1

ROBUST, YET FRAGILE

This is a book about why things bounce back.

In the right circumstances, all kinds of things do: people and communities, businesses and institutions, economies and ecosystems. The place in which you live, the company at which you work, and even, without realizing it, you yourself: Each is resilient, or not, in its own way. And each, in its own way, illuminates one corner of a common reservoir of shared themes and principles. By understanding and embracing these patterns of resilience, we can create a sturdier world, and sturdier selves to live in it.

But before we can understand how things knit themselves back together, we have to understand why they fall apart in the first place. So let's start with a thought experiment:

Imagine, for a moment, that you are a tree farmer, planting a new crop on a vacant plot of earth. Like all farmers, to get the most use out of your land, you will have to contend with any number of foreseeable difficulties: bad weather, drought, changing prices for your product, and, perhaps especially, forest fires.

Fortunately, there are some steps you can take to protect your tree farm against these potential dangers. To lower the risk of fire, for example, you can start by planting your saplings at wide but regular intervals, perhaps 10 meters apart. Then, when your trees mature, none of their canopies will touch, and a spark will have trouble spreading from tree to tree. This is a fairly safe but fairly inefficient planting strategy: Your whole crop will be less likely to burn to the ground, but you will hardly be getting the most productive use out of your land.

Now imagine that, to increase the yield of your farm, you randomly start planting saplings in between this grid of regularly spaced trees. By definition, these new additions will be much closer to their neighbors, and their canopies will grow to touch those of the trees surrounding

them. This will increase your yield, but also your risk: If a fire ignites one of these randomly inserted trees, or an adjoining one, it will have a much higher chance of spreading through its leaves and branches to the connected cluster of trees around it.

At a certain point, continuing to randomly insert trees into the grid will result in all of the trees' canopies touching in a dense megacluster. (This is often achieved when about 60 percent of the available space is filled in.) This, in contrast to your initial strategy, is a tremendously efficient design, at least from a planting and land-use perspective, but it's also very risky: A small spark might have disastrous consequences for the entire crop.

Of course, being a sophisticated arborist, you're unlikely to plant your trees in either a sparse grid or a dense megacluster. Instead, you choose to plant groves of densely planted trees interspersed with *roads*, which not only provide you access to the interior of your property but act as firebreaks, separating one part of the crop from another and insuring the whole from a fire in one of its parts.

These roads are not free: By reducing the area usable for planting, each imposes a cost, so you must be careful where you put them; adding too many roads is just as bad as adding too few. But eventually, with lots of trial and error, and taking into account local variations in the weather, soil, and geography, you may alight upon a near-perfect design for your particular plot—one that maximizes the density of trees while making smart and efficient use of roads to access them. Your tree farm's exceptional design will easily withstand the occasional fire, without ever burning entirely to the ground, all the while providing you with seasonally variable but not wildly gyrating timber returns.

Imagine your horror, then, when one day you discover that much of your perfectly designed plot has been infested by an invasive foreign beetle. This tiny pest, native to another geographic region entirely, stowed away on a shipment from an overseas supplier, then hitchhiked its way to the heart of your tree farm by clinging to your boot. Once inside, it exploited the very features of your design—your carefully placed roads—that were intended to insure against the risks you thought you'd confront.

It is at this moment in our thought experiment that you have painfully discovered that, in systems terms, your tree farm design is *robust-yet-fragile* (or RYF), a term coined by California Institute of Technology

research scientist John Doyle to describe complex systems that are resilient in the face of anticipated dangers (in this case forest fires) but highly susceptible to unanticipated threats (in this case exotic beetles).

On any given day, our news media is filled with real-world versions of this story. Many of the world's most critical systems—from coral reefs and communities to businesses and financial markets—have similar robust-yet-fragile dynamics; they're able to deal capably with a range of normal disruptions but fail spectacularly in the face of rare, unanticipated ones.

As in our tree-planting example, all RYF systems involve critical trade-offs, between efficiency and fragility on the one hand and inefficiency and robustness on the other. A perfectly efficient system, like the densely planted tree farm, is also the most susceptible to calamity; a perfectly robust system, like the sparsely planted tree farm, is too inefficient to be useful. Through countless iterations in their design (whether the designer of a system is a human being or the relentless process of natural selection), RYF systems eventually find a midpoint between these two extremes—an equilibrium that, like our roads-and-groves tree farm design, balances the efficiency and robustness trade-offs particular to the given circumstance.

The complexity of the resulting compensatory system—the network of the roads and groves in the example above—is a by-product of that balancing act. Paradoxically, over time, as the complexity of that compensatory system grows, it becomes a source of fragility itself—approaching a tipping point where even a small disturbance, if it occurs in the right place, can bring the system to its knees. No RYF design can therefore ever be "perfect," because each robustness strategy pursued creates a mirror-image (albeit rare) fragility. In an RYF system, the possibility of "black swans"—low-probability but high-impact events—is engineered in.

The Internet presents a perfect example of this robust-yet-fragile dynamic in action. From its inception as a U.S. military funded project in the 1960s, the Internet was designed to solve a particular problem above all else: to ensure the continuity of communications in the face of disaster. Military leaders at the time were concerned that a preemptive nuclear attack by the Soviets on U.S. telecommunications hubs could disrupt the chain of command—and that their own counterstrike orders might never make it from their command bunkers to their intended

recipients in the missile silos of North Dakota. So they asked the Internet's original engineers to design a system that could sense and automatically divert traffic around the inevitable equipment failures that would accompany any such attack.

The Internet achieves this feat in a simple yet ingenious way: It breaks up every email, web page, and video we transmit into packets of information and forwards them through a labyrinthine network of routers—specialized network computers that are typically redundantly connected to more than one other node on the network. Each router contains a regularly updated routing table, similar to a local train schedule. When a packet of data arrives at a router, this table is consulted and the packet is forwarded in the general direction of its destination. If the best pathway is blocked, congested, or damaged, the routing table is updated accordingly and the packet is diverted along an alternative pathway, where it will meet the next router in its journey, and the process will repeat. A packet containing a typical web search may traverse dozens of Internet routers and links—and be diverted away from multiple congestion points or offline computers—on the seemingly instantaneous trip between your computer and your favorite website.

The highly distributed nature of the routing system ensures that if a malicious hacker were to disrupt a single, randomly chosen computer on the Internet, or even physically blow it up, the network itself would be unlikely to be affected. The routing tables of nearby routers would simply be updated and would send network traffic around the damaged machine. In this way, it's designed to be robust in the face of the anticipated threat of equipment failure.

However, the modern Internet is extremely vulnerable to a form of attack that was unanticipated when it was first invented: the malicious exploitation of the network's open architecture—not to route around damage, but to engorge it with extra, useless information. This is what Internet spammers, computer worms and viruses, botnets, and distributed denial of service attacks do: They flood the network with empty packets of information, often from multiple sources at once. These deluges hijack otherwise beneficial features of the network to congest the system and bring a particular computer, central hub, or even the whole network to a standstill.

These strategies were perfectly illustrated in late 2010, when the secrets-revealing organization WikiLeaks began to divulge its trove of

secret U.S. State Department cables. To protect the organization from anticipated retaliation by the U.S. government, WikiLeaks and its supporters made copies of its matériel—in the form of an encrypted insurance file containing possibly more damaging information—available on thousands of servers across the network. This was far more than the United States could possibly interdict, even if it had had the technical capability and legal authority to do so (which it didn't). Meanwhile, an unaffiliated, loose-knit band of WikiLeaks supporters calling itself Anonymous initiated distributed denial-of-service attacks on the websites of companies that had cut ties with WikiLeaks, briefly knocking the sites of companies like PayPal and MasterCard off line in coordinated cyberprotests.

Both organizations, WikiLeaks and Anonymous, harnessed aspects of the Internet—redundancy and openness—that had once protected the network from the (now-archaic) danger that had motivated its invention in the first place: the threat of a strike by Soviet missiles. Four decades later, their unconventional attacks (at least from the U.S. government's perspective) utilized the very features of the network originally designed to prevent a more conventional one. In the process, the attacking organizations had proven themselves highly resilient: To take down WikiLeaks and stop the assault of Anonymous, the U.S. government would have had to take down the Internet itself, an impossible task.

Doyle points out a dynamic quite similar to this one at work in the human immune system. "Think of the illnesses that plague contemporary human beings: obesity, diabetes, cancer, and autoimmune diseases. These illnesses are malignant expressions of critical control processes of the human body—things like fat accumulation, modulation of our insulin resistance, tissue regeneration, and inflammation that are so basic that most of the time we don't even think about them. These control processes evolved to serve our hunter-gatherer ancestors, who had to store energy between long periods without eating and had to maintain their glucose levels in their brain while fueling their muscles. Such biological processes conferred great robustness on them, but in today's era of high-calorie, junk-food-saturated diets, these very same essential systems are hijacked to promote illness and decay."

To confront the threat of hijacking, Internet engineers add sophisticated software filters to their routers, which analyze incoming

and outgoing packets for telltale signs of malevolent intent. Corporations and individuals install firewall and antivirus software at every level of the network's organization, from the centralized backbones right down to personal laptops. Internet service providers add vast additional capacity to ensure that the network continues to function in the face of such onslaughts.

This collective effort to suffuse the system with distributed intelligence and redundancy may succeed, to varying degrees, at keeping some of these anticipated threats at bay. Yet, even with such actions, potential fragility has not been eliminated; it's simply been moved, to sprout again from another unforeseeable future direction. Worse, like all RYF systems, over time the complexity of these compensatory systems—antivirus software, firewalls—swell until they become a source of potential fragility themselves, as anyone who's ever had an important email accidentally caught in a spam filter knows all too well.

Along the way, paradoxically, the very fact that a robust-yet-fragile system continues to handle commonplace disturbances successfully will often mask an intrinsic fragility at its core, until—surprise!—a tipping point is catastrophically crossed. In the run-up to such an event, everything appears fine, with the system capably absorbing even severe but anticipated disruptions as it was intended to do. The very fact that the system continues to perform in this way conveys a sense of safety. The Internet, for example, continues to function in the face of inevitable equipment failures; our bodies metabolize yet another fast-food meal without going into insulin shock; businesses deal with intermittent booms and busts; the global economy handles shocks of various kinds. And then the critical threshold is breached, often by a stimulus that is itself rather modest, and all hell breaks loose.

When such failures arrive, many people are shocked to discover that these vastly consequential systems have no fallback mechanisms, say, for resolving the bankruptcy of a major financial institution or for the capping of a deep-sea spill.

And in the wake of these catastrophes, we end up resorting to simplified, moralistic narratives, featuring cartoon-like villains, to explain why they happened. In reality, such failures are more often the result of the almost imperceptible accretion of a thousand small, highly distributed decisions—each so narrow in scope as to seem innocuous

—that slowly erode a system's buffer zones and adaptive capacity. A safety auditor cuts one of his targets a break and looks the other way; a politician pressures a regulator on behalf of a constituent for the reduction of a fine; a manager looks to bolster her output by pushing her team to work a couple of extra shifts; a corporate leader decides to put off necessary investments for the future to make the quarterly numbers.

None of these actors is aware of the aggregate impact of his or her choices, as the margin of error imperceptibly narrows and the system they are working within inexorably becomes more brittle. Each, with an imperfect understanding of the whole, is acting rationally, responding to strong social incentives to serve a friend, a constituent, a shareholder in ways that have a significant individual benefit and a low systemic risk. But over time, their decisions slowly change the cultural norms of the system. The lack of consequences stemming from unsafe choices makes higher-risk choices and behaviors seem acceptable. What was the once-in-a-while exception becomes routine. Those who argue on behalf of the older way of doing things are perceived to be fools, paranoids, or party poopers, hopelessly out of touch with the new reality, or worse, enemies of growth who must be silenced. The system as a whole edges silently closer to possible catastrophe, displaying what systems scientists refer to as "self-organized criticality"—moving closer to a critical threshold.

This dynamic—and our first hints about what we might do to improve the resilience of such systems—can be seen in two very different robust-yet-fragile systems that, with new analytical tools, are beginning to powerfully illuminate each other: coral reef ecology and global finance.

OF FISHERIES AND FINANCE

In the 1950s, Jamaica's coral reefs were a thriving picture-postcard example of a Caribbean reef, supporting abundant mixtures of brilliantly colored sponges and feather-shaped octocorals sprouting up from the hard coral base. The reefs were popular habitats for hundreds of varieties of fish, including large predatory species such as sharks, snappers, groupers, and jacks, which were harvested by local fishermen as a reliable and time-honored food source for the island's population.

By the 1970s, things appeared relatively unchanged. In the intervening two decades, however, Jamaica's population had swelled by a third. Local fishermen, struggling to feed the island's growing population, had begun using motorized canoes to harvest not only the predator species, but the smaller, plant-eating fish, such as surgeonfish and parrotfish, as well. The reefs still appeared healthy, however, and most of their habitants seemed to be thriving. Sea urchins, in particular, were flourishing since they no longer had to compete with the plant-eating fish for algae, the mainstay of their diet.

Then, on August 6, 1980, after almost four decades without a major storm, one of the most powerful Caribbean tempests in history, Hurricane Allen, descended on the island and its surrounding reefs. Winds exceeding 175 miles per hour whipped up a forty-foot-high storm surge, which pounded the reefs. Shallow-water corals were devastated; however, their deeper-waters cousins, well below the surface, emerged relatively unscathed. In fact, a few months after the hurricane's strike, substantial coral recruitment—the measure of young corals entering the adult population—was found in the deeper waters. For the next three years, coral cover slowly increased.

The general consensus among marine biologists at the time was that Jamaica's reefs had survived remarkably well after the battering of Hurricane Allen. Data seemed to suggest that the reef system was surviving and, in the deeper waters around the island at least, maybe even thriving.

Then, in 1983, something terrifying happened below the surface of Jamaica's waters. An unidentified pathogen decimated the entire population of Jamaica's long-spined sea urchins. The illness was unprecedented in its lethality and its speed: Within a few days of the onset of symptoms, one observer reported, "on reefs that used to be black with urchins, one could swim for an hour without seeing a single living individual." By February 1984, the sea urchin had been virtually eliminated from its normal habitat—the most extensive and severe mass mortality ever reported for a marine organism.

Coming after the long and slow decline of the other native plant-eating fish species from overfishing, the loss of the urchins proved catastrophic to Jamaica's reefs. With no urchins—or other species—to keep it in check, algae quickly came to dominate every corner of the reef system, eventually covering 92 percent of its surface area and

killing the corals underneath. With the loss of the corals went the remaining fish—reefs that supported hundreds of species for thousands of years were transformed into vacant algal wastelands seemingly overnight.

On a healthy reef, a new pathogen decimating a *single species* (like the urchin) might not have had catastrophic consequences, because an essential reef function—like keeping algae in check—could be performed by more than one species. On the highly compromised Jamaican reef, however, the continued flourishing of the ecosystem as a whole became entirely dependent on a single species continuing to do that job. The loss of the urchins, an otherwise modest trigger, caused the reef to collapse virtually overnight.

Yet, had you asked a marine scientist in 1982 to describe the reefs, you would have received a promising assessment: The reef had proved robust in the face of significant disturbances, ranging from hurricanes to extensive human fishing. There was little evidence of the hidden fragility that was being exacerbated by the slow loss of biodiversity.

Such clarity is available to us in hindsight, but consider the challenges posed by trying to manage such a system at the time. The interactions between the various agents affecting the health of the reef (fish, urchins, algae, corals, and human beings) were imperfectly understood and nonlinear—small changes could lead to big outcomes, and vice versa. Some of the dependencies between these agents were masked: Prior to the loss of the urchins, it was hard to tell if the loss of herbivorous fish was having a significant impact or no impact whatsoever. What's more, recent experience suggested that the system could rebound from a disturbance on the scale of a hurricane. And even a healthy reef can behave in highly variable ways. Fish stocks rise and fall—how would you separate normal variability from the onset of collapse?

Similar questions confront us whenever we try to manage a complex system with highly interdependent parts. Whether we're dealing with fish stocks or financial stocks, to improve the resilience of the system as a whole, we first need measurement tools that take the health of whole systems into account, not just their pieces. At least if we want to keep eating seafood.

• • •

As it happens, in the 1950s, the same decade that our story in Jamaica began, the California sardine industry was also devastated by a collapse. The booming fishing industry had provided the dramatic backdrop for John Steinbeck's novel *Cannery Row*: In the mid-1930s, 790,000 tons of sardines were commercially fished from California's waters. Yet by 1953, the catch had plummeted to less than 15,000 tons—a drop of 98 percent.

