
PART 3: SYSTEMIC RISK IN ECOLOGY AND ENGINEERING

Several fields of engineering and science share with economics a keen concern with systemic risk. Systemic risk is manifested in space shuttle accidents, airplane crashes, the collapse of the New Orleans levees, electrical power blackouts, and the failures of buildings, bridges, and many other engineered systems. Because of these occasional system failures, engineers have more relevant data for the study of systemic risk than do economists. Using these data to conduct retrospective analyses of system problems, engineers have been able to identify and remove some sources of failure (for example, in aircraft). Similarly, epidemiologists and public health experts worry about disease outbreaks and spread, which occasionally reach systemic levels, and they have learned lessons in risk management by studying past epidemics. And ecologists study changes in the state of ecosystems, which may receive less press attention but clearly qualify as systemic developments because they can result in a true regime shift from one equilibrium to another.

There are two ways that one discipline can leverage the experience of another. The first way is by adapting methodologies developed in one field to analyze structures and phenomena in the other field. The examination of the Federal Reserve's Fedwire system in part 4 of this volume exemplifies this mode of intellectual sharing: researchers adapt tools from outside of economics—namely, network theory and graph theory—to learn what insights can be gained by applying them to a problem of systemic behavior in the area of payments. The second way is by sharing insights that are particular to a given field and that, by analogy, might apply to other fields. This is the approach taken in this part of the volume.

USEFUL CONCEPTS FROM ECOLOGY AND ENGINEERING

At the conference, ecologist Simon Levin of Princeton University identified a range of concepts that have proved helpful in understanding complex systems in ecology and that might also apply to financial systems. One useful conceptual model of an ecosystem is a “trophic web,” which represents how species are interconnected. At a coarse level, a trophic web in an ecosystem might be thought of as a set of predator-prey relationships. In this case, sets of differential equations can be successful in modeling the rise and fall of populations as the ecosystem fluctuates around an equilibrium or becomes unstable. More generally, however, “trophic” refers to the flow of energy, so the trophic web for an ecosystem is a framework for representing how the primary source of nutrition (say, sunlight or geothermal vents) is transmitted between levels in the food chain. This interpretation of the trophic web is more applicable to financial systems, in which the interactions are usually less extreme than those in predator-prey relationships; we simply have to interpret “energy” as anything of value that is transmitted through the system. Because of this analogy, it is not surprising that we would find similar, if not identical, phenomena in these two systems, and therefore similar insights might be brought to bear in analyzing them. Complex systems of any sort are characterized by nonlinearities, multiple stable states, hysteresis, contagion, and synchrony, all of which have relevance to the problem of systemic risk.

Nonlinear relationships are a key characteristic of virtually any complex system. They can lead to multiple stable states,

The views expressed in this summary do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

such that the system can exist in one configuration (basin of attraction) for a period of time but then be knocked into a different configuration by a perturbation or shock. This transition can be accompanied by hysteresis, meaning that if the system is to return to its original configuration, it must take a different path. Often, pain and other costs are associated with that recovery pathway.

Nonlinear feedbacks, which can be either positive or negative, can drive a complex system away from a given equilibrium state;¹ the stability of any complex system is determined by the nature of these feedbacks. Feedbacks can result from the low-level processes in the system (for example, the behaviors or individuals in a food chain, traders in a market, or components of an engineered system), from an explicit top-down control system, or from policies enforced

Complex systems of any sort are characterized by nonlinearities, multiple stable states, hysteresis, contagion, and synchrony, all of which have relevance to the problem of systemic risk.

by regulators. Positive feedbacks usually amplify the effect of disturbances, thereby decreasing the stability of steady states. In contrast, we usually think of negative feedbacks as stabilizing. However, that is not always the case, as demonstrated by the suspension bridge over the Tacoma Narrows known as “Galloping Gertie.” The bridge was subject to a negative feedback (a damping) that overcompensated, with the result that a certain wind condition led to escalating oscillations and finally collapse.

Once a system is destabilized, it moves away from the linear regime and can experience nonlinear behaviors such as path dependence (meaning that the next state is dependent on the sequence of events that led to it), sustained oscillations (such as cyclicity in the financial sector), and regime shifts, by which a system moves into an entirely different region of performance, such as the less desirable equilibrium that characterized the Great Depression. However, nonlinear behavior also means that an effective remedy need not require a massive effort, just a well-targeted one.

Another phenomenon common to complex systems is contagion. In ecosystems, contagion is an important part of ecological and epidemiological dynamics, as exemplified by the mechanisms that spread forest fires and disease. In the financial

¹“State” is used here as a shorthand to mean either a single state or a set of dynamically (possibly stochastically) related states in a common basin of attraction, not something static.

sector, contagion manifests itself as cascading losses and increased risk aversion, with the latter leading to herd behavior, funding withdrawals, and a contraction of liquidity. Contagion can be found in two forms in the electric power grid and other complex networks such as road and communications systems. A destabilizing form occurs when the failure of one node (for example, a substation or a bridge) creates a buildup of load on the rest of the system that in turn may lead to a cascade of other failures. But when load switching and rebalancing can effectively redistribute the load, contagion assumes a stabilizing form: it spreads the stress and thereby reduces systemic risk.

Synchrony, another feature shared by some complex systems, is evident when incentives or pressures lead individual actors to fall into step and make similar choices. In nature, one finds benign instances of this phenomenon: some species of fireflies blink synchronously, and flocks of birds and schools of fish can often turn almost as units. However, tight linkages among individuals can also be a cause for concern because they can induce systemic collapse. Conservation biologists have shown considerable interest in the degree of synchrony in species populations: 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. Like contagion, synchrony can lead to systemic risk in the form of a system failure or a sudden jump to a less desirable equilibrium.

Ecosystems, the financial system, and many other complex systems are in fact complex *adaptive* systems, in which collective behaviors emerge from individual actions. In ecosystems, those collective behaviors include the flocking of birds, herding of ruminants, and formation of fish schools. In the world of finance, the Dow Jones Index reflects the integrated effects of many individual decisions, making it an emergent indicator. Many components of the financial system pay attention to these emergent indicators, and what the indicators imply about collective behaviors feeds back to affect individual behavior, but on very different scales of organization and time. Behavioral ecologists have developed some understanding of the principles of collective decision making among animals.²

Complex adaptive systems consist of heterogeneous collections of individual units that interact with one another and thereby influence how the whole system evolves. Often the phenomena that we are interested in are occurring on different scales, and the systems essentially integrate phenomena at multiple scales of space, time, and complexity. The components of the electric power grid (transformers, voltage regulators, generators, relay switches, and so forth), for instance, are nonlinear and have different stochastic

²See, for example, Couzin et al. (2005).

behaviors that might affect only a local neighborhood of the grid, but they interact in ways that can lead to systemic shifts in grid performance, or to failure. Moreover, the observed system performance is actually the integrated result of the grid's behavior along with the behavior of layers of communication, sensing, and control, the fuel supply, human behavior, and the financial transactions that make it function. Clearly, understanding and predicting the performance of a complex adaptive system at that level is a major multiscale and multidisciplinary endeavor.

