

DESPERATELY SEEKING SIGNAL

Just as the residents of L'Aquila, Italy, were preparing for bed on a chilly Sunday evening in April 2009, they felt a pair of tremors, each barely more perceptible than the rumbling of a distant freight train. The first earthquake, which occurred just before 11 p.m. local time, measured 3.9 on the magnitude scale,* a frequency strong enough to rattle nerves and loosen objects but little else. The second was even weaker, a magnitude 3.5; it would not have been powerful enough to awaken a sound sleeper.

But L'Aquila was on edge about earthquakes. The town, which sits in the foothills of the Apennine Mountains and is known for its ski resorts and medieval walls, had been experiencing an unusually large number of them—the two that Sunday were the seventh and eighth of at least magnitude 3 in the span

* News accounts often refer to the Richter scale, named after the Caltech seismologist Charles Richter. In fact, a different and more accurate scale—the moment magnitude scale, developed at Caltech in the late 1970s—is in more common use among seismologists today. The moment magnitude scale is designed to be comparable to the Richter scale—both scales are logarithmic, and a magnitude 8.0 earthquake is a very serious one under either definition. But they are not calculated in quite the same way. The earthquake magnitudes described in this chapter generally refer to the moment magnitude scale.

DESPERATELY SEEKING SIGNAL

of about a week. Small earthquakes are not uncommon in this part of the world, but the rate is normally much less—about one such earthquake every two or three months. These were coming almost one hundred times as often.

Meanwhile, the citizens of a town a mountain pass away, Sulmona, had just survived an earthquake scare of their own. A technician named Giampaolo Giuliani, who worked at Italy's National Institute of Nuclear Physics, claimed to have detected unusually high levels of radon in the area. He theorized this might be a precursor to an earthquake and went so far as to tell Sulmona's mayor that an earthquake would strike the town on the afternoon of March 29. The mayor, impressed by the prediction, ordered vans carrying loudspeakers to drive about town, warning residents of the threat.¹

No earthquake hit Sulmona that day. After the prediction failed, Giuliani was reported to local authorities for *procurato allarme* (bringing about alarm)—in essence, having yelled fire in a crowded theater. He was forced to remove his predictions from the Internet for fear of triggering further panic.

Authorities in L'Aquila told the residents the earthquake swarm* was nothing to worry about; the fault was helpfully discharging energy, explained Bernardo De Bernardinis, the deputy chief of Italy's Civil Protection Department,² reducing the threat of a major earthquake. He agreed with a reporter that they should sit back and enjoy a glass of wine;³ De Bernardinis recommended a local specialty, a Montepulciano.

A major earthquake did hit L'Aquila, however. Measuring at magnitude 6.3, it came at 3:32 a.m. local time on Monday morning. Shaking houses from their foundations, caving in roofs, and turning furniture into projectiles, the quake killed more than 300 residents, left another 65,000 homeless, and caused more than \$16 billion in damage.⁴

What We Do When Our Foundations Are Shaken

L'Aquila ought to have been better prepared. The city sits near a particularly violent type of fault known as a subduction zone, where the African Plate, one of

* Seismologists use the term "earthquake swarm" to refer to a series of small earthquakes.

the eight major tectonic plates that cover the earth's surface, slips slowly and inexorably beneath the Eurasian one. Its first significant earthquake was recorded in 1315, and earthquakes struck again in 1349, 1452, 1461, 1501, 1646, 1703, and 1706;⁵ the most serious one, in 1786, had killed more than 5,000 people. Each time, often on direct order of the pope,⁶ the town was rebuilt and repopulated.

Since then, L'Aquila had tempted fate for more than two centuries. An earthquake hit in 1958, but it was fairly minor—magnitude 5.0⁷—and only the town's oldest residents would have remembered it. The 2009 earthquake was much more powerful. The magnitude scale is logarithmic; a one-point increase in the scale indicates that the energy release has multiplied by thirty-two. Thus, the 2009 earthquake, magnitude 6.3, was about seventy-five times more powerful than the one that had hit L'Aquila in 1958. And it was about 3,000 times more powerful than the tremors—foreshocks to the major earthquake—that L'Aquila had experienced earlier that evening.

Still, while the 2009 earthquake was large by Italian standards, it was barely a hiccup on the global scale. The earthquake that devastated Japan in 2011 measured at magnitude 9.0 or 9.1—almost 11,000 times more powerful. And the largest earthquake recorded since reliable estimates were possible, which hit Chile in 1960 and measured magnitude 9.5, was about 60,000 times stronger than the L'Aquila quake.

Why, then, did L'Aquila—a fairly well-to-do town in a wealthy, industrialized nation—sustain such significant damage? One reason was the city's geology—L'Aquila sits on an ancient lake bed, which tends to amplify the earth's shaking. Mexico City was also built on an ancient lake bed,⁸ and 10,000 were killed there in 1985 from an earthquake whose epicenter was more than two hundred miles away.

But the major reason was simply that the town had become complacent about the seismic danger that lay just fifteen kilometers underground. There was nothing resembling the proper level of earthquake readiness:⁹ building codes, emergency supplies, community drills. Not only were centuries-old buildings leveled by the tremor, but so too were many modern ones, including a wing of a hospital that had been erected as recently as 2000. A little bit of warning would have saved untold lives there.

Had Giampaolo Giuliani provided that warning? In the Italian tabloids, he

had become something of a savant and a martyr. Soft-spoken and disheveled, and often wearing the colors of the local soccer team, he played the role of the humble civil servant or absentminded professor whose insights had been ignored by the scientific establishment. He claimed that he had warned friends and family about the L'Aquila quake and was prevented from telling others only because of the police order against him. He demanded an apology from the authorities—not to him, he said, but to the people of L'Aquila.

Never mind that Giuliani had not actually predicted the earthquake. His prediction had been very specific: Sulmona, not L'Aquila, was at greater risk, and the earthquake would come in March rather than April. In fact, he had suggested to a local newspaper that the danger had passed. "To simplify the concepts," he said before launching into a rambling explanation about the lunar cycle, "the Earth-Moon system has come to visit at perihelion . . . the minimum distance from Earth, and aligned with the planet Venus. . . . I feel I can reassure my fellow citizens because the swarm will be diminishing with the end of March."¹⁰

Perihelion with the planet Venus? Radon gas? What did any of this have to do with earthquakes? And what about Giuliani's failed prediction in Sulmona? It didn't matter. When catastrophe strikes, we look for a signal in the noise—anything that might explain the chaos that we see all around us and bring order to the world again. Giuliani's rambling explanations were the closest thing available.