Two competing hypotheses emerged to explain the collapse: one based on traditional simple overfishing and another on the La Niña currents that cooled California's ocean waters. However, in 2006, after poring over fifty years of data on sardine larvae, George Sugihara, a mathematician and theoretical ecologist at the Scripps Institution of Oceanography, proved that both theories were wrong. The sardine stock had collapsed not because the fishing industry was taking out too many *small* fish, but because it was taking out too many *big* ones: Sugihara found that California's industrialized fishing operations had become so efficient at catching adult sardines that they had significantly changed the age structure of the entire stock. Bereft of adults in 1949 and 1950, the substantially more juvenile sardine population had failed to spawn, and when it encountered additional stressors from the natural world, it flipped into collapse.

Given sufficient time, it's possible that the sardine population could have recovered from this megabust all by itself—such instabilities do occasionally occur naturally in ecosystems from time to time. Unfortunately, fisheries management practices of the time rarely took these kinds of instabilities into account. For much of the twentieth century (and in many places around the world today) the status quo of most fisheries management was based on maximum sustainable yield (MSY), or the largest catch that can be fished from the stock of a species over an indefinite period of time.

MSY is predicated on a linear system with stable equilibrium points. By addressing only one species at a time, and assuming that all other variables were fixed and unchanging, the MSY regime didn't call for moderation from California's commercial fishing fleets when sardine stocks started to decline. Rather, in the face of the intensified economic pressure of a reduced catch, many local fishing operations accelerated their harvesting efforts, pushing the ecosystem past the point of recovery. The fishing system had no breaks.

And MSY is still in widespread use today. It has been implicated in declining fish stocks all over the world: Today 63 percent of all fished species are overfished and in danger, and 29 percent have collapsed outright, meaning they are now, like the sardine, at least 90 percent below their historic maximum catch levels. In 2006, an international research team led by Boris Worm of Dalhousie University in Canada calculated that, if these trends continue, by 2048 all commercial fishing on Earth will cease—there will simply be no more fish in the sea.

To prevent such a catastrophe, Sugihara and others are trying to promote a more holistic alternative to MSY called ecosystem-based fishery management (EBFM). This whole-systems model is rooted in the observation that ecosystems are difficult to model and predict, with constantly changing threshold points that, when exceeded—as in the case of the California sardine population—can cause the system to collapse or restructure. EBFM counteracts this by promoting the maintenance of biodiversity as a central management goal whenever and wherever possible, at every level of organization, from small niches to entire regions of the sea.

Doing so starts with a very different system for measuring what's going on in an ecosystem that's being fished. Under MSY, fisheries managers gathered and analyzed information primarily about the species in the catch and little else; EBFM, on the other hand, requires them to measure the species of fish that are *not* caught as well as the ones that are. Many other factors, referred to as *ecosystem indicators*, need to be gathered and correlated as well, such as coastal upwelling, the pattern of wind and water movement that brings cool, nutrient-rich waters to ocean's surface near the coast, where they provide the diet for plankton, which in turn form the base of the food chain. Critically, EBFM also calls for measurement and correlation of societal trends on land as well as what's happening in the water—merging the social and ecological pictures.

Such a holistic management regime, based on nonlinear, complex systems, has gained traction politically, but implementing EBFM remains challenging: Very few of the fisheries management institutions have the expertise, history, resources, and know-how to turn away from the more straightforward equilibrium models that underpin the maximum sustainable yield regime.

“We like to see the world as consisting of separate parts that can be studied in an isolated, linear way, one piece at a time,” Sugihara explains. “These pieces can then be summed independently to make the whole. Researchers have developed a very large tool kit of analytical methods and statistics to do just this, and it has proven invaluable for studying simple engineered devices. But we persist with these linear tools and models even when systems that interest us are complex and nonlinear. It is a case of looking under the lamppost because the light is better even though we know the keys are lost in the shadows.”

To get fisheries to switch, ecologists are now working to make EBFM principles easier to implement, and to do so they’re borrowing a central concept from finance: the portfolio model. Interacting species in the ecosystem are linked together as a set of correlated assets, and, just as in a financial portfolio, fisheries managers make decisions that balance risks and returns among them. The portfolio method suggests that a particular species should be assessed on the basis of how it contributes to the overall performance of the ecosystem, rather than in isolation—in much the same way that a particular financial stock contributes to the value of a diversified fund. In both cases, there is no intrinsic value to the asset or the species—context is everything.

Managing entire classes of interacting species in this way simplifies what are otherwise very complex and nonlinear dependencies. By focusing on the relationships among categories of species, and not their absolute numbers, a portfolio approach can more easily accommodate complex factors like environmental fluctuations and advances in fishing technologies, even as it makes potential risks to the system much more explicit. And a multispecies portfolio can also be recalibrated according to the changing conditions of the local ecosystem, much like a financial adviser might arrange a portfolio one way for a young, risk-tolerant active investor and very differently when that investor becomes a risk-averse retiree.

The portfolio approach has other benefits as well. In a financial market, investing in groups of assets allows investors to lower the variability of their returns, in exchange for increasing the reliability of those returns. A properly balanced portfolio incorporates various strategies and hedges to smooth out the troughs and the peaks: During boom times, having commodities mixed in with aggressive growth stocks might keep you from generating sufficient returns to buy a

Ferrari, but during a bust it will also keep you out of the poorhouse. Similarly, maintaining a balanced "portfolio" of species allows a fisheries manager to reduce the variability of the annual catch—a major benefit to both fish and fishermen—in exchange for greater stability. In one study of the Chesapeake Bay, ecological economists Martin Smith and Douglas Lipton found that, using just such a portfolio-driven method, over the forty-year period from 1962 to 2003, the bay's fisheries managers could have generated better financial returns from fishing and reduced the season-to-season variability of the catch—a true win-win.

In matters of finance and fisheries, portfolio-driven approaches provide us with choices: If we seek a specific level of return, there are portfolios that can be designed to deliver that return while minimizing the risks; conversely, if we are comfortable with a particular level of risk, there are portfolios that can be designed at that level of risk to maximize our return.

ECOLOGY FOR BANKERS

Even as concepts from financial management begin to inform new strategies for improving the resilience of ecological systems like fisheries, the reverse is also occurring. A new generation of economists and finance experts are beginning to unearth important lessons in ecology for improving the resilience of the global financial economy. These insights are fueling the birth of an entirely new field, called *ecofinance*.

Consider, for a moment, some parallels between the Jamaican coral collapse mentioned earlier and the recent global financial crisis. Just as with the die-off of the sea urchins, the trigger of the financial collapse—Lehman Brothers' $600 billion filing for Chapter 11 bankruptcy in mid-September 2008—was, in the context of a $70 trillion annual global economy, a relatively modest event. Much like the Jamaican reefs' prior survival of Hurricane Allen, the financial markets had robustly weathered a prior decade filled with significant disturbances, including a dot-com bubble, oil shocks, and wars in the Middle East. Yet the decision to let Lehman fail, which we explore in greater detail later in this book, brought about an epidemic of fear and uncertainty that ground the world's entire capital markets to a halt.

Just like the decades-long erosion of biodiversity on the Jamaican reef, the underlying reasons for the financial collapse also involved a number of creeping structural changes to the market that set the stage for catastrophic collapse. And, as we've seen in other RYF systems, like those in the human body or the Internet, when that fragility finally did appear, its effects were amplified by the hijacking of normally beneficial aspects of the global financial system, including its structure, its approach to risk management, its feedback mechanisms, its levels of transparency, and its product innovations.

• • •

In 2008, George Sugihara, the same ecologist who had unearthed the reasons for the California sardine collapse, and two other ecologists, Simon Levin and Robert May, published an essay in *Nature* titled "Ecology for Bankers." They wanted to offer guidance for applying the tenets of holistic, ecosystem-based management to the financial markets. Like a densely planted tree farm, both systems had been managed for increased efficiency at the cost of increasing complexity and fragility. And both had implemented similar risk-management techniques: In the same way that maximum sustainable yield practices were applied to lone species of fish, it was common in the banking sector to apply single-firm risk analysis, looking only at individual banks without ever examining interconnection in the system as a whole. Worse, financial firms tended (and still tend) to calculate risk additively, meaning they take individual measures of risk for each transaction they make and add them up to arrive at a picture of their total potential exposure. This type of modeling makes the financial system seem much more secure than it really is; in fact, just as in marine ecosystems, the nonlinear nature of the financial network is multiplicative: Certain failures multiply the risks of other failures.

"Economics is not typically thought of as a global systems problem," Sugihara explains. "Investment banks are famous for a brand of tunnel vision that focuses risk management at the individual firm level and ignores the difficult and costlier systemic perspective. Monitoring the ecosystem-like network of firms with interlocking balance sheets is not in a corporate risk manager's job description. But ignoring these counterparty obligations and mutual interdependencies is exactly what

prevented financial institutions from seeing and appropriately pricing risk, which in turn dramatically amplified the recent financial crisis."

The essay in *Nature* came on the heels of an earlier, high-level conference in 2006 sponsored by the Federal Reserve Bank of New York, the U.S. National Academies, and the National Research Council to stimulate fresh thinking on systemic risk. Sugihara and Levin introduced the bankers in attendance to the concept of a *trophic web*—the way in which energy and nutrient flows in an ecosystem connect different species.

The basic concept will be familiar from grade-school science class: Aquatic plants convert the energy of the sun into food; the plants are eaten by small fish; the small fish by bigger fish. As fish and plants die, their bodies provide nutrients to smaller organisms and are recycled through the system. Trophic webs map these energy-transfer relationships in detail.

A similar framework can be used to analyze the transfer of value across a financial network. Yet, while creating detailed trophic webs is a common activity in ecology, creating similarly detailed flows is less common for large-scale financial networks. To remedy this, in 2006, the Federal Reserve Bank of New York commissioned a study of the topology of the interbank payment flows inside what's called the Fedwire system—the mechanism that U.S. banks use for transferring funds between them. Fedwire acts as the backbone of the U.S. financial system, processing an astounding average daily volume of almost 500,000 interbank payments, with a daily value of approximately $2.4 trillion (in 2010, the last year for which data was available at the time of writing). These payments can be thought of as the equivalents of energy and nutrients in the financial ecosystem being transferred between different species of institutions.

The sample used for the study involved approximately 700,000 transfers, and more than five thousand banks, taken from a typical day. The picture that emerges is startling: While most banks had a small number of connections, a few hubs had thousands. At the core of the network, just sixty-six banks accounted for 75 percent of the daily value of transfers. Even more telling, the network topology revealed that twenty-five of the biggest banks were completely connected—so intertwined that a failure among any strongly suggested a failure for all, the very definition of "too big to fail."

What's true within the core of the U.S. financial sector is increasingly true at an international level as well. An analysis of the connections between eighteen national financial markets reveals that, over the last two decades, central finance hubs around the world, like London and Hong Kong, have swelled to roughly fourteen times their earlier size; at the same time, the links between them have increased sixfold.

The increased size, connectivity, and volume of capital flowing through the financial system are not intrinsically bad—in fact, in healthier circumstances, they ensured that capital could flow where (and when) it was most needed, generating improved returns and keeping the system liquid. This connectivity was celebrated in the precrash era as a tool for dispersing, not concentrating risks.

The market's densely connected configuration, much like the Internet's, ensured that randomly shutting down one of the innumerable banks in the global system would not be likely to cause systemic problems, because, statistically, the vast majority of banks in the network are at the end of spokes connected to a very limited number of hubs. But flip one of those central hubs (a rare and dangerous occurrence), and you might not only take down the thousands of banks directly connected to it, but the other hubs as well, along with thousands of institutions connected to *them*. Like a game of Jenga, pull a random plank and almost certainly nothing will happen. Pull the wrong one, and the entire robust-yet-fragile edifice comes crashing down.

And that's just what happened: For several decades, a string of moderate and even serious failures was capably handled by the financial system, from the dot-com collapse to oil shocks—each one increasing confidence in the ability of the system as a whole to manage disruption. Then a hub failed. Panic spread throughout the network, and things quickly ground to a halt. As Levin says, "It's not that the hubs were too big to fail, but too *interconnected* to be allowed to fail [italics added]." When calamity struck, there was no way to separate or decouple them from one another.

These interconnections in the financial system didn't just take the normal form of funds flowing between banks to cover daily activities. They were anchored by new, sophisticated debt and insurance derivatives: collateralized debt obligations (CDOs) were the securities that allowed banks to cut up, repackage, and then sell one another the

debt from risky U.S. home mortgages; credit default swaps (CDSs) were the insurance contracts that tied these banks to one another in massive webs of financial interdependence. Together, they were the financial equivalent of nitric acid and glycerin—in small amounts they keep your heart pumping; in large amounts: *boom*.

FINANCIAL CLUSTER BOMBS

A CDO is a financial instrument that can best be understood by imagining a set of wineglasses stacked in a pyramid. When champagne is poured over the pyramid, the glass at the top is filled first, the glasses in the middle are filled next, and those at the bottom are filled last. A CDO was an equivalent financial instrument, but instead of disbursing wine into wineglasses, it poured the monies from mortgage payments into a set of specialized bonds.

To create a CDO, a bank would bundle together a group of mortgages, held by regular U.S. homeowners. Each month, as these homeowners wrote their monthly mortgage checks, the banks would pool these payments together and make payouts to a series of bonds called *tranches* that were stacked up just like the wineglasses. The tranche at the top of the chain, like the glass at the top of the pyramid, got paid first, then the one following it, and so on, until either the tranche at the bottom was paid or the pool of funds was exhausted.

By definition, the top tranche, at the front of the line to be repaid, was the least risky, so it earned both the best rating (AAA) from ratings agencies like Moody's and Standard & Poor's and also the lowest rate of return, perhaps 2 percent. The bottom tranche, on the other hand, was the most risky: If mortgage holders in the CDO pool stopped paying their mortgages, the loss would be felt in that tranche first—so it earned both the lowest rating (say BB) and the highest rate of return, perhaps 10 percent.

So far, so good. But from there, CDO engineering quickly entered into the twilight zone. A bank might, for example, take the lowest-rated tranche (BB) of a particular CDO (let's call it Lucifer) and turn it into its *own* CDO (let's call this bottom tranche Damien). Even though it was comprised of junk, Damien, through the magic of financial engineering, was divided into its own set of tranches, the top one of which was awarded its own triple-A rating. This absurdity masked the fact that the underlying asset upon which it was based—Lucifer's bottom-most BB

tranche—was toxic, high-risk junk. It was like saying “here is the safest house that you can build with this toxic waste” and conveniently forgetting to say the toxic-waste part. Or, if you prefer, like taking a dozen hobbled horses from the glue factory, lining them up according to their speed, and calling the fastest one a Thoroughbred.

When some homeowners originally bundled together under Lucifer stopped being able to pay their mortgages, not only did its BB tranche bondholders lose out, so did those holding *all* of Damien’s tranches—even the ones rated AAA. And there’s the true tragedy: Those holding Damien’s triple-A’s were well-managed, risk-averse pension funds, municipalities, and 401(k) plans.

If CDOs were the financial equivalent of chloresterol-laden junk food mislabeled as health food, the credit default swaps were the mechanism by which the banks ensured that a blockage in any artery would give a heart attack to all. CDSs are insurance contracts, much like the contract that you might sign to insure a car, a home, or your own life. In such arrangements, you make a monthly payment to the insurance company, and, in the event of disaster, it makes you whole. Similarly, these insurance contracts enabled a bank to insure against the loss in value of a stock or bond it might purchase. If Bank A buys $10 million in corporate bonds, and it loses $2 million in value, then the insuring Bank B would make up the difference. In the meantime, Bank A would pay a premium on this insurance to Bank B, much as you pay your health and car insurance premiums.

CDSs had several crucial differences from traditional insurance, however. First, a CDS contract could be traded from investor to investor with no oversight or even regulations ensuring that the insurer had the ability to cover the losses when and if it needed to. By calling the contract a swap, and not insurance, the investors in CDSs were able to avoid the capital reserve requirements and regulatory oversight of the traditional insurance industry. (This was presumably the “innovation” at CDSs’ core.)

Second, CDSs allow firms not only to insure against the possible default of their own investments, but to insure against the possible default of another company’s assets—akin to taking out an insurance policy on your neighbor’s Ferrari. Firms could use swaps as a tool for speculation—to bet a company would fail. This practice was made illegal in traditional insurance markets as far back as the 1700s, before

which it was legal for individuals to buy insurance on British ships that they didn't own, creating—*quelle surprise*—a rash of perfectly seaworthy ships mysteriously sinking to the bottom of the Thames. In its place, Parliament codified the notion of "insurable interest," the requirement that you have an actual economic interest in the asset being insured. It was a concept that reigned unchallenged for two and half centuries—until the rise of the CDS.