The term “complex adaptive system” might leave the impression that the system is adapting and adjusting itself to beneficial effect. What it really means, however, is that some components of the system are adapting and changing, not that the system as a whole is changing in a coordinated way. The adaptation might be in the influenza virus, and its ability to become more effective is not necessarily good for the system as a whole.

A critical attribute of complex adaptive systems that must be properly modeled is path dependence. Imagine rolling a ball down the side of a mountain range. Its path illustrates the natural development of a system. The ball comes to certain decision points where it enters one or another watershed. Once it starts down one path, it is locked into that pathway unless a major perturbation occurs. Thus, the future development of the system is dependent on the path that has been taken—that is, on the history of the system. If there is a major perturbation, however, the system can jump into a new basin of attraction that is conceptually and phenomenologically very different: the system would move from one valley to another.³ This is a regime shift, or system flip, which can be very disruptive. For example, scientists studying ecological systems worry about eutrophication, the over-enrichment of lakes. A system that moves from a healthy oligotrophic lake to a eutrophic lake with large quantities of algae is still a stable system, but the flip is very detrimental for most of the species in the oligotrophic lake. Analogously, a rich land can undergo desertification and become a very different ecosystem.

On a larger scale, ocean circulation patterns can undergo relatively sudden flips. Such flips have occurred in the past and might be triggered again by climate change, but no one knows the likelihood of their recurrence. A qualitative change in ocean circulation patterns—one that altered the topology of the flows—would have major impacts. It would be a regime shift, a shift into a different domain of attraction. Economic markets can go through crashes and recoveries that are also shifts in the basins of attraction. Bank collapses can trigger chain reactions that would represent the same type of shift as a phase transition in physics.

³In economic terms, each valley will have its own rates of saving, interest, employment, productivity, and so forth.

As noted earlier, regime shifts can lead to hysteresis, meaning that the behavior of the system in its recovery phase may be quite different from its behavior in the destruction phase. For example, in the ecological literature, there is considerable interest in the spruce bugworm and other defoliating insects that can completely denude forests of spruce, balsam fir, and other species. After an outbreak of these insects, the system recovers over time, but as the forest quality increases, the bugworm population builds up enough to re-emerge. Once this outbreak occurs, the quality of the forest begins to decline until the system reaches a critical point and collapses. Thus, the system goes through regular periods of outbreak and collapse, each one representing what amounts to a system shift. The fact that the pathway on the way down differs from the pathway on the way up is a hysteresis effect.

Levin pointed out that, unlike systems designed for robustness, complex adaptive systems are systems in which whatever robustness exists has to emerge from the collective properties of the individual units that make up the system; there is no planner or manager whose decisions completely

Ecosystems, the financial system, and many other complex systems are in fact complex adaptive systems, in which collective behaviors emerge from individual actions.

control the system. Therefore, there are no guarantees that things will work well. This leads us to the problem of the global commons, in which we all engage in behaviors based on our own agendas and interests; from these individual behaviors, system properties emerge. For individual organisms, natural selection encourages the development of robust physiological properties. But an ecosystem, banking system, or economic system has not been engineered for robustness.

Collapse in complex adaptive systems is the same as the loss of robustness. If a system is working well, we think of it as robust, whether it is an engineered system, a banking system, or an ecosystem. In various literatures, the terms robustness, resilience, rigidity, and resistance are often used to mean the same thing, although they really describe different components of the system's capacity to function in the presence of internal or external disturbances.

What leads to robustness in complex adaptive systems? There are at least two ways in which a system can be robust in the face of disturbances: by having a rigid design and reliable components, or by having a flexible design that may also include replaceable components. One can see these alternatives

in a stressful marine environment with strong currents. Corals resist the disturbances by being rigid, while kelp withstands the disturbances by being flexible. These are two quite different strategies for responding to the stress of strong currents, and we see the same contrasting strategies in many other systems. Rigidity—sticking with an existing design or decision (think of the Polaroid company and its camera design)—might be the best approach over short periods of time or if the environment is relatively constant. But over longer periods of time or in fluctuating environments, flexibility can prove a more robust approach. In the camera industry, for example, Kodak has continued to change its camera designs and products over the years. Neither the Polaroid nor the Kodak strategy is “right” per se, but each is right over a particular time horizon.

In changing environments, one needs flexibility, whether it is in ecological systems or in banking systems. For example, Levin noted that the flexibility of the influenza virus accounts for its robustness. On the surface of the virus are proteins called surface antigens, in particular haemagglutinin and neuraminidase. The name of a flu strain—say, H5N1 flu—refers to the particular forms of haemagglutinin and neuraminidase associated with that strain, as those proteins change over time. Once a person gets a particular strain of influenza, he or she will never get it again. Individual variants therefore are not very robust; they can be controlled or eradicated by the human immune system if they return. But the influenza virus itself has been around for centuries, maybe millennia, so the virus seen more generally is very robust. It survives because it is adaptive, continually changing its design and its surface proteins.

Therefore, according to Levin, for a system to be robust it must have diversity—analogue to the way the influenza virus is really a family of viruses with variations in their surface proteins—and it must have heterogeneity, so that there is scope for adaptation in the system. For this reason, ecologists attach great importance to biological diversity: even if they do not know what particular species do, the presence of diversity provides a form of insurance. When a system is too homogeneous, it cannot adapt.

Modularity—the degree to which a system can be decoupled into discrete components—also influences robustness. A basic principle in the management of forest fires and epidemics is that if systems are all connected, a perturbation will encounter nothing to stop it from spreading. But when a system is compartmentalized (when firebreaks exist or high-risk parts of a population are vaccinated against an epidemic), then the spread may be contained. Modularity can thus be an important part of robustness if it ensures that an affected component will be isolated from destabilizing feedbacks. However, modularity

often involves a trade-off between local and systemic risk. Because the compartmentalized elements of a system will be less able to withstand some shocks, modularity tends to increase the risk that individual elements will be critically damaged. Although the sacrifice of such elements is assumed to decrease the risk of a calamitous systemic failure, the wrong

In changing environments, one needs flexibility, whether it is in ecological systems or in banking systems.

compartmentalization in financial markets could preclude stabilizing feedbacks, such as mechanisms for replacing lost liquidity, and so could actually increase systemic risk.

Robustness is not the same as stability, which refers to the ability of a system to return to its equilibrium state. It is interesting to note that ecologists have not been able to agree on the relationship between biodiversity and stability. In the 1950s, qualitative arguments led many to believe that biodiversity and stability are positively correlated—for instance, that biodiversity leads to robustness in some macroscopic system properties such as nutrient cycling. But theoretical arguments developed in the 1970s implied that as system complexity or diversity increases, an equilibrium in the relevant system of differential equations is less likely to be asymptotically stable. Some argue that the instability of the system dynamics (in the narrow sense of a stable equilibrium of species densities) is what provides the adaptive capacity to buffer the macroscopic properties: species replace one another, or there are shifts in abundance, and these changes allow the system to adapt to perturbations. Whether diversity increases or decreases stability is an argument over the definition of stability, and it is still being debated.⁴

The lesson that might be inferred is that understanding the behavior of complex adaptive systems requires more than just qualitative analysis and more than just theory. Ecologists have applied alternative mathematical frameworks (for example, interacting particle systems or systems of differential equations), intensive simulations, data-driven analyses, and even experiments in the effort to resolve this issue, and a similar multifaceted effort might be needed to provide policymakers with insights about the root sources of stability in financial systems.

