

In Nature's Casino

By [MICHAEL LEWIS](#)

Published: August 26, 2007

It was Aug. 24, 2005, and New Orleans was still charming. Tropical Depression 12 was spinning from the Bahamas toward Florida, but the chances of an American city's being destroyed by nature were remote, even for one below sea level. An entire industry of weather bookies — scientists who calculate the likelihood of various natural disasters — had in effect set the odds: a storm that destroys \$70 billion of insured property should strike the United States only once every 100 years. New Orleanians had made an art form of ignoring threats far more likely than this; indeed, their carelessness was a big reason they were supposedly more charming than other Americans. And it was true: New Orleanians found pleasure even in oblivion. But in their blindness to certain threats, they could not have been more typically American. From Miami to San Francisco, the nation's priciest real estate now faced [beaches](#) and straddled fault lines; its most vibrant cities occupied its most hazardous land. If, after World War II, you had set out to redistribute wealth to maximize the sums that might be lost to nature, you couldn't have done much better than Americans had done. And virtually no one — not even the weather bookies — fully understood the true odds.



Sasha Bezzubov for The New York Times

Courting destruction in Palm Beach, Fla.

But there was an exception: an American so improbably prepared for the havoc Tropical Depression 12 was about to wreak that he might as well have planned it. His name was John Seo, he was 39 years old and he ran a hedge fund in Westport, Conn., whose chief purpose was to persuade investors to think about catastrophe in the same peculiar way that he did. He had invested nearly a billion dollars of other people's money in buying what are known as "cat bonds." The buyer of a catastrophe bond is effectively selling catastrophe insurance. He puts down his money and will lose it all if some specified bad thing happens within a predetermined number of years: a big hurricane hitting Miami, say, or some insurance company losing more

than \$1 billion on any single natural disaster. In exchange, the cat-bond seller — an insurance company looking to insure itself against extreme losses — pays the buyer a high rate of interest.

Whatever image pops to mind when you hear the phrase “hedge fund manager,” Seo (pronounced so) undermines it. On one hand, he’s the embodiment of what Wall Street has become: quantitative. But he’s quirky. Less interested in money and more interested in ideas than a Wall Street person is meant to be. He inherited not money but math. At the age of 14, in 1950, his mother fled North Korea on foot, walked through live combat, reached the United States and proceeded to become, reportedly, the first Korean woman ever to earn a Ph.D. in mathematics. His father, a South Korean, also came to the United States for his Ph.D. in math and became a professor of economic theory. Two of his three brothers received Ph.D.’s — one in biology, the other in electrical engineering. John took a physics degree from [M.I.T.](#) and applied to Harvard to study for his Ph.D. As a boy, he says, he conceived the idea that he would be a biophysicist, even though he didn’t really know what that meant, because, as he puts it, “I wanted to solve a big problem about life.” He earned his doctorate in biophysics from Harvard in three years, a department record.

His parents had raised him to think, but his thoughts were interrupted once he left Harvard. His wife was pregnant with their second child, and the health plan at Brandeis University, where he had accepted a job, declared her pregnancy a pre-existing condition. He had no money, his parents had no money, and so to cover the costs of childbirth, he accepted a temp job with a Chicago trading firm called O’Connor and Associates. O’Connor had turned a small army of M.I.T. scientists into options traders and made them rich. Seo didn’t want to be rich; he just wanted health insurance. To get it, he agreed to spend eight weeks helping O’Connor price esoteric financial options. When he was done, O’Connor offered him 40 grand and asked him to stay, at a starting salary of \$250,000, 27 times his post-doc teaching salary. “Biophysics was starved for resources,” Seo says. “Finance was hurling resources at problems. It was almost as if I was taking it as a price signal. It was society’s way of saying, Please, will you start solving problems over here?”

His parents, he suspected, would be appalled. They had sacrificed a lot for his academic career. In the late 1980s, if you walked into the Daylight Donuts shop in Dallas, you would have found a sweet-natured Korean woman in her early 50s cheerfully serving up honey-glazed crullers: John’s mom. She had abandoned math for motherhood, and then motherhood for doughnuts, after her most promising son insisted on attending M.I.T. instead of S.M.U., where his tuition would have been free. She needed money, and she got it by buying this doughnut shop and changing the recipe so the glaze didn’t turn soggy. (Revenues tripled.) Whatever frustration she may have felt, she hid, as she did most of her emotions. But when John told her that he was leaving the university for Wall Street, she wept. His father, a hard man to annoy, said, “The devil has come to you as a prostitute and has asked you to lie down with her.”