Finally, CDS contracts were sold and traded privately, or "over the counter." While they added enormously to the risk profile of the institution doing the insuring, they didn't show up on the traditional balance sheet. When the crisis came, nobody knew who owed what to whom and what it meant for anyone's bottom line.

In theory, CDOs and CDSs were originally designed to allow the market to do two things that are quite beneficial: first, to distribute risks to those who were most capable and willing to take them, and second, to allow banks to diversify their portfolios by mixing and matching some of their own activities with one another's. A typical big bank might find itself originating its own mortgages, buying them from others, selling mortgage-backed securities that blended the two, and insuring those of another bank against default. All of this looked, superficially at least, like diversification—a wise strategy for balancing efficiency and robustness.

But, in allowing debt, credit, and risk to be sliced into tranches, packaged, bought, repackaged, sold, and resold, these instruments also made the dependencies between institutions mind-bogglingly complicated. The chain of custody for the underlying assets lengthened to the point of incomprehensibility.

Thus, when the crisis hit, none of the banks could be quite sure if the *other* institutions with which they had contracts might *also* be enmeshed in *other* contracts that left those institutions on the hook in some potentially catastrophic way. This is known as the problem of *counterparty risk*—not the risk that you'll become a deadbeat, or the risk that your partners will become deadbeats, but the risk that some of your partners' partners will become deadbeats.

In the low-risk, precrash world, from an individual bank's perspective, having a contract with another bank—which itself had many other business relationships—was perceived as a good thing: It was a measure of diversification and suggested lower risk. After all,

what's the likelihood of a significant number of contracts in a large portfolio defaulting all at the same time? In the high-risk, postcrash world, however, having a contract with another bank with lots of counterparties was nightmarish: After all, it might be holding a contract with an unknown third party that might, at any moment, blow up in its face. How could you *possibly* determine how creditworthy the other bank was? How could you (or it) know what its obligations were? Or those of its counterparties' counterparties? How could you trust *anybody*?

When the crash came, a sense of transparency disappeared overnight, and with it went the most important variable in the system, trust—a theme we will revisit later in this book.

It didn't help that the derivative contracts themselves were mind-bogglingly complicated. The simplest of these ran some two hundred pages—the most advanced varieties required reading in excess of one *billion* pages. (Reading one page a minute, it would take you slightly more than 1,900 years to read a single contract for one of these products.) Firms had foregone the due diligence and just swallowed the contracts whole. After the crash, figuring out who owed what to whom wasn't just hard, it was impossible. It's unsurprising that the institutions that vaporized amid the destruction—Lehman, Bear Stearns, and AIG Financial Products—had among the largest counterparty exposure. They were attached to an anchor of unknowable size.

The rise of these complex derivatives had another, subtler impact as well. Their promise—of higher returns and dramatically lower risk—proved so intoxicating to financial organizations of all stripes that these firms effectively homogenized their methods of generating revenue and their risk management strategies to take advantage of them. Many different kinds of actors in the market—from commercial banks to hedge funds—started getting into one another's businesses, holding the same classes of assets and liabilities, in the same proportions, seeking out the same goal, in the same way. As each market player began to internalize more and more of the ever-increasing complexity of the market as a whole, market risks became internal risks. This was diversification without a difference. Right under everyone's noses, the complex ecosystem of the global financial markets was turning into a monoculture—of very complicated lemmings.

THE TELLTALE SIGNS OF A SYSTEM FLIP

According to Sugihara, in a complex system, there are telltale warning signs of a critical transition, or system flip, and they were visible in the run-up to the financial crisis. One is a phenomenon called *critical slowing*—the tendency of a system to become unstable near its threshold point. “When a system is under stress, it can be thrown out of equilibrium more easily, and it is slower to recover. Without sufficient recovery time, small perturbations can be amplified until the system is oscillating wildly out of control—even squealing from one stable state to another—like a car being oversteered on an icy patch.”

Just before such destabilization occurs, a system paradoxically may experience synchrony, as agents within it briefly behave in lockstep just before being thrown into chaos. Synchrony can be seen in the brain cells of epileptics, for example, minutes before the onset of a seizure, and it was evident in the financial markets prior to the crash. By the height of the credit boom, from 2004 to 2007, performance across sectors of the financial system was correlated by more than 90 percent, a reflection of the self-similarity of the various market participants. “This was clearly an early indicator of impending danger,” says Sugihara.

“Synchrony is evident when incentives or pressures lead individual actors to fall into step and make similar choices,” adds Levin. “In unsynchronized populations, some individuals thrive while others are in decline; in synchronized populations, a collapse in one place translates into a collapse in all places.”

In a healthy financial system, interconnectivity, risk management, diversification, and product innovation typically act as shock absorbers, dispersing risks and mitigating the impact of inevitable failures. But in the lead-up to the crash, the diversity of the market was slowly eroded, the actions of the market players became synchronized, and the dependencies between them became unknowable. Then, as with the collapse of Jamaica’s sea urchins, the collapse of Lehman introduced a virus of unprecedented lethality and speed that decimated the financial system’s most critical resource: trust. When that happened, the system’s shock absorbers were hijacked and turned into shock amplifiers—spreading the contagion of uncertainty rather than the perception of safety.

In response, the bankers did what seemed rational in such an uncertain situation: They hoarded cash and desperately tried to sell the

depressed assets they had on their books. Yet most of the banks had followed similar business strategies before the crash, and so their most rational responses to it were also almost identical. In a densely connected network filled with clones, both responses, taken en masse, made the situation worse for everyone. Hoarding cash caused a liquidity crisis that made it harder for all banks—including the hoarders—to meet day-to-day obligations. The mass sell-off of depressed assets further accelerated the decline in value of the remaining assets on everyone's books, harming the balance sheets of healthier banks and pulling more institutions into the vortex. In this light, it's easy to see why, even with hundreds of billions in bailouts from the U.S. Treasury, banks were so reluctant to start lending again.

The global financial markets may share various dynamics with robust-yet-fragile ecosystems, but what does it tell us, if anything, about how to avert or lessen the next crisis? After all, doesn't the RYF nature of the market mean that the risk of future calamity is built into the equation?

Increasingly that question is being asked not only by ecofinance theorists like Sugihara and Levin, but by leaders with more formal influence over global financial policy, including an unexpected voice in the world of risk management, working in one of the most traditional financial institutions in the world: Andrew Haldane, the executive director of financial stability for the Bank of England.

Like Levin and Sugihara, Haldane attributes the problems with the precrash financial order to the complexity of the system—the hornet's nest of interconnections between institutions—and the homogeneity of those institutions' business strategies. And his recipe for improving the resilience of the financial network bears striking resemblance to ecologists' prescriptions for ecosystems. "We need more complete, holistic measures of the health of the financial system and the dependencies between various institutions within it; we need to improve communications about it with the public at large in times of impending crisis or system flip, and we need to take steps to improve the financial system's biodiversity."

SEEING THE FINANCIAL WHOLE

Creating more holistic, EBFM-like measures of health for the global financial network begins with collecting a great deal more data. "Today,

risk measurement in financial systems is atomistic—each institution, each node, is judged by itself," says Haldane. "After a systemic failure, this leaves policymakers—and indeed everyone—navigating in a dense fog."

To clear the air, regulators and governments need to be able to see the full number, size, and type of connections, flows, and dependencies between various institutions and markets at a glance. They need financial observatories to continuously replicate the Fedwire study, but on a much larger, more comprehensive, more timely, and more international scale.

Doing so will require building distributed sensor networks for the global financial system and bringing many more transactions into the light of day. This is already starting to happen. In the United States, for example, the package of financial reforms signed into law after the 2008 crash requires that derivatives like credit default swaps be cleared through mechanisms designed to mitigate the problem of counterparty risk. They call these mechanisms centralized counterparty clearinghouses, or CCPs.

In the precrash era, each bank dealt directly and independently with every other, so when the crash came one bank could never be sure if its counterparties had signed more contracts than they could afford to pay. The clearinghouse addresses this murkiness by standing in the middle of every trade, becoming the buyer to each seller and then turning around and becoming the seller to each buyer. If one of the parties goes bankrupt, the CCP covers its side of the contract.

In theory, this achieves several things at once. First, it makes mapping the complicated web of relationships between financial parties much easier, as the clearinghouse sees every relationship between every buyer and every seller in its domain. "Clearinghouses compress the highly dimensional web of financial obligations to a sequence of bilateral relationships with the central counterparty—a simple hub and spoke network. The lengthy chain of relationships is condensed to a single link," says Haldane. Thus, in a crisis, a bank would not have to worry whom its business partners might also have dealings with—it would have dealings only with the clearinghouse. "Provided that link is secure—the hub's resilience is beyond question—counterparty uncertainty is effectively eliminated," he adds.

CCPs also simplify and cut down the thicket of counterbalancing claims between parties; if the First National Bank owes the Second National Bank $50 million, spread across a dozen derivatives contracts, and Second National owes First National a similar amount over a different set of contracts, the CCP can offset the claims through a process called netting. (In the wake of the financial crash, efforts were undertaken to do this manually in the CDS market by “tearing up” redundant, offsetting claims, reducing the volume of contracts by more than 75 percent. A centralized clearing facility would make this process automatic.)

Centralized counterparty clearinghouses also provide a mechanism to ensure that each party can actually pay out its contracts, because the clearinghouse requires each party to set aside sufficient capital to fulfill the contract as if it had to be settled each day—commonly known as marked-to-market accounting. Additionally, clearinghouses make it far less likely that any particular firm will become dangerously overexposed—as AIG and Lehman did—without anyone being able to tell, since the clearinghouse is in a position to see such buildups as they occur.

In addition to using central clearinghouses, other steps toward greater transparency have also recently been mandated in the United States, including the use of an open exchange for such derivatives, akin to a stock exchange. The primary benefit of exchange-based trading is in setting a market price. In the precrash era of direct deals between banks, nobody knew what anyone else was paying for these contracts. By moving to a market-based system, everyone can see what the going rate is, and companies are far less likely to get gouged.

Not everyone is happy with such an arrangement. Making derivatives trades more transparent in turn makes the underlying investment strategies of some investment firms more visible, and many of those strategies have previously depended for their success on being kept private. And Wall Street traders dislike exchange-based trading for another, simpler reason: When nobody knew what anyone else was paying for derivative contracts, they could charge whatever they wanted. The pricing transparency that comes with selling such contracts via exchanges naturally erodes their profits.

Nor is interposing a central counterparty clearinghouse in derivatives trading a panacea—it does not entirely eliminate the fragility in the

financial network, but rather relocates it to the risk-management strategies of the central clearinghouses themselves. If managed well, a CCP can effectively mitigate the risk for a whole market; if not, its very centralized position could give it the starring role the next crisis.

But a centralized counterparty clearinghouse does enable the collection of vastly more information and greater transparency about the activities of the market. This data can be combined with many other sources to inform sophisticated, system-wide measurement tools that, much like EBFM, reveal the connectivity of institutions, not just their size and behavior.

In public health, practitioners often use informal social network mapping to identify and then inoculate super spreaders, those most at risk for initiating or propagating a contagion. So too in finance. “In early 2007, it’s doubtful whether many of the world’s largest financial institutions were more than two or three degrees of separation from AIG,” says Haldane. “And in 1998, it’s unlikely that many of the world’s largest banks were more than one or two degrees of separation from Long-Term Capital Management, originators of the last major crisis. Mapping the links in the financial network might have identified these financial black holes before they swallowed too many planets.”

Yet troublingly, prior to the crash, there was virtually no correlation between the size of the most important financial institutions and how much they held in reserve to cover potential problems. In spite of—or more likely because of—their supposed importance, these critical mega institutions seemed to have held less in reserve as a percentage of assets than their smaller peers. “One explanation is that, because big banks thought that they were diversifying, they thought they could afford greater risk,” says Haldane. “Another explanation is markets allowed these banks to operate less conservatively because of the implicit promise of government support if things turned for the worse. Either way, there was no targeted vaccination of the super spreaders of financial contagion.”

The best way to determine the appropriate levels of inoculation for such super spreaders is to model crises before they occur. This was undertaken in limited form in 2009 when the U.S. Federal Reserve began stress-testing banks that received bailout assistance, to make sure they had sufficient capital to weather difficult conditions such as sustained high levels of unemployment or ongoing home mortgage

defaults. Stress testing gave a snapshot of each bank's health and risk exposure—and in the first round of testing, regulators discovered that more than half of the banks tested required additional capital. But even this limited stress testing is only a once-a-year affair, like an audit, and doesn't analyze how the system as a whole might be impacted if any of its subjects were to collapse. Such tests and simulations are regular occurrences for other network systems like electrical utilities, the military, and the air transportation system—and will need to become a regular feature of the financial system. And not just in a few market hubs, but everywhere.

Still, these are compensatory measures. They contain and minimize the damage of a systemic fragility once it has surfaced, but they can't flip the system back into its preferred state. Is there anything that can help the financial markets better self-regulate? The coral reefs may—once again—hold an intriguing clue.

THE BATFISH AND THE WIR

David Bellwood is obsessed with fish. His friends constantly poke fun at his ability to bring any conversation back around to them. When he was younger, growing up in northern England, he started an aquarium to observe reef fish. The interest propelled him all the way through an undergraduate degree at the University of Bath, where he did an honors project on the aquarium industry and then, later, to a PhD and professorship at James Cook University in Australia, focusing for more than twenty years on what is now his formal area of expertise and informal object of obsession: the parrotfish.

It's easy to see why Bellwood would fixate on this beautiful colored fish. The parrotfish is one of the mainstays of the coral reef system, and it can do some genuinely interesting things: Not only can it change gender—female parrotfish can transform themselves into males when their dominant male leader dies—but certain species of parrotfish have developed the ability to envelop themselves in a transparent cocoon, made from a viscous substance that comes out of an organ in their heads. This homemade sleeping bag disguises the scent of the parrotfish at night, leaving it safely hidden from nocturnal predators.

Of far greater interest to Bellwood, however, is the parrotfish's function on the reef: to consume large quantities of algae extracted from gnawed chunks of coral. Parrotfish are voracious algae eaters—

Bellwood refers to them as “lawn mowers”—and on a healthy reef they play a critical role in regulating the competitive relationship between algae and corals. By cleaning the system of excess algae and enabling new corals to compete for space, parrotfish keep a coral-dominated reef in a continually self-regenerating state and prevent it from flipping into an algae-dominated state. Particularly after the difficult lessons of Jamaica, protecting the parrotfish and other herbivores has become a centerpiece of many reef resilience strategies.

Yet in recent research conducted on the Great Barrier Reef, Bellwood has discovered that the mechanisms that regulate the reef are more complex than marine ecologists had previously thought. His discoveries have significant implications for future reef management, and for other fields as well.

In controlled experiments, Bellwood and other researchers set up large open cages on top of the reefs—each one the size of a small office. In each of the cages they planted dense macroalgae assays, which simulated a reef system fully dominated by algae—what might be described as the parrotfish equivalent of an all-you-can-eat buffet.

Bellwood and his team set up multiple underwater cameras to film what came next. “We thought there was going to be a spectacular circus because we had forty-five species of herbivorous fish in the vicinity, and huge densities of parrotfish. This particular location should have brought a feeding frenzy to the algae.”

The team sat waiting in anticipation for the first few hours. That anticipation soon turned to bafflement. The only thing they could see on the video footage was a thick cluster of algae waving in the current. Not only were there no parrotfish, there were no fish to be seen at all.

“Lights! Camera! Action! And then . . . nothing. Were the fish reproducing or gone somewhere else? I mean, this was three-meter-high, beautiful algae. These are herbivores. These fish should have been having a feast.”

What the team did notice, over the next few weeks, is that the algae were slowly getting *thinner*. “You know how people get older and you can slowly start to see their scalp? Well, the algae were like that. You still had the bits that were reaching to the top but it was getting really thin in the middle. Eventually, there was very, very little on the footage.”

Bellwood was expecting the algae to be gobbled down in a frenzy, primarily by parrotfish, in twenty-four hours. Instead, over a period of

three weeks, the algae were slowly nibbled and eroded until they finally just collapsed—but what was eating them? The results were so curious that Bellwood and the other scientists didn't even have a framework to understand them. All they could do was go back and review the footage again and again.

