3

Risky Business

In each case, the companies and their executives grew rich by taking on excessive risk. In each case, the companies collapsed when these risks turned bad. And in each case, their executives are walking away with millions of dollars while taxpayers are stuck with billions of dollars in costs.

—Henry Waxman, Hearings on the Causes and Effects of the AIG Bailout, October 7, 2008

The chances of a disaster were stunningly large.

On July 18, 1969, two days before Neil Armstrong and Buzz Aldrin first set foot on the moon, a speechwriter in President Nixon's office penned a speech for an all too probable situation:

> Fate has ordained that the men who went to the moon to explore in peace will stay on the moon to rest in peace. These brave men, Neil Armstrong and Edwin Aldrin, know that there is no hope for their recovery.

But they also know that there is hope for mankind in their sacrifice.

These two men are laying down their lives in mankind's most noble goal: the search for truth and understanding. They will be mourned by their families and friends; they will be mourned by their nation; they will be mourned by the people of the world; they will be mourned by a Mother earth that dared send two of her sons into the unknown. In their exploration, they stirred the people of the world to feel as one; in their sacrifice, they bind more tightly the brotherhood of man. . . .

For every human being who looks up at the moon in the nights to come will know that there is some corner of another world that is forever mankind.

It's hard to fathom just how risky the moon missions really were. In the early days of the Apollo program, NASA asked General Electric to calculate the chances of successfully landing men on the moon and bringing them back to earth in one piece. The answer they got was shocking: less than 5 percent. According to the (admittedly crude) numbers, a moon landing was so risky that it would be foolhardy to attempt. If NASA had paid attention to the calculations, the Apollo missions would have had to be scrapped—a horrible political disaster for NASA and a humiliation for the nation. So NASA did what NASA tends to do under such circumstances: it crumpled up the calculations, tossed them in the garbage can, and went ahead with the program anyway. In this particular case, it happened to be the right decision. Even though there was an extremely close call—the Apollo 13 mission came within a hair of killing three astronauts—

the moon landings would never have happened had NASA not disregarded the risks.

NASA engineers were famed for their "can do" bravado. No task was too difficult for the scrappy rocket scientists. The agency sneered at even the most daunting odds. Even if the numbers were dismal and the risks were enormous, NASA administrators disregarded them and plowed onward—with good reason. If Congress got wind of just how risky NASA's human spaceflight projects were, the programs could well have been canceled. It was in NASA's interest to disregard risks to keep their projects alive. On occasion, they even diddled with the numbers to make it look like their rockets were much safer than they actually were.

In 1983, the air force commissioned a study to calculate the risk that the brand-new space shuttle launch system would explode during launch. The study found that there was a dangerously high probability of disaster. As two of the study's authors wrote, "The probability of a solid rocket booster (SRB) failure destroying the shuttle was roughly 1 in 35 based on prior experience with this technology." One in thirty-five was an enormous and unacceptable level of risk. After all, the shuttles were supposed to make hundreds of flights, returning their crews safely every single time. If the shuttle would, on average, lose a crew of seven astronauts once every thirty-five flights, the shuttle program was as good as dead. So NASA disregarded the study, instead deciding "to rely upon its engineering judgment and to use 1 in 100,000 as the SRB failure probability estimate." In other words, NASA simply tossed out the 1-in-35 number and substituted a much more acceptable one—in which you could launch a shuttle every day for decades, totaling thousands upon thousands of launches, and expect not to have a single failure.

On January 28, 1986, a bit more than half a second after *Challenger* left the launch pad, a puff of gray smoke coming from its right solid rocket booster heralded disaster. Nobody knew it at the time, but a small rubber seal in the booster had failed. Fifty-nine seconds into the flight, a small flame erupted from the booster and the conflagration quickly grew out of control. Seventy-three seconds after launch, at an altitude of 46,000 feet, *Challenger* exploded in an enormous yellow-white ball of fire. It had taken only twenty-five shuttle launches before the risks caught up with NASA.

NASA management had deliberately understated the risks of a shuttle flight. Instead of facing the unpleasant reality that the shuttle boosters were risky, the agency decided to engineer a lie that was more acceptable. As physicist Richard Feynman, a member of the *Challenger* investigation panel, put it, "As far as I can tell, 'engineering judgment' means that they're just going to make up numbers!" Instead of performing a genuine assessment of the probability that the shuttle would fail, the management would start with a level of risk that was acceptable and work backward. "It was clear that the numbers . . . were chosen so that when you add everything together, you get 1 in 100,000," Feynman wrote. NASA's risk estimates were complete fictions, and nobody noticed until disaster struck.