A willingness to upset one’s mother is usually a promising first step to a conventional Wall Street career. But Seo soon turned Wall Street into his own private science lab, and his continued interest in deep questions mollified even his father. “Before he got into it, I strongly objected,” Tae Kun Seo says. “But now I think he’s not just grabbing money.” He has watched his son quit one firm to go to work for another, but never for a simple promotion; instead, John has moved to

learn something new. Still, everywhere he goes, he has been drawn to a similar thorny problem: the right price to charge to insure against potential losses from extremely unlikely financial events. “Tail risk,” as it is known to quantitative traders, is where it falls in a bell-shaped probability curve. Tail risk, broadly speaking, is whatever financial cataclysm is believed by markets to have a 1 percent chance or less of happening. In the foreign-exchange market, the tail event might be the dollar falling by one-third in a year; in the bond market, it might be interest rates moving 3 percent in six months; in the stock market, it might be a 30 percent crash. “If there’s been a theme to John’s life,” says his brother Nelson, “it’s pricing tail.”

And if there has been a theme of modern Wall Street, it’s that young men with Ph.D.’s who approach money as science can cause more trouble than a hurricane. John Seo is oddly sympathetic to the complaint. He thinks that much of the academic literature about finance is nonsense, for instance. “These academics couldn’t understand the fact that they couldn’t beat the markets,” he says. “So they just said it was efficient. And, ‘Oh, by the way, here’s a ton of math you don’t understand.’ ” He notes that smart risk-takers with no gift for theory often end up with smart solutions to taking extreme financial risk — answers that often violate the academic theories. (“The markets are usually way ahead of the math.”) He prides himself on his ability to square book smarts with horse sense. As one of his former bosses puts it, “John was known as the man who could price anything, and his pricing felt right to people who didn’t understand his math.”

In the mid-1990s, when Wall Street first noticed money to be made covering the financial risks associated with [hurricanes](#) and earthquakes, it was inevitable that someone would call John Seo to ask him if he could figure out how to make sense of it. Until then, he had specialized in financial, not natural, disasters. But there was a connection between financial catastrophe and natural catastrophe. Both were extreme, both were improbable and both needed to be insured against. The firm that called him was Lehman Brothers, whose offer enticed Seo to quit his job and spend his first year at Lehman learning all he could about the old-fashioned insurance industry.

Right away, he could see the problem with natural catastrophe. An insurance company could function only if it was able to control its exposure to loss. Geico sells auto insurance to more than seven million Americans. No individual car accident can be foreseen, obviously, but the total number of accidents over a large population is amazingly predictable. The company knows from past experience what percentage of the drivers it insures will file claims and how much those claims will cost. The logic of catastrophe is very different: either no one is affected or vast numbers of people are. After an earthquake flattens Tokyo, a Japanese earthquake insurer is in deep trouble: millions of customers file claims. If there were a great number of rich cities scattered across the planet that might plausibly be destroyed by an earthquake, the insurer could spread its exposure to the losses by selling earthquake insurance to all of them. The losses it suffered in Tokyo would be offset by the gains it made from the cities not destroyed by an earthquake. But the financial risk from earthquakes — and hurricanes — is highly concentrated in a few places.

There were insurance problems that were beyond the insurance industry’s means. Yet insurers continued to cover them, sometimes unenthusiastically, sometimes recklessly. Why didn’t

insurance companies see this? Seo wondered, and then found the answer: They hadn't listened closely enough to Karen Clark.

Thirteen years before what would become Tropical Storm Katrina churned toward Florida — on Monday, Aug. 24, 1992 — Karen Clark walked from her Boston office to a nearby Au Bon Pain. Several hours earlier, Hurricane Andrew had struck Florida, and she knew immediately that the event could define her career. Back in 1985, while working for an insurance company, Clark wrote a paper with the unpromising title “A Formal Approach to Catastrophe Risk Assessment in Management.” In it, she made the simple point that insurance companies had no idea how much money they might lose in a single storm. For decades Americans had been lurching toward catastrophe. The 1970s and '80s were unusually free of major storms. At the same time, Americans were cramming themselves and their wealth onto the beach. The insurance industry had been oblivious to the trends and continued to price catastrophic risk just as it always had, by the seat of its pants. The big insurance companies ran up and down the Gulf Coast selling as many policies as they could. No one — not even the supposed experts at Lloyd's of London — had any idea of the scope of new development and the exposure that the insurance industry now had.

To better judge the potential cost of catastrophe, Clark gathered very long-term historical data on hurricanes. “There was all this data that wasn't being used,” she says. “You could take it, and take all the science that also wasn't being used, and you could package it in a model that could spit out numbers companies could use to make decisions. It just seemed like such an obvious thing to do.” She combined the long-term hurricane record with new data on property exposure — building-replacement costs by ZIP code, engineering reports, local building codes, etc. — and wound up with a crude but powerful tool, both for judging the probability of a catastrophe striking any one area and for predicting the losses it might inflict. Then she wrote her paper about it.