Then, during one of their repeated viewings, they spotted something altogether unexpected. From out of the murk, a completely unrelated species, a small black fish with a golden fringe (called a pinnate batfish) appeared on the footage. And then, to everyone's shock, it slowly started feeding on the algae.

"We were stunned for two reasons. To begin with, batfish aren't supposed to eat algae. They are an invertebrate feeder, not an herbivore. Secondly, we could never catch them doing it. As soon as we got in the water, they would swim away. It was like the *Far Side* cartoon where the cows are having a conversation and then the car comes and they all suddenly go back to eating the grass."

Bellwood's research suggests that while the parrotfish act as lawnmowers for the reef, they can do so only when the reef is in the healthy, coral-dominated state. When the system has flipped and algae have taken over, they're no longer able to provide this function. And that's when the batfish—which normally doesn't eat algae—is "deployed" on the reef to correct the imbalance. One prevents a flip, the other reverses it.

Bellwood likens the reef to a golf course. "Under ordinary circumstances, the equipment you need is a lawn mower because it keeps the nice grass down. The coral reef lawn mowers are the parrotfish, surgeonfish, and other species that graze the turf and keep the greens tidy. But if, for some reason, your mowers are broken, the weeds on the golf course start to get big, and soon there is a back garden overgrowth. Now, by the time the lawn mowers show up, they won't work. You need a chainsaw and a bush saw. Surprisingly enough, we discovered that this is how the batfish functions."

The batfish are part of what Bellwood calls a "sleeping functional group," a species or group of species capable of performing a particular functional role—but which do so only under exceptional circumstances.

Bellwood's team was both euphoric and confused by the discovery of these sleeper groups. The good news is that fish exist with the

capability to help reverse a system flip from an algae-dominated state back to a coral-dominated state. The bad news, however, is that in many cases, scientists have no idea which fish they are. Because the resilience function of these species emerges only in exceptional circumstances, no one before Bellwood had ever connected the batfish with coral-algal interactions in more than fifty years of research.

What might it look like for the financial system to have such sleeper functional groups of its own—countercyclical strategies lying dormant within the financial network, stirring only when the system flips, as it did in 2008? Finance may just have found its batfish in tiny Switzerland, in the form of a unique alternative currency called the WIR.

• • •

If you have traveled around Switzerland at any point in the last few decades, you may have noticed stickers in the windows of local shops and businesses: “We Accept WIR.” Perhaps you caught sight of one of WIR’s glossy catalogs displaying a full list of all of its participating businesses. If you stayed in a hotel, the clerk may have even asked you if you wanted to pay through WIR. No doubt she or he soon rescinded the offer upon realizing that you were not a native of Switzerland.

What is this mysterious alternative currency and how does it work?

Bernard Lietaer, a Belgian economist with more than twenty-five years of expertise in currency systems, has been a keen observer of this Swiss alternative currency for many years now. He is not alone. A number of macroeconomic analysts are starting to look more carefully at the way it functions.

WIR—*Wirtschafstring,* or “circle” in German—began in the depths of the Great Depression. In the wake of the stock market crash of 1929, total world trade plummeted—by 20 percent in 1930, another 29 percent in 1931, and another 32 percent in 1932. Unemployment reached 30 million. Wealth that was once all but assured disappeared into thin air, and banks that seemed certain to stand suddenly collapsed, as stocks lost nearly 90 percent of their value.

Switzerland was pulled into the crisis more slowly than some of the other European countries, and it was slower to recover. By 1934, when the United States and Germany were showing faint signs of recovery, Switzerland remained in a malaise. The number of bankruptcies hit

record heights and trade and tourism took deep hits. The Swiss railroad, SBB, carried a deficit double that of the federal budget. Unemployment was the norm, not the exception: In 1934, there was one job opening for every seventy-three official job seekers.

In the midst of all this, sixteen businessmen decided to create a solution for themselves. All of them, along with their clients, had been informed by their banks that their credit lines were no longer open to them; without credit, their businesses were positioned on the brink of bankruptcy. Rather than fail, they decided to set up a complementary form of currency.

Their result, the WIR, is a mutual credit system. A debt in WIR is either reimbursed by bartering in sales with someone else in the network or paid in full in the national currency. Over time, this network expanded to include one-quarter of all the businesses in the country. Today it is a thriving barter network that makes up a well-recognized complementary currency in Switzerland.

An analysis of more than sixty years' worth of data on the WIR by macroeconomist James Stodder has shown that whenever there has been a recession, the volume of business in this unofficial currency has expanded significantly, cushioning the negative impact of lost sales and increased unemployment. Whenever there has been a boom, business in national currency has boomed, and activity in the unofficial currency has dropped proportionally again. In the past, people have attributed the success of the Swiss economy to a national character of pragmatism and thrift. Stodder's study offered unexpected proof that the secret behind the country's legendary stability and economic resilience is the spontaneous countercyclical behavior of this small alternative currency system.

WIR functions precisely like Bellwood's batfish—a latent, just-in-case contingency system that is activated when the economy is at or near a phase shift. Lietaer is an advocate for more complementary currencies that function like the WIR—business to business—in the European Union and the United States. "The substance that circulates in our global economic network—money—is maintained as a monopoly of a single type of currency (bank-debt money, created with interest). Imagine a planetary ecosystem where only one single type of plant or animal is tolerated and officially maintained and where any

manifestation of diversity is eradicated as an inappropriate ‘competitor’ because it is believed it would reduce the efficiency of the whole.”

• • •

The batfish and the WIR illustrate, in very different ways, an essential strategy of many resilient systems: They are embedded countercyclical structures that can respond proportionally, and in the same time signature, to disruptions as they emerge. These structures are part of an essential inventory of diverse tools often found in resilient systems: The batfish is an example of biodiversity, the WIR, of economic diversity. Like all inventory, this diversity imposes a carrying cost. And because such structures are latent in a system—they are only called upon when a crisis emerges—it can be difficult to place a value on them when things are in a more humdrum state of affairs. In a relentless quest for greater systemic efficiency, this diversity can be lost and unmourned—until it’s too late.

As we’ve seen, the fragility and resilience of most systems begins with their structure. The complexity, concentration, and homogeneity of a system can amplify its fragility; the right kinds of simplicity, localism, and diversity can amplify its resilience. The lens of resilience suggests, for example, that what’s needed is a smaller, simpler, more accountable, and more decouplable financial system, with genuinely diverse participants, that is more closely aligned with its original purpose—providing liquidity to organizations and individuals—than with its more recent one, which seemed to be engineering wealth from thin air.

But resilience takes more than just the right structure—it takes the right kinds of processes and practices: measuring the whole health of the system, like EBFM; modeling and stress-testing the system, like the Fedwire study; scanning for emerging disruptions and mobilizing the right, inclusive responses when they strike, like the proposed financial observatories; building in feedback and compensatory systems, like the batfish and the WIR, that help keep the system in check. Most of all, it requires keeping the system dynamic and reconfigurable. As we’ll see next, it is in such dynamism that many systems’ resilience rests.

2

SENSE, SCALE, SWARM

Resilience isn't just found in systems we admire, but sometimes in systems we loathe. Terror networks and even many diseases can be thought of as perversely resilient—they persist and even thrive under sustained and powerful assaults intended to root them out and eliminate them. How do they do it? And how might their strategies for doing so be applied in more positive contexts?

As we'll see, the answers are often rooted in the dynamics of these kinds of entities. For example, terror networks and many microbial infections survive, in part, by modulating their "metabolism" up and down—reducing their operations to near dormancy for long periods, then, when they sense the time is right, scaling up to attack in highly coordinated swarms. Applied in very different circumstances, these sensing, scaling, and swarming tactics hold promise for improving the resilience of systems on which civilization more positively depends.

OPERATION HEMORRHAGE

In November 2010, an unusual document began to circulate in a few less savory corners of the Internet—a digital, full-color magazine that, at a casual glance, might easily have been mistaken for an industry trade publication. And, in a manner of speaking, it was: This was the third issue of *Inspire*, the English-language magazine produced by al-Qaeda in the Arabian Peninsula (AQAP), the branch of the global terror network largely based in Yemen.

Produced by a radicalized American convert to al-Qaeda named Samir Khan, issues of *Inspire* are filled with a mix of propaganda aimed at young American and British citizens: advertisements for jihad, home-brew recipes for terrorist activities (sample headline: "Make a Bomb in the Kitchen of Your Mom"), and fawning profiles of al-Qaeda

leadership. The magazine's writing was naïve, youthful, and earnest, with countercultural flourishes intended to speak to disaffected Western teenagers in a language they understand.

Issue 3 of *Inspire* had a slightly different focus, however, as signaled by its cover, a blurred photo of a UPS cargo plane emblazoned with a bold, singular headline: "$4200."

The dollar figure was a reference to the total cost of an AQAP terrorist plot, entitled Operation Hemorrhage, which was intended to blow up two cargo planes—one operated by UPS, the other by FedEx—carrying packages bound for the United States. As described in exhaustive detail in the magazine's ensuing pages, AQAP bomb makers had carefully hidden high explosives inside two hollowed-out desktop printer cartridges; the cartridges had been connected to detonators and circuit boards lifted from mobile phones, then repackaged with their original printers in a way that made them appear factory sealed. These were then shipped by AQAP operatives from Sanaa, Yemen's capital, addressed to liberal Jewish synagogues in Chicago, President Obama's hometown.

The fictional names of the bomb's recipients, Reynald Krak and Diego Diaz, were derivations of historical figures from the Crusades and the Spanish Inquisition. In a further literary grace note, the bombers also enclosed a copy of Charles Dickens's *Great Expectations*—a reference both to their optimism about the plot's chances of success and a coded message of sorts: Anwar al-Awlaki, AQAP's chief theoretician and religious leader, had become a Dickens fan while in prison in Yemen, when he had been banned from reading Islamic texts. (Al-Awlaki has since been killed in a drone strike, as has Kahn.)

These rhetorical flourishes were hardly the main point of the plot, however. As the issue of *Inspire* goes to lengths to explain, Operation Hemorrhage had two central aims: one strategic and one economic. First, the bombs would act as a sensor: They would pass through the latest air cargo security equipment, providing AQAP with good intelligence regarding the West's explosives detection capabilities. Second, the bombs would be a provocation: The ensuing security fears would force the West to invest billions of dollars in new security procedures:

> From the start our objective was economic. . . . The air freight is a multi-billion dollar industry. . . . For the trade between North America and Europe air cargo is indispensable and to be able to force the West to install stringent security measures sufficient enough to stop our explosive devices would add a heavy economic burden to an already faltering economy.

A few months before, AQAP had prototyped its attack, deploying a bomb of similar design to bring down a UPS cargo plane leaving Dubai International Airport, killing the two pilots aboard. Counterterrorism officials had not registered AQAP's involvement at the time, instead attributing the crash to mechanical failure.

This time, however, Saudi Arabia's intelligence agency caught wind of the attack just as it was coming together and alerted officials in the United Kingdom and Dubai that the packages were being shipped by air from Yemen to the United States by way of Europe. Security officials seized one of the two bombs en route at the East Midlands Airport, in Nottingham, UK, on a UPS plane stopping over on a flight from Cologne to Chicago; the other was found at a FedEx facility in Dubai. Although not independently confirmed, one of the two bombs was defused just seventeen minutes before it was set to explode, according to France's interior minister.

Yet despite the operation's apparent failure, AQAP had deemed the initiative a great success:

> Two Nokia mobiles, $150 each, two HP printers, $300 each, plus shipping, transportation and other miscellaneous expenses add up to a total bill of $4,200. That is all what Operation Hemorrhage cost us. In terms of time it took us three months to plan and execute the operation from beginning to end. On the other hand this supposedly "foiled plot," as some of our enemies would like to call, will without a doubt cost America and other Western countries billions of dollars in new security measures. That is what we call leverage.

Leverage: In the logic of post-9/11 terrorism, it's priceless. AQAP didn't even have to blow up a plane to achieve its aims—all it needed to do

was provoke a costly response.

Indeed, Operation Hemorrhage's value as a media activity may have even outweighed its value as a terrorism activity. In open sourcing and publicizing its methods, AQAP advertised both its own imperviousness and the relative vulnerability of its chosen enemy, all the while encouraging others to emulate its approach. At the same time, it also reinforced its own metanarrative about the conflict and its role within it—engaging in a battle to control the story. That's where those rhetorical flourishes, like the Dickens novel, came in: They're the sorts of sticky, surprising details that stand out and pass virally from person to person, further drawing in potential recruits. Evidence of these tactics' potency already exists: When nine young men in the United Kingdom were arrested in December 2010 and charged with plotting to blow up the London Stock Exchange, Big Ben, and the U.S. embassy, authorities found them carrying two issues of *Inspire* in their pockets.

• • •

In the decade since 9/11, terror networks like AQAP have been on the receiving end of the largest sustained military eradication efforts in modern history. A conventional military force a thousand times its size would have long since been defeated. During the same period, however, even with the recent death of Osama bin Laden, al-Awlaki, and other figureheads, al-Qaeda-affiliated or -inspired chapters have spread to more than sixty countries around the world. Rooted in its networked structure, its ability to modulate its operational metabolism, and its ability to swarm, al-Qaeda's success holds powerful lessons for designing positive forms of systemic resilience in domains far beyond terrorism.

For several decades, Dr. John Arquilla, professor of defense analysis at the U.S. Naval Postgraduate School, has been one of the world's leading thinkers on the resilience tactics of groups like al-Qaeda and the real-world dynamics of netwar, a term he coined in the 1990s with colleague David Ronfeldt to describe what they predicted would become the twenty-first century's dominant mode of conflict.

"Organizational forms profoundly affect a variety of endeavors, whether peaceful or warlike," says Arquilla. "Assembly line–driven organizations, for example, manufacture using certain processes, and their management structures come to echo those processes. The

1950s-era corporate management hierarchy and the military's current command-and-control model are both mirrors of their institutions' basic operations.

"About two decades ago, it became clear that networks were going to enable new organizational structures on the battlefield, and indeed a whole new way of warfare. This would have enormous implications for both the structure and tempo of military operations, because the way you and your enemy organize determines the kind of war you fight."

The first important element of al-Qaeda's success has been the rise of its network franchise model, in which al-Qaeda has become less and less a formal organization and more a global organizing principle and open-source brand. This brand has proved alluring to local groups with very different but overlapping objectives. In exchange for affiliation, such groups became part of a larger movement that amplifies local perceptions of their global power and prestige. Thus, the Salafist Group for Preaching and Combat, based in Algeria, which had previously been focused on overthrowing that government, became al-Qaeda in the Islamic Maghreb; a violent Islamist splinter cell in Indonesia became Al Qaeda in the Malay Archipelago. In addition to providing the larger network with global reach, these various groups also specialized, based on their history, makeup, and unique skill sets. AQAP, for example, emerged as a leading recruiter of Americans to join the global jihad movement—*Inspire* being just one manifestation of their efforts.

This loosely affiliating, self-organizing, networked, and self-specializing dynamic is replicated inside terrorist networks and other asymmetric fighting forces as well as between them. In such arrangements, small groups are not bound together via traditionally strong command-and-control structures, but by ad-hoc, redundant, and informal social connections—less like the Marines and more like a pickup game of basketball. The small scale of the groups within such networks helps them remain agile, while the many-to-many ties in the larger network ensure that even if 10 percent or 20 percent of its membership is eliminated, the network as a whole will continue to function. "How many times have we killed number three in al-Qaeda? In a network, *everyone* is number three," notes Arquilla, dryly.

Terror networks succeed by redefining battlefield success downward—their goal is obviously not to win in traditional military terms, but to embarrass and exhaust their enemies. Battling a traditionally far

superior force to a draw counts as victory. To achieve this, terror groups balance long periods of near dormancy with intermittent high-energy bursts of activity, like Operation Hemorrhage. The purpose of such operations is not merely to wound the enemy, but to incite an ineffective response that will make him appear all the more impotent.

This was a central dynamic in the 2006 Lebanon War, when somewhere between one and two thousand Hezbollah fighters held their own against more than a hundred thousand Israeli troops and a relentless air campaign. Notes Arquilla: “On the first day of the war, Hezbollah launched two hundred rockets; on the last day of the war they launched two hundred rockets. Most Israelis and most people around the world feel that Hezbollah won the war, just by hanging in to the end. And how did they win it? By dividing up into several hundred little firing teams of three to four individuals, marching about, firing off hidden Katyushas, and melting back into the scenery. We call this strategy ‘shoot and scoot.’ And it’s the defining structure of the conflicts in Afghanistan, Iraq—increasingly, everywhere.”

