Summary of presentation by Prof. Massoud Amin and related comment from

New Directions for Understanding Systemic Risk:

A report on a Conference Cosponsored by the Federal Reserve Bank of New York and the National Academy of Sciences

For the NAS book and complete FRBNY report please see:
Economic Policy Review, Federal reserve Bank of New York, Vol. 13, Number 2, Nov. 2007
New Directions for Understanding Systemic Risk, 108 pp, Nat'l Acad. Press, Washington Dc, 2007

The stability of the financial system and the potential for systemic events to alter the functioning of that system have long been important topics for central banks and the related research community. Developments such as increasing industry consolidation, global networking, terrorist threats, and an increasing dependence on computer technologies underscore the importance of this area of research. Recent events, however, including the terrorist attacks of September 11th and the demise of Long Term Capital Management, suggest that existing models of systemic shocks in the financial system may no longer adequately capture the possible channels of propagation and feedback arising from major disturbances. Nor do existing models fully account for the increasing complexity of the financial system's structure, the complete range of financial and information flows, or the endogenous behavior of different agents in the system. Fresh thinking on systemic risk is, therefore, required.

In order to promote a better understanding of systemic risk, the National Academy of Sciences and the Federal Reserve Bank of New York convened a conference in New York in May of 2006 drawing together a broadly interdisciplinary group of scientists, engineers and financial practitioners, ranging from electrical engineers and academic economists to risk analysts and asset managers from major investment banks. The primary purpose of the conference was to promote a cross-disciplinary dialogue in order to examine what possible leverage on the topic of systemic risk could be gained from areas of science not directly related to finance or economics. Accordingly, conference participants from the natural and mathematical sciences and from engineering disciplines drew heavily upon research on complex adaptive systems in order to build a framework both to give some substance and definition to the notion of systemic risk and to point to possible linkages between this research and research on the financial system. Similarly, research economists presented papers that showed how some of these linkages could be leveraged, for example in studies of international trade and, crucially for Federal Reserve policy, in the management of the payments system. Participants from the financial industry also highlighted how thinking on systemic risk and actual systemic events affect trading activities in order to provide a context for the discussion.

Systemic risk as a generic problem

In the world at large, complex systems abound, and concerns about the instability of these systems and their potential for large and potentially catastrophic regime shifts are a dominant social concern. These worries are at the leading edge of many environmental and engineering sciences: for example, in atmospheric science in studies of climate change; for fisheries managers concerned with the sudden collapse of certain economically important fish stocks; as Massoud Amin illustrated, in electrical engineering concerned with preventing disruptions to the North American power grid.

… The commonality of the problem of stability and resilience to shocks in complex systems that these examples point to raises the possibility that approaches to risk management in natural and physical systems could be pertinent to financial risk management. Work presented by Mr. Amin and Mr. Haimes illustrated some of the methods for managing risk in engineering systems, such as “multi-objective trade-off analysis” in which Pareto-optimal actions are derived by considering the subjective probabilities and payoffs associated with different shocks. The methods presented bore some resemblance to methods in financial risk analysis, and much of the subsequent discussion centered around the prospect of integrating methods from various engineering fields into financial risk management.

Prediction and Management of Systemic Failure in the Electric Grid

At the conference, Massoud Amin of the University of Minnesota extended the discussion of risk assessment, modeling and prediction. He described the North American electric power grid, which is another complex system. While it 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 postmortem analysis of major blackouts often shows the root cause to be the failure of one or a few components as the root cause (out of thousands in the portion of the grid ultimately affected), which upsets an equilibrium and leads to a cascade of failures. For example, on August 10, 1996, the North America experienced a major blackout affecting over 7 million customers in 13 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.

Prof. Amin reported that some studies of the August 1996 blackout mentioned above estimated that the blackout could have been avoided if the grid had intelligent controls and was able to shed 0.4% of its load for 30 minutes. Of course, the technologies for recognizing the incipient problem and tailoring a solution are far from obvious.

In an engineered system, like the electric power grid or a telecommunication network, there is indeed the opportunity for control systems, and these can be quite advanced. Prof. Amin reported at the conference on research funded by him in the 1990s at 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 all have to be at the same frequency, 60 hertz, controlling those elements in a coordinated fashion is somewhat analogous to controlling aircraft 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 part of a wing is damaged, which can be done. Accordingly, EPRI’s research is directed toward a control system that has some self-healing capability, meaning it would be able to anticipate disruptive events (identify signals that something important has changed), conduct a real-time assessment of the changing state of the system, determine how close we are to some “edge” in the performance envelope, 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 to control the financial system during disruptions.

