



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
John Kambhau, Scott Weidman, and Neel Krishnan,
Rapporteurs, National Research Council

ISBN: 0-309-10753-9, 120 pages, 6x9, (2007)

**This free PDF was downloaded from:
<http://www.nap.edu/catalog/11914.html>**

Visit the [National Academies Press](#) online, the authoritative source for all books from the [National Academy of Sciences](#), the [National Academy of Engineering](#), the [Institute of Medicine](#), and the [National Research Council](#):

- Download hundreds of free books in PDF
- Read thousands of books online, free
- Sign up to be notified when new books are published
- Purchase printed books
- Purchase PDFs
- Explore with our innovative research tools

Thank you for downloading this free PDF. If you have comments, questions or just want more information about the books published by the National Academies Press, you may contact our customer service department toll-free at 888-624-8373, [visit us online](#), or send an email to comments@nap.edu.

This free book plus thousands more books are available at <http://www.nap.edu>.

Copyright © National Academy of Sciences. Permission is granted for this material to be shared for noncommercial, educational purposes, provided that this notice appears on the reproduced materials, the Web address of the online, full authoritative version is retained, and copies are not altered. To disseminate otherwise or to republish requires written permission from the National Academies Press.

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**

John Kambhu, Scott Weidman, and Neel Krishnan, *Rapporteurs*

NATIONAL RESEARCH COUNCIL
OF THE NATIONAL ACADEMIES

THE NATIONAL ACADEMIES PRESS
Washington, D.C.
www.nap.edu

THE NATIONAL ACADEMIES PRESS 500 Fifth Street, N.W. Washington, DC 20001

NOTICE: The project that is the subject of this report was approved by the Governing Board of the National Research Council, whose members are drawn from the councils of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine.

This study was supported by a contract between the National Academy of Sciences and the Federal Reserve Bank of New York. The views expressed in this report and during the conference "New Directions for Understanding Systemic Risk" represent the authors' own views and do not represent the views of the Federal Reserve Bank of New York, the Federal Reserve System, or the United States government.

Cover Illustration: The background image (in light blue) is adapted, by permission, from Figure 3 of Christopher J. Neely, 2004, "The Federal Reserve Responds to Crises: September 11th Was Not the First," *Federal Reserve Bank of St. Louis Review* 86, no. 2 (March/April): 27-42. The six curves represent daily financial data in the months surrounding the stock market crash of October 19, 1987. The top five curves show the value of the S&P 500 index, the volatility of the New York Stock Exchange as implied by options prices, the yield on 10-year U.S. government bonds, the trade-weighted value of the dollar, and the federal funds rate target against three targets (shown as white dashed lines). The bottom graph shows U.S. official foreign exchange intervention, in purchases of millions of dollars. The vertical line through the six graphs denotes the date of the stock market crash.

International Standard Book Number-13: 978-0-309-10752-5

International Standard Book Number-10: 0-309-10752-0

Copies of this report are available from the National Academies Press, 500 Fifth Street, N.W., Lockbox 285, Washington, DC 20055; (800) 624-6242 or (202) 334-3313 (in the Washington metropolitan area); Internet, <http://www.nap.edu>.

This report is also being republished by the Federal Reserve Bank of New York.

Copyright 2007 by the National Academy of Sciences. All rights reserved.

Printed in the United States of America

THE NATIONAL ACADEMIES

Advisers to the Nation on Science, Engineering, and Medicine

The **National Academy of Sciences** is a private, nonprofit, self-perpetuating society of distinguished scholars engaged in scientific and engineering research, dedicated to the furtherance of science and technology and to their use for the general welfare. Upon the authority of the charter granted to it by the Congress in 1863, the Academy has a mandate that requires it to advise the federal government on scientific and technical matters. Dr. Ralph J. Cicerone is president of the National Academy of Sciences.

The **National Academy of Engineering** was established in 1964, under the charter of the National Academy of Sciences, as a parallel organization of outstanding engineers. It is autonomous in its administration and in the selection of its members, sharing with the National Academy of Sciences the responsibility for advising the federal government. The National Academy of Engineering also sponsors engineering programs aimed at meeting national needs, encourages education and research, and recognizes the superior achievements of engineers. Dr. Charles M. Vest is president of the National Academy of Engineering.

The **Institute of Medicine** was established in 1970 by the National Academy of Sciences to secure the services of eminent members of appropriate professions in the examination of policy matters pertaining to the health of the public. The Institute acts under the responsibility given to the National Academy of Sciences by its congressional charter to be an adviser to the federal government and, upon its own initiative, to identify issues of medical care, research, and education. Dr. Harvey V. Fineberg is president of the Institute of Medicine.

The **National Research Council** was organized by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purposes of furthering knowledge and advising the federal government. Functioning in accordance with general policies determined by the Academy, the Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in providing services to the government, the public, and the scientific and engineering communities. The Council is administered jointly by both Academies and the Institute of Medicine. Dr. Ralph J. Cicerone and Dr. Charles M. Vest are chair and vice chair, respectively, of the National Research Council.

www.national-academies.org

BOARD ON MATHEMATICAL SCIENCES AND THEIR APPLICATIONS

C. DAVID LEVERMORE, University of Maryland, *Chair*
MASSOUD AMIN, University of Minnesota
TANYA STYBLO BEDER, SBCC
MARSHA BERGER, New York University
PHILIP A. BERNSTEIN, Microsoft Corporation
PATRICIA FLATLEY BRENNAN, University of Wisconsin-Madison
PATRICK L. BROCKETT, University of Texas at Austin
DEBRA ELKINS, General Motors R&D Center
LAWRENCE CRAIG EVANS, University of California at Berkeley
JOHN GEWEKE, University of Iowa
DARRYLL HENDRICKS, UBS Investment Bank
JOHN E. HOPCROFT, Cornell University
CHARLES M. LUCAS, AIG (retired)
CHARLES MANSKI, Northwestern University
JOYCE McLAUGHLIN, Rensselaer Polytechnic Institute
JILL P. MESIROV, Massachusetts Institute of Technology
ANDREW M. ODLYZKO, University of Minnesota
JOHN RICE, University of California at Berkeley
STEPHEN M. ROBINSON, University of Wisconsin-Madison
GEORGE SUGIHARA, University of California at San Diego
EDWARD WEGMAN, George Mason University
LAI-SANG YOUNG, New York University

Staff

SCOTT WEIDMAN, Director
NEAL GLASSMAN, Senior Program Officer
BARBARA WRIGHT, Administrative Assistant

For more information on BMSA, see its Web site at <<http://www7.nationalacademies.org/bms/>>, write to BMSA, National Research Council, 500 Fifth Street, N.W., Washington, DC 20001, call at (202) 334-2421, or e-mail at bmsa@nas.edu.

Preface

Guarding against systemic risk in the financial system is a key undertaking for central banks. Defining this type of risk is difficult, but managing it with precision is harder still. Complicating this task is the fact that institutional consolidation, a broadening range of financial products, and greater connectivity among firms have in recent decades materially changed the nature of systemic risk in the financial system.

To stimulate fresh thinking on systemic risk, the Federal Reserve Bank of New York and the National Research Council's Board on Mathematical Sciences and Their Applications held a conference, "New Directions for Understanding Systemic Risk," in May 2006. The main goal of the sessions was to explore parallels between systemic risk in the financial sector and in selected domains of engineering, ecology, and other fields of science. The event attracted more than 100 experts on systemic risk from 22 countries, representing banks, regulators, investment firms, U.S. national laboratories, government agencies, and universities. In addition to bringing together many participants with backgrounds in banking, finance, and economics, the conference broadened the discussions by including the perspectives of mathematicians, statisticians, operations researchers, ecologists, engineers, and physicists.

Although the topic of systemic risk may call to mind the possibility of deliberate attacks, both cyber and terrorist, on the financial system, after careful consideration the conference organizers decided against emphasizing this source of systemic risk. They reasoned that such a focus would downplay the many ways in which systemic risks can arise during the

financial system's normal operations. Analysis of the risks of deliberate attacks might build on the concepts explored in the conference, but it would require additional considerations and tools.

This report was prepared to share some of the insight and excitement generated by the conference and to encourage further cross-disciplinary conversations. It presents no National Research Council recommendations. The views expressed in this report are those of the authors and do not represent the views of the Federal Reserve Bank of New York, the Federal Reserve System, or the U.S. government. We hope that you find it useful and informative.

John Kambhu, Scott Weidman, and Neel Krishnan, *Rapporteurs*

Acknowledgments

The following people planned the conference and contributed in many ways to its success:

Franklin Allen, Wharton School, University of Pennsylvania,
Tanya Styblo Beder, SBCC,
Christine Cumming, Federal Reserve Bank of New York,
Darryll Hendricks, UBS,
Donald L. Kohn, Federal Reserve Board of Governors,
C. David Levermore,* University of Maryland,
Simon Levin,* Princeton University,
Robert H. Litzenberger, Azimuth Trust,
Charles M. Lucas,* AIG Corp. (retired),
Robert M. Oliver, University of California at Berkeley (retired),
M. Elisabeth Paté-Cornell,* Stanford University,
William Pulleyblank, IBM Corp.,
George Sugihara, University of California at San Diego, and
Charles Taylor, Risk Management Association.

Many thanks are given to the Federal Reserve Bank of New York for sponsoring the conference, underwriting the costs of this report, and supporting the involvement of the BMSA.

*These four people were separately appointed by the NRC as a small committee to oversee the conference from the NRC perspective.

