wonder when you see that principles, imagination, and a little mathematics (in a word, the mind) can divine the behavior of the universe. Short of genuine enlightenment, nothing but art comes closer to God.*" It is only understandable that nonphysicists (like, well, financial economists) would want to reach a similar state of rapture. The same relevance and status. Just imagine being able to discover another piece of life's hidden genetic code.

However, finance theorists should humbly recognize that their field does not contain the promise of the potential discovery of immutable, transcendent, immortal truths. Finance is much less pure, much more contaminated, much more vulgar. God's creations are not only strikingly beautiful and chaste, but He also plays with a fair dice. In the universe, the rules don't change in the middle of the game, they are "stationary," thus dependable (i.e., predictable, mathematically tractable). In statistical parlance, the God-given probability distributions are not only knowable

but stable. In the markets, though, things are way messier. Humans are much more treacherous. Untrustworthy. Non-dependable. They don't play fair. They change the rules constantly, without pre-warning. As such, the probability distribution is not only wildly nonstationary (what held in the past does not necessarily hold today, or will hold tomorrow) but basically unknowable. Who can model such a world?

Legendary financial econometrician Andrew Lo, a professor at MIT's Sloan School of Business with hedge fund experience, famously said that in the physical sciences three laws can explain 99 percent of behavior, while in finance, 99 laws explain at best three percent of behavior. Lo is not shy about how he feels regarding the capacities of (traditional, at least) financial theory: "*Neoclassical economics works really well in some areas. But in the markets, neoclassical economists have failed miserably.*"

A big problem for finance theorists is that the markets are an area where dramatically unexpected, dramatically impacting events show a historical tendency to make themselves regularly present. That is, in finance, errors in forecasting and modeling are bound to be made very conspicuous and evident ("How could they miss *that!*"). No place to hide for underperforming economists and mathematicians. If the markets, though still unpredictable, were less subject to Nassim Taleb's famous Black Swans (monstrously unseemly, monstrously consequential occurrences) and behaved more or less smoothly, econometricians and quants may be able to go on toiling away relatively unscrutinized and unquestioned. But, sadly for some, that is not the case here on Earth. Our financial markets are shaped by unpredictable watershed phenomena.

The 1929 Crash, the 1980s Latin American banking crisis, the 1987 Crash, the 1994 bond market meltdown, the 1997 Asian crisis, the 1998 Russian default-LTCM crisis, the 2000 Nasdaq crash, the 2001 Enron bankruptcy, the 2002 WorldCom bankruptcy, and, certainly, the 2007 credit crisis are all extremely impacting events that were not under the prediction radar (keep in mind that with the exception of the first of them, all these debacles occurred at a time when quantitative finance constructs were actively prevalent in the markets and when thousands of academics were spending their days trying to forecast events). We weren't widely warned as to their imminence, as to their inevitability (the *Wall Street Journal* of Monday, October 19, 1987, in a page-one

article made the observation that “*No one is forecasting a crash like that in 1929*”; meanwhile, in the real Wall Street the market was busy that day experiencing its most dramatic one-day massacre ever.

And rightly so, since such “outliers” are bound to be mathematically untamable. It is hard to predict the future existence of something that can’t really be even imagined prospectively nor is represented in the historical data. It is much harder to actually assign rock-solid probabilities to such outcomes (what’s the point of models when you can’t say anything about probabilities?). And what’s really hard is to predict the impact of the outliers, that is, the *expectation* (probability times the associated economic result).

As the many Black Swans that have afflicted the markets show, the real tail events are the final consequences, rather than the (pretty Black Swanly in itself) occurrence that ignited the fuse. The unpredictable Russian default led to the frighteningly system-threatening LTCM meltdown; the unpredictable 2007 credit crisis led to the eye-popping sudden disappearance of the investment banking industry. Foreseeing the igniting Black Swan has proven to be insurmountably challenging; prevising the outcome result is simply not possible, as the markets can literally go to zero, countries can go broke, and banks can melt away into oblivion.

Predicting is relatively easy in Black Swans-devoid “Mediocristan,” where things are boring and outcomes don’t change much (the range of possible uncertain states is very limited). We know that the chance that a U.S. presidential candidate would win more than, say, 85 percent of the vote is predictably insignificant. You could rerun the campaigns over and over and such freakish outcome would never present itself. Or take sports: What is the chance that Roger Federer would lose 50 percent of his tennis matches in the next six months? Barring the Swiss champ being afflicted with some disease, zero. That is, the “Federer asset price” can’t suddenly halve in value. Or consider the odds of finding a 10-foot-tall man? You could reenact the life of the universe several times and still the probability would be insignificant. In Mediocristan, it is actually possible to assign probabilities to things taking place; it is reasonable to discard the extreme as unfathomable.

In “Extremistan,” where the financial markets reside, assets can halve in value (and further) in no time. The rare is not awkward, but frequent (the space of the unknown is amply ample). Outcomes are not enslaved to somewhat stringent constraints, and thus are free to explore the unexplored. In mature democracies, no single party tends to enjoy outlandish

domination. In tennis, top players tend not to lose (too many games) to lesser-ranked peers. But there are no such (granted, nonscientific) rules in the markets. Nothing says that the stock market can't halve next week, or that the value of certain securities can't go to zero. This makes predicting harsh because so many alternative outcomes are possible, including many things that had never happened before. Exciting arenas full of possibilities, like the markets or book sales, are much less tamable than duller (in terms of range of outcomes) arenas like sports, height, or U.S. presidential contests.