No type of catastrophe is more jarring to our sense of order than an earthquake. They quite literally shake our foundations. Whereas hurricanes descend upon us from the heavens and have sometimes been associated with metaphors for God's providence,¹¹ earthquakes come from deep underneath the surface and are more often taken to be signs of His wrath,¹² indifference,¹³ or nonexistence. (The Lisbon earthquake of 1755 was a major spark for the development of secular philosophy.¹⁴) And whereas hurricanes—along with floods, tornadoes, and volcanoes—can often be forecasted in advance, earthquakes have defied centuries of efforts to predict them.

¹⁰ The Japanese word *kamikaze* originally meant "divine wind," referring to typhoons in 1274 and 1281 that had helped to disperse a Mongol invasion.

Magic Toads and the Search for the Holy Grail

Pasadena, California, has long been the world's epicenter for earthquake research. It is home to the California Institute of Technology, where Charles Richter developed his famous logarithmic scale in 1935. The United States Geological Survey (USGS) also has a field office there, where most of its earthquake specialists reside. I traveled there in September 2009 to meet with Dr. Susan Hough, who is one of the USGS's top seismologists and who has written several books about earthquake prediction. She had watched Giuliani's television interviews with suspicion and had written a blistering editorial in the *New York Times*¹⁴ that criticized both Giuliani and the attention paid to him.

Hough's editorial argued that Giuliani's success was merely coincidental. "The public heard about Mr. Giuliani's prediction because it appears to have been borne out," she wrote. "But there are scores of other [incorrect] predictions that the public never hears about."

If you have hundreds of people trying to make forecasts, and there are hundreds of earthquakes per year, inevitably someone is going to get one right. Giuliani's theories about radon gas and lunar cycles had been investigated many times over¹⁵ by credentialed seismologists and had shown little or no ability to predict earthquakes. Giuliani had been lucky: the monkey who typed Shakespeare; the octopus who predicted the World Cup.

Hough's office at the USGS sits near a quiet corner of the Caltech campus where there are more eucalyptus trees than students. She seemed a little road weary when I met her, having just returned from a trip to Turkey where she'd been to study a system of earthquake faults. She has soft features and frizzy hair and her eyes are dark, tired—skeptical. "What's your day job?" she quizzed me a few moments after I greeted her.

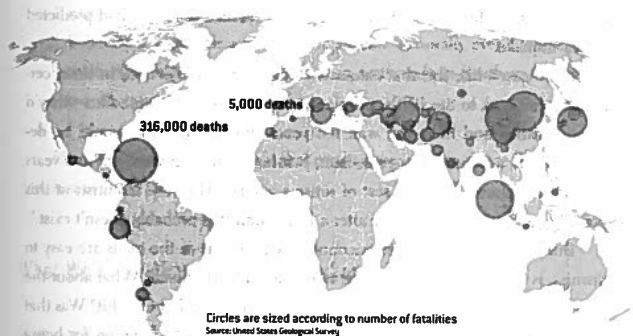
At one point, she pulled a pocket-size globe off her desk, the sort that looks like it was bought at an airport gift shop. She took her index finger and drew a line across the surface of the globe, starting in the Sea of Japan and moving east-southeast.

"They are really concentrated in this belt—stretching from southern China through Greece," Hough explained, referring to the world's most destructive

earthquakes. "It's a complicated earthquake zone, a lot of buildings with vulnerable construction. If you put a big earthquake under Tehran, you could kill a million people."

Indeed, almost all the deadliest earthquakes in modern history (figure 5-1) have occurred along the path that Hough outlined, one which passes through the Cradle of Civilization in the Middle East and through some of the most densely populated regions of the planet, including China and India. Often poor and crowded, these areas lack the luxury to prepare for a once-per-three-hundred-year catastrophe. But the death tolls can be catastrophic when earthquakes hit, stretching into the hundreds of thousands.*

FIGURE 5-1: DEADLIEST EARTHQUAKES SINCE 1900



Earthquakes kill more people than hurricanes, in fact,¹⁶ despite seeming like the rarer phenomenon.¹⁷ Perhaps that is because they are so seldom predicted successfully. Whereas the landfall position of hurricanes can be forecasted at least three times more accurately now than they were even twenty-five years ago, the science of earthquake forecasting seems barely to have evolved since the ninth century A.D., when the Japanese first claimed to be able to an-

* The Haitian earthquake of 2010 was an exception to the pattern geographically, but not in how poverty and lax building standards contribute to immense death and destruction.

ticipate earthquakes by looking at the behavior of catfish.¹⁸ (Cows, pigs, eels, rats, parakeets, seagulls, turtles, goldfish, and snakes have also been reported at various times to behave unusually in advance of an earthquake.)

Kooks like Giuliani are still taken seriously, and not just in the Italian tabloids.¹⁹ The California Earthquake Prediction Council receives hundreds of unsolicited earthquake forecasts per year, most of which, the agency says, “discuss the strange behavior of household pets, intuition, Aunt Agatha’s aching bunions, or other mysterious signs and portents that scientists simply don’t understand.”²⁰ Meanwhile, some of the stuff in academic journals is hard to distinguish from ancient Japanese folklore. A 2010 paper²¹ in a relatively prestigious journal, *The Journal of Zoology*, observed that toads in a pond fifty miles from L’Aquila had stopped spawning five days before the major earthquake there.²² Remarkably, it asserted that this was evidence that they had predicted the earthquake.

It’s research like this that exhausts Hough. “If you look back in time, certainly going back to the 1970s, people would come up with some idea—they’d be optimistic—and then you wait ten years and that method would be debunked,” she told me. “Ten years later, you have a new method and ten years later it’s debunked. You just sort of sense a theme. Most top scientists at this point know better than to chase after a Holy Grail that probably doesn’t exist.”

But while Giuliani’s close encounters with Venus or the toads are easy to dismiss, is there really no way at all to predict an earthquake? What about the swarm of smaller quakes around L’Aquila just before the Big One hit? Was that just a coincidence? The seismological community has a reputation for being very conservative. It was very slow to accept the theory of plate tectonics, for instance²³—the now broadly accepted notion that the shifting of the earth’s continental plates is the primary cause for earthquakes—not adopting it into their canon until the 1960s even though it was proposed in 1912. Had Hough’s skepticism crossed the line into cynicism?

The official position of the USGS is even more emphatic: earthquakes *can* not be predicted. “Neither the USGS nor Caltech nor any other scientists have ever predicted a major earthquake,” the organization’s Web site asserts.²⁴ “They do not know how, and they do not expect to know how any time in the foreseeable future.”