Creating such a control capability for the electric grid required a mixture of tools from dynamical systems, statistical physics, information and communication science, along with research to reduce the computational complexity of the algorithms so they can scale up with the large size of the system being controlled. The EPRI research program has led to working methods that has been applied to a variety of situations, including the electricity infrastructure coupled with telecommunications and the energy markets, cell phone networks on the Internet, and some biological systems. This is a very multiscale challenge: detection of troublesome signals must be done within milliseconds, with some compensatory actions taken automatically, while some load balancing and frequency control on the grid is controlled on a timescale of seconds. At the same time, control functions such as load forecasting and management and generation scheduling take place on a timescale of hours or days. Developing a picture at the atomic level of what is going on in a system and then building up to the macroscale is the challenge that requires multiresolution modeling in both space and time.

Just to give an idea of the complexity of modeling and controlling the electrical grid, Prof. Amin gave some basic facts. In North America, there are more than 15,000 generators, 240,000 miles of high-voltage lines. The overall grid is divided in 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 single element. To determine whether the grid can attain that design goal, we need to simulate the loss of each of 53,000 elements and calculate the effects on each of 50,000 lines, leading to over 2.6 billion cases. The analysis of these systemic risks is very challenging, but it can really make a difference in how to operate the system.

As an additional illustration of the level of detail that can successfully be modeled, Prof. Amin presented an example of a complex model to predict load and demand for DeKalb, Illinois, which is a sizeable market 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 there has been a good deal of research to understand this phenomenon and develop the means to monitor and control it. The models and algorithms are now good enough to simulate the demand by customer type (residential, small commercial, large commercial) on an hour-by-hour basis and attain 99.6-99.7 percent accuracy over the entire year. One value of these predictions is that they enable the power company to proactively dispatch small generators to meet anticipated high demands.

From a broader perspective, any critical national infrastructure typically has many layers and decision-making units and is vulnerable to various types of disturbances. Effective, intelligent, distributed control is required that would enable parts of the constituent networks to remain operational and even automatically reconfigure 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 an elaborate operations control center. But all of these are liable to disruption at the very time when they are most needed (i.e., 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 whatever local resources remain to them in ways consonant with the established global goals to minimize 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 intelligently limiting their messages to only that information necessary to achieve global optimization and facilitate recovery after failure.

If organized in coordination 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-from "horseshoe nail" up to "kingdom"-and perform preprogrammed self-healing actions that require an immediate response. Such simple agents already are embedded in many systems today, such as circuit breakers and fuses as well as diagnostic routines. The observation is that we can definitely account for lose nails and to save the kingdom.

Another key insight came out of analysis of forest fires, which researchers in the one of the six funded consortia found to have similar "failure-cascade" behavior to electric power grids. In a forest fire the spread of a spark into a conflagration depends on how close together are the trees. If there is just one tree in a barren field and it is hit by lightning, it burns but no big blaze results. But if there are many trees and they are close enough together-which is the usual case with trees because Nature is prolific and efficient in using resources-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 it includes a burnable flaw 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 smokejumpers 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 results hold for failures in electric power grids. For power grids, the "one-tree" situation is a case in which every single electric socket had 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 like trees in Nature, electrical systems are designed for efficient use of resources, which means numerous sockets served by a single circuit and multiple circuits for each generator. A failure anywhere on the 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 insufficiently large, the failure bypasses it and continues cascading across the system.

These findings suggest approaches by which the natural barriers in power grids may be made more robust by simple design changes in the configuration of the system, and eventually how small failures might be contained by active smokejumper-like controllers before they grow into large problems. Other research into fundamental theory of complex interactive systems is exploring means of quickly identifying weak links and failures within a system.

Work during the past nine years in this area have developed, among other things, a new vision for the integrated sensing, communications, and control of the power grid. Some of the pertinent issues are why/how to develop protection and control devices for centralized versus decentralized control and issues involving adaptive operation and robustness to various destabilizers. However, instead of performing in Vivo societal tests which can be disruptive, we have performed extensive "wind-tunnel" simulation testing (in Silico) of devices and policies in the context of the whole system along with prediction of unintended consequences of designs and policies to provide a greater understanding of how policies, economic designs and technology might fit into the continental grid, as well as guidance for their effective deployment and operation.