Risks are tricky. We're pretty bad at estimating them. We spend our time worrying about graphic but uncommon events (meteor strikes, child abductions, and shark attacks) when we should really be worrying about—and preventing—more mundane risks (strokes, heart attacks, and diabetes) that are much more likely to cut our lives short. We spend our money chasing after faint hopes of winning the lottery or hitting pay dirt in a get-rich-quick scheme instead of paying off the credit cards that have a serious chance of driving us into ruin. We are terrified of dying in a plane crash but

think nothing about speeding down the highway while talking on a cell phone. We don't have an internal gauge of what behaviors are truly dangerous and what aren't.

In the 1980s, economists Daniel Kahneman (who would later win the Nobel Prize) and Amos Tversky showed how irrational humans can be when confronted with risk. They presented test subjects with a scenario in which they had to make a difficult choice:

> Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed.

The two programs are very different—one is conservative, with a high probability of saving a small number of people, and one is risky, with a small probability of saving a large number of people. The subject has to make a choice about whether to choose the conservative or the risky strategy. But there was a twist. Kahneman and Tversky presented the exact same choice, but with slightly different wordings, to two separate groups of subjects. For the first group of subjects, the wording emphasized saving people from the disease; for the second, the phrasing dwelled on the victims of the disease rather than the survivors.

These differences in wording were purely cosmetic. Mathematically speaking, the two scenarios were identical. If people were logical creatures, the first group of subjects should make the exact same choices as the second group. Yet Kahneman and Tversky found that the wording of the programs made a tremendous difference—it was the phrasing, not the mathematics, that determined how people would behave. When the phrasing emphasized survivors over vic-

tims, subjects voted overwhelmingly—72 percent to 28 percent—to eschew the risky course of action, instead taking the conservative course that saves some patients with high certainty but lets others die. But when the wording spoke of victims rather than survivors, subjects suddenly became less risk-averse; they voted even more overwhelmingly—22 percent to 78 percent—to make a desperate gamble to save lives. The risks were exactly the same in both cases, but people, with their bad sense of risk, couldn't figure that out. The test subjects made their decisions based not on logic but upon how an authority presented the risks to them.

Minor changes in wording can easily make a huge risk seem worth taking or an insignificant risk seem dangerous. As a result, we are vulnerable to manipulation. We can't easily detect when someone is understating or exaggerating risks. We are prey to *risk mismanagement.*

Perhaps more than any other form of proofiness, risk mismanagement means big business. In NASA's case, it meant billions of dollars. By understating the risks of shuttle flights, they got funding from Congress to pursue a disastrous program—one that would kill fourteen astronauts for very little palpable gain. And other entrepreneurs are hoping to follow in NASA's footsteps. However, spaceflight is just a tiny part of the picture. Much of our economy revolves around risk—risk and money go hand in hand. There are entire industries devoted to managing and assessing risk; corporations have figured out how to make incredible amounts of money by measuring risk, packaging it, dividing it, and moving it around. And where there's money to be made, there's risk mismanagement. If you gaze deeply into the center of the ongoing financial crisis, you see risk mismanagement staring back at you. Risk mismanagement is crippling our economy, and along the way making a small handful of malefactors very, very wealthy.

NASA isn't the only game in town when it comes to getting people into space on the back of risk mismanagement. Airline magnate Richard Branson is hard at work trying to snooker private investors just as NASA snookered Congress. Branson is currently running a private spaceflight enterprise, "Virgin Galactic," which within its first five years of operation is supposedly going to launch an estimated three thousand passengers into space. Safely. "Virgin has a detailed understanding of what it takes to manage and operate complex transportation organizations . . . such as Virgin Atlantic Airlines and Virgin Trains which carry millions of passengers each year and have enjoyed superb safety records," brags the Virgin Galactic website. If you believe Virgin, spaceflight will be no riskier than a little jaunt on a private jet.

Hogwash. By comparing spaceflight to plane and train travel, Virgin is effectively underestimating the huge risks you take when you strap yourself to a rocket. It's a very dangerous task to pack enough energy into a cylinder to get you into space—and it's equally dangerous when, falling through the atmosphere, you get that energy back and have to dissipate it away in the form of heat. Throughout the history of spaceflight, about one in a hundred human-carrying rockets has killed its passengers, and that risk seems unlikely to change in the near future.