⁴See National Research Council (2005, pp. 114-5) for a good discussion of this debate. See also Levin (2000, chap. 7).

METHODOLOGIES FOR PREDICTION AND MANAGEMENT

In addition to providing useful concepts for the description and analysis of systems in other disciplines, science and engineering may provide some relevant methodologies for the prediction and management of systemic risk. The rich scientific literature on networks and graph theory, for example, may have some bearing on the management of economic and financial system risk. Networks influence the spread of information, disease, and disturbances, and indeed the spread of effects that can stabilize or destabilize a system. The topology of the network is one of the key factors to study. For instance, are there key nodes in the network whose removal would cause the system to become decoupled? The potential for decoupling might be seen as a vulnerability of the system because it could impair the functioning of the network, but it can also suggest a mode for limiting contagion in that it induces the modularity that is important to robustness. Thus, to control the spread of disease, scientists try to identify those who are super spreaders, the individuals (say, prostitutes or hospital workers) who connect different groups and make the system more likely to exhibit undesirable synchronous effects. More generally, researchers who study the topology of networks and the relationship of that structure to network functionality will consider how the properties of the network affect the spread of money, disease, or information and propagate the spread of disturbances that can cause systems to collapse.

Other scientific research relevant to the management of risk is the literature on the modeling and control of forest fires, the modeling and management of epidemics, and contagious spread more broadly. The whole field of spatial stochastic processes has focused largely on ecological and epidemiological problems. As an example, Levin cited a National Institutes of Health committee he recently chaired that oversaw several agent-based simulations of the potential spread of pandemic influenza in order to identify strategies for controlling that spread. The models developed in this and other research efforts are very computation-intensive. Levin indicated that transferring the techniques from these models to the study of financial systems would not be difficult, both because the parallels were strong and because researchers in the financial sector would be comfortable with the mathematical techniques. The rich literature of epidemic theory, both mathematical and computational, might then be applicable to understanding runs on banks, as long as this approach was properly augmented with knowledge of human behaviors that contribute specifically to bank contagion. Levin suggested that it might also be possible to transfer recent work on social learning to the study of the financial sector.

George Sugihara of the University of California at San Diego expanded on the possibility of rich analogies between ecosystems and financial systems. Perfect parallelism is not required if the goal is merely to stimulate fresh thinking that generates productive hypotheses for research and even policy formation related to financial systems, although empirical corroboration of the analogy is, of course, one way to strengthen its utility.

He pointed out that most ecosystems are innately robust because they are survivors of extreme stress testing. Their existence today sets them apart as the selected survivors of many millions of years of upheaval and perturbation, having withstood continental drift, meteor extinctions, climate fluctuations, and the introduction or evolution of new members. Those that survive show some remarkable constancy in structure that may persist for hundreds of millions of years (for example, the constancy of predator/prey ratios noted in Baumbach, Knoll, and Sepkowski [2002]). Identifying the

The rich literature of epidemic theory, both mathematical and computational, might . . . be applicable to understanding runs on banks, as long as this approach was properly augmented with knowledge of human behaviors that contribute specifically to bank contagion.

common attributes of these diverse systems that have survived rare systemic events could provide clues about which characteristics of complex adaptive systems correlate with a high degree of robustness. These attributes could then be examined as candidate characteristics for lessening systemic risk in other contexts, such as the financial sector. Because experimental stress testing is not feasible in the financial sector, examining such common structural properties of ecosystems should be of interest, and it might help guide policy.

According to Sugihara, recent studies in nonlinear complex systems show rapid and large transitions in state to be common features of many “generic” interconnected dynamic (and cybernetic) systems. Beyond the specific analogy between ecology and economics, certain dynamical behaviors and structural (topological) constraints are common to broad classes of systems. Behaviors and network topologies that are truly generic—as opposed to system-specific—can inform many disciplines. For example, to understand the systemic risk problem, it is useful to know the general properties of complex

systems, particularly the structural ones that promote stability or collapse.

As an example of scientific analysis that can readily be applied to financial systems, Sugihara cited a recent paper in *Science* (Bascompte, Jordano, and Olesen 2006) that examined disassortative networks—networks in which nodes that are in some sense “large” connect with many nodes that are “small,” although the small nodes do not connect to many large nodes. The paper, coming from the field of ecology, focused on the network of pollinators and the plants that they pollinate, but it also dealt more broadly with all networks that are positively reinforcing. The paper showed that the disassortative nature of the pollinator-plant network conveys a great deal of stability—a result, Sugihara suggested, that generalizes to any type of disassortative network, including the network linking U.S. banks to the Fedwire system (see part 4 of this volume). In this case, then, the theoretical analysis of a complex ecological system highlights a characteristic of the financial system that might be essential for stability and therefore worthy of protection.

RISK ASSESSMENT OF EXTREME EVENTS INVOLVING NATIONAL SECURITY

Yacov Haimes of the University of Virginia discussed his work in modeling extreme events, especially those that affect interconnected infrastructures and relate to national security. It is generally impossible to build one single model to represent any such complex system; there are too many cross-cuts and too many ways to examine the processes and effects of a complex system. The analysis of such a system must instead be addressed from multiple perspectives, perhaps hierarchically.

For his approach, Haimes has developed what he calls “hierarchical holograph modeling” (HHM). This method is hierarchical because it includes many different subtopics, such as hardware, software, and organizational influences. He emphasized that the last subtopic must be included in any study of risk because many of the factors that contribute to risk, or follow from extreme events, are organizational problems and human problems. Risk analysis must consider such matters as how well lines of communication function, how much trust exists within a system, and who can share information in a timely and effective way. And, of course, the modern reliance on information technology means that information assurance has also become critical. Haimes calls his method “holographic” because it examines risk from many

different perspectives. For example, in a study conducted for the President’s Commission on Critical Infrastructure Protection, Haimes and his colleagues identified 300 major sources of risk to the U.S. water supply. A good methodology is necessary to structure an analysis encompassing that quantity of information.

This approach to identifying and analyzing extreme events in engineering differs from the approach often used in modeling extreme events in economics and finance. The HHM

Many of the factors that contribute to risk . . . are organizational problems and human problems. Risk analysis must consider such matters as how well lines of communication function [and] how much trust exists within a system.

method starts with an extreme outcome and provides a methodology for exploring what factor or combination of factors would produce that outcome. It is an inverse method in that it works backward from an undesirable outcome to infer what combinations of circumstances could give that result and what the associated probabilities are. In contrast, systemic risk analyses as conducted by financial economists or market practitioners often project forward to infer the ramifications of a hypothesized shock. The two approaches represent different strategies for understanding what factors produce extreme events.

