

those investors included small Norwegian towns, US pension funds, and German regional banks. After all, with an attractive yield, bond insurance, a quasi-governmental AAA stamp of approval, and rising prices, what could go wrong?

Well, everything, really. When already tight credit markets froze in September 2008, prices for these securities collapsed. This further constricted credit, amplifying what had been going on with these securities in the repo market for months. With each bank holding similar assets and liabilities, and as each attempted to rid itself of these assets all at once, prices fell through the floor. But the real surprise, the amplifier, was that the design of these securities, rather than lessening correlation, actually boosted it.

Correlation and Liquidity

In principle, the different tranches of the CDO were isolated from each other. If they went bad, they went bad in reverse order, and it was thought that they couldn't all go bad since different people held different parts of the bond *and* the underlying markets were uncorrelated. Unfortunately, it turned out that the underlying markets were quite strongly correlated. Adding Manhattan to Arizona and Baltimore in a single security *made them correlated*. The sheer volume of cash invested in real estate created one big market in US real estate across the world that became increasingly correlated to equities, particularly to the equities of the banks that were trading real estate. When the income streams of the riskiest (junior) part of the bond dried up as NINJA mortgagors walked away from their debts, investors in the more-secure tranches took fright and dumped their assets as part of the general search for liquidity. What was uncorrelated in theory became extremely correlated in practice. Making matters worse were the CDS (insurance) contracts attached to the CDOs that would pay out if the security defaulted. If these insurance claims happened en masse, the insolvency of the entire system emerged as a distinct possibility. So, when the scope of CDS protection both written by and written on Lehman by firms such as AIG became apparent, not only did the markets take fright, the state for the first time began to see the problem as systemic rather than idiosyncratic, and too big to fail became a reality.

In sum, a web of mispriced risks, with the price set at zero, between the "normal" banking system and the shadow banking system was created through the unseen links between mortgage bonds and CDSs that amplified the grossly underestimated existing correlation between assets. A problem that began in the repo market in 2007 was no longer confined there. The crisis spread globally as investors sought the protection of liquidity but failed to find it. Just as one country's exports depends upon another country's imports, so one bank's liquidity depends upon another bank's willingness to be illiquid. And at that moment, no one wanted to be illiquid.

Note once again that this had nothing to do with the state (beyond the fact that states chose not to regulate derivatives markets—a cause only by omission) or individual morality. The behavior of the whole was not reducible to the sum of its parts. Rather than reducing correlation, these complex assets amplified an already ongoing liquidity crunch that had originated in the repo market months earlier. Too big to fail was the inevitable result of highly levered institutions discovering that all the liquidity in the world really could dry up all at once.

The First Blinder: Tail Risk

So, why didn't anyone see this coming? Queen Elizabeth asked the British economists assembled at the London School of Economics in 2009, who, like analysts everywhere, had failed to see the crisis coming. The answer lies in the way banks measure and manage risk, the third of our seemingly unrelated elements that together generated the crisis and that are quintessentially private-sector, not public-sector failings. Repo runs can start it, and derivatives can amplify it, but to be truly blindsided by a crisis of this magnitude you need to have a theory of risk that denies that catastrophic events can happen in the first place, and then leave it entirely to the self-interested private sector to manage that risk. Unfortunately, almost the entire global financial system worked with just such a theory of risk management.

The first and most basic risk-management technique in finance is called "portfolio diversification," which tries to ensure that your portfolio of assets is not overly exposed to any single source of risk, except by conscious choice. One way to diversify is to try not to buy the same

assets as everyone else. Instead, buy different assets, preferably ones that are uncorrelated or even negatively correlated to other folks' assets.²² The second technique is hedging. Rather than simply rely on passive correlations that are out there in the world to ensure your safety, such as the inverse relationship that typically prevails between the USD and the euro, banks can adopt particular strategies, or trade derivative instruments with specific characteristics, so that the gains from one set of exposures covers (hedges) any losses in another.²³

In principle then, a combination of portfolio diversification and hedging—if appropriately executed in a given market environment—will at the very least keep your investments safe. Think the market will go down? Short sell one asset (profit from a stock price falling by borrowing the stock for a fee, selling it, and then buying it back when its cheaper), and take a long position (buy and hold) in an uncorrelated asset as cover. Want to benefit from the market going up? Use options (the right to buy or sell an asset at a predetermined price) to increase leverage (amplify the bet) while taking a short position as cover. But if this is all it takes to be safe, and to perhaps even make money, why did the banks not see the crisis coming? To answer that question, you need to turn to the trader-turned-philosopher Nassim Nicolas Taleb.