One chance in a hundred might not seem like so much, especially for the rare privilege of becoming an astronaut. But as far as risks go, it's extraordinarily high. For comparison, if today's U.S.

passenger aircraft had a similar failure rate, there would be roughly 275 U.S. plane crashes and 20,000 fatalities every day. A one in a hundred chance of dying every time you set foot on a plane would doom the airline industry; a 1 percent chance of death is simply too risky for any form of transportation to be commercially viable. If the historical failure rate holds, at Virgin Galactic's projected launch rate of one flight per week there would be only a one in three chance that Virgin Galactic goes for two years without a *Challenger*-type disaster. All in all, their chance of getting all three thousand people into space and back home again safely in this (hypothetical) scenario would be about half of 1 percent. People would almost certainly die, sooner rather than later. Even if the company survived the inevitable investigation and embarrassment, it would be just a matter of months before another explosion.

In my opinion, Branson is downplaying the risks, which has helped him convince more than 250 astronaut wannabes to put down $30 million worth of deposits on rides into space. He's also sold politicians and the public on his vision. In 2005, New Mexico politicians started spending tens of millions of dollars—and the governor promised to raise as much as $225 million—to build a spaceport. Two New Mexico counties even passed a sales tax to fund the project. As a smart businessman like Branson probably knows, downplaying risks can be very lucrative. In fact, there are two main ways to mismanage risk for fun and profit. Like NASA or Richard Branson, you can underestimate risks, making something look safer than it actually is. Conversely, you can take something that's mundane and exaggerate its risks, making it loom large in the public's imagination.

Journalists are particularly fond of the latter course. The scarier the story, the bigger the audience. Nothing sells like Armageddon.

Around the turn of the millennium, asteroid scares were all the rage. In 1998, the discovery of a large asteroid gave reporters the opportunity to grab readers with headlines like "October 26, 2028 could be our last day." In 2002, another asteroid gave rise to similar worries: "The world ends on Feb 1 2019 (possibly)." Journalists seemed undeterred when astronomers repeatedly told them that the risk of an actual collision was low. More observations, they said, would be able to pinpoint the asteroids' orbits with greater precision and a smack-up would almost certainly be ruled out—as indeed it was. But the attention-grabbing stories were too good to pass up. In fact, every doomsday scenario, no matter how far-fetched, is guaranteed to get at least some level of attention in the media. Every time physicists start up a new high-energy atom smasher—the Tevatron in 1985, the Relativistic Heavy Ion Collider in 2000, the Large Hadron Collider in 2008—the press chatters on about bizarre claims that the machine will destroy the earth or even the universe. The airwaves are alive with theories that the new machine will create tiny black holes that will swallow the earth, will create a variety of particle known as a "strangelet" that will destroy the planet, or will change the structure of space and time in a way that might annihilate the entire universe. (While no scientist will say that these scenarios are *impossible*—after all, scientific knowledge is uncertain and tentative—every mainstream scientist agrees that they are all very, very improbable.) Yet every time one of these new machines turns on, there are always headlines such as "Is the end nigh? Science experiment could swallow Earth, critics say" and "Physicists fear Big Bang machine could destroy Earth." Nobody seems to care about the science; everyone is fascinated with the prospect of Armageddon, no matter how remote it might be.

Fear sells so well that news organizations occasionally cross

ethical lines to make something appear more risky than it actually is. On November 23, 1986, *60 Minutes* correspondent Ed Bradley reported on allegations that a particular model of car—the Audi 5000—was prone to sudden and unexpected jolts of acceleration even while the driver was pressing on the brake. It was a moving report; a tearful mother told of her horror as her Audi lurched out of control and ran over her son. It was also startling. At one point, the program showed a terrifying sequence that begins with an Audi idling at a dead stop. All of a sudden, the accelerator depresses by itself, and the car is suddenly zooming without a driver. Audi sales plummeted—from about 74,000 cars in 1985 down to a nadir of 12,000. After all, if Ed Bradley was correct, the automobiles were clearly death traps.