In the study for the President’s Commission on Critical Infrastructure Protection, Haimes and his colleagues used HHM as the foundation of an adaptive multiplayer game. Four teams, each with a very different perspective, were assembled in 2005 to develop separate HHMs to learn about the various sources of risk affecting Supervisory Control and Data Acquisition (SCADA) systems. The red team assumed the perspectives of attackers and hackers; the blue team represented the perspectives of SCADA operators and owners; a vendor team embodied the ideas of SCADA developers and vendors; and a policymaker/stakeholder team represented the interests of government and of industry associations. About sixty experts participated in the four teams. Interestingly, because of the teams’ differing perspectives, there was less than 10 percent overlap in the specific risks identified. For instance, several teams identified software and staff training as key risks, but only the policymaker team identified organizational decision making as a potential risk, and only the operators/owners team identified the quality of electrical

infrastructure as a potential risk. This exercise underscores the value of incorporating multiple views and perspectives in efforts to identify sources of risks in complex systems. Team approaches to generating input for risk analysis can be very effective. The key to their success is the mechanism for assimilating the information generated and for anchoring it to concrete evidence. Uncertainty quantification plays a major role in the degree of success of such efforts. The problem most often encountered is that the results are not sufficiently transparent to merit high confidence.

Another study of large-scale risk undertaken by Haimés and his colleagues explored the regional and national economic effects of an attack with a high-altitude electromagnetic pulse (H-EMP).⁵ In an H-EMP attack, an enemy would use a nuclear weapon to inflict systemic damage on the country's electrical and computing infrastructure. Specifically, an atomic bomb would be exploded fifty kilometers above the United States, and most of the damage would be to electronic systems, not people or structures.

Using the inoperability input-output model (an adaptation of Wassily Leontief's input-output model that puts more emphasis on interdependencies), Haimés and his colleagues estimated the percentages of dysfunctionality that would be observed in 485 sectors of the regional economy as a result of an H-EMP attack. These estimates were based on assumptions about the impact of the H-EMP blast on the electrical and computing infrastructure of each sector. As expected, the predicted inoperability effects are not uniform across all sectors, nor are the production losses, which would amount to billions of dollars. By studying the heterogeneous effects of such an event, Haimés explicitly avoids the spatial and sector smoothing that is implicit in some analyses of risk, and draws attention to the varied and localized nature of the economy's vulnerabilities. In this particular case, it was determined that the major impacts sustained by some sectors would nevertheless have a minor effect on the economy per se, and so would not lead to systemic problems. This type of analysis provides policymakers with valuable insights into priorities, highlighting what resources should be protected first or most securely. It can also help illuminate the trade-offs between different recovery strategies, which can be striking.

Presenting another example of a complex analysis of heterogeneous impacts, Haimés described his study of the hypothetical economic impacts of a closure of the Monitor-Merrimac and Hampton Roads bridge-tunnels in southeastern Virginia. That area of Virginia contains a number of military installations, including a major naval base. To understand the

⁵This study was conducted for the Congressional Commission on H-EMP Attacks on the United States.

economic effects, Haimés had to model the driving patterns of many groups of workers and purchasers as they found alternate routes, and the patterns emerging from those models collectively created a picture of the overall system behavior. If these tunnels were destroyed, it would take more than a year to rebuild them, so they are very strategic for Virginia and for national security more generally. This research provides the foundation for choices that prepare us for extreme natural hazards or terrorist attacks and for developing resilience in our interdependent infrastructure and economic systems.

An analysis using an inoperability input-output model revealed that the major sectors whose functioning would be impaired by the closure of the bridge-tunnels would be primary metal manufacturing and textile manufacturing. All the other sectors would be minimally affected. As for the overall economic loss, management services would be affected most, followed by business services and retail trade, while

Team approaches to generating input for risk analysis can be very effective. The key to their success is the mechanism for assimilating the information generated and for anchoring it to concrete evidence. Uncertainty quantification plays a major role in the degree of success of such efforts. The problem most often encountered is that the results are not sufficiently transparent to merit high confidence.

the economic impact in many other sectors would be slight. The analysis shows each sector from different perspectives, producing a broader picture.

In all these risk analyses, Haimés and his colleagues assessed the expected value of outcomes but supplemented that assessment with other information because expected values can be insufficient indicators of risk. Managers and decision makers are often more concerned with the risk attaching to a specific case than with the likelihood of an "average" adverse outcome that may result from all similar risk situations. They are also interested both in the low-frequency, high-damage events—those with major, potentially regime-shifting consequences—and in the more common risks, which dominate the expected value.

Haimes explained how he uses the partitioned multiobjective risk method (PMRM)⁶ to measure and analyze the risk of extreme and catastrophic events by partitioning the probability into several sections, as shown in the following equations:

$$f_2(\cdot) = E[X | X \leq \beta_1] = \frac{\int_0^{\beta_1} xp(x)dx}{\int_0^{\beta_1} p(x)dx}$$

$$f_3(\cdot) = E[X | \beta_1 \leq X \leq \beta_2] = \frac{\int_{\beta_1}^{\beta_2} xp(x)dx}{\int_{\beta_1}^{\beta_2} p(x)dx}$$

$$f_4(\cdot) = E[X | X > \beta_2] = \frac{\int_{\beta_2}^{\infty} xp(x)dx}{\int_{\beta_2}^{\infty} p(x)dx}$$

$$f_5(\cdot) = \frac{\int_0^{\infty} xp(x)dx}{\int_0^{\infty} p(x)dx} = \int_0^{\infty} xp(x)dx.$$

The probabilities displayed in these equations have the following interpretations:


- $f_2(\cdot)$ represents the risk with high probability of exceedance⁷ and low damage, partitioned at β_1 on the damage axis.
- $f_3(\cdot)$ represents the risk with median probability of exceedance and medium damage, partitioned between β_1 and β_2 on the damage axis.
- $f_4(\cdot)$ represents the risk with low probability of exceedance and high damage, partitioned between β_2 and ∞ on the damage axis.
- $f_5(\cdot)$ represents the unconditional (conventional) expected value.

The PMRM can be used to explore trade-offs between the cost of risk management and the potential loss. The chart presents a specific example in which the horizontal axis represents a percentage of electric power capacity at risk and the vertical axis, which is also $f_1(\cdot)$, represents the cost of risk management. Each of the policy options A through D has an

⁶See Asbeck and Haimes (1984).

⁷An exceedance probability (EP) curve specifies the probability that a certain level of loss will be exceeded. If one views the loss as a random variable, the EP is simply the complementary cumulative distribution of the loss.

The Trade-Off between the Cost of Risk Management and Potential Losses



associated cost for risk management and a corresponding loss of functionality. For instance, option A consists of investing significant resources in risk management in order to reduce the likelihood of extreme events. The curve on the left shows the expected value lost, while the curve on the right shows the extreme loss. It is the more meaningful curve.