When not fighting, very little can distinguish unconventional fighters or terrorists from local civilians—particularly where they share strong local, cultural, religious, and even family ties—making them incredibly challenging to detect and eliminate. This ability to scale down rapidly has other benefits as well: It reduces the group’s support requirements, it preserves their future fighting capacity, and it increases the likelihood that their more traditionally organized opposing force might accidentally kill or otherwise offend innocent civilians—mistakes that can be disastrous to the larger effort to frame the narrative of the conflict.

Yet, though dangerous, terrorist cells must at least occasionally “delurk” to complete missions like Operation Hemorrhage. Doing so is extremely dangerous, to be sure, but such operations serve as necessary advertisements for the group’s potency and validate its relevance to potential recruits. More important, they serve to keep the enemy off balance, which is critical to the terrorists’ preferred method of conflict. “By keeping its chosen enemies sporadically and schizophrenically engaged, the global Al Qaeda network hopes we’ll ‘punch ourselves out,’ much like Muhammad Ali did to George Foreman,” says Arquilla.

In the United States, politicians often argue about why al-Qaeda has not staged another 9/11-style spectacular attack on U.S. soil. It may be

that, having induced the United States into a decade-long, trillion-dollar conflict with it and its allies, on its terms, it simply doesn't need to. Our continued, imperfect engagement suits its long-term, rope-a-dope purposes perfectly.

To improve the effectiveness of the attacks they do undertake, terror networks have increasingly adopted a new offensive technique that Arquilla calls "swarming." In this model, small highly distributed teams simultaneously attack nonmilitary targets, overwhelming defenses—if there are any—originally designed to deter a single large aggressor. "In the 2008 Mumbai bombing, just ten attackers in five two-man teams struck at the same time—five different places in one city. They completely overwhelmed a very wealthy and pretty well militarized country—India—and held a major world city hostage for two days, killed two hundred people, and caused untold disruption upon that society. That's the template of what's to come."

LEARNING FROM TB

In their ability to rapidly scale up and down, terror networks bear surprising similarities to another highly resilient but altogether different phenomenon: the pathology of tuberculosis infection in the human body. These similarities suggest both lessons for designing resilient systems and ways of contending with perversely resilient phenomena.

TB is one of the most prevalent diseases on Earth. One in three people carry TB antibodies, meaning they have been exposed at some time in their lifetime, but only a small fraction—about 10 percent of those exposed—will ever become actively sick with the illness. Still, TB's toll is staggering: Every day, about 4,700 people, mostly poor Africans and Asians, die from the disease. After HIV, it remains the most common infectious cause of death in adults worldwide. New infections occur at the rate of about one per second—two new ones occurred while you were reading this sentence.

TB bacteria enter the body through the air, carried in invisible, microscopic droplets to the alveoli, the tiny air sacs of the lungs, where they take up residence. Within approximately six weeks, a newly infected person will develop a small infection, called a primary infection, rarely accompanied by any symptoms.

At a cellular level, however, something truly remarkable is happening. When the TB bacteria arrive in the lung, they are met by

macrophages (literally “big eaters”), the white blood cells at the front lines of the human immune system that are normally responsible for consuming and destroying invading pathogens. In certain cases, macrophages find this task difficult to do, so in a last, kamikaze-like move, they will engulf an unknown invader, coat it in the cellular equivalent of Saran Wrap, and then, on cue, promptly die, taking the pathogen with them. However, in TB, occasionally just the opposite occurs: TB takes over some macrophages, preventing them from dying, and turns them into zombie-like incubators for producing more TB bacteria, which slowly replicate inside until they burst their cellular hosts open and spread to others.

This process isn’t perfect—many TB bacteria are eliminated by untainted macrophages, which sense the distorted chemical signature of their TB-infected counterparts and can turn on them. In more than 90 percent of infected people, the hosts’ immune response is just effective enough to control, but not entirely eliminate, the TB infection, and the disease enters a long period of latency, which can last for years.

“This long latency period is part of the reason TB is so insidious,” says microbiologist Sarah Fortune, a TB researcher at Harvard. “In its latent state, TB doesn’t do very much, which is why our traditional tools against it, antibiotics, are not terribly effective. After all, antibiotics target metabolism, and with latent TB, there isn’t much being metabolized.”

Paradoxically, TB’s long latency is also one of the things that make it so deadly. It allows TB’s human hosts to grow up and reproduce, creating a new generation of people for the disease to spread to. “TB is like a symbiont that occasionally kills you,” says Fortune. “In contrast, illnesses like Ebola are far rarer, precisely because they become deadly so quickly. In a matter of weeks, all the human hosts are gone, and they take the outbreak with them.”

While the precise biomolecular mechanics are still not perfectly understood, the current model of TB’s pathology suggests that during the latent phase of the disease, TB bacteria constantly probe the immune system for weaknesses, much like the AQAP terror cell probing the global security system. “Our reasoned assumption is that during the latent phase, some TB is active, much is inert, and then, some critical threshold is passed, often in an immune system

compromised by other illness such as HIV, alcoholism, or diabetes, and the illness enters its active phase," says Fortune.

At that point, everything changes. In the active phase, the TB infection is aggressively reactivated throughout the body, though the lungs are the preferred location for the illness to strike. Once there, in an eerie parallel to terror networks, the aggressive TB bacteria depend on provoking an overreaction from the immune system in order to spread.

This overreaction is expressed as a caseous granuloma, which forms when a ball of TB bacteria and TB-infected macrophages collect in the lungs, where they are attacked by healthy immune cells. Like cops surrounding a building with a gangster inside, these cells employ a standard immunological strategy: They try to wall off the TB within a fibrous shell.

Tiny granulomas are a normal part of the immune response and form throughout the latent phase of the disease. When they form during the active, aggressive phase, however, things can quickly go haywire. This time, when immune cells surround the ball of TB cells, they die, in turn attracting even more immune cells, which attempt to wall them in. Then *they* die, in a process that repeats over and over again, inflating the granuloma like a balloon.

Inside this expanding shell, the TB bacteria sit amid clusters of now-dead immune cell tissues. Eventually, such granulomas (called "caseous" because of the cottage cheese–like consistency of their dead centers) can grow to be the size of a tennis ball inside the lung. And as they grow, the necrotic material at their core begins to liquefy, forming the perfect environment for TB bacteria to multiply. And the TB does just that: A modest 2-centimeter granuloma alone can contain 100 million active bacteria.

As the granuloma grows, it irritates the lining of the lungs, eventually bursting and spraying the inside of the lungs with the bacteriological equivalent of a swarm attack, hitting many targets at once. "TB depends for its survival on inducing this tactical error on the part of the immune system," says Fortune. "The effort to contain the infection serves only to concentrate and then amplify it."

After that, every time the actively infected person coughs, he or she can release untold millions of invisible TB bacteria into the air. An average infected person can cough up to thirteen infectious doses an

hour; a severely contagious one, up to sixty—before, in about half of all untreated active cases, they finally pass away.

EMBRACING THE SWARM

Both TB and terror networks maintain their resilience amid sustained attack by moderating their metabolism, scaling down to near dormancy for long periods of time, and scaling up to strike when the time is right. Both depend on mechanisms for probing their target's responsiveness and dynamically reorganizing at the right moment. Both spread by provoking an overreaction from their host targets. And both succeed by swarming in coordinated, simultaneous attacks. These similarities are starting to suggest new, biologically inspired metaphors for tackling terror networks.

For Arquilla, the answer to tackling the threat of terror groups is to absorb and mimic their tactics: to meet the network with a network. "The War on Terror is actually the first great war in which we're seeing nations at war with networks. The United States and its allies initially approached things in very traditional ways: massive deployments, overwhelming force, and shock and awe." Like an incomplete course of broad-spectrum antibiotics, these approaches were sufficient for clearing the battlefield but not for holding it. And as they failed to do so, they have slowly, painfully, been replaced with lighter, more agile, targeted and networked approaches.

In an American corollary to Hezbollah's shoot and scoot model, Arquilla describes what finally came to work in the U.S. military's campaign against networks in Iraq, a strategy he calls "outpost and outreach": "Several years into the war, over much resistance, we moved our soldiers from huge bases to hundreds of small outposts thirty to fifty soldiers apiece, dramatically expanding the number of nodes in our physical network." This quickened response times, allowing U.S. troops to respond speedily, flexibly, and proportionally to a threat or problem, within minutes instead of hours. It also encouraged real relationships with local residents, creating connections that generated huge amounts of actionable intelligence.

Complementing this physical reorganization was an outreach and social networking effort that is widely regarded as yielding the greatest tactical success of the war: nurturing the Sunni Awakening movement, a coalition of Sunni tribes that were induced to stop fighting the

Americans and start fighting al-Qaeda's local branch instead. "The combination of the physical network and the social network—that's really what turned the campaign around on the ground. And we couldn't have gotten that to happen without effectively engaging in a battle for the story, with our own counternarrative."

Supported by the success of outpost and outreach, U.S. military forces have also started to embrace swarm tactics of their own, aimed at excising local terror groups whole. Like the slow, persistent efforts used to take down narcotics and Mafia networks, these counterterrorism and counterinsurgency operations are predicated on long wait-and-watch intelligence-gathering efforts, which gradually illuminate latent ties in the terrorist's social network. Then, when as much of the network has been revealed as possible, agile military teams take it down in coordinated, swarming strikes.

Like the use of targeted antibiotics to tackle TB, the timing, speed, simultaneity, and precision of these strikes is important. Because many-to-many networks are naturally resilient, normal military kill or capture strategies are too slow to ever be effective—the loss of individual members is simply routed around by those remaining, like packets routed around downed computers on the Internet.

"This presents a huge and as yet unfinished cultural change for military forces, which have an inherent doctrinal preference for rapid destruction," says Arquilla. "We're slowly realizing that to win, you have to preserve as much information as possible, for as long as possible, before striking. And that means not blowing things up, at least until the last possible minute." In netwar, as in immunology, the timing of force is often more important than its scope.

These efforts, though effective, are still a retrofit of existing military practices. What would it look like for U.S. forces to embrace fully the self-organizing dynamics of its networked adversaries? Arquilla suggests the military may one day adopt the market-driven dynamics of companies like eBay: Instead of issuing pages and pages of operational orders, a commanding general could post tactical objectives to a website and assign points values to them: 100 points for this bridge; 500 for the capture of this town or that enemy combatant.

"The general's various units would log on, see what's listed, and bid for whatever is there. If there was something the general wanted done that wasn't getting done, he could raise the point values. Or lower them.

And, at the end of the day, if there was something that still wasn't done, he might have to give an actual order to someone. He could still do that. But I believe we'd see self-synchronizing campaigns that are faster, more adaptive, and more effective."

And what would the fighting force look like that wages such campaigns? Arquilla envisions an army of swarmers. "I would rest easy if twenty years from now, instead of ten army divisions, we had a hundred clustered small units," he says. Like its enemies, this military would structure itself to scale up and down more quickly. "We have more than two million people right now; I think we could take it down to under four hundred thousand in the active force and have a larger reserve, but just that: a reserve. We could achieve all of this at about half the cost that we're paying today, if we take networks seriously. It's not rocket science. It's network science."

Just as the U.S. military is finding increased resilience and tactical success in embracing some features of its more agile networked adversaries, other big systems are being similarly transformed by these highly distributed, scalable, and networked approaches. Nowhere is this being felt more than in the slow reengineering of the United States' massive electrical grid, and in the planning for the future system that will one day replace it.

THE GRID THAT BREATHES

Referred to by the National Academy of Engineering as "the supreme engineering achievement of the 20th century," the modern grid powering North America is larger and more ubiquitous than the domestic Internet, and just as complex. In the modern grid, electricity generated from all forms pours in from power plants onto an interconnected highway of weblike regional grids that serve areas hundreds or thousands of miles across. This system of systems winds its way from one end of the United States to the other, while crisscrossing back and forth through Canada and Mexico.

Because electrical current follows a path of least resistance and cannot be stored, its journey is not like a car idling at a stoplight, but more like water pouring through a network of gated sluices at light speed. (Most electricity is consumed less than a second after it has been generated.) If there is a traffic jam on one transmission line, electricity is rerouted to its final destination in ways that can be

surprisingly circuitous. When the lights are turned on for dinner in Portland, Oregon, for example, some of the electricity often whizzes from Los Angeles and then takes a detour through Utah before lighting the bulbs up north. When the TV is powered up for *Sunday Night Football* on the East Coast, the electricity might travel, at the speed of light, from Canada and down through Ohio and even Virginia before linking back up to deliver the game in upstate New York.

The system moving all these electrons around is not just technological. It is characterized by tight couplings of human organizations and businesses, a system so big, more than 20 percent of all electrical infrastructure purchases on Earth are used just to keep the North American grid operating. Take a moment to consider the system's vast breadth: According to a 2005 analysis by the Global Environment Fund and the Center for Global Smart Energy, more than 3,200 electric utilities and 2,000 independent power producers serve 120 million residential household customers, 16 million commercial customers, and 700,000 industrial customers. They use the continent's 700,000 miles of high-voltage transmission lines, owned by about two hundred different organizations and valued at more than $160 billion. In turn, more than 5 million miles of medium-voltage distribution lines and 22,000 substations, owned by more than 3,200 organizations and valued at $140 billion, plug into these high-voltage superhighways.

Now catch your breath. Did you take all of that in? Probably not. The system is too complex for most people to comprehend in its entirety. And just describing the connections between the nodes on this vast network captures only one aspect of its performance. For example, the system breathes with the seasons: The electrical conductivity of the grid's physical components—and therefore how much energy needs to be transmitted—changes in hot and cold weather, as their materials sag or contract, like a diaphragm.

The thousands of human beings who make sure the lights stay on have to constantly broker the mission-critical need of humanity for electricity, an electrical production and distribution system of mind-boggling complexity, and the weather. Doing so is an art akin to musical performance, requiring an understanding of the time of day, the day of the year, the geography of the region, major social events (like a rock concert at a local arena), and countless other variables that have nothing to do with the equipment itself.

To make matters worse, a low-probability, high-impact wild-card event like a heat wave or a cold snap can send electrical demand soaring instantly, overwhelming power lines. Similarly, mechanical failure like a downed transmission line or a lightning strike can cause a surge that, if not disposed of correctly by either automatic systems or human operators, can threaten vast regions of the grid itself.

And that's just what happened on August 14, 2003, shortly after 4:00 p.m. Hot summer temperatures in the northern and central regions of the United States prompted more and more people to crank up their air conditioners and fans. In the heat of the summer, several local power lines in northern Ohio started to expand and sag, brushing up against overgrown trees that needed trimming, creating a short circuit and shutting down the power in a cascade of sparks.

On an average day, it would have been a very low-level disturbance, setting off an alarm in a local utility's control room and alerting one of the local operators to the disturbance. Under normal circumstances, that operator would simply reroute the power flows to avoid the site of injury, and a crew of linemen would be dispatched to make the physical repair.

August 14, 2003, however, was not an average day. In addition to the shorted power lines, the alarm software that would have typically notified the local operator of the problem also failed. Unaware of the damage, all of the other operators wheeling power across the regional grid continued to route electricity through the damaged area, forcing the transmission lines to bear an untenable amount of current. Stressed well beyond their capacity, within two hours of the initial short circuit all of the power lines in Ohio cut out entirely.

Just like other robust-yet-fragile systems described earlier, the power outage in Ohio spread from a small triggering event to a systemwide collapse, tripping circuit breakers and sending more and more plants and lines offline. The grid operators were caught completely by surprise and, equipped with outdated and incompatible monitoring equipment, were unable even to see what was happening. Within eight minutes, the blackout affected an estimated 10 million people in Ontario and 45 million people in eight U.S. states. Bigger than 1965, much bigger than 1977, the blackout of 2003 was the largest power outage in North American history. As day turned to night, air and train travel on the East Coast ground to a halt. Cleveland declared a

curfew on all persons under the age of eighteen. In some affected urban areas, the Milky Way became visible for the first time in decades.

All this, triggered by a tree branch and some by-now-familiar factors: systemic complexity, lack of transparency, and a lack of interoperability.

TOWARD A SWARMING GRID

The electrical grid has always had at best incomplete self-monitoring capabilities. In the late 1920s and early 1930s, early analog computers were built using electric components to reproduce miniature electric systems. Each of these early devices—big enough to fill a living room—would simulate the impact of disturbances to the grid. By the 1960s, however, the new digital computers allowed for a more sophisticated system. Most power plants and transmission lines switched over to something called SCADA, an acronym that stands for "supervisory control and data acquisition." These systems collect data from various sensors at a plant or in the lines and then send it off to a central computer that manages and controls it.