It was an illusion, and a relatively implausible one at that. The engine in the Audi 5000 was not strong enough to overpower the brakes—if the driver was pushing hard on the brake, the car would have stayed put even if the engine was going full throttle. (And in all of the accidents in question, the Audi's brakes were in working order.) When the National Highway Traffic Safety Administration studied the issue, they found that the unintended accelerations were caused by drivers pressing the accelerator pedal instead of the brake. In each case, the panicking driver, trying to stop the car, stomped on the accelerator harder and harder, making the Audi zoom out of control. (Which explains why Audi accident investigators occasionally found accelerator pedals that had been bent by the pressure of the driver's foot.) A Canadian investigation found the same thing. The Audis weren't defective; there was nothing mechanically wrong with them. But *60 Minutes* had the unexpected acceleration on tape! What about the eerie footage of the accelerator pedal pressing down on its own and the Audi suddenly zooming out of control? This is

where standard *60 Minutes* fearmongering turns, in my opinion, into bona fide journalistic misconduct. The demonstration seems to have been rigged. A safety "expert" had apparently bored a hole in the Audi's transmission and forced high-pressure air in—it was this high pressure that caused the accelerator to depress "by itself." Bradley never hinted that the Audi had been doctored, and, to all appearances, *60 Minutes* had successfully tricked its audience into thinking that a risk existed when it was just a fiction.*

Exaggerating risks—and the fear these risks cause—is a potent tool, and not just for journalists. Politicians have long known that the public tends to support its leaders in times of crisis, so there's an incentive to make threats seem bigger than they are. According to some critics, the Homeland Security Advisory System—the color-coded five-level terror alert system created shortly after the September 11 attacks—is nothing more than an institutionalized way for the executive branch to manipulate the perceived risk of a terrorist attack. Though it's hard to divine the motivations of the Bush administration, there are some disturbing hints that the alert system was manipulated for political ends. For one thing, the alert level has *never* been relaxed below stage three ("yellow") since the September 11 attacks, indicating an "elevated" risk of terrorist attack. "Elevated" means "higher than usual," so it's completely nonsensical to declare that the everyday state of affairs exposes us to a higher-

* Unfortunately, this sort of hanky-panky isn't unique to *60 Minutes*. In 1992, *Dateline NBC* ran a spectacular segment about the dangers of GM trucks, showing one bursting into flames when struck in the side. There was in fact a design flaw, but *Dateline* had neglected to mention that to ignite the blaze, the TV show had strapped rockets to the truck's gas tank. Unlike *60 Minutes*, *Dateline* was forced to apologize for its actions. Without a hint of irony, Don Hewitt, *60 Minutes*' executive producer, used the occasion to tout his show's superior ethical standards: "If that had happened at *60 Minutes*," he told the *New York Times*, "I'd be looking for a job tomorrow."

than-usual risk of terrorist attack. It's impossible to have an eternally elevated risk, just as it's impossible for all the children in Lake Wobegon to be above average. Another hint comes from indications that political considerations occasionally played a role in raising and lowering the alert level. For example, in 2003, in the lead-up to the invasion of Iraq, the administration admitted relaxing the alert from stage two ("orange") to stage three ("yellow") not because of any change in the level of threat, but so they could elevate it again a few days before the war began. And finally, there's a bit of evidence that is very disturbing even though it should be taken with a big grain of salt. A 2004 study seemed to indicate that every time the government issued a terror warning, the president's approval rate spiked higher. The benefits of scaring the public are all too apparent.*

Nevertheless, if you want to make oodles of money, underestimating risks, not exaggerating them, is where the money is. Richard Branson's millions and NASA's billions are just a drop in the bucket. Savvy businesspeople have figured out how to make themselves very wealthy through risk mismanagement, and their misdeeds are so great that they put the world economy into enormous danger. An unprecedented international economic crisis was born in 2007 because of risk mismanagement.

In the world of finance, risk and reward are flip sides of the same coin. A safe investment—one with extremely little risk—will earn

* The United States is not alone in its mismanagement of terrorist risks. In mid-2009, the United Kingdom relaxed its five-point terrorist threat level from 2 ("Severe") to 3 ("Substantial"). It was the first time since the system was made public in 2006 that the alert level has been that low.

an investor very little money. To make big profits, you have to be willing to bear substantial risks; you have to accept that your gains might never materialize and that you might even lose the money you invested. Successful investors aren't people whose bets always pay off; Warren Buffett regularly loses money when some of the risks he takes don't pan out. Successful investors are the people who have a knack for squeezing the most reward out of the risks they're willing to take. They're the people who maximize the upside, the return on an investment, while keeping the downside, the risks they bear, to a minimum. Conversely, if you want people to give you money—to invest in you—you have to give them a large return with a minimum of risk. The less risk you seem to represent, the more money people are going to give you. This is why risk mismanagement is such big business. If you can understate the risks in your business, if you can hide them or shuffle them around or pawn them off on someone else, you can make billions.