This report has been reviewed in draft form by individuals chosen for their diverse perspectives and technical expertise, in accordance with procedures approved by the National Research Council's Report Review Committee. The purpose of this independent review is to provide candid and critical comments that will assist the institution in making its published report as sound as possible and to ensure that the report meets institutional standards for objectivity, evidence, and responsiveness to the study charge. The review comments and draft manuscript remain confidential to protect the integrity of the deliberative process. We wish to thank the following individuals for their review of this report:

Kenneth J. Arrow, Stanford University,
B. John Garrick, Independent Consultant,
Nigel Goldenfeld, University of Illinois at Urbana-Champaign,
Lord Robert May, University of Oxford,
Christopher A. Sims, Princeton University, and
Richard Zeckhauser, Harvard University.

Although the reviewers listed above have provided many constructive comments and suggestions, they were not asked to endorse the conclusions or recommendations nor did they see the final draft of the report before its release. The review of this report was overseen by Peter J. Bickel of the University of California at Berkeley. Appointed by the National Research Council, he was responsible for making certain that an independent examination of this report was carried out in accordance with institutional procedures and that all review comments were carefully considered. Responsibility for the final content of this report rests entirely with the rapporteurs and the institution.

Contents

1	INTRODUCTION	1
	Transitioning from a Bank-Based to a Market-Based Financial System, 2	
	Systemic Risk as a Generic Problem, 6	
	Systemic Risk in the Financial System, 8	
	Systemic Risk and Regulation, 9	
	The Roles of Policymakers, 11	
	Systemic Risk in Historical Perspective: The Events of 1998, 12	
	Analytical Issues, 16	
	References, 19	
2	CURRENT TRENDS IN ECONOMIC RESEARCH ON SYSTEMIC RISK	20
	Wealth Transfers and Portfolio Constraints, 21	
	Risk and Liquidity in a System Context, 22	
	Market Liquidity and Funding Liquidity, 23	
	Discussion, 25	
	References, 28	
3	SYSTEMIC RISK IN ECOLOGY AND ENGINEERING	29
	Useful Concepts from Ecology and Engineering, 30	
	Methodologies for Prediction and Management, 36	
	Risk Assessment of Extreme Events Involving National Security, 38	

Prediction and Management of Systemic Failure in the
Electric Grid, 43
Analogies in Economics and Finance, 47
Discussion, 49
References, 54

4 THE PAYMENTS SYSTEM AND THE MARKET FOR INTERBANK FUNDS 55

Systemic Dynamics in the Federal Funds Market, 56
Network-Based Modeling of Systemic Risk in the Interbank
Payments System, 59
Discussion, 62
References, 66

5 CONCLUDING OBSERVATIONS 68

Risk Assessment, 68
Modeling, Prediction, and Management, 70
Potential Applications to Policy, 73
Reference, 76

APPENDIXES

A	Conference Program	79
B	Background Paper	82
C	About the Rapporteurs	107

1

Introduction

The stability of the financial system and the potential for systemic events to alter its functioning have long been critical issues for central bankers and researchers. Developments such as securitization and greater tradability of financial instruments, the rise in industry consolidation, growing cross-border financial activity, terrorist threats, and a higher dependence on computer technologies underscore the importance of this research area. Recent events, however, such as the terrorist attacks of September 11, 2001, and the collapse of the hedge fund Long-Term Capital Management (LTCM), suggest that older models of systemic shocks in the financial system may no longer fully capture the possible channels of propagation and feedback arising from major disturbances. Nor can existing models account entirely for the increasing complexity of the financial system, the spectrum of financial and information flows, or the endogenous behavior of different agents in the system. Fresh thinking on systemic risk is therefore required.

Accordingly, in May 2006 a conference cosponsored by the National Academy of Sciences and the Federal Reserve Bank of New York was convened in New York to promote a better understanding of systemic risk. The sessions brought together a broad group of scientists, engineers, economists, and financial market practitioners to engage in a cross-disciplinary examination of systemic risk that could yield insights from the natural and physical sciences useful to researchers in economics and

finance.¹ Presenters from the natural and mathematical sciences and the engineering disciplines provided examples of tools and techniques used to study systemic collapse in interactive systems in nature and engineering. These presentations are summarized in Chapter 3. Similarly, research economists presented studies of systemic risk in cross-border investments, liquidity risk, and the payments system, and these presentations are summarized in Chapters 2 and 4. To provide a context for the discussions, risk managers at large finance institutions described how systemic risk and shocks in the financial system affect trading activities.

TRANSITIONING FROM A BANK-BASED TO A MARKET-BASED FINANCIAL SYSTEM

Financial market practitioners began the conference by highlighting various aspects of systemic risk and systemic events in the financial system. The topics of the presentations ranged from historical systemic episodes, such as the liquidity crisis of 1998 and the failure of LTCM, to risk assessment techniques, such as value-at-risk (VaR) analysis and scenario analysis. Charles Lucas of AIG (since retired), a member of the National Research Council's Board on Mathematical Sciences and Their Applications, introduced the first session by asking the fundamental question: What is systemic risk?

According to Lucas, economists' theoretical understanding of systemic risk stemmed from the experience of the Great Depression and specifically from John Maynard Keynes's interpretation of that experience in *General Theory of Employment, Interest, and Money*. Keynes aimed the formulation of his "general theory" at capturing the dynamics that allowed an economy to transition to an inferior but stable equilibrium, in the process overturning the normal full-employment equilibrium that defined classical models. During the Great Depression, the economy underwent a shock that was sustained by sympathetic movements throughout the financial system—a sequence of events that has come to be called "contagion." Because of policy missteps and a feedback loop with the financial system, the real economy settled into a persistent state of underutilized resources and unemployment. Despite structural changes since that time, the idea of a feedback loop between the financial and real sectors of the economy that leads to an inferior equilibrium with negative consequences for the real economy remains pertinent to current analysis of financial stability.

That system has changed dramatically since the Great Depression, as described in the conference background paper on the evolution of

¹The conference program can be found in Appendix A.

systemic risk.² Though banks still play a large role, many functions that defined their traditional domain are increasingly performed by securities markets and nonbank market participants. For example, hedge funds, private equity groups, and other fund managers now control larger shares of financial capital and take active roles in asset and credit markets. Crises in this more market-based financial system, such as the stock market crash of October 1987 and the market liquidity crisis of 1998, fit a general pattern of rapid decline in the price of some asset or class of assets, leading to a drop in liquidity. The result is contagion, in the form of further sympathetic price declines and a shift in market conditions marked by severely reduced financial market activity and potential negative effects on the real economy. Likewise, although the linkage mechanisms may have changed completely since the Great Depression, and despite the skepticism Keynes's theory received from the research community at the time, this simple model captures the mechanisms underlying the Depression. The endogenous shock in the United States that led to the inferior equilibrium then was a stock market crash followed by a wave of bank runs and loss of liquidity in a feedback mechanism of self-fulfilling prophecy. As Lucas suggested, this comparison offers a basic historical analogy illuminating some of the modern phenomena of systemic risks, such as sudden "regime shifts" in the financial system and the role of feedback mechanisms.

The conference background paper by Hendricks, Kambhu, and Mosser also describes the now well-studied phenomenon of the "bank run." In the classical model, a commercial bank makes illiquid loans on the asset side of its balance sheet, and takes demand deposits on the liabilities side that it is obligated to pay back at any time. In a bank run, even though each depositor would be willing to leave his or her funds on deposit, the belief that other depositors are likely to withdraw theirs causes all rational depositors to try to withdraw their funds as quickly as possible. A run on the bank results, because the bank's loans cannot be liquidated immediately at their full value, leaving the bank with no funds for the last depositors in line. In such a scenario, a run can be triggered by concerns about liquidity even if the bank is otherwise solvent.

Moreover, in this model, self-fulfilling prophecies can make bank runs contagious: If depositors witness a run on one bank, they may believe that runs are more likely to occur on others. This scenario can be attributed to several factors. For example, the issue that sparked a run on one bank, such as excessive loan exposure to real estate or the oil industry, may be

²"Systemic Risk and the Financial System," by Darryll Hendricks, John Kambhu, and Patricia Mosser; the paper can be found in Appendix B of this volume.

perceived to affect other banks, or one or more other banks may have significant interbank exposures to the affected institution.

As the Great Depression revealed, the withdrawal of funding liquidity resulting from bank runs can accentuate economic downturns and generally influence the real economy as lending is curtailed to creditworthy entities.³ Thus, the primary policy approaches to managing financial instability in a bank-oriented financial system—lender-of-last-resort facilities by the central bank, deposit insurance, and banking supervision to ensure credit quality in loan portfolios—were all aimed at preventing or mitigating the effects of these potentially catastrophic withdrawals of funding liquidity from the system. As the relative importance of banks as financial intermediaries has declined with the growth of market-based financial intermediation, market-based systemic events such as the stock market crash of 1987 and the failure of LTCM have shifted the emphasis from *funding liquidity* to *market liquidity*. Moreover, as Federal Reserve Board Governor Donald L. Kohn observed, the Federal Reserve is midway through a long process of adapting its policy tools to this new environment.⁴

As Hendricks, Kambhu, and Mosser describe in their conference background paper (Appendix B of this volume), the shift from a financial system dominated by banks to one dominated by markets has as its hallmark a broadening of the types of activities that banks and other financial intermediaries engage in and the assets that they invest in. The large financial institutions at the core of the system now intermediate the movement of capital in many ways: They assist businesses in the issuance of new stocks and bonds directly to the market (investment banking), they intermediate secondary-market trading of stocks and bonds after issuance on behalf of clients (market making), they lend directly to households and businesses (traditional commercial banking), and they manage asset portfolios on behalf of individuals and institutions (asset management). This latter example has led to the development of new market entities that act as vehicles for household savings, such as mutual funds and pension funds, as well as more leveraged entities, such as hedge funds.