The attention Clark's paper attracted was mostly polite. Two years later, she visited Lloyd's — pregnant with her first child, hauling a Stone Age laptop — and gave a speech to actual risk-takers. In nature's casino, they had set themselves up as the house, and yet they didn't know the odds. They assumed that even the worst catastrophe could generate no more than a few billion dollars in losses, but her model was generating insured losses of more than \$30 billion for a single storm — and these losses were far more likely to occur than they had been in the previous few decades. She projected catastrophic storms from the distant past onto the present-day population and storms from the more recent past onto richer and more populated areas than they had actually hit. (If you reran today the hurricane that struck Miami in 1926, for instance, it would take out not the few hundred million dollars of property it destroyed at the time but \$60 billion to \$100 billion.) “But,” she says, “from their point of view, all of this was just in this computer.”

She spoke for 45 minutes but had no sense that she had been heard. “The room was very quiet,” she says. “No one got up and left. But no one asked questions either. People thought they had already figured it out. They were comfortable with their own subjective judgment.” Of course they were; they had made pots of money the past 20 years insuring against catastrophic storms.

But — and this was her real point — there hadn't been any catastrophic storms! The insurers hadn't been smart. They had been lucky.

Clark soon found herself in a role for which she was, on the surface at least, ill suited: fanatic. "I became obsessed with it," she says. One big player in the insurance industry took closer notice of her work and paid her enough to start a business. Applied Insurance Research, she called it, or A.I.R. Clark hired a few scientists and engineers, and she set to work acquiring more and better data and building better models. But what she really was doing — without quite realizing it — was waiting, waiting for a storm.

Hurricane Andrew made landfall at 5 on a Monday morning. By 9 she knew only the path of the storm and its intensity, but the information enabled her to estimate the losses: \$13 billion, give or take. If builders in South Florida had ignored the building codes and built houses to lower standards, the losses might come in even higher. She faxed the numbers to insurers, then walked to Au Bon Pain. Everything was suddenly more vivid and memorable. She ordered a smoked-turkey and Boursin cheese sandwich on French bread, with lettuce and tomato, and a large Diet Coke. It was a nice sunny day in Boston. She sat outside at a small black table, alone. "It was too stressful to be with other people," she says. "I didn't want to even risk a conversation." She ate in what she describes as "a catatonic state." The scuttlebutt from Lloyd's already had it that losses couldn't possibly exceed \$6 billion, and some thought they were looking at a loss of just a few hundred million. "No one believed it," she says of her estimate. "No one thought it was right. No one said, 'Yeah, \$13 billion sounds like a reasonable number.'" As she ate, she wondered what \$13 billion in losses looked like.

When she returned to the office, her phones were ringing. "People were outraged," she says. "They thought I was crazy." One insurance guy called her, chortling. "A few mobile homes and an Air Force base — how much could it be?" he said.

It took months for the insurers to tote up their losses: \$15.5 billion. (Building codes in South Florida had not been strictly enforced.) Fifteen and a half billion dollars exceeded all of the insurance premiums ever collected in Dade County. Eleven insurance companies went bust. And this wasn't anything like the perfect storm. If it had gone into Miami, it could have bankrupted the whole industry. Clark had been right: the potential financial losses from various catastrophes were too great, and too complicated, to be judged by human intuition. "No one ever called to congratulate me," Clark says, laughing. "But I had a lot of people call and ask to buy the model."

After Hurricane Andrew came a shift in the culture of catastrophe. "This one woman really created the method for valuing this risk," says John Seo. Clark's firm, A.I.R., soon had more than 25 Ph.D.'s on staff and two competitors, Eqecat and Risk Management Solutions. In its Bay Area offices, R.M.S. now houses more than 100 meteorologists, seismologists, oceanographers, physicists, engineers and statisticians, and they didn't stop at hurricanes and earthquakes but moved on to flash floods, wildfires, extreme winter storms, tornadoes, tsunamis and an unpleasant phenomenon delicately known as "extreme mortality," which, more roughly speaking, is the possibility that huge numbers of insured human beings will be killed off by something like a global pandemic.

The models these companies created differed from peril to peril, but they all had one thing in common: they accepted that the past was an imperfect guide to the future. No hurricane has hit the coast of Georgia, for instance, since detailed records have been kept. And so if you relied solely on the past, you would predict that no hurricane ever will hit the Georgia coast. But that makes no sense: the coastline above, in South Carolina, and below, in Florida, has been ravaged by storms. “You are dealing with a physical process,” says Robert Muir-Wood, the chief scientist for R.M.S. “There is no physical reason why Georgia has not been hit. Georgia’s just been lucky.” To evaluate the threat to a Georgia beach house, you need to see through Georgia’s luck. To do this, the R.M.S. modeler creates a history that never happened: he uses what he knows about actual hurricanes, plus what he knows about the forces that create and fuel hurricanes, to invent a 100,000-year history of hurricanes. Real history serves as a guide — it enables him to see, for instance, that the odds of big hurricanes making landfall north of Cape Hatteras are far below the odds of them striking south of Cape Hatteras. It allows him to assign different odds to different stretches of coastline without making the random distinctions that actual hurricanes have made in the last 100 years. Generate a few hundred thousand hurricanes, and you generate not only dozens of massive hurricanes that hit Georgia but also a few that hit, say, Rhode Island.