In the late 1990s, the energy firm Enron was rolling in money. Its officers were millionaires many times over; at one point founder Ken Lay was worth $400 million. Enron was producing enormous profits for its investors, and quickly became a Wall Street darling. In 1995, *Fortune* magazine named it the most innovative company in America. In 1996, it won the title again—and in 1997, 1998, 1999, and 2000. But it was all a façade of risk mismanagement. Lay and his colleagues were making their money, in part, by moving risk from place to place in an attempt to shield it from outside scrutiny. They had created a whole slew of shell corporations (many of which, such as Chewco, Obi-1 Holdings, and Kenobe, had *Star Wars*–inspired names) to hide how risky an investment Enron really was. These corporations assumed much of Enron's debt, taking it off the books. As a result, Enron looked squeaky clean when it was in fact saddled

with billions in debt—and was joined at the hip to dozens of dubious corporations that were on the verge of bankruptcy. When the risk mismanagement came to light, investors fled and the whole house of cards collapsed, taking along with it the savings of many investors who had no clue that Enron was a risky investment.*

If you move money around in clever enough ways, you can camouflage risks and make an absolute mint. Bernard Madoff was worth more than $800 million at his peak—by shuttling money from client to client, he managed to hide just how risky his investment firm really was (and that he was making himself rich by stealing his clients' money). Through his manipulations, Madoff managed to deflect attention for more than a decade—until the markets collapsed, suddenly exposing the fact that his coffers were empty. In making himself rich, he managed to lose more than $50 billion of his clients' money, earning him 150 years in prison.

The Ken Lays and the Bernie Madoffs of the world are able to make their money by hiding risk and moving it from place to place, deceiving their investors. When their criminal enterprises collapse,

* There are supposed to be checks to make sure that this kind of monkey business doesn't happen. Enron was audited by an outside firm, Arthur Andersen, that was supposed to ensure its basic accounting honesty. But such accounting firms have a tremendous conflict of interest; after all, they're paid by the client whose books they're inspecting. As a result, these accounting watchdogs either turn a blind eye to risk mismanagement or face the possibility of losing an important client. More than a decade before the Enron affair, a similar pattern emerged in the savings and loan scandal, which cost taxpayers billions. The example of Silverado Banking, a Texas savings and loan, is typical. As one article put it, "In 1985 Silverado's auditors, Ernst & Whinney, forced the thrift to report $20 million in losses because of problem loans. Silverado's managers weren't pleased with that result, so they got rid of Ernst & Whinney and hired Coopers & Lybrand, which took a more flexible view of the books. In 1986, Silverado reported $15 million in profits and the managers got $2.7 million in bonuses. A year and a half later the enterprise collapsed, at a cost to the government of some $1 billion."

the malefactors rightly get a lot of press attention as they are pilloried and prosecuted. However, as big as these frauds might seem, they're nothing compared to the risk mismanagement that's gnawing away at the heart of our economy. Madoff and Lay are rank amateurs in the proofiness game.

There are more subtle ways to make money off of risk. Just as risk can be hidden or moved from place to place, it can be divvied up and sold. Insurance companies are nothing more than a dumping ground for risk; for an appropriate fee, you can have them assume a risk for you. Afraid of the risk of a house fire? Of hitting someone with an automobile? Of a lawsuit? For enough money, you can convince an insurance company to indemnify you. It's a lucrative business—if you understand your risks.

As an example, imagine that you want to sell fire insurance. You need to have a good estimate for the risk of a house fire (say that there's about a 1 in 250 chance that a typical American home experiences a fire in a year) and how much damage the average house fire causes (say $20,000). These two pieces of information tell you how much money you'll be paying out in claims, and thus how much you need to charge for the insurance. Indeed, given these numbers, you can make a very nice living by charging your customers $100 a year in premiums.

Because the risk of a fire is 1 in 250, for every 10,000 clients you have, you can expect roughly forty of them to have a fire in a given year. Each of those forty fires costs you roughly $20,000—a grand total of $800,000 in damage that your company has to cover. But at the same time, your clients are paying you $1,000,000 in premiums, leaving you with a handsome profit of roughly $200,000. All you do is sit in the middle and watch the money pour in, making the occasional payoff. To make money in this way, though, you really have

to know your risks. If your risk estimates are off, if there are sixty fires in a year instead of forty, you'll have to pay off $1,200,000 in claims—$200,000 more than you collect in premiums, which means that you'll go bankrupt. So you do all you can to keep the risk of fire in your client base—your "risk pool"—as low as possible. You might encourage them to install sprinkler systems. Or you might insure only nonsmokers, as they are at less risk of setting a fire than the general population. Your survival as an insurer depends on understanding your clients' risks and reducing them.