In order to mitigate the risk of systemic failure, the electric grid can be engineered for robustness. Amin presented at the conference an example of intelligent adaptive "islanding," which is a method for blocking contagion. It is possible to simulate the chain of events in a hypothetical major blackout, similar to the one in August 1996 that affected the Western United States. Amin's simulation captured the steady decay in frequency from 60 Hz down below 58 Hz, after which the system would have deteriorated into a blackout. This simulation covers 3.5 seconds of simulated time. Then the simulation was re-run with major lines between Arizona and Southern California taken out to halt the contagion that led to the simulated blackout. This means the Western Interconnect grid was broken into two self-sustaining islands. Amin simulated more than 12,000 cases to stress test the islands, finding that they consistently withstood the damaging contagion. With intelligent islanding (isolation) shortly after a major system disruption, the frequency recovered to close to 60 Hz before a blackout could occur. This example also serves to illustrate the paradigm 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.

Analogies in Economics and Finance

Vince Reinhart from the Federal Reserve Board staff commented on three general ways in which nonlinearity is important to systemic risk. First, he noted that the consequences of events in the financial sector are likely nonlinear. So in designing and enforcing laws and regulations, the goal should not be to minimize the probability of every adverse event, but to lean especially hard against those that have more severe consequences: the risk probabilities have to be weighted by some measure of the welfare or safety attributable to prevention of each serious adverse event. That is the point of the Partitioned Multiobjective Risk Method (PMRM), which is designed to measure and analyze risk of extreme and catastrophic events. In asset-pricing terms, the PMRM uses probabilities in the martingale rather than the physical measure.

Second, some economic processes are self reinforcing. That is, in the runup to a crisis, the size or transmission of some events may be amplified. Margin calls 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 damp spending, that is the financial "accelerator." A lot of this sounds similar to what goes wrong in the power grid when lightening strikes.

A particular example of this is trading activity, which can exhibit very nonlinear behavior. Consider a very simple model, where two people go to a market to exchange. How many resources one person commits to trading depends on how many resources the other person is expected to bring, leading to collective decision-making, which can be a highly nonlinear process where small changes in cost have large changes in overall market activity. Indeed, trading could dry up altogether.

Reinhart offered the possibility of a third type of nonlinearity of particular importance to the financial sector, in that some economic processes are quite dependent on the expectations of the players, which can make prediction very difficult. This implies that there might be multiple equilibriums, and how the market mechanism chooses among them may be tenuous. As a result, randomness and the sequence of events matter, which suggests that the way policy decisions are communicated during the runup to a crisis can have an important influence over how the crisis plays out. It also means that some of the techniques from the physical sciences are not directly transferable—the odds on a 100-year storm don't change because people think it 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 to select a unique equilibrium. In those same models, negative feedback introduces the possibility of multiple equilibriums, which was well known 30 years ago.

Levin observed that there are only coarse options for control of the financial system. Mechanisms such as those within the Federal Reserve try to modify individual incentives and individual behaviors in ways that will be in the collective good. That is a top-down effort to influence some individual behaviors, but we still can't control the spread of panic behavior; we can't always control the system the way we want. The robustness is something that is emergent rather than engineered.

In the financial sector, diversity and heterogeneity translate as multiple institutions and remediation mechanisms, flexibility, and the ability to adapt to change. But some redundancy is, of course, also important in the system.

Discussion

Bob Oliver, past chair of Fair Isaac, noted that Haimes and Amin both had an implicit taxonomy in their risk analysis methodology, in which they ran a risk assessment and then explored risk management. Their talks gave some guidelines of how to go about that linear process. One thing that was missing, which would be really important for central bankers, is how the risk analysis can be turned around to guide a redesign of the architecture and topology of the system under study, as well as the policies that are integral to the system performance. There can be bad systems with good policies, good systems with bad policies, and all combinations in between. How does one get really new insights on those design and architectural questions?

Yacov Haimes suggested that a good way to proceed is to first ask the question of what can go wrong. Looking from many different perspectives (as engineer, economist, social scientist, etc.), one can discover some things that were never 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, assessments of systemic risks must be done through consultations with multiple players, iterating the process, and ultimately converging on a picture of the most important systemic risks.

David Levermore of the University of Maryland pointed out that the large-scale, complex simulations exemplified by the work of Haimes and Amin is only part of the process of understanding systemic risk. In the physical and biological sciences, there is actually an enormous place for very tiny models as well, designed to build understanding. This is comparable in spirit to the work described in Chapter 2, with one possible distinction: in the physical and biological sciences, researchers do not constrain themselves to only those simple models that can be solved analytically. The simple models might have only 3-4 variables, or sometimes just one complicated nonlinear variable, yet 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 is large-scale computing, but one should also note the important studies in dynamical systems on the logistics map, for instance, that were done on very simple computers but still led to enormous insights.

