Part 5: Concluding Observations

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 look extensively at 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 great interest from across many fields for the benefit of the central banks and financial authorities, the financial sector, and the nation's economy more generally? The conference explored this question by focusing on three principal issues associated with catastrophic events in complex systems: risk assessment, modeling and prediction, and mitigation.

Risk Assessment

The economists, central bankers, market practitioners, and scientists and engineers at the conference agreed in large part on key mechanisms that produce instability in large systems. Positive feedback—such as the portfolio insurance and collateral and margin calls that may have played a role in driving the stock market down so dramatically in October 1987—is one such mechanism. Another, synchrony, was mentioned by Simon Levin of Princeton University as possible in any complex adaptive system, sometimes with deleterious consequences, and several conference participants pointed to the increase in systemic vulnerability that can come about when behaviors of various actors become too similar. Charles Taylor of the Risk Management Association amplified this idea in describing how banks' decision making has changed: A number of years ago, while there was a high level of homogeneity in the mix of business taken on by banks, their quantitative methods were less precise and more ad hoc—with some variation in the speed of their responses to events. The result was that individual banks would differ in how they executed processes and how quickly they responded to changes in conditions. Thus, there would be heterogeneity of response to crisis. But now, as the banking system has become more integrated and the time lags have been driven out by efficiency measures, in Taylor's view the system may be evolving in a direction that makes it more fragile in some respects.

One area in which the approaches of financial economists and market practitioners differ from those of engineers such as Yacov Haimes of the University of Virginia and Massoud Amin of the University of Minnesota is in identifying extreme events. The conference background paper[1] and the keynote remarks of Governor Kohn discussed how potential extreme events are identified through stress testing. This procedure 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. Those who commented on stress testing acknowledged that a limitation of this approach is its assumption that behavior in the model does not change dramatically under extreme conditions. This assumption conflicts with what market participants in part 1 of this volume vividly described as the

[1]The background paper can be found in Appendix B.

feeling of regime shift during the events of 1997-98: the Asian currency crisis, the Russian default, and the Long-Term Capital Management collapse.

Part 3 of this volume explains the approaches followed by Haimes and Amin for identifying possible extreme events—for instance, a shutdown of the electric grid—and considering what set of circumstances could produce the failure. Haimes described 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 impact 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, by means of theory or simulation, how events might unfold in response to a given stimulus. Taylor referred 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 affect the potential for catastrophic failure.

> *One area in which the approaches of financial economists and market practitioners differ from those of engineers . . . is in identifying extreme events.*

Central banks over the last two decades have increasingly devoted 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 disrupted. The methodology used to manage risk in engineering may provide insight into means of identifying areas of potential financial instability more systematically. 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, as illustrated in part 2 of this volume, and the research carried out in ecology and engineering, as glimpsed in part 3. Economists were impressed by the quantity and quality of data available to researchers in the examples cited by Levin and by Haimes and Amin.

Research Culture and Directions

Douglas Gale of New York University suggested that the conference brought out "a very striking contrast" between some excellent theoretical research in economics and the pragmatic, holistic modeling of risk in engineered systems. The theoretical research was by young economists who are coming up with new ideas and new concepts for understanding very important phenomena. Although the panel of three talks cannot represent the entire spectrum of economic research, Gale felt it demonstrated the theoretical building blocks that economists use when thinking about problems of financial instability. 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 at understanding general principles of a more generic system. As a means to that end, Haimes stressed that these projects integrate different models, using many different approaches and techniques, rather than just focusing on one model.

In Gale's view, the way economists select their research projects reflects their incentives to pursue that course. Economists certainly know about many of the techniques described in the course of the conference—neural networks, stochastic approximation, dynamical systems, optimal control, and others—and they use them to the extent that they help to accomplish their goals. One can readily imagine adapting the kind of large-scale approaches undertaken by Haimes and Amin to model the financial system. So, one logically asks why academic economists have not pursued that line of research—why they are not using such approaches to provide a foundation for understanding systemic risks. The primary reason is money: In academic economics, in Gale's view, no funding exists for that kind of large-scale research.

The relatively low level of funding for research in economics has had a number of effects on how the discipline is organized. It affects education, promotion and tenure, the publication process, and so on. If, for example, academic economists want to publish in a top journal, an achievement that is very important for their professional recognition and advancement, their papers must normally be about one model and focus on economics rather than other issues. The papers 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.