That's the theory anyhow. Reality, though, is a lot more complicated. Say you don't want to be bothered with all the paperwork and hassle of handling insurance claims and making payoffs to clients. If so, you can pass the risk on to a bigger company, just as your clients passed the risk to you. You could take the bundle of 10,000 insurance contracts and sell them to another firm for, say, $100,000. You make less money than you'd get if you handled claims yourself, but you no longer have to worry about paying out claims. It's pure profit for you. And the big firm is happy too, because it also makes money on the deal—$100,000. That is, it makes money if your risk estimate is correct. If your risk estimate is off and, say, sixty clients burn their homes down, the firm loses a lot of money on the deal ($300,000 to be precise). As for you, the extra fires don't bother you at all! You no longer have a stake in how many clients make a claim on the insurance contracts you sold them. Regardless of what happens to their homes, you still make a profit of $100,000. Even if there's an epidemic of house fires, you can laugh all the way to the bank; you sold a firm bad risk for a nice little profit.

You can probably see where this is going. The moment you sold your insurance contracts to another firm, you stopped caring whether those contracts were risky or not. Your income depends

only on the number of clients you sign up, not on whether or not they burn down their houses and make claims. You no longer have any incentive to maintain a safe risk pool; just the opposite, in fact. If you sell only to nonsmokers, you have fewer clients to choose from, and fewer contracts to sell to the big insurance firm. Instead, you should sell to smokers as well as nonsmokers, the better to drive the number of contracts up. You'll make more money that way. Heck, you can make a profit by selling insurance to fire-eaters, people who try to deep-fry turkeys every Thanksgiving, and serial arsonists. So long as you dump the contracts to an unsuspecting firm before your clients immolate their houses, you'll make money. For the scheme to work, though, you mustn't let the big insurance firm know that you're contaminating the risk pool—you have to pretend that the bundle of contracts is low-risk, when in fact it's absolutely crawling with firebugs. If you do it right, money flows from the large insurance firm into your pocket, all thanks to your clever risk mismanagement.

Instead of "fire insurance," use "mortgages," and all of a sudden you've got a (slightly oversimplified) explanation of the subprime mortgage meltdown in 2007. The risk mismanagement involved selling mortgages to people whose income wasn't sufficient to support one. There are reports of brokers encouraging people to lie on their mortgage application (a federal crime) to allow them to get a nice fat mortgage; in some cases, people didn't even need to prove that they had an income at all. The recipients of these loans were just treading water financially, and were highly likely to default on their loans. This made these so-called subprime mortgages extremely high-risk, but if you were able to dump them on somebody else's balance sheet, you could make a killing.

While times were good and the housing market was rising,

everybody who traded in mortgages and other monetary instruments based upon those mortgages (such as "credit default swaps," which were, functionally, unregulated insurance contracts on these loans) made a fortune. Brokers, banks, and insurance companies passed around the bad risks like a hot potato, creating a complicated web of risk and debt. Every time the mortgage potato circled around, CEOs, corporate officers, managers, traders, and brokers earned themselves millions of dollars in bonuses. By pretending that these high-risk loans were actually low-risk, everybody was making themselves extraordinarily wealthy. But when the economy softened and the housing market began to fall, the good times suddenly came to a screeching halt. People began to default on their mortgages in droves; it's as if an arson craze hit insurance clients all at the same time. There was a giant sucking sound as the bad risk began gulping money away from the firms and other investors who held large amounts of these "toxic assets." Citigroup. Merrill Lynch. Bear Stearns. Lehman Brothers. Morgan Stanley. J. P. Morgan. Freddie Mac. Fannie Mae. AIG. It's hard to know just how much money evaporated—how much of the net worth of these companies was based upon understating the risk of these mortgages—but the damage is in the range of hundreds of billions or even trillions of dollars. This is a mind-boggling sum, not far off from the annual budget of the entire U.S. government. The result was a global economic catastrophe of unrivaled proportions.

During the good times, employees at AIG, Citigroup, and other companies made themselves very, very rich through risk mismanagement. They deliberately underestimated—indeed, ignored—the huge risks of these mortgage-based investments, trading them around and siphoning off money from one another's companies. And when the mortgages blew up, they got to keep their cash. The

companies they worked for, though, were saddled with crushing debt, and looked ready to collapse. The government had to step in to save them, at the cost of trillions of taxpayer dollars.