Earthquakes cannot be predicted? This is a book about prediction, not a book that *makes* predictions, but I’m willing to stick my neck out: I predict that there will be more earthquakes in Japan next year than in New Jersey. And I predict that at some point in the next one hundred years, a major earthquake will hit somewhere in California.²⁵

Both the USGS and I are playing some semantic games. The terms “prediction” and “forecast” are employed differently in different fields; in some cases, they are interchangeable, but other disciplines differentiate them. No field is more sensitive to the distinction than seismology. If you’re speaking with a seismologist:

1. A **prediction** is a definitive and specific statement about when and where an earthquake will strike: *a major earthquake will hit Kyoto, Japan, on June 28.*
2. Whereas a **forecast** is a probabilistic statement, usually over a longer time scale: *there is a 60 percent chance of an earthquake in Southern California over the next thirty years.*

The USGS’s official position is that earthquakes cannot be predicted. They *can*, however, be *forecasted*.

What We Know About How Earthquakes Behave

If you explore the USGS Web site, in fact, you’ll find that it makes lots of tools available to help you *forecast* earthquakes. One particularly neat one is an application that lets you type in the longitude and latitude at any point in the United States; it will estimate the long-term probability of an earthquake there.²⁶ In figure 5-2, I’ve listed the probabilities for earthquakes in a variety of major US cities as provided by the USGS Web site.

We all know that California is very seismically active; the USGS estimates that an earthquake of magnitude 6.8 or higher will hit San Francisco about once every thirty-five years. Many of you will also know that Alaska has many earthquakes—the second largest one in recorded history, magnitude 9.4, hit Anchorage in 1964.

THE SIGNAL AND THE NOISE

FIGURE 5-2. FREQUENCY OF A MAJOR (\geq MAGNITUDE 6.75) EARTHQUAKE WITHIN A 50-MILE RADIUS OF SELECT U.S. CITIES

Anchorage	1 per 30 years
San Francisco	1 per 35 years
Los Angeles	1 per 40 years
Seattle	1 per 150 years
Sacramento	1 per 180 years
San Diego	1 per 190 years
Salt Lake City	1 per 200 years
Portland, OR	1 per 500 years
Charleston, SC	1 per 600 years
Las Vegas	1 per 1,200 years
Memphis	1 per 2,500 years
Phoenix	1 per 7,500 years
New York	1 per 12,000 years
Boston	1 per 15,000 years
Philadelphia	1 per 17,000 years
St. Louis	1 per 23,000 years
Atlanta	1 per 30,000 years
Denver	1 per 40,000 years
Washington, DC	1 per 55,000 years
Chicago	1 per 75,000 years
Houston	1 per 100,000 years
Dallas	1 per 130,000 years
Miami	1 per 140,000 years

But did you know about Charleston, South Carolina? It is seismically active too; indeed, it experienced a magnitude 7.3 earthquake in 1886. The USGS estimates that there will be another big earthquake there about once per six hundred years. If you live in Seattle, you should probably have an earthquake plan ready; it is more earthquake-prone than many parts of California, the USGS says. But you don't need one if you live in Denver, which is a safe distance away from any continental boundaries.

DESPERATELY SEEKING SIGNAL

This seems like an awful lot of very specific and user-friendly information for an organization whose party line is that it is impossible to predict earthquakes. But the USGS's forecasts employ a widely accepted seismological tool called the Gutenberg–Richter law. The theory, developed by Charles Richter and his Caltech colleague Beno Gutenberg in 1944, is derived from empirical statistics about earthquakes. It posits that there is a relatively simple relationship between the magnitude of an earthquake and how often one occurs.

If you compare the frequencies of earthquakes with their magnitudes, you'll find that the number drops off exponentially as the magnitude increases. While there are very few catastrophic earthquakes, there are literally millions of smaller ones—about 1.3 million earthquakes measuring between magnitude 2.0 and magnitude 2.9 around the world every year.²⁷ Most of these earthquakes go undetected—certainly by human beings and often by seismometers.²⁸ However, almost all earthquakes of magnitude 4.5 or greater are recorded today, however remote their location. Figure 5-3a shows the exponential decline in their frequencies, based on actual records of earthquakes from January 1964²⁹ through March 2012.³⁰

It turns out that these earthquakes display a stunning regularity when you graph them in a slightly different way. In figure 5-3b, I've changed the vertical axis—which shows the frequency of earthquakes of different magnitudes—into a logarithmic scale.* Now the earthquakes form what is almost exactly a straight line on the graph. This pattern is characteristic of what is known as a power-law distribution, and it is the relationship that Richter and Gutenberg uncovered.


Something that obeys this distribution has a highly useful property: you can forecast the number of large-scale events from the number of small-scale ones, or vice versa. In the case of earthquakes, it turns out that for every increase of one point in magnitude, an earthquake becomes about ten times less frequent. So, for example, magnitude 6 earthquakes occur ten times more frequently than magnitude 7's, and one hundred times more often than magnitude 8's.

What's more, the Gutenberg–Richter law generally holds across regions of the globe as well as over the whole planet. Suppose, for instance, that we wanted

* Recall that the magnitude scale is already logarithmic, so this is what's technically known as a double-logarithmic plot.

THE SIGNAL AND THE NOISE

FIGURE 5-3A: WORLDWIDE EARTHQUAKE FREQUENCIES, JANUARY 1964–MARCH 2012




to make an earthquake forecast for Tehran, Iran. Fortunately, there hasn't been a catastrophic earthquake there since its seismicity began to be measured. But there have been a number of medium-size ones; between 1960 and 2009, there were about fifteen earthquakes that measured between 5.0 and 5.9 on the magnitude scale in the area surrounding the city.³¹ That works out to about one for every three years. According to the power law that Gutenberg and Richter uncovered, that means that an earthquake measuring between 6.0 and 6.9 should occur about once every thirty years in Tehran.

Furthermore, it follows that an earthquake that measured 7.0 or greater would occur about once every three hundred years near Tehran. This is the earthquake that Susan Hough fears. The Haiti earthquake of 2010, which measured magnitude 7.0 and killed 316,000,³² showed the apocalyptic consequences that earthquakes can produce in the developing world. Iran shares many of Haiti's problems—poverty, lax building codes, political corruption³³—but it is much more densely populated. The USGS estimates, on the basis of high death

DESPERATELY SEEKING SIGNAL

FIGURE 5-3B: WORLDWIDE EARTHQUAKE FREQUENCIES, JANUARY 1964–MARCH 2012, LOGARITHMIC SCALE



tolls from smaller earthquakes in Iran, that between 15 and 30 percent of Tehran's population could die in the event of a catastrophic tremor there.³⁴ Since there are about thirteen million people in Tehran's metro area, that would mean between two and four million fatalities.