Like many protocols, the interconnectedness that came from using such a standard initially resulted in greater efficiency. But over time, as the physical grid and its associated technologies grew in sophistication, SCADA systems proved incapable of capturing the fuller state of the grid and had the effect of slowly eroding, not supporting, its resilience.

According to Phillip Schewe, with the American Institute of Physics, this kind of technology is too old, too outdated for today's challenges. It "does not sense or control nearly enough of the components around the grid. And although it enables some coordination of transmission among utilities, that process is extremely sluggish, much of it still based on telephone calls between human operators at utility control centers, especially during emergencies."

Sometimes even these telephone calls between human operators prove useless because neighboring utilities can use incompatible control protocols. A collection of transcripts gathered on the day of the 2003 blackout illustrates the lack of interoperability and real-time information between the system managers, as in this example, between operators of the Pennsylvania–New Jersey–Maryland and Midwest Independent System Operator regional grids:

PJM: It looks like they lost South Canton–Star 345 line. I was wondering if you can verify flows on the Sammis-Star line for me.

MISO: Well, let's see what I've got. I know that First Energy lost their Juniper line, too.

PJM: Did they?

MISO: And they recently have got that under control here.

PJM: And when did that trip? That might have . . .

MISO: I don't know yet . . .

PJM: And right now I am seeing AEP systems saying Sammis to Star is at 1378 . . .

MISO: Let me see. I have got to try and find it here, if it is possible . . . I see South Canton–Star is open, but now we are getting data of 1199, and I am wondering if it just came after.

PJM: Maybe it did.

The operators were reduced to managing the "greatest engineering triumph of the 20th century" using technology invented in the nineteenth—the landline telephone. Their situational awareness was incomplete, murky, and based on old information.

Simply patching these problems wouldn't make the grid any more secure: Every day, the grid gets larger, more complex, and is tasked with doing things its inventors could scarcely have imagined. And a darker concern lurks: While this blackout has been an accident, what if the next time the malfunction was intentional?

Working on behalf of the power companies, hackers had already discovered security vulnerabilities in SCADA systems: In one simulation, they remotely hacked into the controller typically used by grid operators, surreptitiously recording all of the outbound instructions sent to various generators and monitoring the responses. Then they simply switched the instructions normally sent during the day, when demand is at its peak, and those sent at night, when much of the system is idling, succeeding in crashing the (thankfully simulated) grid in a matter of minutes. A modified version of this approach has now been confirmed in the wild: On November 8, 2011, hackers are suspected of remotely gaining access to a water pump's controller

connected to the Springfield, Illinois, municipal water supply. By sending commands to rapidly turn the pump on and off, they burned it out, destroying it in the process.

Even without the threat of cyberterrorism, the likelihood of repeats of the 2003 blackout have been increasing, driven by the ever-increasing demand for larger power transfers over longer distances, underinvestment in the transmission system, wild-card spikes in demand, and the financial consolidation of power utilities. Then there's the anticipated wave of coming failures among aging electrical transformers, installed in the 1950s, '60s, and '70s and now rapidly approaching the end of their useful lives. Throw into the mix the increasing frequency of climate-change-induced anomalous weather events and the fact that U.S. drivers will soon be plugging hybrid electric vehicles into the grid by the millions, and it becomes clear: If we are ever to ensure that the grid meets the growing needs of the twenty-first century, consistently, safely, and securely, the grid will need to be transformed—but how?

Enter engineers like Massoud Amin, widely considered the father of the smart grid. Amin, a stately Iranian immigrant in his fifties, is now a revered engineering professor at the University of Minnesota. After 9/11, when much of official Washington was concerned about all manner of terrorist attack, Amin was charged with overseeing research into critical infrastructure security and grid operations at the Electric Power Research Institute, a leading Palo Alto–based think tank working on the future of the grid. There, Amin and his team drew from a dizzying array of breakthroughs in fields as diverse as nonlinear dynamical systems, artificial intelligence, game theory, and network theory to develop key concepts for a self-monitoring, self-healing, self-repairing grid.

Amin and his fellow researchers identified three by-now familiar design principles of this grid of the future. The first is real-time monitoring and reaction. In much the same way as economists like Andy Haldane call for the collection of better real-time data about the performance of the global economic network, the grid needs vastly more sensors deployed at every level of its organization, from the core to the periphery. The second principle is anticipation. Current SCADA systems assess isolated bits of information with a thirty-second delay, analogous to driving a car by looking in the rearview mirror. Better monitoring and anticipation would create what Amin refers to, only half

jokingly, as “self-consciousness” in the system: Advanced controllers would perform like a master chess player, modeling supply and demand several steps ahead in a chained link of events. The third design principle is isolation, or decoupling. At the first sign of failure, the grid would break itself down into islands, or isolated entities. Each island would be vulnerable to complete failure, certainly, but the systemwide cascading failures that ultimately led to the blackout of 2003 could be avoided.

Today, as these principles are being implemented, the power grid is undergoing a revolution. Attracted by a financial opportunity that may dwarf that of the Internet, an entire industry has emerged to implement these ideas, at every scale and level of organization, from the individual light switch to the grid as a whole. According to an analysis done by the networking firm Cisco, the resulting technology platform will be one hundred to one thousand times larger than the U.S. domestic Internet, with many trillions of nodes and sensors embedded throughout and millions of low-level software systems to connect them to one another.

Together, these embedded sensors and intelligence systems will imbue the grid with a sense of *proprioception*. Not quite Amin’s self-awareness, “proprioception” refers specifically to “one’s own” perception, or to the positioning of one’s own body in space. Try a little experiment: Shut your eyes and bring both hands above your head. While keeping the fingers of your left hand totally still, use the index fingertip of your right hand to quickly touch your nose. Then, while keeping your eyes closed, use the same right index finger to touch the thumb of your left hand.

For many of us, mastering such a simple task will take a few tries, but we eventually get there—but how? Although we get vital information about our bodies in relationship to space from our sense of sight, we also have a sense of our body’s positioning—where our legs and arms are, for example, at any given moment—even without looking. Special receptors, called *proprioceptors* or *stretch receptors*, are located throughout our bodies and continually relay positional information back to our brains. Our brains, in turn, take stock of the incoming data, reconcile it with what our eyes and other sensory organs tell us, and create a composite, internal sensation of our own body’s orientation in relation to the space around us, even if we can’t physically see our limbs and torso at the moment.

The sensors Amin envisions distributed throughout the grid will go a long way toward imbuing it with an analogous sense. Integrating the signals from these sensors will require something else, however: new protocols for encapsulating and exchanging information inside the system.

As the issues with SCADA systems illustrate, the resilience of systems like the electrical grid is highly sensitive to the design of the information protocols at their core. Protocols are the lingua franca of systems—they define how information is exchanged among their constituent parts. If the underlying protocol is inflexible, or tied too closely to a particular set of hardware and software, it can quickly become obsolete, as the underlying technology ages and is replaced. That's exactly what's happened with SCADA systems; they embody command-and-control principles from the mainframe era.

On the other hand, the Internet's protocol, called TCP/IP, shows how to do it at least partially right. This packet-focused protocol, which we described in the last chapter, has a simple, ingenious, and layered design. To simplify things, imagine an hourglass shape: At the bottom of the hourglass is an ever-expanding menagerie of hardware, from iPods to robots to who-knows-what-next, which connect to the Internet. At the top is a mirror-image, ever-expanding menagerie of software applications, from mobile apps to the future descendant of Twitter, which also connect to the network. In between is the skinny waist of the hourglass: a simple protocol—changing rarely, if ever—for encapsulating information for transport across the network. Write a new application, or create a new Internet-enabled appliance, and if it can speak TCP/IP, it can talk to anything else on the network.

In this way, TCP/IP is like the alphabet—a series of basic forms that can be used to give expression to and make connections between a countless number of ideas, but which itself changes only glacially. When the protocol's authors first invented TCP/IP, most could scarcely imagine a world of iPads and Facebook profiles, but the genius of their system is that they didn't have to—they simply had to ensure that whatever was invented could speak the universal language, and that the universal language could, in turn, transmit whatever they might like to say to one another. (A SCADA system, by analogy, fixes the relationships between the letters, enabling a more limited palette of information and ideas to be exchanged.)

FROM MACRO TO MICRO: SCALE AND SWARM

Embedding sensor networks and the right kinds of extensible protocols into the grid will update an infrastructure originally conceived in the nineteenth century and designed for the twentieth with the management tools of the twenty-first. But the real resilience revolution will arrive when new structural features are added to the grid. The first of these innovations are microgrids—tiny, autonomous, self-sufficient, distributed systems that use locally available energy sources to power a small footprint of a few buildings, homes, or factories.

When the North American power grid was first engineered, the process of making electricity was dirty, dangerous, unsightly, and expensive, and so electricity was generated in large, centralized power plants and shipped to customers who were, at first, dozens, then hundreds, and then thousands of miles away. This is the equivalent of the mainframe era in computing: We all share the resource of the power plant, in much the same way that early computer programmers shared the resources of a room-sized computer.

To drive costs down, and in order to ensure a ready supply of fuel, power plants typically need to be situated near specific kinds of infrastructure, like railroad tracks and rivers. And for both safety and aesthetics, plants ideally need to be some distance from the people who use the energy they produce. As a result, today it's virtually impossible to look at a plug in your home and know where its electricity comes from.

In a grid complemented with microgrids, this situation is inverted. At least some of the energy on the grid is efficiently produced locally, using small-scale wind, solar, and hydropower, for example, and therefore doesn't need be shipped over long distances. In principal, its users need never be connected to a system in which distant failures can adversely affect them.

Such energy would need to be produced cleanly and safely—after all, nobody wants to live with a coal power plant in their backyard, particularly when something goes wrong. And of course the result would have to generate power cheaply—it's not enough to build a grid that delivers all the above but is exorbitant to operate.

If these cost and cleanliness problems can be solved, however, an autonomous, small-scale grid could power every city block, every

home, and perhaps every person individually. The existing grid would be atomized—infinitely replicated in miniature.

No one is more interested in the concept of microgrids than the U.S. military, and it's easy to understand why. In the swarming, networked world of Arquilla's outpost and outreach, there is no front line for troops to be behind—the front line around a forward operating base extends 360 degrees. And yet these bases must be continuously supplied with fuel and matériel, typically via convoys that have to snake their way through extremely hostile territory. These lightly protected caravans of transport vehicles are slow moving and vulnerable: A 2009 study found that U.S. military forces in Afghanistan incur one casualty for every twenty-four fuel resupply convoys. The fully burdened cost of gasoline for bases in-country is a whopping four hundred dollars a gallon; at the peak of the conflict, the Marines ran through 200,000 gallons *a day*.

If their bases could be made self-sufficient, fighting forces would be able to deploy much faster and retain greater autonomy once they had done so. Under an initiative called Net Zero, the U.S. Army is focusing heavily on microgrids that use a wide variety of wind, solar, fuel cells, and other renewable sources to make this possible. Marines in Iraq and Afghanistan are already experimenting with Power Shade, a large tarp covered with flexible solar panels that fits over a standard tent and can quietly power its lighting system without the buzz of a generator that might alert enemy insurgents, and GREENS, or Ground Renewable Expeditionary ENergy System, a larger solar system that can power a platoon's command center. One proposal even suggested that the U.S. forces should encourage Afghan farmers to switch from growing poppies to growing biofuel crops—cutting down on the drug trade, building an indigenous new industry, locally supplying U.S. forces, and undermining the Taliban all at the same time.

In the civilian economy, investors and innovators are betting on microgrids as a major component of the future of energy production—companies like GE and IBM already predict that more than half of American homes will generate some of their own power in the next few decades. As envisioned, microgrids will do more than just distribute power generation: They normalize the relationship that energy consumers have with energy producers. Right now, power travels just one way: from the power station to the grid to the consumer. If microgrids ever reach their apotheosis, consumers will upload their extra reserves of power just as often—if not more frequently—as they

download power from some central source. An Internet of energy would, much like its namesake, be a platform for both individual production and consumption. In this swarm version of the grid, no centralized unit controls the system. Instead, semi-autonomous little units augment the greater whole: a system of systems.

At the far edge of this vision, the renowned MIT chemist Dan Nocera is working on technologies that one day may vaporize the grid altogether—and usher in an era of personal energy—by imbuing the electrical system with a capacity even microgrids lack: storage.

In summer of 2008, Nocera set the energy and chemistry worlds abuzz with a startling announcement: His research team had achieved partial artificial photosynthesis—making fuel directly out of sunlight—with a new chemical catalyst that split water into hydrogen and oxygen. In so doing, they had solved one of the greatest challenges to the large-scale deployment of distributed solar energy: how to store the energy of the sun when it isn't shining. Connect to a box filled with this catalyst, and you can store the energy produced by even a moderately efficient solar cell during the day for intensive use at night.

"In the simplest configuration," Nocera explained to us, "a photovoltaic on your roof will generate the power you need to live when the sun shines. The surplus electricity from the photovoltaic can be fed to a small box filled with our newly discovered water-splitting catalyst to generate hydrogen and oxygen, which are stored locally. At night (or when the sun isn't out), the stored hydrogen and oxygen can be recombined in a fuel cell to give the electricity needed to power your home at night and charge the electric car in your garage while you sleep."

In Nocera's dream of personalized energy, each individual home becomes a solar power grid and a gas station. Because the system is a closed loop—sun and water to hydrogen and oxygen—very little water is actually used up. How much? Nocera's favorite measurement gauge is the pool on the MIT campus where he teaches.

"We need only one-third the amount of water in the MIT pool—an Olympic-size swimming pool—to be converted into hydrogen and oxygen every second, to solve the entire world's energy problems. A lot of new discovery, such as our catalyst, is needed to make this happen, but think of it: only one-third an Olympic-size swimming pool."

Nocera is now the head of the distributed energy storage company Sun Catalytix, where he is working to commercialize and scale his artificial photosynthesis technology. While there's much work to be done, Nocera is confident the dawn of personalized energy—and with it the beginning of genuine energy abundance—could be as little as a decade away. Already, he's formed a partnership with Tata, the Indian superconglomerate, to produce a stand-alone mini power plant that could reinvent rural electricity supply and bring power to the 3 billion people on Earth who don't have it. In such a hypothetical postgrid age, in which we don't subscribe to energy but produce it ourselves, it's not that much greater a leap to imagine a company like Walmart becoming the largest consumer energy company in the world—by selling us packets of energy, or the appliances that generate our own.

Whether such longer-term visions comes to pass, and whether Nocera's technology is ultimately the one to power them, the bottom-up microgrid approach will certainly come to complement the top-down, retrofitted smart grid infrastructure in the years to come. The two visions represent different strategies for endowing the larger electrical system with greater resilience: The latter focuses on imbuing centralized legacy systems with intelligence, the former with imbuing the periphery with greater autonomy and self-sufficiency.

Tying the two together will be newly designed protocols—"skinny waist" replacements for SCADA systems that look much more like the Internet's open TCP/IP protocol (and may indeed be some variant of TCP/IP itself, given its widespread adoption and support). These new protocols will drive greater interoperability and ensure that different local power generation technologies—including those as yet unimagined—can more easily plug and play on the grid. In this way, the smart grid merges aspects of both the energy and information infrastructures.

Embracing these approaches will substantially boost the electrical system's efficiency. The Electric Power Research Institute estimates that the fully realized smart grid will save U.S. consumers $20 billion on their annual energy bill and eliminate roughly 10 percent of national carbon emissions—akin to removing 140 million cars from the roads. Smart grid proponents refer to this reservoir of potential efficiency gains as the "fifth fuel," complementing coal, petroleum, nuclear, and renewables.

Yet the path will not be straightforward. The barriers to the adoption of the smart grid may be as much behavioral as they are technological. In early experiments bringing Amin's design principles of real-time monitoring and anticipation to the part of the grid that touches regular people's lives, in the form of residential smart meters, utilities have run headlong into an inconvenient truth: Many people simply *hate* them.

As originally envisioned, smart meters are an update on the residential electricity meters often found in the basements of people's homes—placed there precisely because they would be out of sight and out of mind. Rather than being read by a human meter reader once a month, smart meters continuously and wirelessly report their energy usage to the customers, allowing them to see their energy use patterns in real time. They also allow the electrical utility to dynamically change the price for electricity—to raise rates when demand is high and lower it when demand is low, pricing signals that are intended to encourage people to change their behavior.