Yet much of that bailout money would wind up in the pockets of those who caused the crisis in the first place. It was almost inevitable—it was an outcome determined by the way humans deal with risk.

The rules that govern the behavior of officers at AIG, Citigroup, and other malefactors in the worldwide economic crisis are the same ones that are causing us to chop down the rain forests, to fill our atmosphere with carbon dioxide, and to overfish the oceans. This phenomenon, known as the *tragedy of the commons*,* is everywhere—even at a friendly dinner.

Imagine that you're at a fancy restaurant with a bunch of acquaintances and you've all decided ahead of time to split the bill evenly among you. Even though the restaurant is pretty expensive, the prix fixe menu isn't too pricy. It gives you a tasty if modest meal, say, a chicken breast, along with a small side salad. It's a wee bit spartan, but it's a very good deal. The other temptations on the menu, such as the lobster, are much more appealing, but their prices are exorbitant. If everybody sticks to the prix fixe menu, the bill will be relatively small. When you divide up the bill, you and your acquaintances will all get a very good deal—a pleasant meal at a cheap price. Alas, that's probably not what happens. More typ-

* So named because the original example concerned the English custom of having a community grassland (the "commons") that livestock owners could use for grazing. As the story goes, in an attempt to capitalize on the free grazing land, people bought more and more livestock and overgrazed the pastures, rendering them useless.

ically, someone at the party has to be a spoilsport and order the lobster. Lobster-boy thinks he's pulling a fast one; since everybody divides the bill equally at the end of the meal, the enormous cost of his lobster is shared among the whole crowd. He personally gets the benefit of eating the expensive entrée, but in the process he drives up the cost of the meal for everybody. Sadly, once lobster-boy puts in his order, all bets are off. Everybody starts ordering the overpriced lobster; even a few people who previously ordered the prix fixe special will change their orders. Perhaps a handful of people show restraint and stick to the cheaper chicken, but it's of no avail. All the lobsters drive the price of the meal up to enormous heights. When the check comes and it's divided up, the cost is eye-popping. This is an example of the tragedy of the commons.

The tragedy of the commons occurs when an individual can take an action that benefits him, yet the negative consequences of that action are diffused—such as when they're divided among a large group of people or when they take a long time to materialize. In situations like these, people act selfishly, getting as much benefit as they can, but as a consequence, we're all worse off. In the restaurant example, lobster orderers benefit from getting the expensive entrée, but the cost is divided among the entire party, making what should have been an affordable, pleasant meal into an outrageously expensive fiasco. Most environmental problems are the result of a tragedy of the commons. Burning fossil fuels, for example, gives the user benefit of cheap, affordable energy, even though each power plant helps warm the planet a little bit—and the consequences of global warming are shared by everybody in the world.

Another example of the tragedy of the commons is all too familiar to urban planners. Since the 1960s, a number of cities around the world have tried to encourage cycling by making bicycles freely

available to anyone who wants to use them. It's a lovely idea; even if you don't use the bicycles yourself, you benefit from the reduction in automobile traffic. Alas, most of the time, these efforts wind up failing because people abuse the bicycles. In 1960s Amsterdam, most of the free bicycles were stolen within weeks. People are no better nowadays. In 2009, Paris's free-bicycle scheme was struggling to survive because people were abusing the bicycles. You can find videos on YouTube where young toughs ride the bicycles down stairs and jump them in skate parks. Vandals destroyed quite a few—smashing them, burning them, and tossing them into the Seine. And many others were simply stolen. A year and a half after the start of the Parisian free-bicycle program, more than half of the 15,000 bicycles were gone. The destructive behavior of individuals destroyed a shared resource. Tragedy of the commons is an immutable fact about society. If the benefits of our actions are divorced, to some extent, from their negative consequences, we're going to take those actions—even if they lead to a very unfortunate outcome for everybody. The tragedy of the commons is a result of human nature, and the moment we ignore human nature we put ourselves at grave risk.

In financial markets, the equivalent concept is known as *moral hazard*—it's what you get when you cross the tragedy of the commons with risk mismanagement. When financial gain is divorced from the risks involved in making that gain, very bad things happen. The subprime mortgage meltdown was a great example of moral hazard in action. The individuals who manufactured and traded those mortgage-backed securities were making oodles of money, but they were insulated from the risks inherent to those mortgages. As a result they misbehaved, misrepresenting risks on a tremendous scale to make money. Even though their actions little by little undermined the security of their companies—and the

stability of the world economy—the party wouldn't stop until the whole economy threatened to come crashing down.