Doug Gale pointed out that, as one of the people who lined up the macroeconomic speakers for the conference, he may have given a biased picture. In actual fact, computational economics is a very large part of economics, and economists typically make a lot of use of data, a point that was echoed by Vince Reinhart. But the idea of modeling an entire system, perhaps even learning how to control it better, is obviously a very large-scale project and is the sort of thing that academic economists are not really ready to take on because of institutional issues. They could follow such a path, but it would require some additional resources. However, 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 break it into understandable and digestible pieces and try to find ways of engineering the system so it is going to be robust without a central control. That does not necessarily require an ability to model the entire system, let alone control the entire model. Economists don't think about this very much, according to Gale. The common research paradigm is to be able to find everything in that model and then study the model rather than the world.

Sugihara noted that sometimes the reliance on simple models, abstracted from reality, can 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 tends to lead 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 overlook it. In studying fisheries, understanding the larger systemic context of an individual species is very important—understanding the web as opposed to the node. In Professor Shin's presentation he explicitly talked about the web of claims and obligations. As people in finance and economics think more in those larger terms, they are going to get a lot more leverage into the reality of the problem. Litzenberger, though, pointed out the value of research such as Milton Friedman's on positive economics, where he spoke about the role of assumptions of modeling in economics. Friedman's model is not meant to capture realistic details, which don't always provide predictive power. Simple models can give a lot of insight and also produce very useful predictions.

Rather than choose between the poles of simple and complex models, several conference participants endorsed the concept of nested hierarchical models. The Sandia-FRBNY collaboration described in

Chapter 4 is a good specific example of what could be accomplished in that direction. George Sugihara suggested the following steps to build on the foundations laid out at the conference:

(i) Make 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 Chapter 2 or agent-based models with simple sets of rules. The latter would include models such as the one proposed by Lux and Marchesi (Nature 1999) which have two simple kinds of behavioral agents, but which can reproduce certain statistical properties of aggregate price series. The work in progress by the FRBNY-Sandia team on building an agent-based model for the Fed Payments Network, is a step in this direction that needs to be constrained more closely by the empirical work they have done. The importance of empirical validation should not be overlooked, and the meaning of the topological patterns uncovered by Soramaki et al. needs to be understood. There is much to be learned from simple models that can inform the systemic risk problem at the most general level.

(ii) Make more complex mechanistic models to complement the simpler ones. This is an ideal to aim for, but 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 (generality) of this across different scales of markets (firm-industry-regional-national-global). Again this needs to be done carefully. Another cautionary experience is from early ecosystem models that appear at first glance to be very complex (they incorporated everyone's favorite variable, so people believed them). Despite the apparent complexity, the overall model dynamic was actually quite simple. Thus the apparent realistic detailed mechanisms were actually a kind of architectural decoration on a basic A-frame that was essentially simple logistic growth (e.g., the FORET model for forests in East Tennessee). Again, this is not to say this cannot be done well, but it needs to be done carefully and is very difficult.

More generally, Levermore noted that some conference speakers seemed to focus on avoiding systemic risk rather than managing the system. In order to quantitatively evaluate risk (a first step toward avoiding it, if that is indeed a realistic goal), we need to be able to model the system to the point where it can be realistically simulated. This was exemplified in Amin's example of testing various "islanding" schemes via simulation. But once we have that level of simulation capability, we are better able to manage the system. The primary benefit of this modeling and simulation capability therefore might not be in avoiding risk, but rather in managing the economy better. For example, if the capability could help craft a regulatory tool that is designed to manage risk, even if this tool is done in such a way as to help the economy run just a fraction of a percent more efficiently, the benefit to the society and to the world would just be enormous and dwarf any cost in developing such a capability. If this capability also helps us to avoid risk that otherwise would be there, that is great, too.

The chapter is not meant to imply that ecology and engineering have overcome all the challenges associated with representing and analyzing complex adaptive systems. Sensing the state of such systems is one ongoing challenge, as is the question of what to measure. Validation of models and verification of software remains a major challenge. There are major computational problems, including how to break models into tractable components. Amin pointed out that self-similar systems can be reduced, but not complex systems like the electrical grid. One can use approximations to decouple complex systems, but it is difficult to analyze the errors thus introduced.