PREDICTION AND MANAGEMENT OF SYSTEMIC FAILURE IN THE ELECTRIC GRID

Massoud Amin of the University of Minnesota extended the discussion of risk assessment, modeling, and prediction by describing past and potential failures in the North American electric power grid, another complex system. While this system might not support multiple equilibria, as ecosystems and financial systems can, it is certainly susceptible to nonlinear amplification of instability, which leads to blackouts. The post mortem analysis of major blackouts often shows the root cause to be the failure of one or a few components (out of thousands in the portion of the grid ultimately affected) that upsets an equilibrium and leads to a cascade of failures. For example, on August 10, 1996, North America experienced a major blackout affecting more than 7 million customers in thirteen states or provinces. It was later determined that the root cause was two transmission faults in Oregon. Ultimately, that modest failure led to power oscillations on the order of 500 megawatts, overwhelming the system's response mechanisms and leading to the blackout.

Amin reported that some studies of the 1996 blackout estimated that it could have been avoided if the grid had intelligent controls and was able to reduce its load by 0.4 percent for thirty minutes. Such studies not only shed light on how to prevent future failures, but also help to clarify what recovery options exist if a similar failure does occur. Recovery is an important part of risk management, and recovery options can be identified by doing a scenario-based quantitative risk assessment in advance. Of course, the technologies for recognizing the incipient problem and tailoring a solution are far from obvious.

Engineered systems such as the electric power grid or a telecommunications network often include advanced control systems that enable recovery. Amin reported on research funded in the 1990s by the Electric Power Research Institute (EPRI) that built on the technology used in control systems for fighter planes. Because a power system includes substations and generators that must operate at the same 60 hertz frequency, controlling those elements in a coordinated fashion is somewhat analogous to controlling planes that are flying in formation. And responding to the loss of one or more components is somewhat analogous to maintaining control of an aircraft when a wing is damaged. Accordingly, EPRI's research was directed toward a control system that would have some self-healing capability—a system, in other words, that could anticipate disruptive events by detecting signals indicating an important change, conduct a real-time assessment of the changing state of the system, determine how close the system is to some “edge” in performance, and remedy or isolate the problem (isolation, sectionalization, and adaptive islanding, which are discussed below). These same sorts of capabilities would be desirable in a system designed to control the financial system during disruptions.

Creating such a control capability for the electric grid requires a mixture of tools from dynamical systems, statistical physics, and information and communication science, as well as research to reduce the computational complexity of the algorithms so they can be scaled up to the large size of the system being controlled.⁸ The electric grid poses a multiscale challenge: troublesome signals must be detected within milliseconds, with certain compensatory actions taken automatically; some load balancing and frequency control on the grid is handled on a timescale of seconds; and control functions such as load forecasting and management or generation scheduling take place on a timescale of hours or days. Identifying at the atomic level what is amiss in a system

⁸Working methods derived from the EPRI research program have been applied in a variety of contexts, including the electricity infrastructure coupled with telecommunications and the energy markets, cell phone networks on the Internet, and some biological systems.

and then responding on a macro-scale requires multiresolution modeling in both space and time.

To convey the complexity of modeling and controlling the electric grid, Amin gave some basic facts. In North America, there are more than 15,000 generators and 240,000 miles of high-voltage lines. The overall grid is divided into several very large interconnected regions, and modeling one of them (which is necessary for understanding the systemic risks) might entail a simulation with 50,000 lines and 3,000 generators. The system is typically designed to withstand the loss of any one of these elements. To determine whether the grid can attain that design goal, we need to simulate the loss of each of the 53,000 elements and calculate the effects on each of the 50,000 lines, leading to more than 2.6 billion cases. Although analysis of these systemic risks is very challenging, the findings can help researchers determine the best way to operate the system.

As an additional illustration of the level of detail that can be successfully simulated, Amin presented a complex model that predicts load and demand for DeKalb, Illinois, a sizable market

In any situation subject to rapid changes, completely centralized control requires multiple, high-data-rate, two-way communication links, a powerful central computing facility, and a sophisticated operations control center. But all of these features are vulnerable to disruption precisely when they are most needed.

with a mixture of commercial and residential customers. Deregulation of the electric system has reduced the correlation between power flow and demand, thus introducing uncertainty into the system, and so a number of researchers have sought new ways to monitor and predict demand. The models and algorithms are now sophisticated enough to simulate the demand by customer type (residential, small commercial, large commercial) on an hour-by-hour basis and attain 99.6 to 99.7 percent accuracy over the entire year. One benefit of these predictions is that they enable power companies to dispatch small generators to meet anticipated high demand.

More broadly, Amin argued that any critical national infrastructure typically has many layers and many decision-making units and is vulnerable to various disturbances. Effective, intelligent, and “distributed control” is required that would enable parts of the constituent networks to remain

operational or even to reconfigure automatically in the event of local failures or threats of failure. In any situation subject to rapid changes, completely centralized control requires multiple, high-data-rate, two-way communication links, a powerful central computing facility, and a sophisticated operations control center. But all of these features are vulnerable to disruption precisely when they are most needed (that is, when the system is stressed by natural disasters, purposeful attack, or unusually high demand).

When failures occur at various locations in such a network, the whole system breaks into isolated “islands,” each of which must then fend for itself. With the intelligence distributed, and the components acting as independent agents, those in each island have the ability to reorganize themselves and make efficient use of the remaining local resources in order to minimize the adverse impact on the overall network. Local controllers will guide the isolated areas to operate independently while preparing them to rejoin the network, without creating unacceptable local conditions either during or after the transition. A network of local controllers can act as a parallel, distributed computer, communicating via microwaves, optical cables, or the power lines themselves, and limiting their messages to only the information necessary to achieve global optimization and facilitate recovery after failure.

If coordinated with the internal structure existing in a complex infrastructure and with the physics specific to the components they control, these agents promise to provide effective local oversight and control without need of excessive communications, supervision, or initial programming. Indeed, they can be used even if human understanding of the complex system in question is incomplete. These agents exist in every local subsystem and perform programmed self-healing actions that can avert a larger failure. Such simple agents are already embedded in many systems today in the form of circuit breakers and fuses as well as diagnostic routines. Echoing the familiar tale of the kingdom that was lost for want of a horseshoe nail, we might say that these agents are like the missing nail: once restored, they can save an entire kingdom.

Another key insight relayed by Amin was drawn from the analysis of forest fires. Researchers in one of the six EPRI-funded consortia found these fires to have “failure-cascade” behavior similar to that of electric power grids. In a forest fire, the transformation of a spark into a conflagration depends on the proximity of the trees to one another. If just one tree in a barren field is hit by lightning, it burns but no big blaze results. But if there are many trees and they are close together, the single lightning strike can result in a forest fire that burns until it reaches a natural barrier such as a rocky ridge, river, or road. If the barrier is narrow enough that a burning tree can fall across it, or if it includes a burnable section, such as a wooden

bridge, the fire jumps the barrier and burns on. It is the role of first-response wild-land firefighters such as smoke jumpers to contain a small fire before it spreads by reinforcing an existing barrier or scraping out a defensible fire line barrier around the original blaze.