Blogger Yaron Koren (yes, *blogger*; if anyone these days would deem such a reference nonrigorous, all I can say is, "Wake up, it's 2008!") believes that the reason we can forecast in Mediocristan but not in Extremistan comes down to the concept of conditional versus independent probabilities: *"In the Mediocristan world of sports, elections, etc., all the factors going into the final outcome are fairly independent of one another: the number of points a team scores in the first half of a game doesn't really affect the number of points they score in the second; whether a person votes for a certain candidate doesn't affect whether their neighbor will vote for that candidate. Thus, for a result to be significantly different from expectations, many things would have to go right (or wrong) independently—enough to make such a result all but impossible. On the other hand, in Extremistan, every event affects every subsequent event. If a book sells a million copies, bookstores begin displaying it prominently; the author gets invited on talk shows to plug it, etc: selling the next million becomes a much easier proposition. Similarly with the price of a stock, or the success of a website, or really most of the other interesting questions in life. . . . So there's a mathematical basis for explaining why the systems that do so well in predicting certain outcomes will fail at all the rest. And why we'll have to remain in the dark about the really important issues."*

It is obviously harder to predict in a world where conditional expectations play a key role. How to tell how one's actions will influence the actions of others? In principle, the act of me buying stock could be seen as increasing the chance of upward prices, as others follow my lead, but it could also cause prices to go down as the market may start to see the stock as overvalued and in need of a correction. It is hard to know when and if the snowball effect will take place, and in which direction. So here you would have two levels of randomness that need to be tamed: first, the original actions by a few people who kick-start a process (will a book get initially sold, will a stock get initially purchased); second, the follow-up by a thundering herd that consolidates the process into a sizeable trend. If

nailing the first one could be tough in itself, deciphering the second one would be truly taxing as the range of possible future paths is expanded, in essence determined by how each individual would react to others' prior actions. That is, in Extremistan things can change much faster and in (apparently) weird directions.

And if theory is not successful at helping us prepare for the events that truly shake our world, what good is it? Perhaps less useful than trying to forecast, say, the winner of a general election by gazing at the stars. In early 2008, when the ravages of the credit crisis were inescapably abundant, Nassim Taleb put it like this: *"If the U.S. Food and Drug Administration monitored the business of financial risk management as rigorously as it monitors drugs, many of these 'scientists' would be arrested for endangering us. We replaced so much experience and common sense with 'models' that work worse than astrology, because they assume that the Black Swan does not exist."*

But, many theorists and indoctrinated outsiders would argue, surely a model, even if somewhat underperforming, is better than nothing. After all, can we afford to walk the markets analytically blind, with no quantitative guide whatsoever? Yes, Taleb says, we can and we should, because *"Trying to model something that escapes modelization is the heart of the problem. . . . Sometimes you need to say, 'No model is better than a faulty model'—like no medicine is better than the advice of an unqualified doctor, and no drug is better than any drug."* The dominance of Black Swans in the markets may make the term "finance theory" somewhat of an oxymoron. When the most important events by far (the 10 largest daily moves in the S&P 500 account for more than half the returns over the past 50 years, the 2007 credit crisis wiped out banks' gains from the previous five years, the major U.S. commercial lenders lost an amount equal to all their previous accumulated profits during the 1980s Latin American debt crisis) cannot be predicted (not just have been consistently non-prophesized, but can't by their very structural nature be presighted), can we really talk of the possibility of a theory?

Financial models suffer from two drawbacks that are particularly acute in a world, like the markets, dominated by rare events. One is that many of the sacred cows in the field assume that the Normal probability distribution reigns supreme, that is, assign negligible chances to asset prices

experiencing wild swings. Another is that the present is being described and the future forecasted through heavy reliance on past historical data.

The summer of 2007 earned a notable place in history on account of the several noteworthy developments that took place under its watch. Tony Blair stepped down as Britain's Prime Minister, Apple released its iPhone, and *The Sopranos* TV show aired its last episode. That holiday period's claim to fame also rests, naturally, on events witnessed in the global financial arena, becoming yet another prominent symbol of the wild tumultuousness that can afflict the markets from time to time (we can all recall stock markets going up by 300 points one day, only to fall by 300 points the following day, only to rise by another 300 points the following, and so on). Summer 2007, and the credit crisis that it witnessed emerge, forever joined a high-profile group of dates (such as September 1998 and October 1987, among many others) that very forcefully show that in the markets the Normal probability distribution does not rule. The actual probability of the extreme is far from negligible.

It is thus with some puzzlement that many readers may receive the news that the assumption of Normality has been a staple of financial theory from its early beginnings (all the Nobels awarded to financial economics are heavily grounded on the Normal assumption; remove such tenet, and the prized theories crumble and crash). Some argue that the reason for this is that it makes mathematical modeling more convenient, as the Normal distribution is quite comfortable to work with. That is, even theorists who may know full well that the markets tend to gyrate wildly with large deviations being the norm, not the rarity, may still borrow from the Normal distribution when concocting analytical constructs. Perhaps their assumption is that other researchers will correct for the deficiency later on, or that pros will learn how to tweak the model so as to make it more attuned to the real world.