Amin noted, in response to a question at the conference, that one can find parts of an engineered system—and presumably in other systems—that are weakly coupled in terms of the dynamics transferred through the system and then approximate those portions with standalone models. That can help us reduce the complexity by dividing and conquering. Alternatively, Prof. Haimes noted that one can essentially decouple the system at the lower level, model the lower-level components or sub-grids, and then impose a higher-level coordination. Sometimes this can be done with even another level of hierarchy. This sort of decomposition is a very effective way of addressing complex systems. Prof. Amin noted that these component models might have very different characteristics, with some being empirical models fit to data, some being physics-based, some being financial, and even some parts that cannot be modeled, such as human behavior and performance. Composing these models is a challenge. To model the electric grid, some of the component models are parametrized so as to provide input to the next level of modeling, using Bayesian analysis. Sensitivity analysis is used to validate the resulting models.

Prof. Amin emphasized the difficulty of identifying meaningful signals from complex systems. For example, when monitoring a large fraction of the U.S. electrical grid, how can we discern whether a perturbation in the system is a natural fluctuation or the signature of a catastrophic failure? Does it reflect a naturally caused phenomenon, perhaps triggered by heat, high humidity, or a high demand in one portion of the grid, or is it actually an attack on the system or the precursor to major disturbance? How close is it to a regime shift or system flip? That can only be addressed with detection systems that can pull up all the data, do data mining, pattern recognition, and then statistical analysis to derive the probability that we were sensing a catastrophic failure or a precursor of one.

This system monitoring problem is exacerbated if sharing of information is limited, as is the case in the banking sector. Charles Taylor asked Prof. Amin how one would monitor and control the reliability of the electrical grid under the assumption that companies did not cooperate with each other but, instead, competed and didn't share information. Amin said that such a situation would lead to a new control mechanism, and the logical question is whether that would stabilize or destabilize the system. He pointed to an EPRI project from the late 1990s, Simulator for Electric Power Industry Agents, which began exploring this case. The analysis was done for four large regions of the United States, and it explored whether one could increase efficiency without diminishing reliability. This concept would need to be scaled up in order to reach a definitive conclusion.[1]

Levin proposed a number of particular challenges facing those who with to better understand systemic risk. For instance, we'd like to be able to develop structure-function relationships—meaning that one could take a snapshot of a system and say something about what its dynamic state is. We don't 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 many complex systems there is an expectation that there are signals that can be observed over a short time scale to tell us when we are getting near a precipice. Other questions include how to overcome the robustness of undesirable configurations, so as to make it easier to

[1] See Amin, Massoud, Restructuring the Electric Enterprise: Simulating the Evolution of the Electric Power Industry with Adaptive Agents, Chapter 3 in *Market Based Pricing of Electricity*, A. Faruqui and M. Crew, eds., Kluwer Academic Publishers, Dec. 2002.

move out of them. How can we get systems out of settings that we think are heading us for trouble and get back from the precipice, and how can we achieve desirable cooperative arrangements?

The tools are available to develop agent-based models of banking systems, in which one builds in rules for individual behavior. Those individuals may be individual people or they may be individual institutions like banks. The models help us understand how individual behaviors become synchronized, become integrated with each other, spread on these networks. Of course there are still a lot of unknowns about those rules, and the gamesmanship and proactive moves are probably more of a factor in the financial sector than they are in ecology or engineered systems. This is just one set of tools, but there are others. George Sugihara has developed an approach to nonlinear forecasting. There is also the work on highly optimized tolerance that John Doyle and Jean Carlson have been developing in California, in which they basically use a genetic algorithm, a neural network approach 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 just as you would train a chess playing program, they modify these systems until these systems become more tolerant to the disturbances to which they are exposed. So that is a way, even when one can't solve the mathematics, that one can improve the structure of systems. The difficulty with these approaches, as Doyle and Carlson point out, is that systems become robust yet fragile in their terminology, meaning systems that are engineered or have evolved to be tolerant to a particular set of disturbances often do so at the expense of their response to other classes of disturbances, something that we have to be careful about in the design of systems.[2]

[2] See, e.g., T. Zhou, J. M. Carlson and J. Doyle, Mutation, specialization, and hypersensitivity in highly optimized tolerance, *Proceedings of the National Academy of Sciences* 99:2049-2054. 2002. and J. M. Carlson and J. Doyle, Complexity and robustness, *Proceedings of the National Academy of Sciences* 99 suppl. 1:2538-2545. 2002.

5
WHAT HAS BEEN LEARNED?

Complex systems abound, and many different disciplines are concerned with understanding catastrophic change in such systems. People who study atmospheric science are very interested in precipitous climate change, people in ecology are very interested in so-called regime shifts and precipitous ecological change, engineers design complex systems so as to lessen the risk of catastrophic failures. What opportunities exist to leverage this concern from across many different fields for the benefit of the central banks and financial authorities, the financial sector, and the nation's economy more generally? We focus on three principal areas: risk assessment, modeling and prediction, and mitigation.