The companies’ models disagreed here and there, but on one point they spoke with a single voice: four natural perils had outgrown the insurers’ ability to insure them — U.S. hurricane, California earthquake, European winter storm and Japanese earthquake. The insurance industry was prepared to lose \$30 billion in a single event, once every 10 years. The models showed that a sole hurricane in Florida wouldn’t have to work too hard to create \$100 billion in losses. There were concentrations of wealth in the world that defied the logic of insurance. And most of them were in America.

The more John Seo looked into the insurance industry, the more it seemed to be teetering at the edge of ruin. This had happened once before, in 1842, when the city of Hamburg burned to the ground and bankrupted the entire German insurance industry many times over. Out of the ashes was born a new industry, called reinsurance. The point of reinsurance was to take on the risk that the insurance industry couldn’t dilute through diversification — say, the risk of an entire city burning to the ground or being wiped off the map by a storm. The old insurance companies would still sell policies to the individual residents of Hamburg. But they would turn around and hand some of the premiums they collected to Cologne Re (short for reinsurance) in exchange for taking on losses over a certain amount. Cologne Re would protect itself by diversifying at a higher level — by selling catastrophic fire insurance to lots of other towns.

But by their very nature, the big catastrophic risks of the early 21st century couldn’t be diversified away. Wealth had become far too concentrated in a handful of extraordinarily treacherous places. The only way to handle them was to spread them widely, and the only way to do that was to get them out of the insurance industry and onto Wall Street. Today, the global stock markets are estimated at \$59 trillion. A 1 percent drop in the markets — not an unusual event — causes \$590 billion in losses. The losses caused by even the biggest natural disaster would be a drop in the bucket to the broader capital markets. “If you could take a Magnitude 8 earthquake and distribute its shock across the planet, no one would feel it,” Seo says. “The same principle applies here.” That’s where catastrophe bonds came in: they were the ideal mechanism

for dissipating the potential losses to State Farm, Allstate and the other insurers by extending them to the broader markets.

Karen Clark's model was, for Seo, the starting point. When he first stumbled upon it and the other companies' models, he found them "guilty until proven innocent," as he puts it. "I could see the uncertainty in them," he says, "just by looking at the different numbers they generated for the same storm." When they run numbers to see what would happen if the 1926 Miami hurricane hit the city today, A.I.R. puts the losses at \$80 billion, R.M.S. at \$106 billion and Eqecat at \$63 billion. They can't all be right. But they didn't need to be exactly right, just sort of right, and the more he poked around inside them, the more he felt they were better than good enough to underpin financial decisions. They enabled you to get a handle on the risk as best you could while acknowledging that you would never know it exactly. And after all, how accurate were the models that forecast the likelihood that Enron would collapse? Next to what Wall Street investors tried to predict every day, natural disasters seemed almost stable. "In the financial markets, you have to care what other people think, even if what they think is screwed up," Seo says. "Crowd dynamics build on each other. But these things — hurricanes, earthquakes — don't exhibit crowd behavior. There's a real underlying risk you have to understand. You have to be a value investor."

The models were necessary but insufficient. True, they gave you a rough sense of the expected financial losses, but they said nothing about the rewards. Financial markets exist only as long as investors feel the odds are stacked in their favor. Investors — unlike roulette players — can honestly expect to make a gain (their share in the profits of productive enterprise). But how big a gain? How should the payout vary, from government bonds to blue-chip stocks to subprime mortgages? The rewards in each market tended to vary with investors' moods, but those in catastrophe insurance were just incredibly volatile. Hurricane insurance rates would skyrocket after a big storm, then settle back down. This wouldn't do: if big investors were going to be persuaded to take billions of dollars in catastrophic risk, they would need to feel there was some reason in the pricing of that risk. "The market," as Seo puts it, "needs an acceptable mode of failure."

In the spring of 2001, to the surprise of his colleagues, Seo left his big Wall Street firm and opened a hedge fund — which, he announced, wouldn't charge its investors the standard 2 percent of assets and 20 percent of returns but a lower, flat fee. "It was quixotic," says Paul Puleo, a former executive at Lehman who worked with Seo. "He quits this high-paying job to basically open a business in his garage in a market that doesn't exist." Seo opened his new shop with his younger brother Nelson and then brought in their older brother, Michael. (His third brother, Scott, had studied astrophysics but decided that "there was no future in astrophysics" and eventually turned himself into an ophthalmologist.) Seo named his firm Fermat Capital Management, after one of his intellectual heroes. "I had once read the letters between Pierre de Fermat and Blaise Pascal," he wrote in a recent e-mail message. "From my father I had learned that most great mathematicians were nasty guys and total jerks (check out Isaac Newton . . . extra nasty guy), but when I read the Fermat-Pascal letters, you could see that Fermat was an exception to the stereotype . . . truly a noble person. I loved his character and found that his way of analyzing profitless games of chance (probability theory) was the key to understanding how to analyze profitable games of chance (investment theory)."