What the Gutenberg–Richter law does not tell us anything about is *when* the earthquake would strike. (Nor does it suggest that Tehran is “due” for an earthquake if it hasn't experienced one recently.) Countries like Iran and Haiti do not have the luxury of making contingency plans for a once-every-three-hundred-year event. The earthquake forecasts produced using the Gutenberg–Richter law provide for a good general guide to the hazard in an area. But like weather forecasts determined from statistical records alone (it rains 35 percent of the time in London in March), they don't always translate into actionable intelligence (should I carry an umbrella?). Geological time scales occupy centuries or entire millennia; human life spans are measured in years.

The Temptation of the Signal

What seismologists are really interested in—what Susan Hough calls the “Holy Grail” of seismology—are *time-dependent* forecasts, those in which the probability of an earthquake is not assumed to be constant across time.

Even seismologists who are skeptical of the possibility of making time-dependent earthquake forecasts acknowledge that there are some patterns in the earthquake distribution. The most obvious is the presence of aftershocks. Large earthquakes are almost always followed by dozens or even thousands of aftershocks (the 2011 earthquake in Japan produced at least 1,200 of them). These aftershocks follow a somewhat predictable pattern.³⁵ Aftershocks are more likely to occur immediately after an earthquake than days later, and more likely to occur days later than weeks after the fact.


This, however, is not terribly helpful when it comes to saving lives. This is because aftershocks, by definition, are always less powerful than the initial earthquake. Usually, if a particular fault produces a sufficiently powerful earthquake, there will be a few aftershocks and then that’ll be the end of the fireworks for a while. This isn’t always the case, however. For example, the incredibly powerful earthquake that hit the New Madrid Fault on the Missouri-Tennessee border on December 16, 1811, evaluated by seismologists as magnitude 8.2, was followed just six hours later by another shock of about the same magnitude. And the fault was not yet quiesced: the December 16 quakes were succeeded by another magnitude 8.1 earthquake on January 23, and then yet another, even more powerful 8.3 earthquake on February 7. Which ones were the foreshocks? Which ones were the aftershocks? Any interpretation is about as useless as any other.

The question, of course, is whether we can predict earthquakes *before* the fact: can we tell the foreshocks and aftershocks apart in advance? When we look at data that shows the distribution of earthquakes across time and space, it tempts us with the possibility that there might be some signal in the noise.

Figure 5-4a, for instance, shows the distribution of earthquakes near L’Aquila³⁶ from 2006 until the magnitude 6.3 earthquake hit in 2009.³⁷ All the data in this chart, except the large black circle that indicates the main earth-

quake, shows earthquakes that occurred before the main shock. In the case of L’Aquila, there does seem to be a discernible pattern. A big cluster of earthquakes, measuring up to about magnitude 4, occurred just before the main shock in early 2009—much higher than the background rate of seismic activity in the area.

FIGURE 5-4A: EARTHQUAKES NEAR L’AQUILA, ITALY
JANUARY 1, 2006–APRIL 6, 2009




A more debatable case is the Japan earthquake of 2011. When we make one of these plots for the Tōhoku region (figure 5-4b), we see, first of all, that it is much more seismically active than Italy. But are there patterns in the timing of the earthquakes there? There seem to be some; for instance, there is a cluster of earthquakes measuring between magnitude 5.5 and magnitude 7.0 in mid-2008. These, however, did not precipitate a larger earthquake. But we do see an especially large foreshock, magnitude 7.5, on March 9, 2011, preceding the magnitude 9.1 Tōhoku earthquake³⁸ by about fifty hours.

Only about half of major earthquakes are preceded by discernible foreshocks,³⁹ however. Haiti’s was not (figure 5-4c). Instrumentation is not very good in most parts of the Caribbean, so we don’t have records of magnitude 2 and 3 earthquakes, but seismometers in the United States and other areas should be able to pick up anything that registers at 4 or higher. The last time there had been even a magnitude 4 earthquake in the area was in 2005. For years

THE SIGNAL AND THE NOISE


FIGURE 5-4B: EARTHQUAKES NEAR TÔHOKU, JAPAN
JANUARY 1, 2006–MARCH 11, 2011



before the magnitude 7.0 earthquake hit in 2010. There was just no warning at all.

Complicating matters further are false alarms—periods of increased seismic activity that never result in a major tremor. One case well known to seismologists is a series of smaller earthquakes near Reno, Nevada, in early 2008.


FIGURE 5-4C: EARTHQUAKES NEAR LÉOGÂNE, HAITI
JANUARY 1, 2000–JANUARY 12, 2010



DESPERATELY SEEKING SIGNAL

The Reno earthquake swarm looks a lot like the one we saw before L'Aquila in 2009. But it never amounted to anything much; the largest earthquake in the series was just magnitude 5.0 and no major earthquake followed.

FIGURE 5-4D: EARTHQUAKES NEAR RENO, NEVADA
JANUARY 1, 2006–DECEMBER 31, 2011



This is just a taste of the maddening array of data that seismologists observe. It seems to exist in a purgatory state—not quite random and not quite predictable. Perhaps that would imply that we could at least get halfway there and make some progress in forecasting earthquakes—even if we can never get to hard-and-fast predictions. But the historical record of attempts to predict earthquakes is one of almost complete failure.

A Parade of Failed Forecasts

Hough's 2009 book, *Predicting the Unpredictable: The Tumultuous Science of Earthquake Prediction*, is a history of efforts to predict earthquakes, and is as damning to that enterprise as Phil Tetlock's study was to political pundits. There just seems to have been no progress at all, and there have been many false alarms.