Whatever the case, it seems appropriately hard to approve of theories that assume the existence of a Platonic financial universe where the vast majority of events resemble the average and where the probability of extreme deviations is deemed to be negligible. According to the Normal distribution, events that move more than three "*standard deviations*" (the conventionally accepted measure for risk and volatility in finance, itself valid as statistical tool only under the Normality assumption) from the mean should not happen. And yet the markets are almost regularly

displaying behavior that is far, far crazier. We all still remember the complaints by David Viniar, Goldman Sachs CFO, on August 2007 as the credit crisis was starting to break loose: “*We were seeing things that were 25-standard deviation moves, several days in a row.*” In a Normal world such happenstance is utterly impossible. The universe isn’t old enough to accommodate such small probability. Obviously, we must not live under Normality.

And that summer’s travails, though particularly detectable, are not, again, by any means the only high-profile instance of the markets yelling out loud, “Don’t call us Normal!” The posterior, even wilder, events in the credit, equity, and interest markets implacably bear witness, but there is also plenty of historical precedent. During the European Exchange Rate Mechanism debacle in 1992 (whereby Europe’s system of officially managed currency rates collapsed), 50-standard deviation moves in interest rates were witnessed, while 1987’s Black Monday was a 20-standard deviation (or 20-sigma) event. During the summer 1998 convolutions that eventually brought down giant Long Term Capital Management, 15-plus sigma deviations became the norm. Plenty of smaller (yet still sensationally non-Normal) similar gyrations have been observed in finance. So-called “one in a million years” events have been experienced, several times, by people whose age is way below one million years. Which one is wrong, the real world or the model? If you said the model, you got it right. The real probability distribution has fat (not thin) tails that grant extreme events the weighty weight that they deserve. In finance, rare events are not that rare. When it comes to wheeling and dealing in the markets, we are not Normal.

In fact, it seems to me that the assumption of Normality when it comes to the financial markets shows strong correlation with that other famously misguided imposition, namely Prohibition in 1920s America. Just like Prohibition forbade regular folks from (legally) drowning down their sorrows, Normality “forbids” investors from taking the markets beyond certain levels. Such probabilistic assumption denies individuals the capacity to cross certain lines, explore certain territories, discover certain realities. It is, thus, a very constraining assumption. A tyrannical one, you might say. Just like with any form of reactionary totalitarianism, individuals are judged to be of limited capabilities, requiring pre-set, centrally imposed, stringent regulations. They are assumed to be unable to reach beyond certain limits, forever confined to a restricted existence.

A financial theory world ruled by Normality is a world where humans (the only ones that can move a market) are prohibited from realizing their full potential, where they are caged in a dull universe of severely reduced possibilities, where freedom is only a word. Perhaps Prohibition isn't the only historical parallel with the Normality assumption after all. Do I hear *ism*? To all those eager to break free from the Normality dictatorship, history may provide a comforting message. Prohibition and *ism* eventually did, of course, spectacularly fail. Why? Simply put, because they run dramatically counter to human nature. People (generally) want to drink. People (usually) want to be free. If one tries to set artificial limits on humans' natural desires, ambitions, and capabilities, the eventual end result is bound to be one of failure.

Real-life markets show us with astounding regularity that investors (who, despite occasional evidence to the contrary, are all too human) also want to be free. They want to spread their wings and be able to explore any possible price level, no matter how remote, no matter how inaccessible, no matter how unthinkable. They want to realize their full potential, and invariably do so. The theoretical straightjacket imposed by Normality seems as much at odds with humanity as were Prohibition and *ism*. Inhumanely unrealistic. Inhumanely unworkable. And yet all too tempting for certain freedom-denying technocrats.

Interestingly, the unavoidably obvious presence of non-Normal markets may be the direct result of a widespread belief in normal markets on the part of investors and speculators (conventional theories might have played a part in building such expectations). Bluntly stated, people's belief in the absence of rare events will eventually cause the rare event to take place. Rare events must always be unexpected, otherwise they would not occur. The assumption of Normality will make people take actions that will render the actual distribution non-Normal. The religion of thin tails will deliver the paganism of fat tails. Outliers are created by people who don't believe in outliers.

The message from a Normal distribution is that waters will be, for the most part, quite calm. No significant storms on the horizon. It is a comforting message for those considering the possibility of sailing through the marketplace. After all, most people would not dive into a market if they expect that at the end of the road there will be a crash. The possibility of a crash must be deemed negligible if a market rally is to sustain momentum. As more and more Normality-believing

investors join the bandwagon, the discarding of a crash as a viable event becomes more widespread and, in effect, conventional wisdom. The more investors join a booming market, the more normal the investment looks to others. The Normality assumption becomes not just the rationale for entering into the market, but an end in itself. The market reaches a point where people are not buying any specific asset. They are buying Normality (i.e., the complete absence of nasty surprises). If they expected something else, they wouldn't have joined the party.

But with every new inflow of cash into the market, the chances of a rare event go up. The more participants in the market, the more chances that someone, somewhere would react negatively to a new development (such as corporate losses, an accounting scandal, or disappointing economic figures), would panic, and would liquidate as a result, prompting other investors to panic, liquidate, and so on all the way to a crash. In essence, when faced with the unexpected presence of the unexpected, Normality-believers will tremble and exacerbate the downfall. They never believed in outliers until they experienced one, and their reaction gives strength to the outlier, making it stronger, fattening the tails. The non-Normal distribution is thus unavoidably born, a testament to people's wildly changing trading habits. As long as people continue to not expect rare events to occur, rare events will inevitably take place.