Similar outcomes can be observed for failures in electric power grids. For power grids, the “one-tree” situation is one in which every single electric socket has a dedicated wire connecting it to a dedicated generator. A lightning strike on any wire would take out that one circuit and no more. But the efficient use of resources argues against such a system, and instead favors one in which numerous sockets are served by a single circuit and there are multiple circuits for each generator. A failure anywhere in such a system causes additional failures until a barrier—a surge protector or circuit breaker, say—is reached. If the barrier does not function properly or is an insufficient impediment, the failure bypasses it and continues cascading across the system.

These findings suggest risk management approaches in which the natural barriers in power grids may be made more robust by simple design changes, or in which small failures might be contained by active smoke-jumper-like controllers before the failures grow into large problems. Other research

Echoing the familiar tale of the kingdom that was lost for want of a horseshoe nail, we might say that [independent] agents are like the missing nail: once restored, they can save an entire kingdom.

into the fundamental theory of complex interactive systems is exploring methods of quickly detecting weak links and failures within a system. Phased risk assessments have been very helpful in this regard. That is, experience indicates the value of performing “coarse-grained” risk assessments to identify important contributors. Rather than considering fifty initiating events for crisis scenarios, one might collapse them into five or six key events, and then focus on what is most important.

According to Amin, work over the past nine years in this area has led to a new vision for the integrated sensing, communications, and control of the power grid. Some of the pertinent issues are why and how to develop protection and containment devices for centralized as opposed to decentralized control and questions involving adaptive operation and the resistance to various destabilizers. In researching these issues, EPRI has refrained from conducting “in vivo” societal tests, which can be disruptive, and has instead

performed extensive simulation testing (in silico) of devices and policies in the context of the whole system. The EPRI simulations have produced a greater understanding of how policies, economic designs, and technology might fit into the continental grid (while exposing some unintended consequences of possible designs and policies), and provided guidance on the effective deployment and operation of these resources.

To mitigate the risk of systemic failure, the electric grid can be engineered for robustness. Amin presented an example of intelligent adaptive “islanding,” which is a method for blocking contagion. His results were based on a simulation of a hypothetical major blackout similar to the August 1996 blackout in the western United States. The simulation results he displayed captured the steady decay in frequency from 60 hertz to less than 58 hertz, after which the system would have deteriorated into a blackout. This simulation covered 3.5 seconds of simulated time. Then the simulation was re-run with major power lines eliminated between Arizona and Southern California to halt the contagion that led to the simulated blackout. As a result, the Western Interconnect grid was broken into two self-sustaining islands. Amin simulated more than 12,000 cases to stress-test the islands, and found that they consistently withstood the damaging contagion. With intelligent islanding (isolation) shortly after a major system disruption, the frequency recovered to close to 60 hertz before a blackout could occur.

This example also illustrates the practice in some engineering risk analysis of identifying undesirable outcomes first and then developing the fault trees and associated probabilities that could lead to those outcomes. The engineering community extensively employs both inductive reasoning (the event-tree thought process) and deductive reasoning (the fault tree) in its risk assessments. The most common approach is to use the event tree to structure the scenarios and fault trees to quantify the split fractions of the event-tree branch points.

ANALOGIES IN ECONOMICS AND FINANCE

Vincent Reinhart of the Board of Governors of the Federal Reserve System commented on three general forms of nonlinearity that are important to systemic risk. First, he noted that the consequences of events in the financial sector are likely nonlinear. Therefore, in designing and enforcing laws and regulations, the goal should not be to minimize the probability of every adverse event, but to guard against those that have

more severe consequences: In other words, the risk probabilities have to be weighted by some measure of the welfare gain that would arise from the prevention of each serious adverse event. That is the point of the partitioned multiobjective risk method, which—as we saw earlier—is designed to measure and analyze the risk of extreme and catastrophic events.

In a second form of nonlinearity, some economic processes are self-reinforcing. That is, in the run-up to a crisis, the size or transmission of some events may be amplified. Margin calls

If intermediaries restrict the availability of credit and therefore weaken spending, that action becomes the “financial accelerator.” These self-reinforcing effects are similar to those that can occur in the power grid when lightning strikes.

may cause selling that forces prices down more sharply, leading to a “fire sale.” Concerns about collateral values or an uncertain stock of capital may reduce arbitrage. If intermediaries restrict the availability of credit and therefore weaken spending, that action becomes the “financial accelerator.”⁹ These self-reinforcing effects are similar to those that can occur in the power grid when lightning strikes.

Trading activity can exhibit this second form of nonlinearity. Consider a simple model in which two people go to a market to trade. The amount of resources that one person commits to trading depends on the amount that the other person is expected to bring. This situation leads to collective decision making, which can be a highly nonlinear process in which small changes in cost bring about large changes in overall market activity. Indeed, trading could dry up altogether.

The third form of nonlinearity described by Reinhart was the dependence of some economic processes on the expectations of the players. This dependence can make prediction very difficult and implies that there might be multiple equilibria. How the market mechanism chooses among these equilibrium outcomes may be unclear. As a result, randomness and the sequence of events matter, suggesting that the way policy decisions are communicated during the run-up to a crisis can have an important influence on how the crisis

⁹The term “financial accelerator” refers to how endogenous developments in credit markets can amplify shocks in an economy. See Bernanke, Gertler, and Gilchrist (1996).

plays out. It also means that some techniques from the physical sciences are not directly transferable to economic and financial risk—the odds on a 100-year storm do not change because people think that such a storm has become more likely.

Reinhart also noted that, in a simple economic model, positive feedback can be destabilizing. But if one introduces an asset that is priced in a forward-looking manner, positive feedback is a mechanism for selecting a unique equilibrium. In those same models, negative feedback introduces the possibility of multiple equilibria—as was well known thirty years ago.

Levin observed that, in contrast to management of the electric power grid, there are only coarse or indirect options for control of the financial system. The tools available to policymakers—such as those used by central banks—are designed to modify individual incentives and individual behaviors in ways that will support the collective good. Such top-down efforts to influence individual behaviors can often be effective, but it is still difficult to control the spread of panic behavior or to manage financial crises in an optimal way. Within the financial system, robustness is something that emerges; it cannot be engineered.

Levin also noted that the key determinants of robustness—diversity and heterogeneity—are the same for biological, engineered, and financial systems. The influenza virus is robust because it takes on diverse forms; the analogue in the financial sector is the variety of institutions and remediation mechanisms, which makes the financial system more resistant to large-scale failures. In both cases, the system is able to adapt to change. But some redundancy—the ability of one component to perform another’s function—is, of course, also important in these systems. Otherwise, the chance loss of one component could be catastrophic.

DISCUSSION

Robert Oliver of the University of California at Berkeley noted that both Haimés and Amin had an implicit taxonomy in their risk analysis methodology: they first ran a risk assessment and then explored risk management. Their talks gave some guidelines for carrying out that linear process. However, those talks did not illustrate how engineers also turn around risk analyses to guide redesigns of system architectures and topology and of the policies that are integral to system performance. Since that process could be of value to central bankers, Oliver asked for comments on how one might reach new insights on those design and architectural questions.