Risk Assessment

The economists, central bankers and market practitioners and the scientists and engineers at the conference agreed in large part on key mechanisms that produce destructive instability in large systems. Positive feedbacks, such as the portfolio insurance and collateral and margin calls that drove the stock market down so dramatically in October 1987, are one such mechanism. Another such mechanism, synchrony, was mentioned by Simon Levin as possible in any complex adaptive system, sometimes with deleterious consequences, and several conference participants mentioned the increase in systemic risk that can come about when behaviors of various actors become too similar. Charles Taylor amplified this by contrasting banks' decision-making in the past and today: a number of years ago, while there was a high level of homogeneity in the mix of business taken on by banks, their access to information was less and their quantitative methods were less precise and more ad hoc. The result was that individual banks would differ in how they executed processes and how quickly they responded to changes in conditions. So there would have been heterogeneity of response to crisis. But now, as the banking system has become more integrated and the time lags driven out by efficiency measures, the system may be evolving in a direction that makes it more fragile in some respects.

One area in which financial economists' and market practitioners' approach differs from that of the contributions from engineering by Yacov Haimes and Massoud Amin was in identifying extreme events. The background paper, Don Kohn in his remarks and the ECB in a paper submitted to the conference all discussed how potential extreme events are identified through stress testing. This involves developing a model of an economic or market process, applying extreme values from the distribution of the drivers of the model, and examining the output. They acknowledge that a limitation of this approach is that it assumes that behavior in the model does not change dramatically under extreme conditions. That is in contrast to what market participants in Chapter 1 vividly described as the feeling of regime shift during the Asian/Russian/LTCM crisis of 1998.

Yacov Haimes and Massoud Amin outlined in Chapter 3 an approach that involves identifying possible extreme events—e.g., the electrical grid shuts down—and considering what set of circumstances could produce the failure. Haimes describes a systematic process using small models and arranging factors in a hierarchy that probes what failures, mechanisms and regime shifts in what combination might lead to catastrophic failure. This paradigm of identifying a range of possible bad outcomes (risks) and backtracking to estimate their probabilities and identify options for reducing their likelihood or lessening their impacts is a common one in engineering. It is in contrast to the paradigm in which a given set of conditions is stipulated and then one explores, via theory or simulation, how events might unfold in response to a given stimulus. Charles Taylor refers to the former paradigm as "looking through the wrong end of the telescope."

While Haimes's process inevitably involves intuition and judgment, the data-rich environment in which his methods are applied grounds his modeling sufficiently so that one can draw meaningful inferences, even if they are not susceptible to classical statistical tests. For example, this method can be used to refine estimates of unconditional and conditional probabilities and correlations as well as the measurement of impacts. These estimates allow the analyst to make informed judgments about factors that could trigger systemic collapse. The stacking, if not necessarily nesting, of models in tiers also allows the analyst to assess how behavioral changes during a regime shift impact the potential for catastrophic failure.

Central banks over the last two decades have devoted increasing resources to research and analysis of financial stability. A major purpose of these efforts is to identify potential triggers of instability: events, as well as market, policy, or institutional mechanisms that can generate instability or propagate it once the financial system is perturbed. The methodology used to manage risk in engineering may provide insights about how to more systematically identify areas of potential financial instability. Central banks may have an interest in evaluating these methodologies.

Modeling, Prediction, and Management

The conference generated lively discussion of differences in the approach to research in economics described in Chapter 2 and the research carried out in ecology and engineering as described in Chapter 3. Economists were impressed by the quantity and quality of data available to researchers in the examples cited by Levin, Haimes and Amin.

Research Culture and Directions

Doug Gale suggested that the conference brought out "a very striking contrast" between some excellent theoretical research on economics and pragmatic, holistic modeling of risk in engineered systems. The theoretical research was by young economists who are coming up with new ideas, new concepts for understanding very important phenomena. Although the panel of three talks can't represent the entire spectrum of economic research, he feels it demonstrated the building blocks, which are representative of the way in which economists go about their business. The engineering research by Amin and Haimes represents a very different approach. They engaged in very large-scale projects—comprehensive, holistic modeling of risk phenomena using real data—that aim at realism and at prediction and control of particular systems rather than understanding general principles of a more generic system. As a means to that end, Haimes stressed that these projects integrate different models, using a lot of different approaches and techniques, rather than just focusing on one model.