Lima, Peru

One of the more infamous cases involved a geophysicist named Brian Brady, who had a Ph.D. from MIT and worked at Colorado School of Mines. Brady asserted that a magnitude 9.2 earthquake—one of the largest in recorded history—would hit Lima, Peru, in 1981.⁴⁰ His prediction initially had a fair amount of support in the seismological community—an early version of it had been coauthored with a USGS scientist. But as the theory became more elaborate—Brady would eventually invoke everything from the rock bursts he had observed in his studies of mines to Einstein's theory of relativity in support of it—colleagues had started telling him that theory was beyond their understanding;⁴¹ a polite way of saying that he was nuts. Eventually, he predicted that the magnitude 9.2 earthquake would be just one in a spectacular series in Peru, culminating in a magnitude 9.9 earthquake, the largest in recorded history, in August 1981.⁴²

The prediction was leaked to the Peruvian media and terrified the population; this serious-seeming American scientist was sure their capital city would be in ruins. Their fear only intensified when it was reported that the Peruvian Red Cross had requested 100,000 body bags to prepare for the disaster. Tourism and property values declined,⁴³ and the U.S. government eventually dispatched a team of scientists and diplomats to Peru in an effort to calm nerves. It made front-page news when there was no Great Peruvian Earthquake in 1981 (or even a minor one).

Parkfield, California

If Lima had provided a warning that false alarms can extract a substantial psychological and economic cost on the population, it did not stop seismologists from seeking out the Holy Grail. While Brady had been something of a lone wolf, there were cases when earthquake prediction had much more explicit backing from the USGS and the rest of the seismological community. These efforts did not go so well either.

Among the most studied seismic zones in the world is Parkfield, California, which sits along the San Andreas Fault somewhere between Fresno, Bakers-

field, and the next exit with an In-N-Out Burger. There had been earthquakes in Parkfield at what seemed to be regular intervals about twenty-two years apart: in 1857, 1881, 1901, 1922, 1934, and 1966. A USGS-sponsored paper⁴⁴ projected the trend forward and predicted with 95 percent confidence that there would be another such earthquake at some point between 1983 and 1993, most likely in 1988. The next significant earthquake to hit Parkfield did not occur until 2004, however, well outside of the prediction window.

Apart from being wrong, the Parkfield prediction also seemed to reinforce a popular misconception about earthquakes: that they come at regular intervals and that a region can be "due" for one if it hasn't experienced an earthquake in some time. Earthquakes result from a buildup of stress along fault lines. It might follow that the stress builds up until it is discharged, like a geyser erupting with boiling water, relieving the stress and resetting the process.

But the fault system is complex: regions like California are associated with multiple faults, and each fault has its own branches and tributaries. When an earthquake does strike, it may relieve the stress on one portion of a fault, but it can transfer it along to neighboring faults, or even to some faraway portion of the same fault.⁴⁵ Moreover, the stress on a fault is hard to observe directly—until an earthquake hits.

What this means is that if San Francisco is forecasted to have a major earthquake every thirty-five years, it does not imply that these will be spaced out evenly (as in 1900, 1935, 1970). It's safer to assume there is a 1 in 35 chance of an earthquake occurring every year, and that this rate does not change much over time regardless of how long it has been since the last one.

Mojave Desert, California

The Brady and Parkfield fiascoes seemed to suppress efforts at earthquake prediction for some time. But they came back with a vengeance in the 2000s, when newer and seemingly more statistically driven methods of earthquake prediction became the rage.

One such method was put forward by Vladimir Keilis-Borok, a Russian-born mathematical geophysicist who is now in his late eighties and teaches at UCLA. Keilis-Borok had done much to advance the theory of how earthquakes

formed and first achieved notoriety in 1986 when, at a summit meeting in Reykjavík with Mikhail Gorbachev, President Reagan was handed a slip of paper predicting a major earthquake in the United States within the next five years, an event later interpreted to be the Loma Prieta quake that struck San Francisco in 1989.⁴⁶

In 2004, Keilis-Borok and his team claimed to have made a “major breakthrough” in earthquake prediction.⁴⁷ By identifying patterns from smaller earthquakes in a given region, they said, they were able to predict large ones. The methods that Keilis-Borok applied to identify these patterns were elaborate and opaque,⁴⁸ representing past earthquakes with a series of eight equations, each of which was applied in combination with the others at all conceivable intervals of time and space. But, the team said, their method had correctly predicted 2003 earthquakes in San Simeon, California, and Hokkaido, Japan.

Whether the San Simeon and Hokkaido predictions were publicly communicated ahead of time remains unclear;⁴⁹ a search of the Lexis-Nexis database of newspapers reveals no mention of them in 2003.⁵⁰ When we are evaluating the success of a forecasting method, it is crucial to keep “retrodictions” and predictions separate; predicting the past is an oxymoron and obviously should not be counted among successes.⁵¹

By January 2004, however, Keilis-Borok had gone very public with another prediction:⁵² an earthquake measuring at least magnitude 6.4 would hit an area of the Mojave Desert in Southern California at some point within the subsequent nine months. The prediction began to attract widespread attention; Keilis-Borok was featured in the pages of *Discover* magazine, the *Los Angeles Times*, and a dozen or so other mainstream publications. Someone from Governor Schwarzenegger’s office called; an emergency panel was convened. Even the famously skeptical USGS was willing to give some credit; their Web site conceded that “the work of the Keilis-Borok team is a legitimate approach to earthquake prediction research.”⁵³

But no major earthquake hit the Mojave Desert that year, and indeed, almost a decade later, none has. The Keilis-Borok team has continued to make predictions about earthquakes in California, Italy, and Japan but with little success: a 2010 analysis found three hits but twenty-three misses among predictions that they had clearly enunciated ahead of time.⁵⁴

Sumatra, Indonesia

There is another type of error, in which an earthquake of a given magnitude is deemed unlikely or impossible in a region—and then it happens. David Bowman, a former student of Keilis-Borok who is now the chair of the Department of Geological Sciences at Cal State Fullerton, had redoubled his efforts at earthquake prediction after the Great Sumatra Earthquake of 2004, the devastating magnitude 9.2 disaster that produced a tsunami and killed 230,000 people. Bowman’s technique, like Keilis-Borok’s, was highly mathematically driven and used medium-size earthquakes to predict major ones.⁵⁵ However, it was more elegant and ambitious, proposing a theory called accelerated moment release that attempted to quantify the amount of stress at different points in a fault system. In contrast to Keilis-Borok’s approach, Bowman’s system allowed him to forecast the likelihood of an earthquake along any portion of a fault; thus, he was not just predicting where earthquakes would hit, but also where they were unlikely to occur.

Bowman and his team did achieve some initial success; the massive aftershock in Sumatra in March 2005, measuring magnitude 8.6, had its epicenter in an area his method identified as high-risk. However, a paper that he published in 2006⁵⁶ also suggested that there was a particularly low risk of an earthquake on another portion of the fault, in the Indian Ocean adjacent to the Indonesian province of Bengkulu. Just a year later, in September 2007, a series of earthquakes hit exactly that area, culminating in a magnitude 8.5. Fortunately, the earthquakes occurred far enough offshore that fatalities were light, but it was devastating to Bowman’s theory.