Let's now tackle the issue of historical data. Back to Taleb for this: *"In the beginning, when I knew close to nothing about econometrics, I wondered whether the time series reflecting the activity of people now dead or retired should matter for predicting the future. Econometricians who knew a lot more than I did about these matters asked no such question; this hinted that it was in all likelihood a stupid inquiry. . . . I am now convinced that, perhaps, most of econometrics could be useless—much of what financial statisticians know would not be worth knowing."* Polemic stuff, no doubt. If anything, the field of "financial econometrics" seems stronger than ever, with prominent academics and academic institutions devoting lots of attention to it, and with one of its inventors actually receiving the Nobel Prize just a few years ago. And yet it seems hard to disagree with Taleb. As the famed trader-turned-philosopher says, when it comes to the financial markets econometric analysis is bound to be less than relevant.

At its core, econometrics is an attempt to forecast the future based on what happened in the past. As every former and present economics student worldwide can attest, this exercise can involve extremely complex statistical and mathematical maneuvers. It is no exaggeration to say that the proliferation of econometrics has been a decisive factor behind the outrageously excessive formalization of economic theory in the past decades.

Lately, econometrics has found its way into financial research. Past market data is used to predict future market movements, through the use of funky models with increasingly funkier names such as GARCH, EGARCH, AARCH, APARCH, FIGARCH, STARCH, TARCH, SQGARCH, and CESGARCH. But intelligently designed as these tools surely are, it is not easy to become a believer. Simple old-fashioned common sense ruthlessly dictates that past information should not be very useful in predicting the future of financial markets.

Why? Among other reasons (like the fact that you will never be able to capture all the variables that affect decision making), because, as Taleb apparently simplistically though innovatively, insightfully points out, we would be trying to predict what current financial players are going to do based on what ancient players did in the past. In the markets, prices move for one reason only: human action. If more humans decide to buy than to sell, prices will go up. If more humans decide to sell than to buy, prices will go down. Clearly, each human being has his own, independent, decision-making capabilities. The financial prices of a certain historical period would be the result of the actions taken by those individuals active in the market at that time. Those prices thus reflect the average consensual decisions of the players who happened to be around, given the relevant circumstances then present.

Econometricians would try to use those prices to forecast the prices of several periods later. The problem is that many of those individuals originally involved in setting the prices included in the time series used in the analysis would by now be either dead or no longer active in the market. Econometricians would in fact be borrowing from inactive brains, attempting to predict the decision-making process of a group of independently thinking individuals from the decision-making processes of a different group of independently thinking individuals who are no longer around. Why should Peter's particular stock pickings 20 years ago matter for predicting Paul's particular stock pickings today, particularly

since Peter has been retired in the Bahamas for the last decade? It might be sensible to use data from Peter's past actions to predict *Peter's* current actions, but it definitely looks a bit suspect to use that data as a predictor for the actions of another, different, unrelated human being.

What financial econometricians are trying to do is akin to predicting the number of goals to be scored by a soccer player next season by looking at the time series of goals historically scored by his club. As any soccer fan would tell you, it would be weird to try to infer anything relevant from goals-scored data that includes players who no longer play. There are simply different people involved. Since goals (like stock, bond, or commodity prices) are all about people, it seems truly far-fetched to assume that historical time series can tell me anything about the future, no matter how complex the techniques involved. The fact that the legendary George Best managed to score 180 goals in his time at Manchester United in the 1960s tells us absolutely nothing about the scoring capabilities of today's striker Wayne Rooney.

The famous LTCM story can help us understand why borrowing from Peter to predict Paul's actions (i.e., trusting that older situations with different people present under different circumstances can provide a reliable guide as to the future) does not look like a winning proposition. LTCM took big bets in fixed-income and equity markets, making the data-backed assumption that markets return to normality. The fund had constructed money machines that would cash in big when such return predictably took place. Based on the historical evidence, such structure seemed flawless. Nothing could go wrong.

However, LTCM forgot that there was something new in the picture that distorted everything so much that it made past references useless: LTCM itself. Historical data did not reflect the existence of such a giant fund taking such giant positions in a few specific markets. It couldn't, of course, because such a giant had not existed until now. LTCM's actions had changed the game and the probability distribution because now everybody else's actions depended on LTCM's. An LTCM-less past could not be a reliable guide to an LTCM-dominated present. LTCM's boss John Meriwether put it best: *"The hurricane is not more or less likely to hit because hurricane insurance has been written. In financial markets this is not true. The more people write financial insurance, the more likely it is that disaster will happen because the people who know you have sold the insurance can make it happen."*

The data-backed return to normalcy unraveled because, once LTCM suffered a bit of trouble, the rest of the market began to trade against their portfolio, that is, began to bet against normality. As a result, things became more, not less, abnormal than ever. Courtesy of the uniquely unique presence of an entity like LTCM at that very precise point in time. New people and new circumstances can render old-timers hopelessly irrelevant.

The widespread presence of quantitative investors (or quant funds), now a ubiquitous element of the markets, may be a particularly acute case of what could be deemed the “new kids on the block” phenomenon, and of how under such situations past data becomes extra unreliable. Quant punters (in more or less intense fashion) are in the habit of employing very advanced technological and scientific tools in the quest for making money. They also tend to (a la LTCM) take a hard statistical look at the historical rearview mirror as guidance for position taking and risk measurement.