The conference background paper also explains how a securities-market-based financial system works best when capital markets are liquid. In this context, liquidity refers to tradability: Markets are liquid when any individual trade is unlikely to have a major effect on the asset price because large numbers of willing traders are on the buy and sell sides of

³See the conference background paper described in footnote 2.

⁴Governor Kohn's observations, as reported in this summary, are based on his conference remarks, "The Evolving Nature of the Financial System: Financial Crises and the Role of the Central Bank."

the market. Liquidity normally rests on a number of foundations; foremost are market making, trading, and arbitrage. Market makers buy and sell securities out of inventory they maintain to meet customer demand, thereby providing intertemporal liquidity to smooth out short-run imbalances in supply and demand. Traders contribute to market liquidity by trading on bets that prices will converge to long-run fundamental levels. This activity speeds the convergence of prices to fundamental levels and provides stability to the market. Systemic shocks occur when one of these foundations is compromised.

Market-oriented crises tend to begin with a large asset price decline that becomes self-sustaining. Normally, when asset prices drop sharply, investors step up to buy assets that have declined sufficiently—an action that largely prevents market stress from worsening. This type of stabilizing correction is natural for a well-functioning, efficient asset market. In systemic crises, however, investors and traders are either unable or unwilling to step in, perhaps because their losses have limited their risk-taking and market-making capacity or because a structural failure in, say, the payments or settlement system has made such a step difficult. As prices decline, more market participants either sell from a change in their risk appetite or are forced to sell by a tightening of financing constraints, and prices are pushed down. Like the self-fulfilling-prophecy aspect of the bank-run model, this sequence of events can be self-reinforcing as market participants retire to the sidelines.

Market-based systemic crises are often characterized by a coordination failure: A wide cross section of market participants simultaneously decide to reduce risk taking and effectively refrain from financial activities, such as trading stocks, issuing debt and equity, and lending. While no one institution is necessarily insolvent or illiquid, each firm reduces its activity and risk to protect capital. In aggregate, the firms' actions combine to reduce financial market activity severely as asset prices fall, possibly harming the real economy in the process as the provision of financial services to otherwise creditworthy entities is curtailed and declines in asset prices impact firms' balance sheets. As Governor Kohn explained, the stock market crash of 1987 followed this pattern: simultaneous efforts to reduce equity market exposures were followed by a broad pullback in all risk taking.

In a market-based systemic crisis, as in the bank-run model, the actions and beliefs of individual participants across the financial system can combine to disrupt the entire system, even though the great majority of institutions are not at risk of collapse.⁵ When cast in these terms, the notion of systemic risk in the financial system bears a strong resemblance

⁵See the conference background paper described in footnote 2.

to the dynamics of many complex adaptive systems in the physical world. Many of the features of complex systems described by conference participants from the natural and mathematical sciences are clearly present in the financial system. For example, Simon Levin of Princeton University cited nonlinearities, multiple stable states, hysteresis, contagion, and synchrony as features common to all complex adaptive systems. These features are evident in models of financial crises: (1) contagion is seen in the self-reinforcing character of price declines and transmission of liquidity shocks across institutions; (2) multiple stable states and hysteresis can appear in the move to an inferior but stable equilibrium; and (3) nonlinearities in expectations and investment decisions can lead to sharp changes in the volatility and covariation of asset prices in an apparent regime shift, as discussed by risk managers below in this introduction. This commonality between financial and other complex adaptive systems points to the broad social importance of the study of systemic events.

SYSTEMIC RISK AS A GENERIC PROBLEM

In the world at large, complex systems abound. Accordingly, the instability of these systems and their potential for large and potentially catastrophic regime shifts are a dominant social concern—and one of high importance to many environmental and engineering sciences. For example, atmospheric scientists examine such questions in the context of climate change, as do fishery managers concerned with the sudden collapse of certain economically important fish stocks. As the presentations by Massoud Amin of the University of Minnesota and Yacov Haimes of the University of Virginia made clear, engineers grapple with similar issues to prevent disruptions to the North American power grid and to analyze for government entities the wider economic effects of terrorist attacks.

The ubiquity of such problems across so many fields suggests the possibility of finding common principles at work. As George Sugihara of the University of California at San Diego explained, engineers and public health professionals alike may be interested in how actions to address high-frequency but low-amplitude events, such as small floods or small outbreaks of disease, might predispose a system to low-frequency but very-high-amplitude disturbances. For instance, overuse of antibiotics to combat small-scale outbreaks of disease can lead to high-consequence outbreaks of antibiotic-resistant illnesses. Or, in an example from recent experience, the construction of levees in New Orleans to protect against intermediate-strength storm surges led to the higher-consequence damage from the lower-probability Hurricane Katrina.

Recent studies that have identified many common characteristics of nonlinear complex adaptive systems in the physical world may point to

a tentative vocabulary of systemic risk. A key concept that can be used to describe the process of adverse systemic change in both ecology and finance is the tendency toward a rapid and large transition from one stable state to another, possibly less favorable, state—what one might call a regime shift. Levin cited this phenomenon in the eutrophication of bodies of water, in which a shock, such as excessive heat, can lead to overenrichment of the water with nutrients, resulting in excessive growth of some organisms and a depletion of oxygen that is damaging to other populations. The new state of the body of water is a new stable equilibrium. Somewhat analogously, in financial markets, an exogenous change in the economic environment can lead to new profit opportunities in certain assets that attract capital and, if investment in the assets is excessive, to an asset price bubble vulnerable to a change in investor confidence. If a shock triggers a collapse of asset prices, there is a risk of a broader contraction if the normal self-correcting features of markets fail to work. Absent those self-corrections, the flight to quality by investors seeking safe assets could become a self-sustaining transition to a state with lower levels of credit and real economic activity.

In Levin's terminology, in both situations some shock leads to coordinated behavior within the system, a process known as "synchrony"—excessive growth of nutrients in the first example and excessive investment in an asset price bubble in the second. This synchronized behavior leads to reinforcing feedbacks, causing the initial shock to spread and cause contagion. Under the combined effect of the shock and contagion, a system makes a transition, or regime shift, from a stable state to an inferior stable state while shedding energy so that it cannot readily recover its original state, a process known as "hysteresis." Levin explained that much of the research on complex systems in the natural world has focused on the properties of robustness and resilience to shocks that either can prevent regime shifts and hysteresis from taking place, or can lead to recovery if they occur.

The commonality of stability and resilience to shocks in complex systems suggests that approaches to risk management in natural and physical systems could be pertinent to financial risk management. Amin and Haimes each 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 semblance to those used in financial risk analysis, and much of the subsequent conference discussion centered on the prospect of adapting methods from various engineering fields to financial risk management. Adaptation is clearly necessary because the range of behavior in financial markets is not mirrored in, say, the behaviors of

humans operating a complex engineered system. A risk analysis of an engineered system can assume that the people involved are attempting to fulfill their roles, which are relatively defined, and share a common objective. In contrast, in the financial system traders and investors operate in a competitive environment and might change their roles and behaviors opportunistically and creatively.

SYSTEMIC RISK IN THE FINANCIAL SYSTEM

Systemic risk in the financial system is difficult to define precisely. Although a literature exists on financial crises and systemic risk, a range of views can be found on what constitutes systemic risk.⁶

An adage among traders is that, in times of crisis, everything is correlated. Though conference participants did not share a consensus on the *definition* of systemic risk, the descriptions of systemic events by risk managers at the conference reflected this view. For example, Thomas Daula of Morgan Stanley described systemic events as regime shifts in which periods of extreme volatility combine with losses of liquidity to produce solvency risk. These crisis periods, according to Daula, “are characterized by very sharp increases in correlations and, therefore, they look and feel a lot like a regime shift—and a regime shift where you are moving from a normal regime, where there are relatively low correlations amongst financial markets, to a different regime, where you have extremely high volatility and a sharp spike in correlation.”

Under such regime shifts, the normal assumptions culled from historical experience that guide day-to-day trading break down. As D. Wilson Ervin of Credit Suisse observed in regard to the Russian default and the collapse of LTCM: “The most memorable part of this episode were the questions around fundamental issues that were normally unquestioned in day-to-day activities, about the reliability of your counterparties, about how markets would work under various circumstances, about whether liquidity would be there under certain circumstances.” In the presence of such uncertainty and market panic, traders can tend toward herd movement as they attempt to avoid losses—what the literature refers to as “phase locking”—and the normal mechanisms of price determination can break down. According to Robert Litzenberger of Azimuth Trust: “What happens is, in what we might refer to as crisis periods or liquidation periods . . . prices are generated internally by the market microstructure.

⁶For a central banking perspective, see Brimmer (1989); Bernanke and Gertler (1990, pp. 87-114) describe links between financial distress and the real economy. For a recent paper on systemic risk, see Chan et al. (2006).

Trades that were previously uncorrelated become correlated because they are being liquidated at the same time.”

In the tentative vocabulary of systemic risk suggested above, the self-reinforcing uncertainty and market panic that can characterize a systemic episode are a clear example of contagion. The jump in correlations appearing at the onset of a systemic event can in turn be seen as an example of self-reinforcing feedback and synchrony. Furthermore, the transition from a normally functioning market to one in which prices are generated by the internal market microstructure is accompanied by widespread and simultaneous liquidations. Financing constraints and the loss of liquidity make a return to the pre-crisis state very difficult—an asymmetrical transition and example of hysteresis. Thus, the notion of systemic risk, which financial market participants are at least viscerally acquainted with, can be worked into the framework of complex systems research from other fields.