In theory, this signaling should allow customers to trade changes in behavior for lower electrical bills, and it should help utilities better manage supply and demand, smoothing peaks and valleys. A more stable grid is, in principle, not only more efficient and more cost effective for consumers, it's one onto which utilities can more reliably plug in more renewable sources, lowering carbon emissions—a huge potential win-win.

But aligning supply and demand requires people to do something new: to start carefully attending to *when* they use electricity by making choices about at what time, say, their refrigerator de-ices, or time shifting to nonpeak hours when they do a load of laundry. A less variable load on the grid requires more variability among consumers—by getting some of them to use electricity countercyclically. The combination of both financial incentives and penalties designed to encourage this, combined with the real-time nature of the system, have made some early customers of smart meters feel like they're being constantly watched by Big Brother and charged for making incorrect decisions.

Unfortunately, in some rollouts of the first generation of smart meters, utilities did a nonexistent-to-terrible job of explaining the system, spurring vocal opposition groups to form that challenged the meters on grounds ranging from privacy and cost to health and safety.

“I used to think PG&E was just incompetent,” said one early and involuntary recipient of Pacific Gas and Electric’s smart meters, after the company decided to roll them out concurrently with a price increase. “Now I think they’re incompetent and *spying on me*.”

To counter such hostile reactions, many utilities are turning not to additional technology, but to a mix of behavioral economics, clear communications, and, of all things, the lowly smiley face. They’re doing so by embracing a consumer-focused energy conservation platform developed by a company called Opower.

While searching for a business plan for their new energy-focused start-up, Opower’s founders, Dan Yates and Alex Laskey, realized that even if electricity consumers were interested in conserving energy (something most say they support), the task was usually rendered impossible by the indecipherable, poorly targeted cookie-cutter bills they received from their utility. Here was a classic example of a poorly designed feedback loop—bills filled with irrelevant, nonactionable, and untimely information that gave consumers little sense of how their behaviors affected the system and no incentives for change.

Then, in a chance meeting, Yates and Laskey were introduced to the work of behavioral scientist Robert Cialdini, a renowned emeritus marketing professor from Arizona State University, who had spent a thirty-year career studying the science of persuasion, in particular the role of social norms in getting people to change their behavior. Cialdini had authored hugely popular books on persuasion and marketing and enhanced his scholarly research with covert fieldwork, enlisting for stints as a used car salesman, telemarketer, and other jobs that depended on persuading strangers.

One of Cialdini’s central areas of research was *social proof*—the concept that people are influenced more by their immediate neighbors’ behavior than by any other form of reward.

In one classic study of the power of social proof, Cialdini and his colleagues conducted an experiment with the ubiquitous towel reuse programs found in hotels, where guests are asked to keep their towels from being laundered in order to conserve water. In the bathrooms of a hotel in the American Southwest, they randomly placed three variations of the standard towel reuse sign, one focused on environmental protection, one focused on cooperation with the hotel, and a third focused on the behavior of other guests:

HELP SAVE THE ENVIRONMENT

You can show your respect for nature and help save the environment by reusing your towels during your stay.

PARTNER WITH US TO HELP SAVE THE ENVIRONMENT

In exchange for your participation in this program, we at the hotel will donate a percentage of the energy savings to a nonprofit environmental protection organization. The environment deserves our combined efforts. You can join us by reusing your towels during your stay.

JOIN YOUR FELLOW GUESTS IN HELPING TO SAVE THE ENVIRONMENT

Almost 75% of guests who are asked to participate in our new resource savings program do help by using their towels more than once. You can join your fellow guests to help save the environment by reusing your towels during your stay.

The first and second signs generated almost identical responses—38 percent and 36 percent compliance rates, respectively—but the last was significantly more effective: Almost one in two of the guests (48 percent) reused their towels. (To ensure consistency, Cialdini's team checked only rooms with single guests, and only on the first night of their stay.) Here was concrete proof that, no matter what people said they believed in, actual conservation behavior was strongly based on one's beliefs about what others around you were doing.

Based on the success of studies like this one, Cialdini, who became Opower's chief scientist, agreed to help Yates and Laskey build a system that uses sophisticated data analytics and clear, usable communications, informed by social proof, to encourage consumers to conserve electricity. To do that, the company first collects vast datasets on consumer electricity usage from the utility, combines them with data about the weather and other events, and then applies sophisticated algorithms that disaggregate the portion of a consumer's electrical usage that is related to heating and cooling, one-time events, and the running of appliances like refrigerators that are on all the time.

Then Opower uses this data to deliver a customized energy report, written in plain English, to the utility's customers that gives them a sense of how they're doing—not in the abstract, but in comparison to a tiny pool of about one hundred of their closest neighbors living in similarly sized homes, along with some concrete, customized steps they can do to improve.

Customers who are using energy more efficiently than the composite average of their neighbors are rewarded with a smiley face on their report; those in the top 20 percent are rewarded with two smileys.

And that's it. No bonus offers. No financial incentives. Just the right information, modest incentives, concrete recommendations, lots of follow-up . . . and smiley faces.

The social reward centers of the brain are so exquisitely sensitive to social norms that that's all it takes. Opower claims that 85 percent of households which receive its reports, and the various communications which follow them, modestly cut their power consumption—evidence that the normative behavioral approach works across age, income levels, and level of education. Independently, utilities that have already deployed the system report initial average energy savings from 2 to 3 percent among their customers—at a tiny fraction of the cost of doing things like retrofitting buildings or changing light bulbs. At their current scale (30 million households and rising, as of this writing), the company claims it will soon help customers save more energy than is produced by the entire U.S. solar industry.

The larger point is this: For all of the capacities that a new swarming, dynamic, smart, and resilient grid enables, people are still the linchpin to its long-term resilience. A grid that does not include, empower, and properly motivate them to participate is hardly smart at all.

• • •

Several themes emerge from this discussion of resilience dynamics in action. Build a system with modular component parts, arranged in networks and connected by open, "skinny waist" protocols; imbue the parts with distributed intelligence; and give people the right information and incentives, and you can create the circumstances in which resilience can flourish. The system can begin to sense its own state, sense the state of its environment, anticipate disruption, dynamically

modulate its scale in response, and localize or decouple its operations when needed.

Such systems embrace an ethic of decentralization and shared control—swarming—so that no single entity is absolutely in charge. But neither are they utterly anarchic. They do not do away with all centralized authority, but instead balance it with the right kinds of local empowerment and self-sufficiency.

Nor do these systems achieve their equilibrium in a maximally dispersed fog of connections—in fact, quite the opposite. As we'll see next, the pattern that emerges in many resilient systems is a kind of distributed, diverse density that finds its closest analogy in the city, the reef, and the garden. This is the principle of clustering.

9

BRINGING RESILIENCE HOME

In our tour of resilience patterns, we have crossed a range of geographies, disciplines, contexts, and ideas. We've seen some of the ways systems can become fragile and break. We've examined how their connectedness can create feedback loops that (often invisibly) either amplify or erode their resilience. We've explored how resilience strategies from one domain might find application in another. And we've glimpsed how individuals, groups, and communities can bolster their resilience by embracing a warm zone of connectedness, collaboration, and diversity.

It turns out that Goldilocks had it right all along: Resilience is often found in having just the *right* amounts of these properties—being connected, but not too connected; being diverse, but not too diverse; being able to couple with other systems when it helps, but also being able to decouple from them when it hurts. The picture that emerges is one of strategic looseness, an intentional stance of both fluidity (of strategies, structures, and actions) and fixedness (of values and purpose).

The problem, of course, is that many of the aspects of a given system's resilience are defiantly context specific. The particular approach that makes one organization more resilient in a given situation may make another more fragile. (Note the word "more"; there are no absolutes in resilience, no binaries, just measures of more and less.) Every resilient solution is rooted in its setting and not necessarily a sure path to others' success.

How, then, do we put these principles into action?

MAPPING FRAGILITIES, THRESHOLDS, AND FEEDBACK LOOPS

While there's no single recipe for every circumstance, every journey toward greater resilience begins with continuous, inclusive, and honest efforts to seek out fragilities, thresholds, and feedback loops of a system—grasping its holistic nature, identifying its potential sources of vulnerability, determining the directionality of its feedback loops, mapping its critical thresholds, and understanding, as best we can, the consequences of breaching them. Doing so calls us to greater mindfulness—assessment, without judgment, of the world as it truly is. And this is just as true for organizations and communities as it is for people.

As we've seen, fragilities can take many forms: chronic, long-term challenges like persistent poverty; or the erosion of social mobility; or increased susceptibility to environmental shocks; or the vulnerability of supply chains and infrastructure to climate disruptions. Fragilities can emerge when a corporate culture slowly migrates to an inappropriate level of risk tolerance; or when effective governance wanes; or when a loss of cognitive diversity leads to groupthink; or when a loss of biodiversity devastates an ecosystem. Fragilities can be rooted in the impossibly complex Gordian knots that sometimes tie systems together, as we saw in the tortilla riots, or in the brittleness of a system that is unable to cooperate when it counts. In an organization or a community, fragility can arise from a lack of trust among community members or employees, from a deep resistance to change, or from the slow erosion of psychological resilience. And fragilities don't even have to be negative, per se: They can take the form of a "golden handcuffs" dependency on an external subsidy, or on a single mode of highly profitable but undiversified operations.

All civilizations face their fragilities. Many residents of the world's wealthiest nations, particularly Americans, have felt fortunate to live through a period largely insulated from shocks and disruptions. This "vacation from history" enabled many to become accustomed to living at the efficient-but-fragile end of the robust-yet-fragile continuum. In a world temporarily devoid of consequences, the slow erosion and increasing inelasticity of the country's political, financial, socioeconomic, and ecological systems scarcely seemed to matter. Now that a new, more volatile chapter has begun, those now-compromised systems have flipped from being engines of resilience to sources of fragility themselves.

All of these vulnerabilities are distinct and require different responses, yet all have the same effect: They make their host systems less adaptive and increasingly robust-yet-fragile—capable of dealing with anticipated disruption but likely to collapse in the face of an unorthodox challenge.

In spite of this, surprisingly few communities or organizations have any kind of structure in place to think broadly and proactively about the fragilities and potential disruptions that confront them. This has to change. Today, it's unthinkable for an organization of any meaningful size not to continuously monitor its financial or supply chain risk; soon, it will be equally unthinkable not to scan for a broader array of potential disruptions, from environmental issues such as carbon, water, energy, and climate risks to internal cultural factors like levels of cooperation and trust, and social issues such as the health and well-being of the communities these organizations operate in.

Even so, it's important to remember that seeking fragilities does not guarantee finding or eliminating them. Surprises are by definition inevitable and unforeseeable, but seeking out their potential sources is the first step toward adopting the open, ready stance on which resilient responses depend.

This is not to say that resilience is solely about risk management—in the decades to come it's also going to become a significant driver of innovation in its own right. Developing new technologies and services that help enterprises decouple themselves from scarce ecological resources or that help customers contend with unpredictable shocks is going to be big business, worth many billions, if not trillions of dollars.

Consider Kilimo Salama, an innovative agricultural microinsurance program developed in rural Kenya, and one example of a resilience strategy with significant future economic prospects. Created by the Syngenta Foundation for Sustainable Agriculture, UAP Insurance, and Kenyan mobile operator Safaricom, Kilimo Salama (which means "safe farming" in Swahili) insures tens of thousands of smallholder farmers, who may plant as little as a single acre of corn, against the possibility of severe weather events such as drought and excessive rain. These are precisely the kinds of events that are likely to become more frequent in a future increasingly influenced by climate change.

Kilimo Salama's customers are often near-subsistence farmers, whose labor costs are very low but whose adaptive capacity in the face

of bad weather is almost nonexistent: One badly timed drought can be catastrophic. To mitigate this risk, Kilimo Salama enables farmers to purchase very small insurance contracts on each bag of seed they buy, at the point of sale, for 5 percent of the cost. Contracts are purchased using M-Pesa, a mobile payments platform deployed ubiquitously throughout Kenya, on the farmer's own mobile phone.

Once they've registered a farmer, Kilimo Salama uses a series of distributed, solar-powered wireless weather stations to monitor climate patterns around the farmer's land throughout the growing season. If there's too much rain, or too little, at the wrong time, the farmers get a payout automatically via their mobile phones, covering the cost of their seeds.

The innovations here are numerous: First, using automated weather sensors means that one of the costliest aspects of delivering insurance —having to check manually if an adverse event has actually affected a particular farm—is effectively eliminated. "Farm visits make sense if you have a large-scale farm, because the cost of the visit can be rolled into the policy," says Rose Goslinga, Kilimo Salama's founder. "But for a small-scale one-acre farm, the math just doesn't add up—and this has been among the main reasons insurance hasn't been available. By using wireless sensing stations, we don't have to visit the farm, which transforms the economics."

Similarly, by linking the weather sensors directly to the M-Pesa accounts of the customers, the entire claims process is eliminated: If the weather turns bad, payment is automatically made via a mechanism the farmers know and trust, much faster than if a claims adjuster had to be involved. Better still, the information that these sensors collect is used to develop insights on relevant regional climate trends. In a positive feedback loop, these are then delivered back to farmers via targeted, timely text messages helping them to become better farmers even during good years.

Not only does microinsurance insure against loss, it lets people make investments they wouldn't feel comfortable making otherwise, because they can hedge them. Goslinga points to one farmer she knows who took a chance on a higher-yielding, more expensive variety of corn only because he could ensure it with Kilimo Salama. His yield increased 150 percent.

Other innovations are less technical but no less important to the success of the program. Developing Kilimo Salama required Goslinga and her colleagues to figure out how to use the technology to bundle and service nontraditional customers in a nontraditional way, how to restructure the insurance supply chain accordingly (all the way up to the reinsurers, who insure the insurance company), and—most important—how to get skeptical farmers to sign on.

The resulting service is not just an important part of the future of climate resilience, it's going to be an important part of the future of the global insurance industry itself. The MicroInsurance Centre estimates the market for microinsurance products of all kinds to be at least 1 billion people, less than 3 percent of whom are currently served. With opportunity like that at hand, the Kilimo Salama model will soon be exported, not just to other countries in the Global South, but, in a reversal of received stereotypes about the flow of innovation, also to countries in the Global North. In this century, resilience will not just be a social good, but a global market opportunity from which developed nations have as much to learn as they do to teach.

EMBRACING ADHOCRACY

As we've seen again and again, wherever we find social resilience, we rarely find just big, formal institutions at work. Instead, we often find a rich stew made up of bits and pieces of public and private organizations, informal social networks, government agencies, individuals, social innovators, and technology platforms, all working together in highly provisional, spontaneous, and self-organized ways. Since each disruption and circumstance is unique, there can be no prefabricated organizational chart for the players—indeed, one of their first tasks is to create it.

There's a name for this mode of organization, which was first popularized in the 1970s by the futurist Alvin Toffler and management theorist Henry Mintzberg: It's called an *adhocracy*. It's characterized by informal team roles, limited focus on standard operating procedures, deep improvisation, rapid cycles, selective decentralization, the empowerment of specialist teams, and a general intolerance of bureaucracy. In the digital age, an adhocracy can be put together in a plug-and-play, Lego-like way, well suited in fast-moving, fluid

circumstances where you don't know what you'll need next. If it were a musical genre, adhocracy would be jazz.

A successful adhocracy was, for example, at the heart of the success of Mission 4636 initiative in Haiti, in which dozens of self-coordinated collaborators—some as large as the International Red Cross and the U.S. Marines, others as small as an individual graduate student volunteer—worked together with nothing more than shared purpose and good software holding them together. Adhocracy also underwrites the success of CeaseFire's violence interrupters, who are able to bring together the right array of people and resources, at the right time, to cool down a shooting. In these cases and others like them, the resilient response is improvisational and provisional, established at the speed circumstances dictate.

In a larger sense, adhocracy is what all of the lessons, insights, and strategies in this book are really about. Whether expanding the cognitive diversity of a team, like Greg Fontenot and the leaders of Red Team University; or weaving diverse social networks together like Noah Idechong in Palau; or facilitating cooperation among traditional antagonists like William Ury and his colleagues at Abraham Path; or co-opting the swarming tactics of al-Qaeda in the Arabian Peninsula; or promoting eBay-like approaches to managing complex campaigns as espoused by John Arquilla; or creating complementary currencies like the WIR; or building the adaptive capacity of individuals through mindfulness training like Elissa Epel and her colleagues—all help to ensure that an adhocracy can emerge and thrive when it's needed.