The problem with this is that, of course, never before had so many smart scientists and computer geniuses coincided in the markets, often playing exactly the same type of investment games, often armed with billions of monetary units in ammunition. So when they look at past data (notwithstanding how extremely sophisticated the lenses may be) they don't find themselves. They couldn't, because they weren't there. In the case of certain specific strategies maybe one or two early pioneers had been going at it as far as two or three decades ago, but nothing remotely close to today's reality, both in terms of the crowded number of players, the size of their wallets, and the technological prowess. Today's quant funds are employing computers and mathematical models that simply did not exist until quite recently. When the data captures neither the people nor the tools, it is impossible to confidently borrow much from such ancient wisdom.

Thus, the presence of quanty folks using cutting-edge modernish technologies to trade and with a habit to statistically analyze the past may be akin to a catch-22 situation (on occasions, at least): The historical guide is rendered faulty by the very current (and not past) existence of those folks. Take risk. A quant fund may measure the riskiness of a strategy by back-testing its past performance and building devices such as Value at Risk. But those numbers would not account for the aliveness of that fund and its contemporary siblings, clouding the picture by not reflecting the possibility, for instance, that a liquidation by one member

of the quant family (for whatever reason) would trigger further sell-offs among its brainy peers, thus rendering huge losses for the strategy overall. Something like this, of course, is what happened during August 2007, when quant troubles unraveled global equity markets and, as a consequence, the returns of quant plays themselves.

In a way, all this could be labeled as “Econometrics against Econometrics”; as new complex analytical techniques are devised and applied in the markets, older strategies and strategists are rendered obsolete, and past data is irredeemably condemned to not displaying evidence of the current methodologies at play. As this process progresses, it becomes additionally hard to predict the future based on the past. If Peter and Paul engaged on simpler, fundamentals-driven analysis for picking their portfolios, the fact that they inhabit different times may be less of an issue (though still naturally a big one). It might be somewhat reasonable to draw some useful lessons given that the tools (essentially, reading the newspapers) were so similar. But not so in quantland, where the tools can be drastically renovated from one point to the next.

Besides different, dead, or retired people, using historical data suffers from another, simple problem: How far do I go back into the past, and how sure can I be that such selected past period encompasses all the possible events that can take place in the forecasted period? This is where the Black Swan issue makes a heavy presence. If on October 1, 1987, someone had used 50 years of data to try to predict the behavior of U.S. stock markets, the sample evidence would have dictated that there was no chance in hell that the market may drop 25 percent on a single day. Armed with that information you might, say, have confidently sold out-of-the-money puts on the S&P 500. If the data was right, you could make a boatload of premium money, safe in the knowledge that you would never be exercised. Three weeks later, by October 20, you would of course have been wiped out from the previous day’s quite real 23 percent meltdown on Wall Street.

A Black Swan is by definition something that has never (or very rarely) happened before, making the probabilistic detection of the events that most alter our financial environment through naively looking at past data a pretty hopeless task. Black Swans (or “tail events”) are prospectively incomprehensible; we only fully understand what our imagination missed after the fact. It is simply too much to ask of past data that it should consistently contain warning evidence of phenomena for which there is little or no reliable precedent.

Andrew Lo lists the following as one of his field's most important unanswered questions: "*What is the best way to measure the likelihood of rare events and manage such risks if, by definition, there are so few events in the historical record?*" Nassim Taleb's ready answer would be: "We can't." There is no way for us to conceive of a cold probability figure that can be nonastrologically assigned to a Black Swan taking place. Think about it. Someone asks you for your estimate that Wall Street would tumble by 30 percent next week or that the dollar would be worth as much as the euro by next month, what can you say? Zero chances? Not really, because the fearsome tail event lurks in the darkness of the financial world, always omnipresent, never discardable. But if not zero, then how much? How much probability would you have assigned on March 1, 2008 (share price at \$80) to the event that Bear Stearns (founded in 1923, and a historical bastion of the American financial establishment) would be gone within three weeks? The fact that in its 85-year history Bear had never gone under before made for a hard prediction through the use of historical data. The Black Swan is not probabilistically discernable, and no amount of econometric complexity seems likely to change things.

A very illuminating and rabidly current example of the limited power of past data (either because it can't contain the predictably one-of-a-kind financial Black Swan or because the selected sample period is particularly deficient) are the widely publicized failings of Value at Risk (VaR) models during the credit crisis. VaR is a regulators-sanctioned, industrywide-employed risk measurement tool that aims to describe expected maximum losses (within a certain confidence interval) from a financial position or conglomerate of positions, based on historical data and statistical assumptions (mostly the prevalence of the Normal distribution). VaR models disturbingly failed to predict the monstrous subprime-related losses that have afflicted banks and others. The numbers it had been churning pre-mayhem had been way too low, way too comforting, way too unworrisome. That is, VaR (which outputs are regularly followed by senior management and disclosed in public) provided a picture of tranquility right before the world went crazy. Why?

Simple, the markets had gone through a prolonged calm phase before the summer of 2007, and thus the data (banks tend to use one to five years of historical evidence) described nothing but serenity. According to the most revered risk measure, there was nothing to lose sleep over. This placid message may have endowed financial executives with a refreshing, statistically-backed, "scientifically" reinforced sense of confidence

(VaR had worked too well for a long while). Akin to the captain of the *Titanic* accelerating the pace because recent records showed no presence of large icebergs in that part of the ocean.