In Gale's view, economists select their research projects the way they do because there are incentives for them to pursue that course. They certainly know about a lot of the wide range of techniques mentioned in the course of the conference—neural networks, stochastic approximation, dynamical systems, optimal control, and so on—and they use them to the extent that it helps them do what they want to do. One can readily imagine adapting the kind of large-scale approaches undertaken by Haimes and Amin to model, predict, and control the financial system. So one logically asks why academic economists have not pursued that line of research, why they are not providing a foundation for understanding systemic risks. The primary reason is money: in academic economics, there just isn't funding for that kind of large-scale research.

The relatively low level of funding for research in economics has had a number of effects on the way in which the discipline is organized. It affects education, promotion and tenure, the publication process, and so on. If, for example, an academic economist wants to publish a paper in a top journal (which is very important for their professional recognition and advancement), that paper must normally be about one model and focus on economics rather than other issues. The paper typically must include a methodological innovation. The prestigious journals would not be interested in research that consists of applying well-known techniques or models to some very practical problem. In contrast, engineers (and scientists in some applied fields) have broader latitude in the type of research they pursue and the roles for which they are rewarded, in part because of a wider array of funding sources. While some engineering research is geared solely toward scholarly publications, other work (even by the same individuals) might consist of studies that inform very pragmatic decisions. The premier honorary society for engineers in the United States, the National Academy of Engineering, includes a mix of those who have advanced the academic foundations of their field and those who have advanced the profession in other ways, perhaps as founders or managers of major enterprises. As a result of these different cultures, economists end up working in very small teams on what are to some extent theoretical, as opposed to very practical, problems. Even when doing empirical work, what is called applied economics, or doing something which is ostensibly addressing issues of regulation or optimal policy, generally speaking economists do not have incentives to produce something that can be immediately applied in the real world. Economists are looking for insight, and that is a very different kind of activity. Gale said he could imagine a role for research into systemic risk, and he thinks it would very exciting.

There was some discussion at the conference about the level of resources devoted to understanding systemic risk, with several conference participants observing that the amount spent on studying systemic risk is a miniscule fraction of the amount spent on understanding and managing the risks of individual entities. Gale noted that a prerequisite for significant change in the types of research economists conduct is a large-scale shift in funding for the discipline. The need is not just to provide money for particular studies on the financial system or systemic risk, but changing an entire discipline, which means changing incentives across the field.

Reinhart raised the possibility that change could occur through revisiting scholarly work that had been overlooked by the profession. In that connection, Reinhart quoted work in 1992 by Simon Levin, "A popular fascination of theorists in all disciplines, because of the potential for mechanistic understanding, has been with systems in which the dynamics at one level can be understood as the collective behavior of aggregates of similar units." That is an appealing mechanism, if it were true. But it is not true for the financial system or an economy as a whole. The economy is a network of heterogeneous, not similar, agents. Instead of transmission lines, transformers, and switches, financial markets have individual investors with different strategies, market makers, brokers, market utilities, and beta providers. Economists have known for 30 years that heterogeneity cannot be assumed away: in *Micro Motives and Macro Behavior*, Nobel laureate Tom Schelling provided many examples of how individual behaviors produced clustering and self organization. This conference is evidence that the lure of a more mechanistic model is loosening.

Impact on Policymaking

David Levermore of the University of Maryland suggested that the ultimate benefit of the new directions suggested by the conference may not be so much for managing risk, which is an important component of course, but for understanding the economy better. Better models of systemic risk can incorporate and build on the theory and intuition of central bankers and economists and refine them through additional quantitative insight. For example, in designing a regulation that currently might affect all institutions of a certain type, future policymaking might include gradations such that maybe only big institutions are touched one way and smaller institutions are relatively unencumbered because their health doesn't constitute a systemic risk. Having that greater degree of latitude will allow policymakers to be more creative and productive. Vince Reinhart noted that such a tiered system is already emerging as a result of the Basel II Accord.

Charles Taylor of the Risk Management Association added fine-tuning to Levermore’s suggestion. The financial system has an enormous tendency to pursue efficiency, and so policymakers do not need better models geared toward system efficiency. Nor is there much reason for policymakers to worry about improving models of institutional risk management, because there are very strong incentives for institutions to manage their risks well. The public policy gap is in understanding how systems could evolve so as to be more robust to tail events. As Reinhart noted, though, we simply don’t have much data about tail events, by definition.

Litzenberger amplified that point. When we attempt to implement risk models for catastrophic periods, we want objective measures based in some way on historical data. But if the data pertains to just one event, that is a scenario analysis, and there is no statistical reliability with respect to the estimation. That is a major problem we face when we use sophisticated empirical techniques with very limited data to fully model the system. When we try to extend this thinking beyond the Fed funds system, with its good data, when we start going to the broader financial system, we could get very muddled. He compared the situation to econometric models of the U.S. economy that he learned about in graduate school. They were impressive, but in truth they never predicted very well, and many people eventually became very disillusioned with some of those models.