Haimés suggested that a good way to proceed is to ask first what can go wrong. Looking from many different perspectives (as engineer, economist, social scientist, and so forth), one can discover some things that have never been expected to go wrong. To identify systemic risks, one has to look at everything. Since no one can really capture all of the relevant perspectives, systemic risks must be assessed through consultations with multiple players, which ultimately converge on a picture of the most important risks.

David Levermore of the University of Maryland pointed out that large-scale, complex simulations as exemplified by the work of Haimés and Amin are only part of the process of analyzing systemic risk. In the physical and biological sciences,

To identify systemic risks, one has to look at everything. Since no one can really capture all of the relevant perspectives, systemic risks must be assessed through consultations with multiple players, which ultimately converge on a picture of the most important risks.

very tiny models, designed to build understanding, also play an important role.¹⁰ These models are comparable in spirit to the work described in part 2 of this volume, with one possible distinction: in the physical and biological sciences, researchers do not limit themselves to only those simple models that can be solved analytically. The simple models might have only three or four variables, or sometimes just one complicated nonlinear variable, but still be complex enough to preclude analytic solution. Thus, research in the physical and biological sciences might rely more on computation than is the case in macroeconomics. Some of this research entails large-scale computing, but one should also note that studies yielding highly significant insights, such as the studies in dynamical systems on the logistics map, have been undertaken on very simple computers.

Douglas Gale of New York University observed that in helping to identify speakers for the conference, he had looked for those who would discuss the theoretical research being done on financial stability. This emphasis may have given a biased picture of current research in economics. In fact, Gale noted, computational economics is a very large part of

¹⁰See, for example, May (2004, pp. 790-3) and Keeling et al. (2003). Both papers also illustrate the possible limitations of simple models.

economics, and economists typically make great use of data, a point that was echoed by Reinhart. But an effort to model an entire system, with the aim of learning how to control it better, is a very large-scale project and one that academic economists will not readily take on because of the way the profession is organized and financed. They could follow such a path, but it would require additional resources. Moreover, Gale expressed some doubts about whether a large-scale computational approach is the right way to look at a system. Instead, it might be more fruitful to divide that system into understandable and digestible pieces and then find ways of engineering the system to ensure its robustness without a central control. Such an undertaking would not require an ability to model the entire system, still less an ability to control the full model.

Sugihara noted that the reliance on simple models, abstracted from reality, can sometimes have misleading consequences. For instance, the ideal gas laws, which are a mainstay of the physical sciences, assume a certain kind of functional form that often invites researchers to fit a scattering of points to that form. But the reason for the scatter might be quite important, and simplistic laws can lead researchers to

An effort to model an entire system, with the aim of learning how to control it better, is a very large-scale project and one that academic economists will not readily take on because of the way the profession is organized and financed.

overlook it. In the study of fisheries, understanding the larger systemic context of an individual species—the web as opposed to the node—is very important. The presentation by Hyun Song Shin of Princeton University, Sugihara noted, explicitly addressed the web of claims and obligations. As researchers and policymakers in finance and economics continue to think in those larger terms, they are going to reach a fuller understanding of the reality of the problem. Robert Litzenberger of Azimuth Trust, however, pointed out the value of abstraction in research, citing Milton Friedman’s paper on positive economics, which assigns an important role to assumptions and modeling. In Friedman’s view, assumption allows the economist to abstract from the things that are less important in order to focus on the key variables. The economist’s model is not meant to offer realistic description, which can fail to have predictive power. Simple models can provide considerable insight and also produce very useful predictions. The ultimate test of an assumption is its predictive power.

Rather than choose between the extremes of simple and complex models, several conference participants endorsed the concept of nested hierarchical models. The collaboration between Morten Bech of the Federal Reserve Bank of New York, Walter Beyeler and Robert Glass of Sandia National Laboratories, and Kimmo Soramäki of the European Central Bank, described in part 4 of this volume, is a good example of what could be accomplished in that direction. Pursuing the notion of combining different types of models, Sugihara suggested the following steps to build on the foundations laid at the conference:

- Devise minimal (simple) models first to see how much real variation in the data can be explained. Examples might be Shin’s model of leverage, presented in part 2 of this volume, or agent-based models with simple sets of rules. The latter would include models that can reproduce certain statistical properties of aggregate price series, such as the model proposed by Lux and Marchesi (1999). The work in progress by Bech, Beyeler, Glass, and Soramäki on an agent-based model for the Fedwire payments network is a step in this direction. The importance of empirical validation should not be overlooked, and the meaning of the topological patterns uncovered by Bech and his collaborators needs to be understood. There is much to be learned from simple models that can elucidate the systemic risk problem at the most general level.
- Create more complex, mechanistic models to complement the simpler ones. This task aims for the ideal, and it needs to be done carefully and in tandem with the simple models. Nonlinearities in functional relationships fix the scale of the model mechanisms (aggregation problem) and can hinder the applicability of those models across different market scales: firm, industry, regional, national, and global. The difficulty of developing complex models is exemplified by the early efforts to develop ecosystem models. These models appeared to be very complex because they incorporated many variables. But the overall model behavior was essentially simple logistic growth: much of the apparent complexity did not add real insight. While the ecosystem models provide a note of caution, it is nevertheless the case that complex models *can* be built well.

Taking a broader view, Levermore noted that some conference speakers seemed to focus on avoiding systemic risk rather than managing the system. To evaluate risk quantitatively (a first step toward avoiding it, if that is indeed a realistic goal), one must be able to model the system to the point where it can be plausibly simulated. An example was Amin’s practice of testing various “islanding” schemes through simulation. Once that level of simulation capability is achieved, managing the system becomes easier. The primary benefit of

this modeling and simulation capability, then, might not lie in avoiding risk but in managing the economy more effectively. For example, if the capability could help craft a regulatory tool designed to manage risk, even if that tool could help the economy run only a fraction of a percent more efficiently, the benefit to society would be enormous, easily dwarfing any cost in developing such a capability. If this capability also helped us to *avoid* risk, it would be better still.

This discussion is not meant to imply that ecology and engineering have overcome all the difficulties associated with representing and analyzing complex adaptive systems. Assessing the state of such systems is an ongoing challenge,

We do not know how to anticipate the collapse of a system by looking at it and recognizing that something is not right. Are there ways to look at trends in the stock market and know when a collapse is coming? In the view of many observers, complex systems produce signals that will tell us when we are approaching a precipice.

as is determining precisely what to measure. The validation of models and verification of software remain major challenges. Computational problems—including how to decouple models into tractable components—are also a continuing source of concern. Amin pointed out that self-similar systems can be reduced, but complex systems such as the electric grid cannot. Researchers can use approximations to decouple complex systems, but it is difficult to analyze the errors thus introduced.