While conference participants from the financial industry agreed on the “look” and “feel” of systemic episodes, there was some diversity of opinion on the more academic question of what actually constitutes systemic risk or a systemic event. Darryll Hendricks of UBS compared systemic risk with the Loch Ness monster: People claim that it exists or must exist, but nobody can point to a definitive episode. As Hendricks noted, most definitions of systemic events involve a transmission of shocks from the financial sector to the real economy—for instance, disruptions in credit provision as well as a propagation mechanism such as self-reinforcing feedback.

Therefore, is a systemic event simply one that creates externalities? Most would probably agree that this is too low a threshold for classifying an event as systemic. Lucas put forward the idea that a proper definition of systemic risk would involve transition from a stable equilibrium to some inferior but stable equilibrium, as explained above. This idea accords well with the regime-shift characterization used by financial industry participants. However, questions remain about how to characterize these equilibria. Was an event systemic because a disturbance propagated across diverse actors through self-reinforcing feedback, or did some policy mistake form the common cause of the disturbance to all actors, such as insufficient liquidity provision during the Great Depression?

SYSTEMIC RISK AND REGULATION

With the range of opinions on the proper identification of systemic risk, it is natural to wonder why the definition is so important. The 1987 stock market crash and the 1998 nexus of the Russian default, the failure of LTCM, and the resulting liquidity crisis were episodes of systemic

magnitude propagated and sustained by self-reinforcing mechanisms in the financial sector, and these episodes had potential consequences for the real economy. The definition is important, Daula explained, because regulation to ameliorate systemic risk constitutes a tax, and therefore a clear understanding of the risks is needed for the most protection at the lowest potential cost. Regulation is a tax in the sense that direct expenditures are required to comply with regulatory directives, and potential costs imposed in terms of efficiency losses in the allocation of capital. For an example of the latter, consider capital requirements that vary by the nature of one's business or the assets on one's balance sheet; they can create a wedge between market prices absent regulation and actual market prices. One could also wonder if a particular regulatory regime has a cost effect on the banking sector's industrial organization. For example, do certain forms of capital requirements—or, more generally, the costs of compliance with regulatory regimes—encourage consolidation?

Given these considerations, the importance of a sound method for identifying systemic risk becomes obvious. Without it, policymakers face a strong incentive to build expansive regulatory regimes capable of influencing practices that may or may not truly reduce systemic risk, because the potentially disastrous consequences of a real systemic event would justify the costs of such regulation. As Governor Kohn stated: "The natural inclination is to take more intrusive actions that minimize the risks of immediate disruption, and this inclination is probably exacerbated by ignorance and uncertainty." He explained that too much regulation could harm efficiency or generate moral hazard as market participants begin to take regulators' corrective measures for granted and increase risk taking. For example, they may fail to engage in adequate due diligence when extending credit or fail to maintain adequate capital for the risks they undertake. Further, to borrow from a theme raised by Levin, excessive regulation could introduce rigidities that may limit the natural flexibility of markets to respond to shocks.

A detailed understanding of what constitutes systemic risk is therefore important to forming a regulatory regime that balances costs and benefits. Indeed, in all the roles policymakers fill in preventing systemic events and mitigating systemic risk, a proper analytical framework is crucial for defining the correct scope and mode of action. For central bankers in particular, a clear method for identifying systemic risk and the onset of systemic events is critical for decision making on whether and how to intervene.

THE ROLES OF POLICYMAKERS

The Federal Reserve's role in setting monetary policy gives it the ability to mitigate the consequences of systemic events by easing access to liquidity. After the August 1998 financial market turmoil associated with the Russian loan default and the subsequent collapse of LTCM, for example, the Federal Open Market Committee lowered the target federal funds rate to soften the effects of "increasing weakness in foreign economies and of less accommodative financial conditions domestically."⁷ Other policymaking roles assumed by the Federal Reserve—services provider, bank supervisor and regulator, and crisis manager—also help to position it to mitigate systemic events. As a financial services provider, the Federal Reserve, through its Fedwire system, is the backbone of the interbank payments system. The conference presentation on Fedwire, discussed in Chapter 4 of this report, highlighted how important this role is in tempering the effects of a crisis. Referring to the hours after the attacks of September 11, 2001, the study highlighted how infrastructure disruptions and the resulting payments miscoordination threatened to seriously disrupt the payments system. In response, the Federal Reserve extended the operating hours of Fedwire and increased liquidity provision by using the discount window and open market operations, actions that significantly reduced the impact of the disruption. This episode exemplifies the Federal Reserve's role as crisis manager.

Governor Kohn remarked that the Federal Reserve and other regulatory agencies, as banking supervisors, can do much to reduce systemic risk by maintaining a healthy banking system. Collective efforts of regulators and the private sector to enhance market discipline, improve risk management practices, and strengthen the clearing and settlement systems reduce the likelihood that a sharp change in asset prices or questions about a major market participant will lead to a systemic financial crisis. In today's market-dominated financial system, banks still have a large role to play in financing traders' securities positions and in clearing and settling trades in their brokerage activities. In their role as providers and conduits of liquidity, healthy banks can be bulwarks against the propagation of financial turmoil.⁸

As a crisis manager, the Federal Reserve can avert many problems by monitoring conditions and identifying risks as they arise. Indeed, as Governor Kohn explained, a common element in the Federal Reserve's response to both the 1987 stock market crisis and the 1998 liquidity crisis

⁷Federal Open Market Committee Statement, September 29, 1998 (<<http://www.federalreserve.gov/boarddocs/press/general/1998/19980929/>>).

⁸For analysis of banks as liquidity providers in a crisis, see Saldenber and Strahan (1999) and Gatev, Schuermann, and Strahan (2006).

was its public acknowledgment that a crisis was under way. In announcing a crisis and articulating its response, the Federal Reserve reassured market participants that it was working to mitigate the systemic effects of the crisis; such reassurance can go a long way toward encouraging a return to risk taking. In both episodes, the Federal Reserve also used open market operations to ease reserve market conditions and the stance of monetary policy, monitored the flow of credit through the financial system, and worked with lenders to emphasize their collective interest in avoiding a credit gridlock.

The Federal Reserve's actions relied on an early determination of the potential systemic effects of the two events. Largely as a result of these actions, neither the 1987 event nor the 1998 episode led to a disruption in real economic activity.

SYSTEMIC RISK IN HISTORICAL PERSPECTIVE: THE EVENTS OF 1998

Governor Kohn observed that the Federal Reserve has been involved in a long process of adapting its tools to the market-dominated financial system that is still emerging today. Accordingly, the 1987 stock market crash and the 1998 liquidity crisis are natural case studies to examine when defining systemic risk, as neither was triggered by the bank-run phenomenon that characterized many systemic problems in the nineteenth century.

The events of 1987 and 1998 had many common elements. First, both began with sharp movements in asset prices that were exacerbated by market conditions—portfolio insurance in 1987 and the closing out of positions in 1998. Second, market participants became highly uncertain about the dynamics of the market, the true value of assets, and the future movement of asset prices. In terms of the regime-shift scenario described earlier, events outpaced the standard risk management systems, which had been based on historical data and experience. Third, large and rapid price movements called into question the creditworthiness of counterparties, which could no longer be judged by now-obsolete financial statements. Fourth, the decline in asset prices decreased wealth and raised the cost of capital, developments that threatened to reduce both consumption and investment in the real economy.

Although the 1987 and 1998 events shared many features, conference participants tended to focus on the more recent 1998 episode because financial institutions, instruments, and practices then were more similar to the way they are today. The potential negative effects on the real economy and the systemic character of this episode were highlighted by the withdrawal of investors from the commercial paper market.

As noted by Ervin, the events of 1998 were catalyzed by the Russian default. In 1998, Russia was in a precarious position as a fledgling democracy attempting to transition to a market-based economy. It had a high dependence on energy exports at a time when the price of oil was dropping, a massive trade deficit, an unsustainable pegged exchange rate, and a large government budget deficit. It was also financed mainly by short-term debt. Despite a large loan package in July 1998 from the International Monetary Fund, a sustained reversal in market sentiment led the Russian government to announce in August of that year that it would default on short-term local-currency debt. The result was disastrous: Many Russian counterparties failed, and liquidity in Russian instruments dried up.⁹

Investor losses were estimated to be on the order of \$100 billion. As Ervin explained, every working assumption about the Russian market came into doubt: the rules, the participants, the prices, the functioning of markets, even the legal system. This was surely a systemic crisis for Russia. Moreover, it threatened to become a systemic crisis for the international financial system when the market turmoil affected a particular hedge fund, LTCM, and the liquidity of core markets in the financial system.

In the mid-to-late 1990s, LTCM was a very large and well-known hedge fund, both highly leveraged and highly successful. Its primary investment strategy centered on finding arbitrage opportunities or near-arbitrage opportunities in which the market seemed to be out of line with long-term economic fundamentals—a trading strategy based on the idea that certain pricing gaps will close over some (potentially long) period of time. As part of its trading strategy, Ervin explained, LTCM would tend to buy older, illiquid Treasury bonds, then short-sell current, on-the-run Treasury bonds and eke out a small yield differential between the two. The goal was to capture a small yield differential in relative asset prices, allowing LTCM to earn steady returns as relative prices converged to fair values while the fund avoided directional risk. At the time at least, this strategy was considered somewhat state-of-the-art, and many financial entities attempted to emulate it, if not to mirror LTCM's positions outright.

The primary problem with LTCM's strategy was that, as a relative-value trader, it was very exposed to liquidity shocks and correlation assumptions, even if the fundamentals underlying its positions were correct. In the days following the Russian default, there was a large flight to quality in developed markets that caused credit spreads to widen sharply. Interestingly, this trend was not limited to U.S. corporate debt; interest rate swap spreads—an indication of the credit conditions of international

⁹This discussion draws heavily on Ervin's conference presentation.

banks—also widened sharply. As this shock rippled through the financial system during August, it also began to affect equity markets, and the Dow dropped 357 points on August 27 and a further 512 points on August 31. Implied volatility in prices of equity options also increased substantially, more than doubling its pre-crisis levels.