When the risk finally caught up with the securities and the whole risk-mismanagement scheme collapsed, these individuals tended to hold on to their wealth. Some of them even managed to make lots of cash even after the meltdown, thanks to another incarnation of moral hazard.

AIG, a massive insurance company, was at the heart of the meltdown, with tens of billions of dollars in bad mortgage risk. By the time the economic crisis came about, AIG had swollen to gigantic proportions. It had a trillion dollars in assets, and a finger in almost every big financial institution in the nation. If it collapsed, it would send tremors through all the banks and the investment firms in the United States; it might even cause some of them to collapse as well. People would panic, and potentially our entire economic infrastructure could be shaken to rubble in a matter of days. Economists call this potential *systemic risk*, and it was the ace in AIG's hand—by using that risk, AIG could ensure that the government would keep shoveling money into its coffers.

Because the government couldn't let AIG collapse for fear of destroying the economy, company officials realized that they had carte blanche to go wild and take whatever risks they wanted. AIG officials were largely insulated from the consequences of their actions, because they suspected that their company wouldn't be allowed to fail no matter what risks they took. This gave AIG's managers the green light for risk mismanagement on an enormous scale. They could deceive their investors about the risks that they were taking, misuse their assets, siphon off money through bonuses, and gamble on risky investments in hopes of reaping a windfall. If things turned

sour, the government would have to step in and bail them out. Congress had no choice but to subsidize AIG's risk-taking.

Things did, of course, turn sour with AIG's risky investments, and the government duly stepped in. On September 16, 2008, the feds bailed it out with an $85 billion loan. "Nothing made me more frustrated, more angry than having to intervene, particularly in a couple of cases where taking wild bets had forced these companies close to bankruptcy," Federal Reserve chairman Ben Bernanke admitted in mid-2009. But because of the fear of triggering an economic disaster, "I had to hold my nose and stop those firms from failing."

Unfortunately, the bailout didn't stop the misbehavior. After all, AIG was still too big to fail—so long as the government was afraid of systemic risk, AIG officials had no reason to act honorably and try to put the company on a sound financial footing. Within days of the bailout, AIG executives were spending nearly half a million dollars of the bailout money on a retreat to a posh resort in California, where they indulged themselves with spa treatments, lavish banquets, and of course, plenty of golf—$7,000 in greens fees. Lawmakers were livid. But that anger wasn't enough to stop them from giving the company an additional $38 billion in October and then another $40 billion in November. Edward Liddy, the CEO of AIG, promised to behave. "We are tightening our belt," he told a reporter in October. "Just as the American consumer, the American taxpayer is tightening their belt, we are doing the same thing. But we're not stopping at one notch; we're going three and four and five notches." Apparently, AIG belts are constructed somewhat differently, because within a few months, Liddy was explaining to Congress why he felt it necessary to use the bailout money to pay

AIG employees $160 million in "retention" bonuses. The top bonus was a whopping $6.4 million. Seventy-three people got more than $1 million each—and eleven of them had already left the company, so the "retention" bonus could hardly be expected to work as advertised. It was unbelievable misbehavior, embarrassing Congress and the new president.

All companies that are too big to fail—Citigroup, General Motors, Fannie Mae, Merrill Lynch—are swimming in moral hazard. Once they know that they won't be allowed to collapse, it's almost guaranteed that they will fill their own pockets while passing the consequences of their risky behavior on to the taxpayers.*

Every few years, tremendous cases of risk mismanagement extract money from average citizens and put it into the hands of the wealthy. Whether it's the savings and loan scandal or Enron or the subprime mortgage crisis, the end is always the same. The people who are willing to lie about risks make themselves very rich, and the taxpayer suffers the consequences. Even if one or two of the malefactors wind up in jail, there are always many more who made themselves much better off at others' expense and never suffered any serious consequences.

Risk management is the form of proofiness that hits the pocketbook most directly. However, other forms of proofiness can have consequences that are even more grave. They can undermine the press, deny us our vote, put us in jail, and sap the strength of our democracy. Mere thievery pales by comparison.

* As I was finishing up this manuscript, the hemorrhage showed no signs of ending. The newspapers were expressing outrage about how Citigroup had announced that it would give its twenty-five senior executives enormous bonuses, including one worth a mind-blowing $98 million, making AIG look tightfisted.