Engineers as well as scientists in some applied fields have more latitude [than economists] in the types of research they can pursue and the roles for which they are rewarded, in part because a wider array of funding sources exists.

In contrast, engineers as well as scientists in some applied fields have more latitude in the types of research they can pursue and the roles for which they are rewarded, in part because a wider array of funding sources exists. 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. Economists, operating in a very different culture, end up working in small teams on what are to some extent theoretical, as opposed to practical, problems. Even when conducting empirical studies—as in applied economics—or when addressing issues of regulation or optimal policy, economists generally do not have incentives to produce work that can be immediately applied. Economists are looking for insight, and that is a very different kind of activity. Gale indicated that he could imagine a role for research into systemic risk, one that would be very exciting.

Some discussion centered on 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 type 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 to change an entire discipline, which means changing incentives across the field.

Vincent Reinhart of the Board of Governors of the Federal Reserve System raised the possibility that change could occur through revisiting scholarly work that had been overlooked by the profession. In that connection, he quoted from work by Levin (1992): "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 market makers, brokers, market utilities, beta providers,[2] and individual investors with different strategies. Economists have known for thirty years that heterogeneity cannot be assumed away: In *Micro Motives and Macro Behavior*, Nobel Laureate Thomas 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 waning.

The Role of Data

Reinhart suggested that the difference between the research style of economists and that of engineers and physical scientists (at least as demonstrated at the conference) might revolve around data and computing power. As noted in part 3 of this volume, there is more of a tradition of data sharing, and more nonproprietary data with which to work, in engineering and the physical and life sciences. As economists gain access to large data sets—opening up the possibility of seeing redacted data on individual transactions and individual behavior, as exemplified by the Fedwire projects described in part 4—economists and financial economists will be driven to cooperate more. To the extent that economic researchers start developing more complex models to represent the heterogeneity of economic agents and combining them with large data sets—for instance,

[2]Beta providers are investors whose trading drives the prices of related assets to converge toward their normal relationships when prices diverge.

of individual transactions in markets—their work will likely become more computational, as has been the pattern in much of the natural sciences.

In studies of systemic risk in the financial sector, key data are transaction prices, transaction volumes and timing, financial institution position and exposure measures, and economic and other news. In centrally organized exchange markets, such as the New York Stock Exchange (NYSE), good data on prices, volumes, and timing are collected and could be used in research. In over-the-counter markets, where transactions are arranged between institutions and are not recorded centrally, electronic quotation and trading systems have improved the availability of price information. But a preponderance of information required to study systemic risk at some scale remains the proprietary information of financial institutions.

The central bankers, regulators, and economists were impressed by the cooperative arrangements in the electrical power generating industry for sharing proprietary information used in researching and managing systemic stability and the insight gained from using detailed data. As risk management and financial analysis have advanced over the last two decades, financial institutions have developed large databases of financial information. While financial firms are unlikely to share very recent data, the proprietary value of information in detailed financial institution data may decay fairly quickly, given the rapidity with which market conditions and market opportunities fade. If financial institutions share central bankers and regulators' interest in risk management tools, the examples of data sharing from other industries might be helpful in demonstrating the benefit of even a modest information-sharing effort.

Potential Applications to Policy

The conference also compared sources of robustness in financial and economic systems with those in ecological and engineering systems and considered the implications for mitigation. Several participants agreed that there is a need for more research into robustness strategies in preventing systemic events and for more analysis of the implications for policy responses when such events occur.

The sessions revealed some important differences in approaches to regime shift and hysteresis, with implications for mitigation. Charles Lucas of AIG (since retired), a member of the National Academy's Board on Mathematical Sciences and Their Applications opened the conference with a discussion of the dramatic and deleterious regime shift that occurred in the wake of the financial crisis of the late 1920s and early 1930s: the shift from the booming, but troubled, 1920s to the Great Depression. Economists considering systemic risk have wrestled with the questions of when a financial disturbance can or will lead to macroeconomic effects, when those macroeconomic impacts represent a new equilibrium for the economy, whether the shift to a new and inferior equilibrium is the result of financial disturbance or policy errors, and what sort of hysteresis—resistance to a return to the previous equilibrium state—exists.

While one response to financial crisis might be to shut markets down, under the implicit assumption that they are not strong enough to withstand the shock, financial economists and financial authorities generally recommend that markets remain open—a view based on their trust in the flexibility of markets.