… In responding to systemic risk, one needs to think about the timeframe over which policy is expected to work. Reinhart speculated that the presence of portfolio insurance and dynamic hedging in 1987 might have been a market mechanism that tended to amplify the down trend. It is not obvious what a central bank could do in that event; the market was falling, and you just can't step into that process. The central bank was able to remind commercial banks that down streaming funds to investment banks would be a good thing, and it provided assurances about the availability of liquidity. The market went up subsequently, and the economy performed generally well despite the destruction of wealth associated with the initial stock price decline.

Reinhart asserted that quick action is the right thing to do, but there is nowhere near as much research to inform crisis management as there is to understand crisis propagation. He thought it would be appropriate to apply the sophistication of the work seen in the conference to crisis management as well. It is noteworthy that the research reported on by Haimes and Amin explicitly included exploration of policy alternatives, while the economic research reported at the conference was not geared toward informing policy. This might signal a difference in disciplinary culture, since in both cases the researchers were modeling systemic behaviors.

References

Amin and Schewe, "Preventing Blackouts," Scientific American, pp. 60-67, www.Sciam.com, May 2007

Amin and Gellings, "The North American power delivery system: Balancing market restructuring and environmental economics with infrastructure security," *Energy,* Vol. 31, Issues 6-7, pp. 967-999, May-June 2006

Amin and Wollenberg "Toward a Smart Grid," IEEE Power and Energy Magazine, Vol.3, No 5, pp. 34-38, Sept/Oct. 2005

Amin, "Energy Infrastructure Defense Systems," *Proceedings of the IEEE*, Vol. 93, Number 5, pp. 861-875, May 2005

Amin, "Restructuring the Electric Enterprise: Simulating the Evolution of the Electric Power Industry with Adaptive Agents," Chapter 3, pp. 27-50, in *Electricity Pricing in Transition,* Ahmad Faruqui and Kelly Eakin (editors), Kluwer Academic Publishers, September 2002

Amin, "National Infrastructures as Complex Interactive Networks," chapter 14 in: *Automation, Control, and Complexity: An Integrated Approach*, T. Samad & J. Weyrauch (Eds.), pp. 263-286, John Wiley and Sons Ltd., NY, March 2000

Amin, "Toward Self-Healing Infrastructure Systems," *IEEE Computer Magazine*, pp. 44-53, Vol. 33, No. 8, Aug. 2000

Amin, "Toward Self-Healing Energy Infrastructure Systems," *IEEE Computer Applications in Power*, pp. 20-28, Vol. 14, No. 1, January 2001

Amin, "Modeling and Control of Electric Power Systems and Markets," *IEEE Control Systems Magazine*, pp. 20-25, Vol. 20, No. 4, Aug. 2000

Amin and Ballard, "Defining New Markets for Intelligent Agents," *IEEE IT Professional*, pp. 29-35, Vol. 2, No. 4, July/Aug. 2000

Special Issue of Proceedings of the IEEE on Energy Infrastructure Defense Systems, (Guest editor: Amin), Vol. 93, Number 5, pp. 855-1059, May 2005

Special issues of IEEE Control Systems Magazine on Control of Complex Networks, (Guest editor: Amin), Vol. 21, No. 6, December 2001 and Vol. 22, No. 1, February 2002

Special issue of IEEE Control Systems Magazine on Power Systems and Markets, (Guest editor: Amin), pp. 20-90, Vol. 20, Number 4, Aug. 2000

Network, Control, Communications and Computing Technologies in Intelligent Transportation Systems, (Guest co-editors: Amin, Garcia-Ortiz and Wootton), *Mathematical and Computer Modeling*, Elsevier Science Ltd, Vol. 22, No. 4-7, 454 pp. , Aug.-Oct. 1995

Amin, "Electricity," in *Digital Infrastructures: Enabling Civil and Environmental Systems through Information Technology*, R. Zimmerman and T. Horan (Editors), Chapter 7, pp. 116-140, July 2004

Amin, "Balancing Market Priorities with Security Issues: Interconnected System Operations and Control under the Restructured Electricity Enterprise," IEEE Power and Energy Magazine, Vol.2, No 4, pp. 30-38, July/August 2004

Starr and Amin, "Global Transition Dynamics: Unfolding the Full Social Implications of National Decision Pathways," 11 pp., submitted to the President of the US National Academy of Engineering, September 2003