Four years later, Seo's hedge fund still faced two problems. The smaller one was that investors were occasionally slow to see the appeal of an investment whose first name was catastrophe. As one investor put it, "My boss won't let me buy bonds that I have to watch the Weather Channel to follow." That objection doesn't worry Seo much. "Investors who object to cat-bond investing usually say that it's just gambling," he says. "But the more mature guys say: 'That's what investing is. But it's gambling with the odds in your favor.' "

His bigger problem was that insurance companies still didn't fully understand their predicament: they had \$500 billion in exposure to catastrophe but had sold only about \$5 billion of cat bonds — a fifth of them to him. Still, he could see their unease in their prices: hurricane- and earthquake-insurance premiums bounced around madly from year to year. Right after Andrew, the entire industry quintupled its prices; a few tranquil years later, prices were back down nearly to where they had been before the storm. Financial markets bounced around wildly too, of course, but in the financial markets, the underlying risks (corporate earnings, people's moods) were volatile. The risk in natural-disaster insurance was real, physical and, in principle, quantifiable, and from year to year it did not change much, if at all. In effect, the insurers weren't insuring against disaster; they were only pretending to take the risk, without actually doing so, and billing their customers retroactively for whatever losses they incurred. At the same time, they were quietly sneaking away from catastrophe. Before the 1994 Northridge earthquake, more than a third of California homeowners had quake insurance; right after, the insurers fled the market, so that fewer than 15 percent of California homeowners have earthquakes in their policies today.

The market was broken: people on fault lines and beachfronts were stuck either paying far too much for their insurance or with no real coverage except the vague and corrupting hope that, in a crisis, the government would bail them out. A potentially huge, socially beneficial market was moments from birth. All it needed was a push from nature. And so on Aug. 24, 2005, John Seo was waiting, waiting for a storm. And here it came.

Wall Street is a machine for turning information nobody cares about into information people can get rich from. Back when banks lent people money to buy homes and then sat around waiting for interest payments, no one thought to explore how quickly homeowners would refinance their mortgages if interest rates fell. But then Wall Street created a market in mortgage bonds, and the trader with better information about how and when people refinance made a killing. There's now a giant subindustry to analyze the inner financial life of the American homeowner.

Catastrophe bonds do something even odder: they financialize storms. Once there's a market for cat bonds, there's money to be made, even as a storm strikes, in marginally better weathermen. For instance, before the 2005 hurricane season, a Bermuda cat-bond hedge fund called Nephila found a team of oceanographers in Rhode Island called Accurate Environmental Forecasting, whose forecasts of hurricane seasons had been surprisingly good. Nephila rented the company's services and traded bonds on the back of its reports. "They kind of chuckle at what we do," says a Nephila founder, Frank Majors. "The fact that we're making \$10 million bets on whether Charley is going to hit Tampa or not. It made them a little nervous at first. We told them not to worry about what we're going to do with the information. Just give it to us."

As Katrina bore down on New Orleans, a cat bond named Kamp Re, issued by the insurance company Zurich, was suddenly at risk. If Zurich lost more than \$1.2 billion on a single hurricane in about a two-year period, investors would lose all their money. If Zurich represented about 3 percent of the U.S. insurance market — that is, it was on the hook for about 3 percent of the losses — a hurricane would need to inflict about \$40 billion in damage to trigger the default. Since no event as big as this had ever happened, it was hard to say just how likely it was to happen. According to R.M.S., there was a 1.08 percent chance that Kamp Re bond holders would lose all their money — assuming the scientists really understood the odds. The deal had been a success. One of its biggest buyers was John Seo.

As Katrina spun, the players in nature's casino gathered around the table. When the storm jogged east and struck not New Orleans directly but the less populated, and less wealthy, coastline between Louisiana and Mississippi, they all had the same reaction — relief — but Hemant Shah felt a special relief. Shah is one of the founders of R.M.S., and he was at that moment driving to catch a flight from San Francisco to New York, where he hoped to speak at a conference devoted to predicting terrorism. When he saw Katrina miss New Orleans, he said to himself, O.K., it's big, but it's not catastrophic, and he boarded his plane.

As he flew across the country, R.M.S. and its competitors replicated Katrina inside their computers in much the same way that Karen Clark had once replicated Hurricane Andrew. Just hours after landfall, all three firms sent clients in the insurance industry their best estimates of financial losses: R.M.S. put them at \$10 billion to \$25 billion; Eqecat called for a range between \$9 billion and \$16 billion; Clark's A.I.R. had a range of \$12.7 billion to \$26.5 billion. Big, as Shah said, but not catastrophic. Traders who had underwritten Kamp Re took calls from an investor at a Japanese bank in London. Cheered by Katrina's path, the fellow was looking to buy some Kamp Re bonds. The traders found another investor eager to unload his Kamp Re holdings. The London investor bought \$10 million of Kamp Re at a price of \$94.