These events were preceded by the decision by Salomon Brothers to close its bond arbitrage group in the spring and summer of 1998. The departure of such a large bond trader potentially left the market with less liquidity than it would have had, because of both the liquidation of Salomon's very large positions in the months prior and the absence of a large player whose trading otherwise would have contributed to market liquidity. Vincent Reinhart of the Federal Reserve's Board of Governors and Litzenberger pointed to the Salomon decision as creating the initial market stresses that escalated with the Russian default and the emergence of LTCM's financial problems.

The mechanisms that led to the subsequent fall of LTCM highlight aspects of the nonlinearities and reinforcing feedbacks cited in the conference's discussion of a securities-market-based financial system. LTCM had a strategy of targeting the volatility of the Standard and Poor's index as a type of risk control mechanism, using it as a benchmark against which to assess the value-at-risk of its own positions. VaR analysis is a widespread portfolio-management strategy that calculates the maximum potential loss over a certain time period, given a specified level of confidence. Historical volatilities are used to form a VaR estimate, Litzenberger explained. ARCH and GARCH methods (autoregressive conditional heteroskedasticity and generalized autoregressive conditional heteroskedasticity, respectively) are common tools for obtaining a volatility estimate based on historical volatility and covariance. The usual way of employing these estimation techniques at the time was to consider historical volatilities, not just volatilities during hypothetical crisis states.

The pressure on LTCM's position caused by the liquidation of Salomon's bond arbitrage group was combined with pressure from the widening of credit spreads following the Russian default. As VaR models reflecting the spike in price volatility indicated higher risk, market participants began to liquidate their positions defensively. This reaction illustrates the concept of reinforcing feedbacks: As volatility increased, market participants reasoned that risk had also increased, so they began to liquidate those positions, a step that in turn led to further elevations in volatility and more decisions to liquidate. The process also illustrates the importance of linkages and nonlinearities in systemic events: Even though Russian instruments were a small proportion of LTCM's overall portfolio, market participants began to question their own rationale for holding other, non-Russian positions that LTCM also held. Thus, they

began liquidating those positions in anticipation of liquidation spilling over into other markets, and in this way a seemingly small disturbance propagated quickly.

As Litzenberger explained, in the 1997 run-up to LTCM's failure, the arbitrage market was marked by high liquidity and low volatility. Under these conditions, to maintain a target risk profile (for example, VaR) when volatility was low, traders such as LTCM would leverage their positions. Recall that LTCM maintained a strategy of targeting its risk taking on the volatility of the Standard and Poor's index; the fund's response to the situation in 1997 was essentially to add leverage by returning a substantial portion of capital to its investors. This strategy was consistent with attempting to maintain profitability when trading opportunities were harder to find as trades were mean-reverting faster. According to Litzenberger, this would have been an entirely reasonable strategy if conditions in 1997 had constituted a steady state. However, the liquidation of Salomon's bond arbitrage group and the Russian default combined to disrupt this steady state and cause a considerable rise in volatility. The subsequent apparent increase in risk triggered widescale liquidations, as the assumptions underlying these positions came into serious doubt. The resulting pressure on LTCM led it to send investors a letter on September 2 asking for more capital. Just three weeks later, the firm was taken over by its creditor banks to enable orderly liquidation of the fund's positions.

Ervin explained that, from the viewpoint of derivatives trading desks, the events leading up to the collapse of LTCM resembled a regime shift: At times, trading in U.S. dollar interest rate swaps dried up completely; pricing for typical bonds, such as investment-grade mortgages, widened to the point that one could not get a price, or at least a real price. "During this period, people simply didn't have confidence they understood what was going on. They weren't sure they understood the new rules of the game, who would survive, and how they should play," he said.

Beyond the breakdown of trading models based on historical correlations, systemic events have a psychological character, as all the rules seem to collapse and participants enter into a state of high uncertainty about their counterparties. According to Governor Kohn, this points to a crucial role for policymakers:

Heightened uncertainty is the key characteristic of episodes of financial instability. The central bank may not have any more information than market participants do. In economic models, based on such uncertainties, it is the central bank's willingness to act in the face of uncertainty that differentiates it from other market participants and gives it a positive role to play during financial crises.

This role must be buttressed by a clearer understanding of the fundamental dynamics underlying the securities-market-based financial system; yet many obstacles to this ideal remain, both for market participants seeking to insulate themselves from the effects of crises and especially for regulators seeking to prevent them. Among these obstacles are the difficulty in simulating financial crises, the lack of historical episodes to study, and—crucially for entities such as the Federal Reserve—hindrances to the types of data sharing among market participants and regulators that would allow central banks to act with certainty during systemic crises.

ANALYTICAL ISSUES

In a bank-dominated financial system, Governor Kohn observed, it is much easier to gather the information necessary to regulate effectively against the possibility of systemic disruptions. In such a context, the critical information would come from fellow bank regulators with which the Federal Reserve had been working and from banks the regulators had been examining. However, in a more market-dominated context, in which many financial institutions have a presence in many cross-border business lines, obtaining the information on counterparty exposure and risks necessary to develop cogent analyses and to inform decision making in a possible crisis requires widespread cooperation across disparate entities. Moreover, in many instances, market participants may regard this information as proprietary. Scant availability of data and inadequate data sharing present challenges for regulators and market participants alike.

Governor Kohn remarked that, as the prime source of constraint on potential crisis-causing behavior, market discipline through vigilance among private parties is always preferable to regulatory dictates. For market discipline to be effective, however, counterparties must have a clear understanding of each other's risk profile. This often requires them to share proprietary information, and confidentiality agreements between counterparties may be necessary to ensure comfort.

He acknowledged that market participants may be wary of sharing proprietary information. However, information sharing can greatly increase the probability that credit will continue to flow during systemic disruptions, resulting in a lower probability of a sustained systemic disruption, a reduced need for government intervention, and enhanced financial stability without moral hazard.

Governor Kohn added that, in a market-based system, sound risk management by *all* market participants is essential to protect against the risk of a low-probability—or “tail”—event causing a financial crisis. For example, the bringing together of practitioners in risk management policy groups can potentially lead to improved reporting of risk information to

counterparties and allow best practices to be transferred across market participants with respect to valuation, exposure measurement, limit setting, and internal checks and balances. Indeed, a lesson drawn from the 1998 crisis by the President's Working Group on Financial Markets (1999) was how weakness in risk management and counterparty credit discipline enabled a firm to acquire large leveraged positions of a size that could magnify the effects of negative events.

Governor Kohn described how financial regulators, through supervision, can promote market discipline and sound risk management. The regulatory capital framework proposed in Basel II would (1) require the largest internationally active banking organizations to enhance measurement and management of their credit and operational risks, (2) prescribe a rigorous methodology for entities to assess overall capital adequacy in relation to their risk profile, and (3) require entities to disclose publicly information about their risk profile.

Sound risk management practices among market participants rely heavily on sophisticated analytical methods that present challenges beyond limited data availability and information sharing. Discussions among economic researchers, financial market practitioners, and members of the engineering and natural sciences fields pointed to the considerable differences between the financial system and other complex systems. Among these differences is the inability to conduct or observe natural experiments on systemic crises in the financial system because crisis occurrences are too infrequent. Another difference is the role of human behavior in the financial system and the nonlinearities and anticipatory behavior it can introduce, a factor largely missing in studies of complex systems in engineering or the physical sciences. The presentations by financial market participants addressed these issues in discussions of scenario analysis.

Scenario analysis, as Daula explained, is the primary tool that market participants use to examine the risks posed by systemic events. Aside from being what it implies, scenario analysis was defined by Daula in economic terms: It starts with a particular scenario about the economy and then defines a general equilibrium, inferring the conditional expectations of all the consequences of that scenario for markets around the globe and their relative prices. He identified three ways of specifying the scenario. The first is to look at historical episodes. This approach has the advantage of being grounded in an actual event; the drawback is that changes in market structure since the chosen episode can lessen the predictive power of the analysis. A second approach is to fashion a purely hypothetical event. This has the advantage of allowing one to match the scenario to the particular market structure at the time; the obvious drawback is the difficulty knowing with certainty whether the hypothetical event is at all

likely or whether the analysis performed accurately reflects how the event would actually unfold. The third approach, which addresses some of the pitfalls of the previous two, is to use a hybrid, mixing in something that may have occurred in the past in a slightly different context and analyzing how it may play out in today's context; conditional expectations for changes in market structure are adjusted along the way.

It is difficult to choose the optimal scenario to analyze. Daula suggested a method that pointed to some possible interdisciplinary linkages between this type of financial research and engineering approaches to systemic risk: Choose a set of scenarios broad enough to span collectively the types of market fluctuation likely to be encountered. If the scenarios selected are sufficiently broad, common elements may emerge. Daula emphasized, though, that this type of exercise may result in unlikely scenarios. Providing an example, he noted that one often-considered scenario is a monetary crisis in a reserve currency such as the U.S. dollar, an event that arguably has not occurred in thirty years. Incorporating extreme tail events such as these would address Litzenberger's concern that many quantitative risk management approaches rely too heavily on data from relatively benign periods and thus allow history to grant a false sense of security.