Taleb's Black Swans and Fat-Tailed Worlds

A common refrain when the crisis first hit was that no one *could* have seen it coming. It was the financial equivalent of the meteor that wiped out the dinosaurs. All the diversification and hedging strategies that were supposed to keep banks from blowing up were, as David Viniar, the chief financial officer of Goldman Sachs put it, blindsided by “25 standard deviation moves, several days in a row.”²⁴ This is similar to the “ten sigma event” claim reportedly made by John Meriweather when his hedge fund, Long Term Capital Management (LTCM), blew up in 1998.²⁵

What these sigmas refer to is the number of standard deviations from the mean of a probability distribution at which an outcome will, probabilistically speaking, occur, with each higher sigma (number) being increasingly less likely than the last. According to Mr. Viniar, what happened in 2008 was “comparable to winning the lottery 21 or

22 times in a row.”²⁶ LTCM's ten sigma in 1998 was, likewise, an event that should have occurred roughly three times in the life of the universe. That these two events happened a mere nine years apart shows us that such claims are nonsense. It also tells us why Nassim Taleb has a huge problem with the idea of risk management and financial engineering.

Claims about sigmas typically refer to a “normal-distributed” probability distribution. The shape of the distribution is important. If the shape is “normal,” it conforms to what is called a Gaussian distribution, the classic bell curve, where most of the action is in the middle of the distribution, and less action is likely to occur the further you go out into the tails (see figure 1.1).

To understand why this is important, imagine that we have sampled the height of 10,000 randomly selected adults. We find out that most people are between five and six feet tall, that far fewer people are either seven feet or three feet tall, and that no one in our sample is outside that range. Knowing this, we can figure out the probability of any one person of a given size being close to the mean of the distribution. Under a normal distribution, a one-sigma deviation means that there is a 68 percent chance that person is close to the mean height. Two




Figure 2.1 The “Normal” Distribution of probable events

sigmas translates into a 95 percent chance of being close to the mean, and so on, out into the (very) thin tails, where no one is ever eight feet tall. As the numbers get bigger, the probability of encountering someone of such an extreme size gets exponentially smaller. The chance that someone will fall completely outside the sample becomes so unlikely that you can basically forget about it.

Change the variable from height to default probability, and you can see how such a way of thinking about the likelihood of future events could be of great use to banks as they tried to risk-adjust their portfolios and positions. The piece of technology that allowed banks to do this is known as Value at Risk (VaR) analysis, which is part of a larger class of mathematical models designed to help banks manage risk. What VaR does is generate a figure (a VaR number) for how much a firm can win or lose on an individual trade. By summing VaR numbers, one can estimate a firm's total exposure. Consider the following example.

What was the worst that could have happened to the US housing market in 2008? As in the height example, the answer depends on a data sample that calibrates the model. Prior to 2007, the worst downturn firms had data on was the result of the mortgage defaults in Texas in the 1980s, when houses lost 40 percent of their value. Take this data as the parameter limit, or how far out the tail goes before the sigma becomes too large to be imaginable, and you will conclude, given the assumption of a normal distribution of events, that the probability that *all* the mortgages in your portfolio would lose more than 40 percent of their value at once is ridiculously small. So small, in fact, that you can ignore it. Indeed, the probability that all your mortgage bonds will go bad or that a very large bank will go bust is absurdly small, ten sigma or more, again, *so long as you think that the probability distribution you face is normally distributed*. Your VaR number, once calculated, would reflect this.

Nassim Taleb never bought into this line of thinking. He had been a critic of VaR models as far back as 1997, arguing that they systematically underestimated the probability of high-impact, low-probability events. He argued that the thin tails of the Gaussian worked for height but not for finance, where the tails were "fat." The probabilities associated with fat tails do not get exponentially smaller, so outlier events

are much more frequent than your model allows you to imagine. This is why ten-sigma events actually happen nine years apart.

Taleb's 2006 book *The Black Swan*, published before the crisis, turned these criticisms of VaR into a full-blown attack on the way banks and governments think about risk. Taleb essentially asked the question, what would happen if you ran into an eight-and-a-half-foot-tall person having sampled 10,000 people who were shorter? You might, given that we have never run into such a person, say with confidence that she doesn't exist. Meeting her would be a ten-sigma event. Taleb would bet against you, and you would lose, because in finance there is no way of knowing that you will not run into the equivalent of an eight-and-a-half foot-tall person.