John Seo just watched. For the past four years, he and his brothers had made money at such moments as this: "live" cat trading, it's called. A few investors would inevitably become jittery and sell their cat bonds at big discounts, what with the Weather Channel all hysteria all the time. ("The worst place to go if you're taking risks," says one cat-bond investor, "is the Weather Channel. They're just screaming all the time.") But entering the 2005 hurricane season, the Seo brothers had reconsidered their habit of buying in a storm. "The word had gotten out that buying in the storm was the smart thing to do," Seo says. "And we were afraid our past successes would give us an irrational interest in buying. Everything's all fuzzy in these events. And when things are fuzzy, your brain gives you an excuse to push the envelope. So we adopted a policy, before the season, of staying out of the market."

A few hours later, Hemant Shah's plane landed in New York. Shah turned on his BlackBerry and discovered that the New Orleans levees had broken: much of the city would soon be underwater. "My first reaction," Shah says, "was, Uh-oh, we have a problem." In the imaginary 100,000-year history of hurricanes that R.M.S. had in its computers, no hypothetical storm that struck so far from New Orleans had ever caused the levees to fail. The models, like the intuition they replaced, had a blind spot.

The Kamp Re bonds collapsed, the price dropping from the mid-90s to the low 20s. A few weeks later, an announcement from Zurich American made it clear that the investors in Kamp Re wouldn't be getting any money back, and Kamp Re's price fell from \$20 to 10 cents. But then the real trouble started: R.M.S., the modeling company, declared that it was rethinking the whole subject of hurricane risk. Since 1995, scientists had noted a distinct uptick in hurricane activity in the North Atlantic Basin. The uptick had been ignorable because the storms had not been making landfall. But between July 2004 and the end of 2005, seven of history's most expensive hurricanes had struck the American coast, leaving behind 5.5 million insurance claims and \$81 billion in insured losses. The rise in hurricane size and frequency was no longer ignorable. R.M.S. convened a panel of scientists. The scientists agreed that unusually warm sea-surface temperatures were causing unusually ferocious and frequent storms. The root cause might be [global warming](#) or merely the routine ups and downs of temperatures in the North Atlantic Basin. On cause they failed to agree. On consequence they were united. At the beginning of August 2005, R.M.S. had judged a Katrina-size catastrophe to be a once-in-40-years event. Seven months later, the company pegged it as a once-in-20-years event. The risk had doubled.

It had been just 13 years since Karen Clark's model swept the industry, but the entire catastrophe risk-taking industry now lived at the mercy of these modelers. The scientists were, in effect, the new odds-makers. It was as if the casino owner had walked up to his roulette table, seen a pile of chips on 00 and announced that 00 would no longer pay 36:1 but would henceforth pay only 18:1. The agencies that rated the insurance companies — S & P, Moodys, etc. — relied on the scientists to evaluate their exposure. When the scientists increased the likelihood of catastrophic storms, S & P and Moodys demanded that the insurance companies raise more capital to cover their suddenly more probable losses. And so in addition to the more than \$40 billion they had lost in Katrina, the insurance companies, by edict of the ratings agencies, needed to raise \$82 billion from their shareholders just to keep their investment-grade rating. And suddenly they weren't so eager to expose themselves to losses from hurricanes.

John Seo felt differently. Katrina had cost him millions. But at the same time, in a funny way, it had vindicated his ideas about catastrophe. He had lost only what he had expected to lose. He had found an acceptable mode of failure.

As a boy, John Seo learned everything he could about the Titanic. "It was considered unsinkable because it had a hull of 16 chambers," he says. The chambers were stacked back to front. If the ship hit something head on, the object might puncture the front chamber, but it would likely have to puncture at least three more to sink the ship. "They probably said, What are the odds of four chambers going?" he says. "There might have been a one-in-a-hundred chance of puncturing a single chamber, but the odds of puncturing four chambers, they probably thought of as one in a million. That's because they thought of them as independent chambers. And the chambers might have been independent if the first officer hadn't gambled at the last minute and swerved. By swerving, the iceberg went down the side of the ship. If the officer had taken it head on, he might have killed a passenger or two, but the ship might not have sunk. The mistake was to turn. Often people associate action with lowering risk or controlling risk, but experience shows more often than not that by taking action you only make the risk worse."

The Titanic offered another lesson for the investor in catastrophe: the threats that seem to us the most remote are those we know the least about. Catastrophe risk is fundamentally different from normal risk. It deals with events so rare that experience doesn't help you much to predict them. How do you use history to judge the likelihood of a pandemic killing off 1 in every 200 Americans? You can't. It has happened only once. (The Spanish flu epidemic of 1918.) You lack information. You don't know what you don't know. The further out into the tail you go — the less probable the event — the greater the uncertainty. The greater the uncertainty, the more an investor should be paid to live with it.