Between a Rock and a Hard Place

After the model’s failure in 2007, Bowman did something that forecasters very rarely do. Rather than blame the failure on bad luck (his model had allowed for some possibility of an earthquake near Bengkulu, just not a high one), he reexamined his model and decided his approach to predicting earthquakes was fundamentally flawed—and gave up on it.

"I'm a failed predictor," Bowman told me in 2010. "I did a bold and stupid thing—I made a testable prediction. That's what we're supposed to do, but it can bite you when you're wrong."

Bowman's idea had been to identify the root causes of earthquakes—stress accumulating along a fault line—and formulate predictions from there. In fact, he wanted to understand how stress was changing and evolving throughout the entire system; his approach was motivated by chaos theory.

Chaos theory is a demon that can be tamed—weather forecasters did so, at least in part. But weather forecasters have a much better theoretical understanding of the earth's atmosphere than seismologists do of the earth's crust. They know, more or less, how weather works, right down to the molecular level. Seismologists don't have that advantage.

"It's easy for climate systems," Bowman reflected. "If they want to see what's happening in the atmosphere, they just have to look up. We're looking at rock. Most events occur at a depth of fifteen kilometers underground. We don't have a hope of drilling down there, realistically—sci-fi movies aside. That's the fundamental problem. There's no way to directly measure the stress."

Without that theoretical understanding, seismologists have to resort to purely statistical methods to predict earthquakes. You can create a statistical variable called "stress" in your model, as Bowman tried to do. But since there's no way to measure it directly, that variable is still just expressed as a mathematical function of past earthquakes. Bowman thinks that purely statistical approaches like these are unlikely to work. "The data set is incredibly noisy," he says. "There's not enough to do anything statistically significant in testing hypotheses."

What happens in systems with noisy data and underdeveloped theory—like earthquake prediction and parts of economics and political science—is a two-step process. First, people start to mistake the noise for a signal. Second, this noise pollutes journals, blogs, and news accounts with false alarms, undermining good science and setting back our ability to understand how the system actually works.

Overfitting: The Most Important Scientific Problem You've Never Heard Of

In statistics, the name given to the act of mistaking noise for a signal is *overfitting*.

Suppose that you're some sort of petty criminal and I'm your boss. I deputize you to figure out a good method for picking combination locks of the sort you might find in a middle school—maybe we want to steal everybody's lunch money. I want an approach that will give us a high probability of picking a lock anywhere and anytime. I give you three locks to practice on—a red one, a black one, and a blue one.

After experimenting with the locks for a few days, you come back and tell me that you've discovered a foolproof solution. If the lock is red, you say, the combination is 27-12-31. If it's black, use the numbers 44-14-19. And if it's blue, it's 10-3-32.

I'd tell you that you've completely failed in your mission. You've clearly figured out how to open these three particular locks. But you haven't done anything to advance our theory of lock-picking—to give us some hope of picking them when we don't know the combination in advance. I'd have been interested in knowing, say, whether there was a good type of paper clip for picking these locks, or some sort of mechanical flaw we can exploit. Or failing that, if there's some trick to detect the combination: maybe certain types of numbers are used more often than others? You've given me an *overly specific* solution to a *general* problem. This is overfitting, and it leads to worse predictions.


The name overfitting comes from the way that statistical models are "fit" to match past observations. The fit can be too loose—this is called underfitting—in which case you will not be capturing as much of the signal as you could. Or it can be too tight—an overfit model—which means that you're fitting the noise in the data rather than discovering its underlying structure. The latter error is much more common in practice.

To see how this works, let's give ourselves an advantage that we'll almost never have in real life: we'll know *exactly* what the real data is supposed to look

THE SIGNAL AND THE NOISE

like. In figure 5-5, I've drawn a smooth parabolic curve, which peaks in the middle and trails off near the ends. This could represent any sort of real-world data that you might like: as we saw in chapter 3, for instance, it represents a pretty good description of how baseball players perform as they age, since they are better in the middle of their careers than at the end or the beginning.

FIGURE 5-5: TRUE DISTRIBUTION OF DATA




However, we do not get to observe this underlying relationship directly. Instead, it manifests itself through a series of individual data points and we have to infer the pattern from those. Moreover, these data points are affected by idiosyncratic circumstances—so there is some signal, but there is also some noise. In figure 5-5, I've plotted one hundred data points, represented by circles and triangles. This looks to be enough to detect the signal through the noise. Although there is some randomness in the data, it's pretty clear that they follow our curve.

What happens, however, when we have a more limited amount of data, as will usually be the case in real life? Then we have more potential to get ourselves in trouble. In figure 5-6a, I've limited us to about twenty-five of our one hundred observations. How would you connect these dots?


DESPERATELY SEEKING SIGNAL

FIGURE 5-6A: LIMITED SAMPLE OF DATA




Knowing what the real pattern is supposed to be, of course, you'll still be inclined to fit the points with some kind of curve shape. Indeed, modeling this data with a simple mathematical expression called a quadratic equation does a very good job of re-creating the true relationship (figure 5-6b).

FIGURE 5-6B: WELL-FIT MODEL



When we don't know the Platonic ideal for our data, however, sometimes we get greedy. Figure 5-6c represents an example of this: an overfit model. In figure 5-6c, we've devised a complex function⁵⁷ that chases down every outlying data point, weaving up and down implausibly as it tries to connect the dots. This moves us further away from the true relationship and will lead to worse predictions.

FIGURE 5-6C: OVERFIT MODEL



This seems like an easy mistake to avoid, and it would be if only we were omniscient and always knew about the underlying structure of the data. In almost all real-world applications, however, we have to work by induction, inferring the structure from the available evidence. You are most likely to overfit a model when the data is limited and noisy and when your understanding of the fundamental relationships is poor; both circumstances apply in earthquake forecasting.

If we either don't know or don't care about the truth of the relationship, there are lots of reasons why we may be prone to overfitting the model. One is that the overfit model will score better according to most of the statistical tests that forecasters use. A commonly used test is to measure how much of the variability in the data is accounted for by our model. According to this test, the overfit model (figure 5-6c) explains 85 percent of the variance, making it

"better" than the properly fit one (figure 5-6b), which explains 56 percent. But the overfit model scores those extra points in essence by cheating—by fitting noise rather than signal. It actually does a much worse job of explaining the real world.⁵⁸

As obvious as this might seem when explained in this way, many forecasters completely ignore this problem. The wide array of statistical methods available to researchers enables them to be no less fanciful—and no more scientific—than a child finding animal patterns in clouds.⁵⁹ "With four parameters I can fit an elephant," the mathematician John von Neumann once said of this problem.⁶⁰ "And with five I can make him wiggle his trunk."