Key here is the issue of observational experience. If you haven't been around for a third of the life of the universe (ten sigma), then how can you know what is possible over that time period? It's the assumed distribution that tells you what is possible, not your experience. To return to the height example, just because your model estimates that an eight-foot-tall person does not exist, it doesn't follow that she doesn't actually exist and that you will not run into her. In Taleb's example, all swans were white until Europeans went to Australia and found black swans. Their exhaustive, multiyear, multisite sample of all known swans had convinced Europeans that all swans were white—until they were not. Nothing in their prior sample, no matter how complete it was, could have told them that a black swan was coming. How, then, do you hedge against risks that are not in your sample? How can you know that which is unknowable until it happens? The answer is, basically, you can't, and if you think you can, you are setting yourself up for a fall.

Counting the Bullets

One way to think about the problem is to imagine playing a game of Russian roulette. Most people would prefer not to play this game when offered the option because the risk-reward ratio is too high, which is correct if we assume the classic "one bullet and six chambers in the gun" setup. But what if I have information you don't have, derived from a mathematical model called Brains at Risk (BaR), that

tells me the gun has over a billion chambers and only one bullet, and that I can figure out where the bullet is by sampling (pulling the trigger millions of times)? Let's also say that each time I pull the trigger I receive \$100. I click once and am \$100 richer, so I click again. By lunchtime, I am a millionaire and have grown confident. In terms of estimating the risk that I face, each "click without a bang" is a piece of information about the probability distribution. As I sample (click) more, I grow more confident about the shape of the distribution. I think that I am generating a more accurate prediction of where the bullet is with each piece of information (click), right up until the moment when I blow my brains out. I have just run into a black swan: a low probability (given the sample and assumed distribution) and (very) high-impact event.

VaR and associated techniques sample the past to predict the future, and from this information we derive theories about the way the future should play out based on our expectations of the probability distribution rather than our actual experiences in the world. We also assume that more information is better than less information, regardless of how it's generated, and therefore believe that the more we sample, the more we converge upon the world "as it really is." But we do not actually do so. Rather, we are assuming far more stability than is warranted—simply because the gun hasn't gone off yet. As *The Black Swan* tells us, we get hit by events that our sample could not have warned us about right at our point of maximum certainty that such events will not happen.

What VaR and similar models make us forget is that we do not see the generators of reality (the number of chambers in the gun), only their outcomes (the clicks of the trigger), and as a consequence we massively underestimate the payoffs we face, most of which are decidedly negative. We think we see the generators, what causes things, but we do not. Instead, we have theories about what causes things and we act upon those theories, which activities, as the Russian roulette example shows, tend to end both abruptly and badly.

Let's apply the black swan idea to risk management in banks in 2008. Consider a data set comprising returns to the Western banking system. If you were to take a monthly time series average of financial sector profitability from June 1947 until June 2007, you could talk with

some degree of accuracy about the mean rate of return, the "standard" deviation, and all the rest, until June 2007. But if you include returns from July 2007 through December 2008, you will have included an outlier so large that it will blow your earlier historical measures out of the water. Nothing in your VaR or other analyses can tell you that this event is coming. The risk is in the tail, not the middle of the distribution, and it's massive. Like the proverbial drunk looking for his keys only under the lamppost, we are drawn to see "normal" distributions in decidedly nonnormal worlds because that is where we find the light.

So, part of the reason no one saw the crisis coming lay within the very models that the banks used to see things coming. Such models see the future only as a normally distributed replication of the past. This makes big, random, game-changing events impossible to foresee, when in fact they are all too common. Such technologies give us, as Taleb says, the illusion of control. We thought that we were diversified and hedged. We thought that we were taking few risks, when they were in fact mounting exponentially, just below the surface, ready to blow up. This is why the events of 2007 and 2008 seemed to participants to be ten and even twenty-five sigma, but what it really shows is that the models used were worse than inaccurate. As Andy Haldane of the Bank of England put it, "these models were both very precise and very wrong."²⁷ Add tail-risk-blind management techniques to a derivative-amplified and leverage-enhanced run on the repo market, and you end up with one heck of a multi-trillion-dollar mess. Not only did we not see it coming, we didn't see it coming because we didn't think it was possible in the first place.

Note once again how none of this has anything to do with the state's spending habits or individual morality. The causes are once again systemic and arise out of the interaction of the parts to produce an outcome that is irreducible to them. Why would people have such faith in a technology that hides risk rather than measures it? To answer that question, we need to address the deepest cause of the crisis—the other reason no one saw it coming: the theories of a generation of economic thinkers who only ever saw markets as good and the state as bad, which takes us back to economics as a morality tale, albeit of a different type.