Most damning for VaR, firms that ended up doing very badly out of the crisis had lower VaR figures pre-crisis than those who ended up faring relatively well (of course, many times this can be explained by widely differing balance sheet sizes; but the point remains that the largest losses, by far, corresponded to those that were reporting lowish VaR numbers). Merrill Lynch (which posted a Q4 2007 loss of almost \$10 billion, its largest ever) had a much lower Q3 VaR than Goldman Sachs (which, uniquely among Wall Street peers, reported record earnings in Q4). Bear Stearns, which eventual fate does not require clarifications, disclosed a Q3 VaR (average daily VaR of some \$30 million) five times lower than Goldman's. After the fact, Merrill seemed to be convinced that something had not gone quite right with the mathematical risk monitors: "*VaR, stress tests and other risk measures significantly underestimated the magnitude of actual loss from the unprecedented credit market environment,*" said Merrill's Q3 filing with the U.S. Securities and Exchange Commission. "*In the past, these AAA CDO securities had never experienced a significant loss in value.*"

As the crisis progressed and intensified, VaR kept underperforming. While the theoretical expected maximum loss churned out from the model increased as market turbulence hit the roof, it wasn't even in the vicinity of a close reflection of the carnage that was about to ensue (among other reasons, because more recently incorporated data only slowly starts to modify the picture, getting lost in a vast sea of historical, non-mayhem-containing past evidence). Bear's daily VaR was still a lowly \$60 million just days before the firm disappeared and \$8 billion of value melted away. The number of VaR *exceptions* (the number of days when actual trading losses exceeded theoretical losses) reached outrageous levels at most financial institutions, an irreverent admission of the mechanism's utter failures as a risk radar. It's not just that VaR underperformed so savagely. The wound was so sore because the thing failed when guidance was most acutely needed. The past was inexcusably misleading when it mattered the most.

The very disappointing performance of VaR during the crisis has fueled a debate as to how it should be modified so as to prevent similar failings going forward. Some argue that longer data samples should be used, so as to have a better chance of capturing extreme events. But

others defend the opposite tactic, saying that a shorter window would act as a faster warning signal. Anyways, the problems run deeper, are more structural than the mere arbitrary selection of a certain time window. For one, VaR is calculated in terms of sigma, our familiar standard deviation parameter, which, once more, only works as measure of dispersion if we assume Normality (naturally, such probabilistic assumption works heavily in the direction of rendering unrealistically low risk numbers). Secondly, VaR may have been condemned to a hopeless task by its academic, quanta, and regulatory sponsors: the past is simply not a reliable guide to the future when there are humans around doing mischief. After all, it's not as if the 2007 crisis signified VaR's first-ever crisis of confidence.

Less than a decade ago, this glorified risk alerter went through another very painful period, again letting the financial world down. During the 1998 crisis (coming on the heels of the prior year's Asian crisis) banks experienced several acute VaR exceptions (i.e., true losses turned out to be way higher than those forecasted by the model). In one study of U.S. banks, some institutions were found to have had up to three and even five exceptions during the August–October 1998 period, when the model (at 99 percent confidence level) would predict only one exception out of every 100 working days. And not only that, the exceptions (as during the credit crisis) were quite large, more than two standard deviations beyond VaR in some cases, and more than seven sigmas in another. Under the Normal distribution, the probability of a loss just one standard deviation beyond 99 percent is virtually zero. So banks' actual P&L suffered much bigger losses, much more often than VaR had warned about. Interestingly (or, as some may have it, scarily), VaR quite possibly contributed decisively to the Russian default transforming into a pronounced tailspin for global markets, and the system-threatening LTCM collapse. Trying to predict the future based on historical data may not just be utterly impractical, but actually pretty dangerous. But that's another story that we reserve for later.

While failing to predict Black Swans is certainly a big indictment on finance theory (failing at warning when the big bad wolf is at the door), one could generously argue that, well, Black Swans are so utterly unexpected that perhaps theorists could be somehow excused (okay, may counter-punch some theory abhorers, but then don't allow financial

economists and quantitative analysts to go around saying that they can mathematically tame the markets; without Black Swans, there are no markets). But even after having been granted such generosity, theory would still present a shaky report card. It turns out that economists are also lacking when it comes to predicting the small stuff, those regular market and economic movements that, while important, are not likely to cause panic-inducing tremors. That is, it is not only that headline-grabbing crises are not being predicted, but even non-crisis-caliber, seminormal changes in key variables are consistently being widely mistargeted.

We are all familiar with how off-the-mark predictions of future GDP growth, inflation, unemployment, exchange rates, stock markets, or interest rates have traditionally been. After all, there is a whole huge market (called the government bond market) that thrives on such unreliability. Traders make bets on bonds based on their perception as to the future levels of variables such as those listed above, or on their perception as to their peers' perceptions as to those future levels, which would be affected by currently available forecasts. If such forecasts were invariably right, little money would be made in the market and activity would dry up. It is thus tempting to conclude that one of the reasons for the bond market's extraordinary liquidity is that pros have little faith in economic predictions.