The approach of collectively analyzing a broad range of scenarios may allow for linkages with optimization methodologies from engineering fields such as operations research. Presenting one possible inroad, Haimes offered an overview of the "partitioned multi-objective risk method," in which systems are analyzed both for the low-frequency but high-cost events and for the high-frequency but low-cost events. Drawing on such mathematical work from engineering fields may also enable one to analyze how certain attempts to increase a system's resilience and robustness may actually predispose the system to low-frequency but high-damage events. Needless to say, any such analysis must be very careful in its assumptions of probability distributions.¹⁰

Researchers and policymakers face many challenges in arriving at a better understanding of systemic risk in our evolving securities-market-dominated financial system. Market participants and regulators face a dual problem: They must determine the factors that can trigger contagion, the prospect for sudden regime shifts, and the potential for hysteresis; they must also craft policies that strengthen resilience to the threat of systemic events in a way that neither predisposes the system to even larger disruptions nor imposes unjustifiable costs on market participants.

¹⁰This topic is discussed in Goldenfeld and Kadanoff (1999).

REFERENCES

- Bernanke, B., and M. Gertler. 1990. "Financial Fragility and Economic Performance." *Quarterly Journal of Economics* 105, no. 1 (February): 87-114.
- Brimmer, A. F. 1989. "Distinguished Lecture on Economics in Government: Central Banking and Systemic Risks in Capital Markets." *Journal of Economic Perspectives* 3, no. 2 (spring): 3-16.
- Chan, N., M. Getmansky, S. M. Haas, and A. W. Lo. 2006. "Systemic Risk and Hedge Funds." In M. Carey and R. Stulz, eds., *The Risks of Financial Institutions*. Chicago: University of Chicago Press.
- Gatev, E., T. Schuermann, and P. E. Strahan. 2006. "How Do Banks Manage Liquidity Risk? Evidence from the Equity and Deposit Markets in the Fall of 1998." In M. Carey and R. Stulz, eds., *The Risks of Financial Institutions*, 105-30. Chicago: University of Chicago Press.
- Goldenfeld, N., and L. Kadanoff. 1999. "Simple Lessons from Complexity." *Science* 284, no. 5411 (April): 87-9.
- Keynes, J. M. 1936. *General Theory of Employment, Interest, and Money*. Harcourt, Brace and World/Harbinger (1964).
- President's Working Group on Financial Markets. 1999. "Hedge Funds, Leverage, and the Lessons of Long-Term Capital Management." Washington, D.C.: U.S. Department of the Treasury, April.
- Saidenberg, M. R., and P. E. Strahan. 1999. "Are Banks Still Important for Financing Large Businesses?" Federal Reserve Bank of New York *Current Issues in Economics and Finance* 5, no. 12 (August).

2

Current Trends in Economic Research on Systemic Risk

A conference session on current research directions featured three papers examining market-based crises—crises in which financial institutions are affected by shocks that propagate through asset prices and market liquidity.¹ In these crisis models, shocks affect financial institutions through the prices of securities that the institutions hold in common—not through chains of connections between institutions, as in a payments network.

While market-oriented models of financial crises differ from the traditional bank-oriented models in the way shocks are propagated, they share with bank models the possibility of multiple equilibria and transitions driven by positive feedback. Thus, a shock can cause a transition from a normal state to a crisis state from which the system need not recover endogenously. Indeed, the models often feature path-dependent behavior in which the transition out of a crisis state entails a path different from the one leading to the crisis and may require some form of external intervention. These characteristics of market-based models—and the dynamics of the models more generally—are the subject of the three papers presented.

¹The large literature on systemic risk and financial crisis cannot be represented in any set of three papers. The papers in this session of the conference were selected to illustrate current thinking about financial crises that propagate through securities markets (for example, the bond and stock markets). Further, the conference organizers sought out analytical or theoretical papers that would show the conceptual underpinning of the literature on financial crises; empirical analyses of financial crises were not included.

As the discussion that followed the presentations made clear, the papers open some potentially productive new avenues for research. More insight is needed into how financial markets recover from crisis states and what policies or regulatory regimes would speed that recovery and contribute to a more robust financial system.² A related issue that merits further research is the trade-off in risk management practices between the objective of limiting risk *ex ante* and the effects of risk management constraints in the midst of a crisis. For instance, mark-to-market accounting is a risk management practice that makes trading performance transparent and prevents managers and traders from concealing losses while trying to gamble their way out of losing positions.³ Further, marking to market the value of trading positions combined with risk management loss limits that force a closeout of a losing position can prevent a loss from becoming large enough to bring down a firm. (Some bank failures and catastrophic investment fund losses are attributable to the failure to adhere to this basic risk management discipline.) However, as the papers presented suggest, the collective and mechanical exercise of such discipline on a widespread scale after a large market shock can create the type of liquidity spiral that leads to a market crisis.

WEALTH TRANSFERS AND PORTFOLIO CONSTRAINTS

The first paper, by Anna Pavlova of the London Business School and Roberto Rigobon of MIT (presented by Rigobon), examined the transmission of shocks between countries with cross-border trade and investment. Pavlova and Rigobon (2006) began studying this issue after they uncovered a divergence of views on a simple question: Would it be good for the stock market in the United States if the dollar depreciated? They found that the answer depended on whether the initial shock was a supply or a demand shock and also on the effects of wealth redistribution arising from the changes in the relative prices of goods and financial assets. The presentation focused on how a shock plays out in the real side of the economy and in the financial system and how the two sectors interact through the effects of wealth redistribution.

The paper highlights the ways in which financial market imperfections and institutional features of the financial system affect the transmission of shocks across countries. The model presented has a center country

²For examples of research on these issues, see Allen and Gale (1994, 2005) and Holmstrom and Tirole (1998).

³Mark-to-market accounting requires that the value of an investment, which might vary over the period for which it is held, be assigned the current market price of such an investment.

and two peripheral countries; significantly, it also includes a constraint on the center country's financial sector that can be interpreted as a risk management constraint on that country's investors—for instance, a constraint against concentration risk. With this model, Pavlova and Rigobon seek to understand how the exchange rates, interest rates, and stock markets in the three countries evolve in response to shocks. Is there comovement in asset prices of the peripheral countries and, if so, does it depend on the tightness of the constraint? The analysis uses a general equilibrium framework that illuminates the role of wealth redistribution in the transmission of shocks.⁴

In the model, the constraint creates a common risk factor or covariation in stock prices and terms of trade (the exchange rate). In the presence of shocks, the portfolio constraint leads to wealth transfers that create comovement among the terms of trade and stock prices in the peripheral countries, while reducing the comovement between the stock markets of the center country and the peripheral countries. These results are consistent with empirical findings documenting contagion among the stock prices and exchange rates of countries belonging to the same asset class (for example, emerging markets). One of the model's implications for policy is that during a crisis, interventions that relax the portfolio constraint in the center country's financial system could be a more effective response to a systemic crisis than providing assistance to the country suffering the initial shock. The alleviation of the constraint short-circuits the wealth transfers that transmit the shock to others, reducing the likelihood of contagion.

RISK AND LIQUIDITY IN A SYSTEM CONTEXT

Hyun Song Shin of Princeton University examined how liquidity shocks can propagate through the linkages between balance sheets of financial institutions and securities prices. The starting point of Shin's (2006) analysis is the fact that most of the assets on the balance sheets of financial institutions are claims against other parties. This fact leads to interesting and possibly complex interrelationships in which asset prices can fluctuate together. How creditworthy one party's liabilities are depends on the strength of the assets on its balance sheet, which in turn depends on the creditworthiness of other parties' liabilities, and so on.

In Shin's analysis, the financial system is a system of interlinked balance sheets. An objective of the study is to analyze fluctuations in appar-

⁴In a general equilibrium analysis, all decision makers behave optimally relative to others (subject to constraints such as budget limitations), and supply and demand in all markets are in balance at the equilibrium prices.

ent risk appetites that arise endogenously from solvency constraints and financial institutions' interlinked balance sheets. In the model, all assets are marked to market, and economic agents are assumed to be risk neutral so that the analyst can observe how asset prices respond to the liquidity effects arising from market participants' interlinked balance sheets, rather than to changes in risk preferences or risk aversion.

In the model, the market value of each firm's debt depends on the value of the firm's assets. Since some of these assets are the debt of other firms, linkages arise in the value of the debt of all the firms. An equilibrium is a fixed point of these asset value equations. With the addition to the model of a target leverage ratio determined by, for instance, a risk management constraint, financial institutions will shrink or expand their balance sheets in response to shocks to their capital—actions that will set off liquidity drains and lending booms. In this model, supply and demand curves have counterintuitive shapes, and a fall in prices can actually increase the supply of assets. In such a case, a negative shock to bank capital raises a bank's leverage ratio above its target; to reduce leverage, the bank must sell assets. These sales depress prices even more, causing a further negative shock to all banks' capital and setting in motion additional asset sales and a downward spiral in asset prices.

A policy-related implication of this analysis is the potential for feedback effects to arise from mark-to-market accounting. Now that a much wider range of assets can be marked to market, will such an accounting convention enhance stability or undermine it? Accounting is absolutely crucial for thinking about incentive problems because gains and losses are recognized on the balance sheet, and it is the unit of account that drives decisions.

In thinking about systemic risk, Shin considers the difference between domino effects and price effects. In domino scenarios, shocks propagate between banks through the payments system or through cascading defaults. Price effects, however, can propagate shocks even when no balance sheet or payment linkages exist. Further, price effects operate even in the absence of large players. Price changes are a lightning rod that coordinates expectations and actions and that affects the system through the similarity of positions across firms regardless of firm size or the lack of direct linkages between the firms.