The converse is also true: When systems are structurally overconnected—as they were in the 2008 banking crisis; or when responses fail to represent a diversity of genuine stakeholders, as they did at 33 Liberty Street; or when interventions are bureaucratically imposed on communities rather than developed with them, as in the case of the Bangladeshi wells, there is no space for an adhocracy to germinate.

This is not to say that formal institutions don't have a role in resilience—they certainly do. But when we focus too strongly on them as the *sole* actors in response to a disruption, we don't just ignore, but can actually smother the opportunities for these kinds of successful, improvisational approaches to emerge. Unfortunately, in fields ranging from international development and diplomacy to disaster recovery,

that's exactly where our instincts take us: We usually bias in the direction of the bureaucracy, rather than away from it. What's needed is an approach that complements these silos of excellence and works in the white spaces between them, where resilience (and social innovation) is so often found. That's what resilient organizations, and their translational leaders, do: They create the opportunity, connectivity, permission, and encouragement for people to meet in the white spaces. The leadership imperative in such circumstances is centered on influence and coordination, not command and control.

THE FIERCE URGENCY OF DATA

Adhocracies thrive on data. And by a stroke of fantastic luck, we're currently witnessing the global birth of an adhocracy of data—a global revolution that, for the first time, empowers organizations with the capacity to collect and correlate widely distributed real-time information about the way many critical systems are performing. This kind of open data will play a central role in resilience strategies for years to come.

A perfect example is given by FLOW, a powerful new data reporting system developed by the international water and sanitation organization Water for People, which is used for measuring the success of water and sanitation projects over time.

"Today, we rightly celebrate the day when a community gets water for the first time," says Ned Breslin, Water for People's CEO and a longtime veteran of global water projects. "But the way many water and sanitation agencies operate, it's as if this is the end of the story. It's not—it's just Day One. Unfortunately, due to poor planning and a lack of community ownership, within a few years of being completed, up to 60 percent of the water projects in Africa and Asia are broken or abandoned. I've walked right past broken wells dug five or ten years ago to build new ones."

To tackle this problem, Water for People developed FLOW, or Field Level Operations Watch, software that allows field-workers to use their mobile phones to map the health of water access points—sort of an Ushahidi for water. The system provides at-a-glance reports on thousands of health and sanitation points in a given region—not only their operational status, but also the level of community support and oversight—and injects a new level of transparency, efficiency, accountability, and resilience into water and sanitation efforts. Imagine

how differently things might have turned out in Bangladesh had such a platform existed. Or in New Orleans after Katrina, for that matter.

Mining the data exhaust from platforms like FLOW and Ushahidi can tell us a lot about where systemic fragilities are emerging and where interventions aren't working—often in a matter of minutes, rather than months. By making this kind of data externally available (with appropriate safeguards for privacy and anonymity) it can be remixed and correlated to produce dramatically more valuable insights, far beyond what might be feasible for any one organization to collect. That in turn helps separate an early, weak signal of a genuine disruption from simple background noise, and gives us more time to act upon it.

In the future, this data will increasingly be the fodder for sophisticated, predictive machine-learning algorithms that will not only help identify where weak signals are occurring, but where they will likely occur, helping us not only to react faster, but preemptively.

That's the theory behind EpidemicIQ, an initiative to track the global emergence of public health epidemics worldwide through the collection and correlation of literally billions of social media data points every day. The system monitors countless Twitter feeds, Facebook updates, mobile text messages, blog posts, local news reports, and the like, in thousands of languages around the world, in real time. A sophisticated artificial intelligence system ingests and sorts through this gargantuan mountain of data, seeking to separate possibly relevant reports from irrelevant ones. Once potential hits are identified, they're sent to human beings for verification, including, ingeniously, online players of the video game *FarmVille* on Facebook. (Players are offered virtual goods in exchange for correctly identifying whether a particular social media post is actually about a disease outbreak.) These confirmed hits are then sent on to a pool of subject-matter experts who analyze them using proven epidemiological models. This is an adhocracy of a different sort: of artificial intelligence software, untrained human beings, and skilled experts, all seamlessly collaborating to find digital needles in vast digital haystacks.

And it works. During a recent outbreak of Ebola in Uganda, EpidemicIQ was tracking cases five days faster than the Centers for Disease Control and eight days ahead of the World Health Organization. But the ultimate promise of such systems may be in telling us where the pattern might go next, before it does. Cross-

referenced with data about human and animal migration, transportation systems, trade networks, and countless other variables, a future fast-moving contagion might be arrested by selectively decoupling and islanding parts of various global systems without inducing mass panic. Imagine how a similar system of financial observatories might help regulators like Andy Haldane in the next economic crisis.

Such possibilities are a big part of the reason global organizations like the UN and the World Bank have made commitments to opening the data about their development programs around the world. It then becomes possible to correlate outcomes from, say, the UN's food relief efforts with the quality of water access points from Water for People and the development programs funded by the bank. This is also why global companies like Nike, GE, and IBM are creating open data initiatives covering many aspects of their operations. They're betting that the competitive cost incurred by becoming more transparent will be dwarfed by the value of understanding how their own operations relate to other critical global trends.

These approaches also hold promise for gluing governments and informal community networks together for greater resilience. An effort called Open311, for example, seeks to create a standardized interface to the 311 public information service used in many American cities like New York, so that a resident could, for example, automatically report a pothole, receive a text alert when its filled, and map a route to work that is as pothole free as possible. Such a gateway could be used for far more than reporting when something is broken; it could also collect subjective data on how residents feel about the safety and livability of their neighborhood, or residents' thoughts on what to do with a vacant lot, or how to guide volunteer efforts during a crisis to the places they are most needed. These efforts expand resilience by broadening the possibilities of cooperation, collaboration, imagination, and responsiveness among the city's residents. In so doing, they expand the varieties of informal governance far beyond what the formal government could do on its own—once again, by embracing adhocracy. And for organizations of all types there is a powerful lesson here: Resilience benefits accrue to organizations that prioritize the collection, collation, presentation, and sharing of data.

This is not to say that the open data revolution is without risks or downsides: Just as correlations can be made to sense and respond to trouble, they can be used by those seeking to create trouble as well.

Some governments—and not just despots and dictators—are already finding open data tools an irresistible tool for tracking citizens in ways that may violate their constitutional liberties. And it's not just governments, but criminals and terrorists who are also getting into the game: After Mexican drug cartels successfully silenced their traditional media by killing journalists who reported on their activities, they took to reverse engineering the identities of bloggers and social media users as well, recently hanging several of their mutilated corpses from a bridge with a sign reading, "This is what happens to people who post funny things on the Internet. Pay attention."

Equally disturbing, during the Mumbai hotel terrorist attacks mentioned in chapter 2, the terrorists on the ground remained in constant satellite phone contact with a distant coordinator, who looked up the online identities of hostages; when he confirmed that a hostage was, for example, an American or Jewish, his compatriots shot him. In a further perverse echo of Mission 4636, the same coordinator scanned Twitter postings from the terrified populace and used it to improve the attackers situational awareness and body count; indeed, the open data was so valuable to the terrorists, the Mumbai government asked the populace to stop tweeting live updates about the attack.

There are no simple answers here. Crisis mappers and those involved in the humanitarian community have begun to discuss ways to improve security and to implement codes of conduct and shared principles for harm minimization, but these will always be in tension with the need to ensure that platforms for sharing open data are as accessible and participatory as possible. Ultimately, like many aspects of our modern information society, it comes down to a value judgment: that while the risks of potential abuse are real, they are, in the main, vastly outweighed by the benefits.

REHEARSING THE FUTURE

Even with robust data in hand, healthy communities can't be expected to anticipate *all* of the surprising, nonlinear ways in which systems behave when they're disturbed, any more than the experts do. What's needed is an inclusive way of thinking about various possible, probable, and preferable futures and their implications. Here again, new tools and technologies promise to help communities and

organizations map critical thresholds, rehearse the consequences of disruption before they occur, and make better choices.

One pioneering example is a scenario-planning platform called Marine InVEST, developed by a consortium called the Natural Capital Project. Like an ecologically focused mash-up of Google Maps and the video game *SimCity*, the Marine InVEST software allows coastal communities to develop a sophisticated perspective of the many interactions between the land and the sea in their particular geographic region. Starting with a spatial map of the coast and aquascape, ecologists add science-driven models of local biodiversity and species distribution, ecosystem services, fishing, industrial activity, shipping, recreation, tourism, and the like. The resulting map illuminates the real financial value of the services that the ocean currently provides and makes clear how various activities constrain and shape one another. It also translates between denominations of value in the system—revealing the underlying exchange rate of dollars, fishing stocks, biodiversity, tourists served, miles of coastline preserved, etc.

Many of the connections between these systems are multidimensional. For example, adding fish-farming pens to a system might increase a local company's revenue, but it may also affect wild fisheries, as pens are a perfect breeding ground for parasites that affect penned and wild species alike. At the same time, adding such pens can alter the water quality by increasing effluents, which in turn can affect how much of a coastal system can be used for recreation.

By making visible these kinds of connections, Marine InVEST enables onshore communities to generate plausible answers to real-world what-if questions: What impact might a specific increase in commercial fishing have on nonfished species? How would installing an ocean-energy system affect aquatic recreation and tourism? What's the future cost (and benefit) of protecting a certain amount of shoreline—not just in dollars, but in biodiversity? How much additional activity is safe before important variables are pushed close to a critical threshold?

Understanding the answers helps the community holistically balance competing economic, ecological, and social interests while avoiding the critical thresholds that could flip the system into a less desirable state—just as ecosystems-based fisheries management (EBFM) approaches call for.

Picking a specific approach is no mean feat as, in the real world, some individual stakeholders will inevitably be shortchanged to some degree: Fewer fish will be caught than could be caught; fewer tourists will be accommodated than could be. No plan is without its compromises, and getting people to live with less-than-peak efficiency isn't a software problem—it's a political one.

To address this, the Marine InVEST team works closely with the community to build consensus around the goals, trade-offs, and costs of a given scenario. "Without this engagement, the results are meaningless," says Anne Guerry, one of the project's scientists. "On the other hand, when the community members have scenarios in hand it changes the nature of their conversation. It lets stakeholders play things out in a particular direction, see the risks and rewards of a particular scenario and how it might get into dangerous territory and fail, but without the real-world consequences of doing so."

In a manner reminiscent of Red Team University, this form of serious play encourages community members to think through the complex and less-than-obvious second- and third-order implications of a given set of choices. And it often reveals as much about the nature of relationships between the stakeholders as it does about the implications of the scenario itself. "In many ways, these renewed relationships and engagement processes *are* the software's deliverable," says Guerry. "The models get better every time we engage with nonscientists and"—echoing Noah Idechong's work in Palau—"so do their relationships with one another."

Even a tool as sophisticated as Marine InVEST won't capture every possible source or consequence of disruption. There will always be surprises, and nothing can force people to make the "right" decisions. But getting communities to reflect together about the implications of such scenarios builds resilience nonetheless: A community that learns to discuss one possible disruption is better prepared to deal with any possible disruption.

LEARN FROM RESILIENT PLACES

It's important to state again here that the path to resilience doesn't mean the elimination of every disturbance. As we noted earlier, artificial, prolonged stability can itself be a sign of fragility—an indication that a system may be conserving too many of its resources,

like a forest overdue for a prescribed burn. Weathering the occasional disruption is one of the most important ways systems learn, by turning the spotlight on potentially severe fragilities without causing the system to flip completely into a degraded state. They also serve as mechanisms of creative destruction (to borrow Schumpeter's famous phrase) that can clear a path for new arrangements—in our political system, our communities, our ecosystems, our infrastructure, and our economies alike. Modest, regular disruptions also amplify the internal diversity of the system—ensuring that some part of it is continuously seeding, growing, maturing, dying, and fertilizing the whole.

Over time, this messy cycle shapes the culture of resilient places and communities and the beliefs and values of the people who live in them. This may be one reason that social resilience is often found in places that have experienced deep challenges, like the Gulf Coast, the South Side of Chicago, or Detroit. Their all-too-routine, painful experiences of disruption build a deep cultural memory of resilience.

That's the lesson of the last story in this book: the story of Hancock Bank and its remarkable response to Hurricane Katrina.

Thanks to a culture of extensive disaster planning, the 112-year-old bank, based in Gulfport, Mississippi, was able to get its essential computer operations back up and running just three days after the storm ransacked its gleaming seventeen-story headquarters.

But on the ground, in the wake of the storm, that scarcely seemed to matter. Ninety of the bank's 103 branches had been either obliterated or severely damaged; electricity was wiped out over the entire region; police and firefighters were occupied with disaster recovery and unable to provide banking protection; and many customers had lost everything, including their IDs and checkbooks. Amid the chaos, it was impossible for Hancock's staff even to know who their customers were, not to mention how much money they had on deposit with the bank. And yet, with credit-card systems offline, citizens needed cash more urgently than ever before, and all of the local sources—Hancock's and its competitors' alike—were devastated.

So what to do?

At the height of the calamity, the bank's executives went back to the institution's founding charter, more than a century old, and noticed that it focused entirely on serving people and taking care of communities.

The word "profit" never appeared.

So Hancock Bank's executives did something remarkable—they enlarged the tribe. While the winds were still subsiding, Hancock employees stood in front of forty of the branches knocked offline, operating from card tables, under tarps, and out of mobile homes, and offered two hundred dollars in cash to anyone who would sign a slip of paper with his or her name, residence, and Social Security number. Not just Hancock customers—anyone. No ID, no problem. Much of this money was recovered from the ruins of local casinos and literally laundered and ironed by branch workers before being handed out.

By the time the disaster was over, the company had put more than $42 million into the hands of local residents—customers and noncustomers alike—with nothing more than Post-it Note IOUs as a record. This helped keep the local economy going at a time when it faced months of near-certain paralysis.

Hancock's incredible act of trust paid off handsomely for the bank: Within months, thirteen thousand new accounts were opened at Hancock and deposits had risen by $1.5 billion; within three years, all but $200,000 worth of those initial $200 loans—99.5 percent—had been paid back.

At the center of this remarkable story was a clear, shared vision of the bank as a long-term social institution first, and a short-term profit-making institution second. The bank had obviously prepared for disasters, but no one had contemplated handing out millions to strangers, two hundred dollars at a time. It was Hancock's extensive rehearsal for future disruption; its strong, shared social values; its trust in its community; and its empowered adhocracy of employees and stakeholders that fed its ability to rapidly flip, like the batfish or the WIR, into a completely new mode of behavior, in which the normal rules of operation were suspended and both the bank and the community could be restored.

And no regulator had to tell them to do it.

• • •

In so many of the stories we've explored about resilience, we see a great confluence of factors coming together—the right systems and structures, the right technologies and information, the right kinds of community empowerment, and the right values and habits of mind.

When these are joined, multiple victories ensue, no matter the scale or context. Neighbors who work together to design and plant an urban farm in a vacant lot, for example, will certainly go some of the way toward reversing urban blight and improving their food security, health, and nutrition. They'll save a little money, get some exercise, and suck a little carbon dioxide out of the atmosphere all at the same time. While they're at it, they may amplify their own sense of self-reliance and adaptive capacity and build the relationships that may ensure their community's resilience in the face of some crisis yet to come.

Even our own thoughts play a role here, not only in our own resilience, but in others' as well. The work of researchers like Richard Davidson, Elissa Epel, Clifford Saron, Amishi Jha, and others shows us paths to improving our own resilience through reflective practice and the discovery of greater meaning in our lives. And Gary Slutkin shows us how such habits of mind can be contagious, for good or ill. Tie these threads together, and you have the first links in a chain that connects your personal resilience to that of your social circle, your community, the place you work and live, and out across the world. What you choose to believe, the mental practices you cultivate, and how you respond to disruption truly shape the whole. Resilience can radiate out from within.

The journey toward resilience is the great moral quest of our age. It is the lens with which we must necessarily adjust our relationships to one another, to our communities and institutions, and to our planet. Even so, we must remember that there are no finish lines here and no silver bullets. Resilience is always, perhaps maddeningly, provisional, and its insistence toward holism, longer-term thinking, and less-than-peak efficiency represent real political challenges. Many efforts to achieve it will fail, and even a wildly successful effort to boost it will fade, as new forces of change are brought to bear on a system. Resilience must continuously be refreshed and recommitted to. Every effort at resilience buys us not certainty, but another day, another chance.

Every day is Day One.