In this regard, Amin noted that if one can find parts of an engineered system—and presumably parts of other systems—that are weakly coupled in terms of the dynamics transferred through the system, then one can approximate those portions with standalone models. Such a strategy essentially reduces the complexity by dividing and conquering. These component models might assume a variety of forms: some might be empirical models fit to data, others might be physics-based or financial, and still others might include elements—such as human behavior and performance—that cannot be modeled. Haines observed that, as an alternative strategy, one can decouple the system at the lower level, model the lower level components or subgrids, and then impose a higher level coordination. In some cases, this can even be done with an

additional level of hierarchy. This type of decomposition is a very effective way of addressing complex systems. In either case, aggregating (composing) the outputs of these component models into an overall picture is very challenging. To model the electric grid, for example, researchers have parametrized some of the component models so as to provide input to the next level of modeling, using Bayesian analysis. Sensitivity analysis is used to validate the resulting models.

Amin emphasized the difficulty of identifying meaningful signals from complex systems. For example, when monitoring a large fraction of the U.S. electric grid, how can we discern whether a perturbation in the system is a natural fluctuation or a sign of catastrophic failure? Is it a naturally caused phenomenon, perhaps triggered by heat, high humidity, or strong demand in one portion of the grid, or is it actually an attack on the system or the precursor to a major disturbance? How close is it to a regime shift or system flip? These questions can be addressed only with detection systems that can call up all the data and perform data mining, pattern recognition, and statistical analysis to derive the probability that a catastrophic failure is either developing or occurring now.

This system-monitoring problem is exacerbated if the sharing of information is limited, as it is in the banking sector. Charles Taylor of the Risk Management Association asked Amin how one would monitor and control the reliability of the electric grid under the assumption that companies did not cooperate with each other but instead competed and did not share information. Amin said that such a situation would lead to a new control mechanism, and the logical question would be whether that mechanism would stabilize or destabilize the system. He pointed to a project undertaken by the Electric Power Research Institute in the late 1990s—the Simulator for Electric Power Industry Agents—which addressed such a case. The analysis, applied to four large regions of the United States, explored whether one could increase efficiency without diminishing reliability.¹¹ This preliminary analysis would need to be carried out with more data and realism in order to reach a definitive conclusion.

Levin identified particular challenges facing those who wish to understand systemic risk more fully. For instance, we would like to be able to develop structure-function relationships—meaning that one could take a snapshot of a system and infer something about its dynamic state. We do not know how to anticipate the collapse of a system by looking at it and recognizing that something is not right. Are there ways to look at trends in the stock market and know when a collapse is coming? In the view of many observers, complex systems produce signals that will tell us when we are approaching a

¹¹See Amin (2002).

precipice. But the unfolding of market disruptions is affected greatly by confidence, herding, and other behaviors that are not mirrored in risk assessments for complex engineered systems. Other questions include, How do we overcome the robustness of undesirable configurations, so as to make it easier to move out of them? How can we get systems out of potentially problematic settings, and how can we achieve desirable cooperative arrangements?

The tools are available to develop agent-based models of banking systems—models in which one builds in rules for the behavior of individual people or institutions. These models help us understand how individual behaviors become synchronized or integrated with one another and how they spread through the financial sector. Of course, there are many unknowns about these rules, and the gamesmanship and proactive moves probably figure more importantly in the financial sector than in ecology or engineered systems. This is just one set of tools, but there are others. Sugihara has developed an approach to nonlinear forecasting. John Doyle

of the California Institute of Technology and Jean Carlson of the University of California at Santa Barbara have done work on highly optimized tolerance in which they use a genetic algorithm to evolve the properties of systems. They consider a variety of systems with particular structures and feedback properties, expose them to perturbations, observe their recovery, and then—in the same way that one might “train” a chess-playing program—modify these systems until they become more tolerant of the disturbances to which they are exposed. Doyle and Carlson’s strategy offers a way to improve the structure of systems when the mathematics cannot be solved. Nevertheless, as the authors themselves point out, their approach does have a drawback: Systems that are engineered or have evolved to be tolerant of a particular set of disturbances often do so at the expense of their response to other classes of disturbances. Such systems are at once robust and fragile—an outcome that policymakers and researchers might wish to guard against as they seek better ways to manage risk and avert systemic failures.¹²

¹²See, for example, Zhou, Carlson, and Doyle (2002) and Carlson and Doyle (2002).

REFERENCES

- Amin, M.* 2002. "Restructuring the Electric Enterprise: Simulating the Evolution of the Electric Power Industry with Intelligent Adaptive Agents." In A. Faruqui and K. Eakin, eds., *ELECTRICITY PRICING IN TRANSITION*, 27-50. Boston: Kluwer Academic Publishers.
- Asbeck, E., and Y. Y. Haimes.* 1984. "The Partitioned Multiobjective Risk Method." *LARGE SCALE SYSTEMS* 6, no. 1: 13-38.
- Bascompte, J., P. Jordano, and J. M. Olesen.* 2006. "Asymmetric Coevolutionary Networks Facilitate Biodiversity Maintenance." *SCIENCE* 312, no. 5772 (April 21): 431-3.
- Baumbach, R., A. Knoll, and J. Sepkowski, Jr.* 2002. "Anatomical and Ecological Constraints on Phanerozoic Animal Diversity in the Marine Realm." *PROCEEDINGS OF THE NATIONAL ACADEMY OF SCIENCES* 99: 6854-9.
- Bernanke, B., M. Gertler, and S. Gilchrist.* 1996. "The Financial Accelerator and the Flight to Quality." *REVIEW OF ECONOMICS AND STATISTICS* 78, no. 1 (February): 1-15.
- Carlson, J. M., and J. Doyle.* 2002. "Complexity and Robustness." *PROCEEDINGS OF THE NATIONAL ACADEMY OF SCIENCES* 99, suppl. 1: 2538-45.
- Couzin, I. D., J. Krause, N. R. Franks, and S. A. Levin.* 2005. "Effective Leadership and Decision-Making in Animal Groups on the Move." *NATURE* 433, no. 7025 (February 3): 513-6.
- Keeling, M. J., M. E. J. Woolhouse, R. M. May, G. Davies, and B. T. Grenfell.* 2003. "Modelling Vaccination Strategies against Foot-and-Mouth Disease." *NATURE* 421, no. 6919 (January 9): 136-42.
- Levin, S.* 2000. *FRAGILE DOMINION*. New York: Perseus.
- Lux, T., and M. Marchesi.* 1999. "Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market." *NATURE* 397, no. 6719 (February 11): 498-500.
- May, R. M.* 2004. "Uses and Abuses of Mathematics in Biology." *SCIENCE* 303, no. 5659 (February 6): 790-3.
- National Research Council.* 2005. *MATHEMATICS AND 21ST CENTURY BIOLOGY*. Washington, D.C.: National Academies Press.
- Zhou, T., J. M. Carlson, and J. Doyle.* 2002. "Mutation, Specialization, and Hypersensitivity in Highly Optimized Tolerance." *PROCEEDINGS OF THE NATIONAL ACADEMY OF SCIENCES* 99: 2049-54.

The views expressed in this summary do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.