MARKET LIQUIDITY AND FUNDING LIQUIDITY

In the session's last paper, Markus Brunnermeier of Princeton University (presenter) and Lasse Pedersen of New York University explored the relationship between market liquidity and funding liquidity, giving particular attention to how they interact through risk management practices at

financial institutions. Market liquidity is the ease of trading an asset and is asset-specific, while funding liquidity is the availability of funds and is agent- or borrower-specific. Brunnermeier and Pedersen's (2006) paper links the two liquidity concepts by arguing that they are mutually reinforcing: when funding liquidity is abundant, traders have the resources to finance trading positions that smooth out price shocks, and markets will be liquid. This process is self-reinforcing because liquid markets are less volatile and assets become better collateral—conditions that lead to a relaxation of funding constraints on trading activity. This feedback loop is what Brunnermeier and Pedersen set out to study.

They construct a model that would explain four stylized facts about market liquidity. The first fact is the most important one for the systemic risk question—the sudden loss or fragility of liquidity. Second is the commonality of liquidity and the way market liquidity comoves across different assets. Third is the apparent correlation between liquidity and volatility: whenever volatility is high, liquidity is low. The last is the flight-to-quality phenomenon, whereby traders flock to low-volatility securities when their capital is eroded, causing the liquidity of riskier assets to deteriorate.

In the model, a market liquidity shock is defined as the price deviation from the fair value of an asset. To examine endogenous illiquidity effects, the researchers assume that offsetting liquidity shocks exist: thus, in the initial period, a liquidity shock causes the price to deviate from fair value and, in the subsequent period, an offsetting shock occurs that restores the price to its initial fair value.⁵ In addition to liquidity shocks, a source of risk in the model is a fundamental shock that changes the fair value of the asset. Traders in the model buy and sell securities in an attempt to profit from the liquidity shocks and, in so doing, provide liquidity to the market. This liquidity provision is risky, however, because of the fundamental shocks that change the fair value of the asset. Traders are constrained by their net worth and need to finance their trading positions subject to a margin or “haircut” on the amount they can borrow, where the margin is a credit risk mitigation device imposed by the lender and is determined by the volatility of the fundamental value of the asset. The traders face funding liquidity risk because a fall in their net worth or a rise in the margin required for trading positions may deprive them of funds needed for trading.

In this model, the relationship between the margin requirement and

⁵Liquidity shocks are price shocks that are unrelated to fundamental value. For example, an investor may sell bonds to meet a need for cash, placing downward pressure on the bond price; at a different moment, an investor who has experienced a cash windfall may buy bonds, producing an opposite effect on the price.

the asset's price and volatility will influence whether equilibrium outcomes with fragile market liquidity and illiquidity spirals occur. Trader losses from price shocks can lead to self-perpetuating falls in market liquidity as trading is endogenously curtailed because of the difficulty of funding the margin required for trading positions.

DISCUSSION

Herdlike Behavior and Incentives for Contrarian Trading Strategies

The three presentations summarized in this chapter highlighted the positive feedback effects that produce herdlike behavior in markets, and the subsequent discussion focused in part on means of encouraging heterogeneous investment strategies to counter such behavior. Investors who sit on the sidelines during boom times will not be weakened by the inevitable downturn and will be well positioned to profit by entering the market to buy assets at distressed prices. Such contrarian investment behavior would mitigate the sort of systemic collapse that was analyzed in the papers presented. A number of conference participants asked what incentives for this type of stabilizing behavior fund managers have. Would fund managers who were content to hold cash and low-yielding liquid assets when the markets were flourishing be able to convince their investors to stay with them when everyone else was earning tremendous profits riding the upside of a bubble? Which investors are willing to earn very little in anticipation of realizing high returns by purchasing undervalued assets after a market crash?

If it is costly to hold liquid assets in order to be a buyer and to provide liquidity in a market crash, why *would* anyone choose to do it? In an equilibrium analysis that accounts for the incentives to sit on the sidelines in a boom, the market crash must be big enough to assure liquidity providers that they will earn sufficient profits buying at distressed prices to compensate them for forgone profits. So, in the absence of government or central bank intervention, the paradox is that the inducement to adopt contrarian investment strategies is greater when the severity of the crash is greater.⁶

The conference participants discussed the role the central bank or government might play in encouraging the sort of contrarian behavior that would stabilize failing markets. Collateralized lending by the central bank could be one way to short-circuit the feedback in asset prices and distress-driven selling of those assets; investors could acquire liquidity

⁶Allen and Gale (1994, 2005) study these issues.

by borrowing against assets instead of selling them.⁷ However, the type of assets that investors might want to offer as collateral could be different from the asset types normally used as collateral when borrowing from the central bank—especially in a situation in which investors' best assets have already been used in collateralized borrowing from the markets. Further, there could also be incentive effects—such as moral hazard—that change behavior in boom times in undesirable ways. If investors anticipate that illiquidity would be mitigated in a crash, they may have even more reason to ignore the risks in an emerging price bubble.

Another policy option mentioned in the discussion in this session would be to change reserve requirements and capital requirements to counteract the positive feedback effects—that is, to raise requirements in boom times and lower them in bad times. Alternatively, when markets are prospering, banks could be required to increase their liquid asset holdings so that they can provide liquidity more effectively when markets fail. The problem here, of course, is that these requirements act like a tax on these institutions, and taxes are always unpopular and would place the institutions at a disadvantage relative to other market participants—at least in the good times.

The Range of Economic Models in the Study of Systemic Risk

Participants in the session also discussed the types of models used to study systemic risk and commented on the challenges and trade-offs researchers face in developing their models. One type of model is the falling domino model. When applied to data on the linkages among banks through interbank loans and exposures in the payments system, for example, the model is used to study how cascading losses following the collapse of a bank propagate through the banking or payments system. In such an event, what would happen to other banks and how would liquidity in the payments system be affected? Another type of model takes into account the optimal behavior of market participants in analyses of their response to shocks. These models can be general equilibrium or game-theoretic models: The former look at the interaction between financial asset markets through, say, investors' portfolio choices; The latter examine strategic interaction between economic agents in which agents act in anticipation of how others will behave. In addition, the models can be either comparative static models or dynamic models: The former analyze

⁷Examples of such liquidity provision are the discount window lending facilities at central banks that provide emergency liquidity to banks, and the repo options that the Federal Reserve made available to nonbanks to address concerns about liquidity shocks associated with the Y2K vulnerability in computing systems.

differences between the pre-shock and post-shock equilibrium states of the financial system, while the latter examine what occurs in the transition from one equilibrium state to the other.

The work by Pavlova and Rigobon is representative of the current literature on international crises involving exchange rates and cross-border shocks to financial systems and economic activity. The studies by Shin and by Brunnermeier and Pedersen are illustrative of the models that look at feedback effects to clarify the interactions between market prices and the behavior of financial institutions. These papers highlight the importance of the financial system's institutional features—mark-to-market accounting, margin requirements in trading, and risk management constraints more generally—to an understanding of systemic risk. The papers are stripped-down approaches examining the equilibrium of a system of price determination equations to simplify the analysis of feedback effects. Adding to the analysis a consideration of heterogeneity among investment strategies, as in the discussion above, increases the complexity of the effort considerably. For instance, one could step back and ask how investors would choose their initial portfolios if they anticipated the feedback effects and linked sequences of events in possible future scenarios. Or one could ask what incentives or compensation arrangements would motivate an investor or fund manager to act on that anticipation.

The challenge in these and other models is the trade-off between analytical tractability and realism. Given the current state of the art, significant simplification and abstraction are required to build models that can be used to answer practical questions. Yet the simplicity of a model by its nature means that potentially important factors can be missed. Indeed, a key goal of the conference was to determine whether there are modeling techniques in other disciplines that can deal with complexity yet still keep sight of the important features of the system under study.

Adequacy of Buffers Against Systemic Shocks in the Financial System

A third discussion topic that drew considerable interest was whether competitive pressures and risk management practices are undermining the robustness of the financial system. More sophisticated methods of assessing collateral and margin requirements in the financing of trading positions may be lowering the overall margin and collateral amounts held against these exposures. For instance, the use of portfolio margining allows the netting and offsetting of positions and results in a lower margin on posted collateral. Certainly, the technique has advantages: Netting of margin across gaining and losing positions in a portfolio can alleviate the liquidity shocks from margin-driven selling of the losing

position, reducing the positive feedback effects analyzed above. At the same time, however, portfolio margining reduces the amount of overall margin, resulting in a smaller cushion if correlated shocks occur simultaneously across the whole range of margined investments.

A critical risk management issue here is the treatment of correlation assumptions in determining margin amounts for a portfolio of diverse assets. Correlations among asset prices can change radically in a crisis. A conference participant observed that a truly sophisticated risk manager would set portfolio margin requirements that take into account how those correlations can change in a crisis, and not look myopically at the average correlations of the last three years. Whether such an approach would be rewarded, however, brings us back to the earlier discussion of incentives and contrarian behavior: Do risk managers have meaningful incentives to use conservative portfolio margin requirements when their competitors are basing their margins on optimistic assumptions about correlations of margined positions?

REFERENCES

- Allen, F., and D. Gale. 1994. "Limited Market Participation and Volatility of Asset Prices." *American Economic Review* 84, no. 4 (September): 933-55.
- _____. 2005. "From Cash-in-the-Market Pricing to Financial Fragility." *Journal of the European Economic Association* 3, no. 2-3 (April-May): 535-46. Papers and Proceedings of the 19th Annual Congress of the European Economic Association.
- Brunnermeier, M. K., and L. H. Pedersen. 2006. "Market Liquidity and Funding Liquidity." Unpublished paper, Princeton University, November.
- Holmstrom, B., and J. Tirole. 1998. "Private and Public Supply of Liquidity." *Journal of Political Economy* 106, no. 1 (February): 1-40.
- Pavlova, A., and R. Rigobon. 2006. "The Role of Portfolio Constraints in the International Propagation of Shocks." Unpublished paper, MIT Sloan, December.
- Shin, H. S. 2006. "Risk and Liquidity in a System Context." Bank for International Settlements Working Paper no. 212, August.