Overfitting represents a double whammy: it makes our model look *better* on paper but perform *worse* in the real world. Because of the latter trait, an overfit model eventually will get its comeuppance if and when it is used to make real predictions. Because of the former, it may look superficially more impressive until then, claiming to make very accurate and newsworthy predictions and to represent an advance over previously applied techniques. This may make it easier to get the model published in an academic journal or to sell to a client, crowding out more honest models from the marketplace. But if the model is fitting noise, it has the potential to hurt the science.

As you may have guessed, something like Keilis-Borok's earthquake model was badly overfit. It applied an incredibly complicated array of equations to noisy data. And it paid the price—getting just three of its twenty-three predictions correct. David Bowman recognized that he had similar problems and pulled the plug on his model.

To be clear, these mistakes are usually honest ones. To borrow the title of another book, they play into our tendency to be fooled by randomness. We may even grow quite attached to the idiosyncrasies in our model. We may, without even realizing it, work backward to generate persuasive-sounding theories that rationalize them, and these will often fool our friends and colleagues as well as ourselves. Michael Babyak, who has written extensively on this problem,⁶⁰ puts

⁵⁸ If you feed a computer a string of coin tosses (a random mix of 1's and 0's representing heads and tails), and then test out statistical parameters to try to fit a pattern-matching model, eventually it will think it can call 60 percent or 70 percent or (if you include enough variables) 100 percent of coin flips correctly. All this is artificial, of course; over the long run, it will call exactly 50 percent of coin flips correctly, no more and no less.

the dilemma this way: "In science, we seek to balance curiosity with skepticism." This is a case of our curiosity getting the better of us.

An Overfit Model of Japan?

Our tendency to mistake noise for signal can occasionally produce some dire real-world consequences. Japan, despite being extremely seismically active, was largely unprepared for its devastating 2011 earthquake. The Fukushima nuclear reactor was built to withstand a magnitude 8.6 earthquake,⁶¹ but not a 9.1. Archaeological evidence⁶² is suggestive of historic tsunamis on the scale of the 130-foot waves that the 2011 earthquake produced, but these cases were apparently forgotten or ignored.

A magnitude 9.1 earthquake is an incredibly rare event in any part of the world: nobody should have been predicting it to the exact decade, let alone the exact date. In Japan, however, some scientists and central planners dismissed the possibility out of hand. This may reflect a case of overfitting.

In figure 5-7a, I've plotted the historical frequencies of earthquakes near the 2011 epicenter in Japan.⁶¹ The data includes everything up through but not including the magnitude 9.1 earthquake on March 11. You'll see that the relationship almost follows the straight-line pattern that Gutenberg and Richter's method predicts. However, at about magnitude 7.5, there is a kink in the graph. There had been no earthquakes as large as a magnitude 8.0 in the region since 1964, and so the curve seems to bend down accordingly.

So how to connect the dots? If you go strictly by the Gutenberg-Richter law, ignoring the kink in the graph, you should still follow the straight line, as in figure 5-7b. Alternatively, you could go by what seismologists call a characteristic fit (figure 5-7c), which just means that it is descriptive of the historical frequencies of the earthquake in that area. In this case, that would mean that you took the kink in the historical data to be real—meaning, you thought there was some good reason why earthquakes larger than about magnitude 7.6 were unlikely to occur in the region.

FIGURE 5-7A: TŌHOKU, JAPAN EARTHQUAKE FREQUENCIES
JANUARY 1, 1964–MARCH 10, 2011




FIGURE 5-7B: TŌHOKU, JAPAN EARTHQUAKE FREQUENCIES
GUTENBERG-RICHTER FIT





FIGURE 5-7C. TŌHOKU, JAPAN EARTHQUAKE FREQUENCIES
CHARACTERISTIC FIT



Here is another example where an innocuous-seeming choice of assumptions will yield radically distinct conclusions—in this case, about the probability of a magnitude 9 earthquake in this part of Japan. The characteristic fit suggests that such an earthquake was nearly impossible—it implies that one might occur about every 13,000 years. The Gutenberg–Richter estimate, on the other hand, was that you’d get one such earthquake every three hundred years. That’s infrequent but hardly impossible—a tangible enough risk that a wealthy nation like Japan might be able to prepare for it.⁶⁴

The characteristic fit matched the recent historical record from Tōhoku a bit more snugly. But as we’ve learned, this type of pattern-matching is *not* always a good thing—it could imply an overfit model, in which case it will do a worse job of matching the true relationship.

In this case, an overfit model would dramatically underestimate the likelihood of a catastrophic earthquake in the area. The problem with the characteristic fit is that it relied on an incredibly weak signal. As I mentioned, there had been no earthquake of magnitude 8 or higher in this region in the forty-five years or so prior to Tōhoku. However, these are rare events to begin with: the

Gutenberg–Richter law posits that they might occur only about once per thirty years in this area. It’s not very hard at all for a once-per-thirty-year event to fail to occur in a forty-five-year window,⁶⁵ no more so than a .300 hitter having a bad day at the plate and going 0-for-5.⁶⁶ Meanwhile, there were quite a few earthquakes with magnitudes in the mid- to high 7’s in this part of Japan. When such earthquakes had occurred in other parts of the world, they had almost always suggested the potential for larger ones. What justification was there to think that Tōhoku would be a special case?

Actually, seismologists in Japan and elsewhere came up with a few rationalizations for that. They suggested, for instance, that the particular composition of the seafloor in the region, which is old and relatively cool and dense, might prohibit the formation of such large earthquakes.⁶⁷ Some seismologists observed that, before 2004, no magnitude 9 earthquake had occurred in a region with that type of seafloor.

This was about like concluding that it was impossible for anyone from Pennsylvania to win the Powerball jackpot because no one had done so in the past three weeks. Magnitude 9 earthquakes, like lottery winners, are few and far between. Before 2004, in fact, only three of them had occurred in recorded history anywhere in the world. This wasn’t nearly enough data to support such highly specific conclusions about the exact circumstances under which they might occur. Nor was Japan the first failure of such a theory; a similar one had been advanced about Sumatra⁶⁸ at a time when it had experienced lots of magnitude-7 earthquakes⁶⁹ but nothing stronger. Then the Great Sumatra Earthquake, magnitude 9.2,⁷⁰ hit in December 2004.