The financial markets, or, at any rate, the arcane corner of Wall Street that dealt exclusively with highly unlikely financial events, had figured this out. The traders who sold insurance against extreme market collapses — the tail risks — all tended to charge exactly the same price, between four and five times their expected losses. Expected loss could be defined like this: Say an investor wanted to buy \$1 billion of insurance for a year against a once-in-100-years stock-market crash. The expected loss would be 1 in 100, 1 percent of \$1 billion: \$10 million. The insurance would thus cost \$40 million to \$50 million. The pattern held across Wall Street. The trader at Lehman Brothers who priced stock-market-crash insurance didn't know the trader at Harvard Management who priced the insurance against drastic interest-rate changes, and he didn't know the trader at O'Connor and Associates who priced the insurance against the dollar's losing a third of its value. But their idea of a fair premium for insurance against financial disaster suggested they were reading the same books on the subject — only there were no books. "The reigning theory is that the taste for risk is as arbitrary as the value of a painting," Seo says. "But if this is so, why are these preferences so consistent across markets?"

Seo thought, Maybe risk is not like art. Maybe there is some deep rule that governs it. And maybe the market is groping its way to that rule all by itself.

Intuitively what the market was doing made sense. Highly improbable events were especially unsettling. The person who insured others against an unlikely event faced not only the problem of judging its likelihood; even if he knew how often it would occur, he didn't know when it would occur. Even if you had complete certainty that a U.S. stock-market crash happened just once every 25 years, you still didn't know which year. If you had set up a business to sell crash insurance in January 1987, you would have been bankrupted by the crash in October; on the other hand, if you had gone into the business in 1988, you would have gotten rich. There was no justice in it. The catastrophic risk-taker was a bit like a card counter at the blackjack table allowed to play only a few hands: yes, the odds are in his favor, but he doesn't always get to play long enough for the odds to determine the outcome.

The uncertainty in these extreme, remote market risks meant that the person who took them should be paid more to do so. But how much more? Extreme events were treated on Wall Street as freak outliers that bore no relation to other, more normal events. There was a striking consistency in the pricing of these risks across Wall Street, but there was no hard logic under them: it was all being done by feel.

The logic is what Seo stumbled upon back in 2000 at Lehman Brothers after someone handed him a weird option to price. An industrial company had called Lehman with a problem. It

operated factories in Japan and California, both near fault lines. It could handle one of the two being shut down by an earthquake, but not both at the same time. Could Lehman Brothers quote a price for an option that would pay the company \$10 million if both Japan and California suffered earthquakes in the same year? Lehman turned to its employee with a reputation for being able to price anything. And Seo thought it over. The earthquakes that the industrial company was worried about were not all that improbable: roughly once-a-decade events. A sloppy solution would be simply to call an insurance company and buy \$10 million in coverage for the Japanese quake and then another \$10 million in coverage for the California quake; the going rate was \$2 million for each policy. “If I had been lazy, I could have just quoted \$4 million for the premium,” he says. “It would have been obnoxious to do so, but traders have been known to do it.” If either quake happened, but not both, he would have a windfall gain of \$10 million. (One of his policies would pay him \$10 million, but he would not be required to pay anything to the quake-fearing corporation, since it would get paid only if both earthquakes occurred.)

But there was a better solution. He needed to buy the California quake insurance for \$2 million, its market price, but only if the Japanese quake happened in the same year. All Seo had to do, then, was buy enough Japanese quake insurance so that if the Japanese quake occurred, he could afford to pay the insurance company for his \$10 million California insurance policy: \$2 million. In other words, he didn’t need \$10 million of Japanese quake insurance; he needed only \$2 million. The cost of that was a mere \$400,000. For that sum, he could insure the manufacturing company against its strange risk at little risk to himself. Anything he charged above \$400,000 was pure profit for Lehman Brothers.

And that was that, except it wasn’t. He saw something. Each risk by itself was not unusual: the quakes being insured against were once-a-decade events. But since each earthquake had a 1-in-10 chance of happening in a year, the chances that both of them would occur were far more remote: 1 in 100 (10 percent of 10 percent). When you combined these more ordinary risks, you simulated extremely unlikely ones. “What I noticed, after the fact, is that this exotic option’s price was special,” he says. “It was related to tail pricing.” The risk of catastrophe wasn’t some freak outlier with no connection to more mainstream risks. It bore a fixed relationship to those risks. Indeed, one way of thinking about natural catastrophes was as a combination of more likely events.

Thus the hunches of Wall Street professionals found vindication in Seo’s arithmetic. The expected loss of the more ordinary risk of a single earthquake was \$1 million (a 10 percent chance of a \$10 million loss). The insurance cost \$2 million, or twice the expected loss. The expected loss of the remote combined risk was \$100,000 (a 1 percent chance of a \$10 million loss). But the insurance cost \$400,000: four times the expected loss. All those practical traders who were pricing tail risk at roughly four times the expected losses had been on to something. “Here I saw the beginnings of a market mechanism that directly links 1-in-10-year risk pricing to 1-in-100-year risk pricing,” Seo says. The intuitive reason that extreme, remote risk should be more highly priced than normal everyday risk was “a happy agreement between human psychological perception and hard mathematical logic.”