The effects of some financial disturbances are seen as salutary by many economists and central bankers, leading to improved risk management and a better long-term allocation of resources, at least in some sectors. The banking problems of the early 1990s and the failure of Barings in 1995 have been widely cited as precipitating substantial innovation that improved credit and counterparty risk management.

Many economists cite the resilience of financial markets in handling disturbances, even long-run disturbances, principally through the effectiveness of the price mechanism, but also by creating new markets and contractual and institutional arrangements. Even if prices fall very sharply, revaluation of assets and liabilities, if allowed to occur, often results in markets finding a new equilibrium after transactions resume. That process may take weeks, as it did after the 1987 stock market crash, when even though prices rebounded sharply the next day, overall trading activity and international equity capital flows took about ten weeks to recover to normal levels. Or it may take longer, as it did after banks began writing down their real estate loans and selling them off in the early 1990s. Thus, the potential for regime shift and subsequent hysteresis as a result of systemic events in financial markets is to some extent offset by the flexibility and resilience of the markets in assessing and responding to systemic shocks.

Consistent with Levin's discussion of rigidity and flexibility as strategies to create robustness in ecosystems, financial systems appear to possess flexibility as a key bulwark of robustness. One challenge for financial markets is that the underlying infrastructure that manages the flow of transactions may have some inherent rigidity because of its legacy technologies and reliance on scale and network economies; another question is whether the flexibility of some activities is reduced by consolidation.

Rigidity and flexibility are opposite, but equally valid, strategies to achieve robustness: a system can either be strong enough to resist disturbances or it can be flexible enough to "bend" to them. These two strategies also map to differing perspectives on policymakers' appropriate response to financial disturbances. While one response to financial crisis might be to shut markets down, under the implicit assumption that they are not strong enough to withstand the shock, financial economists and financial authorities generally recommend that markets remain open—a view based on their trust in the flexibility of markets. There are circumstances in which markets have been suspended: In the immediate aftermath of the destruction of September 11, 2001, the equity markets remained closed for four days; the NYSE instituted circuit breakers for trading after the October 19, 1987, stock market crash; and banking holidays are sometimes declared during major weather events.

In responding to systemic risk, monetary and financial authorities need to think about the time frame 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 downtrend. It is not obvious what a central bank could do in that event; the market was falling, and the central bank could not just step into that process. It was able to remind commercial banks that downstreaming funds to investment banks would be a good thing, and it provided assurances about the availability of liquidity. The markets were kept open, trading resumed, and the markets rose subsequently; 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 step to take, but there is not nearly as much research available to inform crisis management as there is to understand crisis propagation. He thought it would be appropriate to apply the sophistication of the work presented at the conference to crisis management as well.

David Levermore of the University of Maryland suggested that the ultimate benefit of the new directions suggested by the conference might not apply so much to managing risk, which is an important component of course, but to understanding the economy better. Improved 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 redesigning a regulation that currently affects all institutions of a certain type, future policymakers might include gradations, such that perhaps only large institutions are affected while smaller institutions are relatively unencumbered because their health does not constitute a systemic risk. Having that greater degree of latitude will allow policymakers to be more creative and productive. Reinhart noted that such a tiered system is already emerging as a result of the Basel II Accord on bank capital requirements.

Taylor added that the public policy objective is to understand how systems can evolve so as to be more robust to tail events. As Reinhart noted, though, we simply do not have much data on tail events, by definition. Robert Litzenberger of Azimuth Trust 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 pertain to just one event, then that is a scenario analysis, and there is no statistical reliability with respect to its assessment. That is a major problem we face when we use sophisticated empirical techniques with very limited data to model the system fully. When we try to extend this thinking beyond the Fedwire system, with its good data, to the broader financial system, we run out of the data that would be needed if the models are to make useful predictions. Litzenberger compared the situation with that of econometric models of the U.S. economy that he studied in graduate school. They were impressive, but in truth they never predicted very well, and many researchers eventually became disillusioned with some of those models. To arrive at a better understanding of systemic risk and to improve risk management tools and policies, the discussion pointed to the immense potential value from developing rich data sets of financial information, financial asset prices, and institutions' risks and earnings.

References

Levin, S. A. 1992. "The Problem of Pattern and Scale in Ecology." Ecology 73, no. 6 (December): 1943-67.

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.