The Gutenberg–Richter law would not have predicted the exact timing of the Sumatra or Japan earthquakes, but it would have allowed for their possibility.⁷¹ So far, it has held up remarkably well when a great many more elaborate attempts at earthquake prediction have failed.

The Limits of Earthquakes and Our Knowledge of Them

The very large earthquakes of recent years are causing seismologists to rethink what the upper bounds of earthquakes might be. If you look at figure 5.3b,

which accounts for all earthquakes since 1964 (including Sumatra and Tōhoku) it now forms a nearly straight line though all the data points. A decade ago, you would have detected more of a kink in the graph (as in the Tōhoku chart in figure 5-7a). What this meant is that there were slightly fewer megaquakes than the Gutenberg–Richter law predicted. But recently we have been catching up.

Because they occur so rarely, it will take centuries to know what the true rate of magnitude 9 earthquakes is. It will take even longer to know whether earthquakes larger than magnitude 9.5 are possible. Hough told me that there may be some fundamental constraints on earthquake size from the geography of fault systems. If the largest continuous string of faults in the world ruptured together—everything from Tierra Del Fuego at the southern tip of South America all the way up through the Aleutians in Alaska—a magnitude 10 is about what you'd get, she said. But it is hard to know for sure.

Even if we had a thousand years of reliable seismological records, however, it might be that we would not get all that far. It may be that there are intrinsic limits on the predictability of earthquakes.

Earthquakes may be an inherently *complex* process. The theory of complexity that the late physicist Per Bak and others developed is different from chaos theory, although the two are often lumped together. Instead, the theory suggests that very simple things can behave in strange and mysterious ways when they interact with one another.

Bak's favorite example was that of a sandpile on a beach. If you drop another grain of sand onto the pile (what could be simpler than a grain of sand?), it can actually do one of three things. Depending on the shape and size of the pile, it might stay more or less where it lands, or it might cascade gently down the small hill toward the bottom of the pile. Or it might do something else: if the pile is too steep, it could destabilize the entire system and trigger a sand avalanche. Complex systems seem to have this property, with large periods of apparent stasis marked by sudden and catastrophic failures. These processes may not literally be random, but they are so irreducibly complex (right down to the last grain of sand) that it just won't be possible to predict them beyond a certain level.

The Beauty of the Noise

And yet complex processes produce order and beauty when you zoom out and look at them from enough distance. I use the terms signal and noise very loosely in this book, but they originally come from electrical engineering. There are different types of noise that engineers recognize—all of them are random, but they follow different underlying probability distributions. If you listen to true white noise, which is produced by random bursts of sound over a uniform distribution of frequencies, it is sibilant and somewhat abrasive. The type of noise associated with complex systems, called Brownian noise, is more soothing and sounds almost like rushing water.⁷²

Meanwhile, the same tectonic forces that carve fault lines beneath the earth's surface also carve breathtaking mountains, fertile valleys, and handsome coastlines. What that means is that people will probably never stop living in them, despite the seismic danger.

Science on Trial

In a final irony of the LAquila earthquake, a group of seven scientists and public officials were quite literally put on trial for manslaughter in 2011.⁷³ Prosecutors from the city of LAquila alleged that they had failed to adequately notify the public about the risk of a Big One after the earthquake swarm there.

The trial was obviously ridiculous, but is there anything the scientists could have done better? Probably there was; there is fairly clear evidence that the risk of a major earthquake increases substantially—perhaps temporarily becoming one hundred to five hundred times higher than its baseline rate⁷⁴—following an earthquake swarm. The risk was nevertheless extremely low—most earthquake swarms do not produce major quakes—but it was not quite right to imply that everything was normal and that people should sit down and have a glass of wine.

This book takes the view that the first duty of a forecaster is always fealty to

THE SIGNAL AND THE NOISE

the truth of the forecast. Politics, broadly defined, can get in the way of that. The seismological community is still scarred by the failed predictions in Lima and Parkfield, and by having to compete against the likes of Giuliani. This complicates their incentives and distracts them from their mission. Bad and irresponsible predictions can drive out good ones.

Hough is probably right that the Holy Grail of earthquake prediction will never be attained. Even if individual seismologists are behaving responsibly, we nevertheless have the collective output of the discipline to evaluate, which together constitutes thousands of hypotheses about earthquake predictability. The track record suggests that most of these hypotheses have failed and that magic-bullet approaches to earthquake prediction just aren't likely to work.

However, the track record of science as a whole is a remarkable one; that is also a clear signal. It is probably safe to conclude that the same method attempted over and over with little variation is unlikely to yield different results. But science often produces "unpredictable" breakthroughs.

One area in which seismologists have made some progress is in the case of very short term earthquake forecasts, as might have been relevant in L'Aquila. Next to the Gutenberg-Richter law, the knowledge that major earthquakes essentially always produce aftershocks is the most widely accepted finding in the discipline. Some seismologists I spoke with, like John Rundle of UC Davis and Tom Jordan of the University of Southern California, are concentrating more on these near-term forecasts and increasingly take the view that they should be communicated clearly and completely to the public.

Jordan's research, for instance, suggests that aftershocks sometimes move in a predictable geographic direction along a fault line. If they are moving in the direction of a population center, they can potentially be more threatening to life and property even if they are becoming less powerful. For instance, the magnitude 5.8 earthquake in Christchurch, New Zealand, in 2011, which killed 185, was an aftershock of a 7.0 earthquake that occurred in September 2010 in a remote part of the country.⁷⁵ When it comes to aftershocks, there is clearly a lot of signal, so this may be the more natural place to focus.

Finally, technology is always pushing forward. Recent efforts by NASA and by Rundle to measure fault stress through remote sensing systems like GPS satellites have shown some promise.⁷⁶ Although the efforts are crude for the

DESPERATELY SEEKING SIGNAL

time being, there is potential to increase the amount of data at seismologists' disposal and get them closer to understanding the root causes of earthquakes.

These methods may eventually produce some forward progress. If success in earthquake prediction has been almost nonexistent for millennia, the same was true for weather forecasting until about forty years ago. Or it may be that as we develop our understanding of complexity theory—itsself a very new branch of science—we may come to a more emphatic conclusion that earthquakes are not really predictable at all.

Either way, there will probably be some failed predictions first. As the memory of our mistakes fades, the signal will again seem to shimmer over the horizon. Parched for prediction we will pursue it, even if it is a mirage.