Seo’s math — which soon left middle school for graduate school — served two purposes: to describe this universal rule about the pricing of risk and to persuade investors that there was a

deeper, hidden logic to investing in catastrophe. They could have some sense of what the price of the risk should be. It was an extraordinary idea: that catastrophe might be fair.

Then came Katrina. The reaction to the storm has put a fine point on Americans' risk disorientation. The single biggest issue in Florida's 2006 governor's race, for instance, was the price of insurance. The Republican, Charlie Crist, got himself elected on the strength of his promise to reduce Floridians' home-insurance rates by creating a state-subsidized pool of \$28 billion in catastrophe insurance coverage. "Florida took this notion of spreading this risk and turned it on its head," says one former state insurance commissioner. "They said, 'We're going to take all this risk ourselves.'" The state sold its citizens catastrophe insurance at roughly one-sixth the market rates, thus encouraging them to live in riskier places than they would if they had to pay what the market charged (and in the bargain, the state subsidized the well-to-do who live near the beach at the expense of the less-well-to-do who don't). But if all the models are correct, \$28 billion might not cover even one serious storm. The disaster waiting to happen in Florida grows bigger by the day, but for a man running for governor of Florida, ignoring it is a political no-brainer. If he's lucky — if no big storms hit in his term — he looks like the genius who saved Floridians billions in catastrophic-risk premiums. If he's unlucky, he bankrupts Florida and all hell breaks loose, but he can shake down the federal government to cover some of the losses.

Louisiana's politicians are usually quicker than most to seize upon shrewd politics that generate terrible social policy, but in this case they could not afford to. Louisiana cannot generate and preserve wealth without insurance, and it cannot obtain insurance except at the market price. But that price remains a mystery. Billions of dollars in insurance settlements — received by local businesses and homeowners as payouts on their pre-Katrina policies — bloat New Orleans banks and brokerage houses. The money isn't moving because the people are paralyzed. It's as if they have been forced to shoot craps without knowing the odds. Businesses are finding it harder than ever to buy insurance, and homeowners are getting letters from Allstate, State Farm and the others telling them that their long relationship must now come to an end. "I've been in the business 45 years," says a New Orleans insurance broker named Happy Crusel, "and I've never seen anything remotely like this." An entire city is now being reshaped by an invisible force: the price of catastrophic risk. But it's the wrong price.

Insurance companies, John Seo says, are charging customers too much — or avoiding their customers altogether — instead of sharing their risk with others, like himself, who would be glad to take it. New Orleans, as a result, is slower than it otherwise would be to rebuild. "The insurance companies are basically running away from society," he says. "What they need to do is take the risk and kick it up to us." They need to spread it as widely as possible across the investment world and, in the process, minimize the cost of insuring potential losses from catastrophes.

But this, too, is happening. The people on Wall Street who specialize in cat bonds now view Katrina as the single most important thing that ever happened to their business: overnight it went from a tiny backwater to a \$14 billion market, and it is now stretching and straining to grow. In March of this year, a single insurer, Allstate, announced its intention to sell \$4 billion in catastrophe bonds. A \$14 billion market is a trivial sum next to the half-trillion or so dollars that the insurance industry stands to lose from megacatastrophes and next to the additional trillions of

dollars worth of property that has gone uninsured in the places most likely to be destroyed by nature, like California, because the insurance is so expensive. But there are all around John Seo signs of a shift in the culture of catastrophe. “It has all the features of providential action,” he says. “It’s like all the actions of man and nature serve to grow the cat-bond market.”

When Katrina struck and his Kamp Re bonds collapsed — from \$100 to 0 — Seo was able to view his loss with detachment. The models had badly underestimated the risk, but it was in the nature of extreme risk that the prediction of it would sometimes be mistaken. “The important thing is that the money wasn’t lost in an unearned manner,” he says, by which he means that it wasn’t lost dishonestly or even unwisely or in what his community of investors would consider a professionally unacceptable manner. Investors will endure losses as long as they come in the context of a game they perceive as basically fair, which is why they don’t abandon the stock market after a crash. “That’s all I need to know,” Seo says. “That’s all my clients need to know.” Actually, he goes even further: “I would be embarrassed if we had a big event and our loss wasn’t commensurate with it. It would mean that we didn’t serve society. We failed society.”

Seo’s returns in 2005 were only slightly positive, compared with the roughly 10 to 12 percent he had been delivering, but the demand for his services boomed. He now controls \$2 billion, or more than twice what he had before the most costly natural disaster in history. Big investors weren’t scared off by Katrina. Just the reverse. It has led many of them to turn to Seo and others like him to make money from catastrophe. And they probably will. But what interests Seo more is what might happen in the bargain, that the financial consequences of catastrophe will be turned into something they have never been: boringly normal.

Michael Lewis is a contributing writer. The paperback edition of his book “The Blind Side: Evolution of a Game” will be published next month.