Long before the appearance of the Black Swan concept (and of economists-basher-in-chief Taleb) in the scene, economic forecasting and modeling had already received plenty of negative praise. For instance, celebrity Harvard professor and author John Kenneth Galbraith once uttered that "*The only function of economic forecasting is to make astrology respectable.*" Paul Ormerod, a leading UK forecaster, published in the late 1990s a polemic tome called *The Death of Economics* where he offered that "*The record of economists in understanding and forecasting the economy at the macro-level is not especially impressive. Indeed, uncharitable writers might be inclined to describe it as appalling. . . . The Japanese recession, by far the deepest since the war, was not predicted. Neither the strength of the recovery in America in the second half of 1992 nor the slowdown of the recession in Germany was foreseen by the models.*" A 2001 paper commissioned by the U.S. Federal Reserve's Board of Governors opened by openly stating that "*Economists have never had much luck in forecasting asset prices in general or exchange rates in particular. . . .*" In 1994, *Wall Street Journal* economics editor Alfred Malabre's very readable book "Lost Prophets" reflected on the dreadful forecasts

he had witnessed during his career: *“In late September 1969, a bare three months before a recession actually began, I conducted a survey for the Journal. The headline of my article carrying the survey results read ‘Most Economists Doubt Recession Will Occur.’ The consensus forecast for the year ahead was that overall economic activity would rise slightly more than 5%. . . . In fact, GNP fell in the final quarter of 1969. . . . the measure continued to drop through much of 1970. . . . Most forecasters, having incorrectly signaled the start of a recession in 1968, now compounded their error by predicting recession-free growth at the very time a recession was setting in. Credentials seemed to matter little.”* And so on.

In 1999, Washington, D.C.-based Heritage Foundation conducted a study on the forecasting ability of the International Monetary Fund’s famous twice-yearly economic projections (presented as part of the IMF’s World Economic Outlook). Given the IMF’s global clout and, crucially, its deep bench of analytically oriented PhDs (though it must be said that the Fund seems to have of late embarked on a PhDs-dismissing strategy), it could be reasonably argued that the accuracy of its forecasts can serve us well for the purpose of analyzing the general reliability of the “science” of economic predicting.

Heritage’s study looked at the IMF’s forecasts for 1971–1998 for both industrial and developing countries (as a libertarian über-American institution, the Foundation was trying to attest the real effectiveness of the IMF, to which the U.S. government had just allocated several billions of dollars). The findings showed that while overall forecasting performance was quite decent when it came to developed nations (it was understandably paltrier in the case of less developed ones), IMF economists had consistently missed out on key “turning points,” including Latin American hyperinflation in the 1980s (inflation forecasts made mistakes in the hundreds of percentage points), industrial growth slowdown in the mid-1990s (with across-the-board overoptimistic projections in the 1–2 percent range), and Japan’s economic crisis in the 1990s (persistent overestimation). In other words, some of the world’s most applauded econometricians had completely missed the most decisive events.

In 2001, the Riksbank (Central Bank of Sweden) conducted its own study, testing the predicting abilities of a very large sample of forecaster-wanna-bes, including investment banks, corporates, rating agencies, and universities (all presumably employing some type of quantitative estimator). The evidence showed that, for the 1990–2001 period, crystal ballers erred annual GDP predictions (taking into account both upside

and downside errors), on average, by 1.20 percent (U.S.), 1.60 percent (Japan), or 0.93 percent (Germany), and in similar inflation forecasts by 0.55 percent (U.S.), 0.48 percent (Japan), or 0.61 percent (Germany). Interestingly, the best names (in principle, those able to hire the most renowned forecasters) did not seem to perform better. Once more, turning points were not foreseen (until they had happened).

Perhaps the most illustrating analysis of how sophisticated quantitative methods fare when it comes to guessing the future is the one conducted during 1979–2000 by Spyros Makridakis and Michele Hibon (also mentioned in Taleb's *The Black Swan*), professors at the highly prestigious French business school INSEAD. They essentially conducted a forecasting competition among econometricians, focusing on business and economic time series. The goal, of course, was to see how accurate the methods proved to be. The first such test, taking place in 1979, yielded the surprising conclusion that simple methods outperformed sophisticated ones. This was not well received by the econometric intelligentsia. To respond to such criticisms, Makridakis and Hibon launched the so-called M-Competition in 1982, increasing the number of time series and of methods and, crucially, having many other experts conduct their own forecasts using their preferred instruments. The empirical results did not vary. Statistically complex tools do not perform better, in spite of their technical prowess. Such strong empirical evidence seems to have been ignored by theoretical econometricians, who have unveiled themselves to be extremely hostile to such verification exercises (they obviously don't want the world to know the results). Rather, econometricians, Makridakis and Hibon offer, have concentrated on developing yet more abstruse models without regard for the ability of those models to more accurately predict real-life data.

Faced with such an unfriendly environment, the INSEAD professors decided in 2000 to embark on a final attempt to settle the accuracy issue, through the M3-Competition which included yet more experts, yet more methods, yet more series. The shocking conclusion? Statistically complex methods do not necessarily produce more accurate forecasts. Makridakis and Hibon conclude with what may seem to many as unfettered common sense, but may be considered hostile fire inside many an ivory tower office: *"Pure theory and elaborate methods are of little practical value unless they can contribute to improve the accuracy of post-sample predictions. . . . [T]he time has come to accept this finding so that pragmatic ways can be found to*

improve predictions. . . . [T]hose criticizing Competitions, and empirical studies in general, should stop doing so and instead concentrate their efforts on explaining the anomalies between theory and practice.”

In this chapter, we have provided several arguments that seem to reinforce the view that *finance theory* may be considered an oxymoronic term. The preeminent presence of unchartable humans, the dominant weight of unimaginable and unprecedented Black Swans, and the limited explanatory power of historical data all indicate that financial markets are doubtfully tamable through mathematical wizardry, no matter how complex (in fact, more complexity may result in even less reliable results). So why do economists and quants persist? And should we care that they do? The first question will be (tentatively) answered in the next couple of chapters. We will tackle the second question now.

Should we care that a few hundred professors and financial engineers choose to spend their days trying to apply mathematical models to the practical-only discipline of finance? Many of us may deem the effort doomed from the start and hopelessly hopeless but, really, what's it to us if a bunch of strangers have chosen to embark down that unseemly road? If they have been lucky enough to find university deans and trading floor honchos willing to finance (extremely generously in some cases) such a lifestyle, then more power to them, right?

Not only that. Perhaps quantitative finance research should actually be encouraged (again, provided that others are the ones having to spend their days immersed in stochastic calculus, numerical methods, and time series analysis). Think about it. Wouldn't it be nice if the hurly-burly of the markets could really be accurately synthesized through a few equations and theoretical dogmas? Certainly, few potential discoveries appear as temptingly attractive as that one. Playing God in this case would not be entirely blasphemous: The end goal would be highly beneficial (if you like reduced volatility and “fairer” asset prices), and the techniques used are quite decorously intelligent. Shouldn't we actively encourage brilliant mathematicians from Carnegie Mellon University, Goldman Sachs, or Standard & Poor's to focus their talents on building models that aim to unlock, and tame once and for all, the markets' DNA? The possible upside seems obviously grand. Should we deny financial economists the triumph in store?

Or, rather, should we protect them from themselves, and, more importantly, protect ourselves from them? It is hard to deny that thanks to the quantification of finance, market practices and research have benefited from the arrival of previously unsuspected characters (mostly from the physical sciences) who have contributed magnificently. The contributions to the finance arena of people like Emanuel Derman, Paul Wilmott, or Steven Shreve might have never happened had analytics and modeling not gained greatly enhanced relevance inside trading floors. So the theoretization of finance has produced tangible benefits, at least when it comes to the quality of human capital.

But there is a darker side. First, economists and mathematicians may be missing out on a lot by embracing abstruseness. Not only would all that time spent solving equations may be later revealed as wasted in a sea of inapplicability, but the opportunity cost from not focusing on the other stuff could be taxing. Why be involved in finance at all if you are going to turn a blind eye to the all-exciting, real-world aspect of the markets? Why hide yourself behind an imaginary, self-concocted Platonic universe, when the real version is so unmissably sexy? Renowned journalist John Cassidy once wrote that 1996's Economics Nobel winner William Vickrey apparently refused to be judged by the mathematical scheming that had earned him the prize (at age 82) in the first place, and instead insisted on being known for his ideas for solving practical problems, like subway reform and the budget deficit. Vickrey in fact dismissed his rewarded theoretical contributions as "*one of my digressions into abstract economics. . . . At best, of minor importance in terms of human welfare.*" Today's financial theorists should act now and try to hedge themselves from experiencing their "Vickrey moment." Unlimited formalism might win you trophies (and gainful employment), but it may not gain you (real) relevance, possibly not even to your own eyes.

Worse, the theories may end up causing harm. In this scenario, the models would be able to claim lots of practical relevance, but of the wrong kind. This has of course happened aplenty in general economics. The applications of Karl Marx's ideas left lots to be desired in terms of human welfare. Keynesianism has been accused of several ills, including unbearable stagflation. Milton Friedman and the Chicago School may have contributed to supporting dictators in power. Reaganite "Supply Siders" created monstrous fiscal deficits. The subfield of financial markets has also been hampered by theories gone wild. While there are precise,

specific instances that can be highlighted, the potential malaise may be more structurally ingrained. That is, because of the proliferation and gradual acceptance (whether real, perceived, or “faked”) of financial theory, the threat that the misuse of the latter implies may be both systematic and hard to eliminate. A key theme that would make a presence throughout this book is that theories may provide a false, misplaced, inadequate sense of confidence and of quantifiable certainty, thus blinding pros to the dimension of the actual risks and encouraging (and excusing) forays into dangerous places. All this math-enabled deceit could end up very badly.

But quite possibly the most potentially harmful effect of the theories would be not so much that they provide faulty guidance based on the illusion of understanding (bad as this would be in itself, naturally), but rather that their prevalence and acceptance would do away with another possible and historically useful source of guidance, namely human intuition and wisdom. When mathematics and statistics take over as decision-making tools, the mind (with its treasure of accumulated experiences and battle scars) may be relegated to a relevance-lacking backseat role. The numbers soullessly churned out from the computer, not the softer sapience of traders and other players, become the key deciding factor. Decades of folk wisdom, passed down through generations of market warriors, may be irredeemably lost, all in the name of the scientification of that which may not be subject to being scientificized in the first place. This is important. The most profoundly insightful and informationally relevant source of financial intelligence (i.e., human experience, intuition, and oral traditions) may be entirely thrown to waste in exchange for the dominance of quantitative tools that present highly doubtful real-life credentials.

In sum, finance theory may present the world with a double threat. Not just the potential for dangerously faulty mathematically charged steering, but also the excreting of that most traditional and primordial of counselors: the experience-honed human gut feeling.