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 Chapter 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 chapter.

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, 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

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.

control system, or from policies enforced 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 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 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

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

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.

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 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

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

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—analogueous 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 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

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

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.

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 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 Chapter 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 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, Haines 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 Haines 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 Wasily Leontief's input-output model that puts more emphasis on interdependencies), Haines 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

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

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 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 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

$$f_2(\cdot) = E[X | X \leq \beta_1] = \frac{\int_0^{\beta_1} xp(x)dx}{\int_0^{\beta_1} p(x)dx}$$
$$f_3(\cdot) = E[X | \beta_1 \leq X \leq \beta_2] = \frac{\int_{\beta_1}^{\beta_2} xp(x)dx}{\int_{\beta_1}^{\beta_2} p(x)dx}$$
$$f_4(\cdot) = E[X | X > \beta_2] = \frac{\int_{\beta_2}^{\infty} xp(x)dx}{\int_{\beta_2}^{\infty} p(x)dx}$$
$$f_5(\cdot) = \frac{\int_0^{\infty} xp(x)dx}{\int_0^{\infty} p(x)dx} = \int_0^{\infty} xp(x)dx$$

FIGURE 3.1 Partitioning the risk of extreme and catastrophic events.

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 equations in Figure 3.1.

The probabilities displayed in the equations in Figure 3.1 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.

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.

- $f_3(\cdot)$ 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. Figure 3.2 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 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.

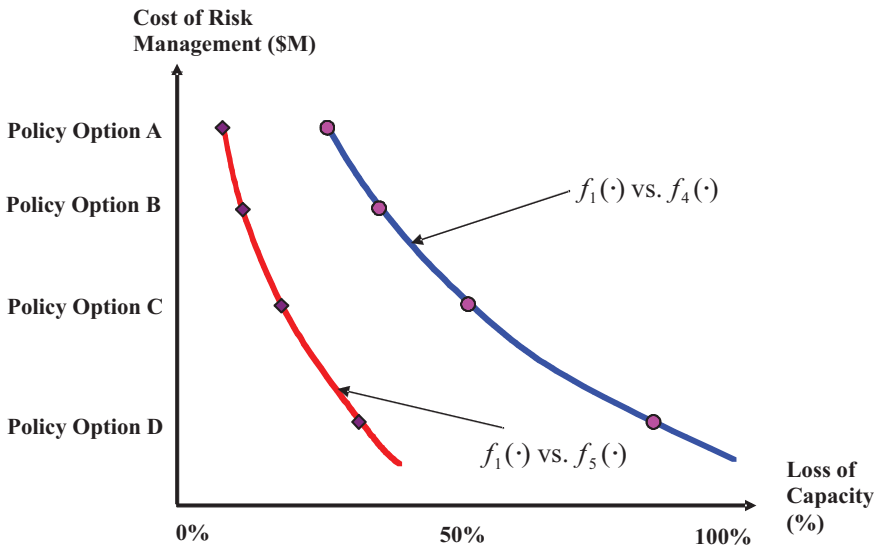


FIGURE 3.2 Notional depiction of the trade-off between the cost of risk management and the potential loss.

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 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 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

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.

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 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 multi-objective 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 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 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;

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

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, very tiny models, designed to build understanding, also play an important role.¹⁰ These models are comparable in spirit to the work described in Chapter 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,

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

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 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 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 Fed-

eral Reserve Bank of New York, Walter Beyeler and Robert Glass of Sandia National Laboratories, and Kimmo Soramäki of the European Central Bank, described in Chapter 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 Chapter 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, 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. Haimes 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 parameterized 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 infor-

mation 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 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

¹¹See Amin (2002).

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

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. “Modeling 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.: The 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.

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

4

The Payments System and the Market for Interbank Funds

The two papers in a session on systemic issues in the federal funds market examined the interbank payments system and the market for borrowing funds used to settle interbank payments. Both analyses were based on data on payments made through the Federal Reserve's Fedwire system.

Fedwire—the nation's primary interbank payments system—is a real-time gross settlement system through which payments are settled individually and with finality. More than 9,500 participants use Fedwire to send and receive time-critical and/or large-value payments on behalf of corporate, financial, and individual clients as well as to settle positions with other participants stemming from payments in other systems, securities settlement, and interbank loans. In the first half of 2006, an average of 525,000 Fedwire payments were made each day, with the daily value of funds transferred averaging \$2.1 trillion.

Complementing the Fedwire payments system is the U.S. federal funds market, where banks borrow funds to settle the payments they make over Fedwire. Fed funds are the bank reserves on deposit at the Federal Reserve used to settle payments between banks. The fed funds market plays a critical role in allocating liquidity in the financial system as well as in supporting banks' ability to finalize the settlement of interbank payment obligations.

The first paper, by Adam Ashcraft of the Federal Reserve Bank of New York and Darrell Duffie of Stanford University, explored whether trading frictions in the fed funds market affect the reallocation of excess reserves to banks requiring funds to complete their payments. Ashcraft

and Duffie found that fed funds trading is driven partially by individual banks' precautionary targeting of balances, and that this targeting contributes to systemic stability. The banks' trading of funds mitigates the risk of overconcentration of reserves in some banks, and contributes to the liquidity of the fed funds market and keeps price volatility relatively low. The second paper, by Morten Bech of the New York Fed, Walter Beyeler and Robert Glass of Sandia National Laboratories, and Kimmo Soramäki of the European Central Bank (who did not present), analyzed the network structure of interbank payments and developed a model for a payments system within such a network. The paper concluded that a liquidity market allows a payments system to achieve a specified level of performance with much less liquidity than would otherwise be required—a finding that sheds light on the trade-offs between liquidity within the payments system and friction within the liquidity market. The combination of network topology and bank behavior within a model, the study also found, is critical for analysis of systemic risk in payments systems. According to the authors, it is not sufficient just to understand the topology of the network; knowledge of the processes operating in that topology, such as bank behavior, is equally important.

SYSTEMIC DYNAMICS IN THE FEDERAL FUNDS MARKET

Ashcraft and Duffie analyzed how allocational frictions affect trading in the federal funds market. They also considered whether these frictions could lead to systemic risk in the form of "gridlock," in which individual financial institutions fail to transfer balances quickly to counterparties as they wait for the counterparties to transfer balances to them. Gridlock creates a self-fulfilling slowdown in the efficient reallocation of excess balances.

Like every over-the-counter market, the fed funds market is subject to allocational frictions because trading is executed through bilateral negotiation.¹ These frictions can be any sources of transaction costs or delays in identifying suitable counterparties, negotiating trades, or executing trades, and they can impact market efficiency. Existing theories of trading dynamics in over-the-counter markets have focused on "search" frictions, whereby traders locate each other with delays, to some extent by trial and error, and negotiate prices that depend in part on the difficulty of finding suitable alternative counterparties. Prices also reflect the relative benefits of making a trade immediately rather than later (and with a newly found counterparty) for each of the two counterparties. As frictions increase and

¹In the fed funds market, brokers reduce but do not eliminate the allocational frictions.

the matching of suitable pairs of counterparties becomes more difficult, a trader with an urgent need to transact has relatively less “leverage” during a bilateral negotiation, and this condition is reflected in the contracted price. The efficiency of allocation and the effect of search frictions on pricing are among the main concerns of the various theoretical studies. Yet there is little *empirical* work on these aspects of the microstructure of over-the-counter markets.

Ashcraft and Duffie broaden this relatively small body of empirical research as well as add a unique dimension to it. Their analysis of allocational frictions uses transaction-level data from Fedwire, and the majority of their work focuses on the top 100 institutions by payment, or “send,” volume for business days in 2005.² This data set permits the construction of real-time balances for each institution and allows for the tracking of the sender and receiver of payments and loans for every minute of the day. The authors identified a particular send as a loan, as opposed to another form of payment, by analyzing the terms of payments in the reverse direction between the same two counterparties on the next business day. (Fed funds loans are for overnight repayment.)

Ashcraft and Duffie documented evidence of federal funds trading being driven partially by individual banks’ precautionary targeting of balances. Banks are motivated to end each day with non-negative balances in their reserve accounts at the Federal Reserve because overnight overdrafts are not permitted except in special circumstances.³ These institutions are also motivated to end each day with relatively small balances, in part because the Federal Reserve does not pay interest on overnight balances and in part because the institutions have other ways to meet their reserve requirements over their two-week maintenance periods, such as by holding currency in large ATM networks and by sweeping funds in reservable accounts into nonreservable accounts.

By targeting its balances, a bank contributes to systemic stability. When its balances are larger than they typically are at a particular time of the day, a bank has an incentive to trade, and especially to lend, so as to reduce the balances. Ashcraft and Duffie showed that, empirically, banks indeed act consistently in response to this incentive. Likewise, when balances are low, banks trade (in particular, borrow) on average so as to raise them. This self-interested balance targeting at the individual bank level

²The researchers protected the confidentiality of the institutions by removing firm-specific details and conducting all data-related work within the Federal Reserve Bank of New York.

³The Federal Reserve’s discount window is available, but at terms that make it preferable to achieve non-negative balances through fed funds trading with other banks before the end of the day.

promotes systemic stability.⁴ It mitigates the risk of overconcentration of reserves in some banks and underallocation in others. Balance targeting reduces the risk of gridlock, and plays a role in keeping the federal funds market liquid and funds rate volatility relatively low.

Drawing on their discussions with fed funds traders, Ashcraft and Duffie reported that fed funds trading is significantly more sensitive to reserve balances in the last hour of the day. For example, at some large banks, fed funds traders responsible for targeting a small, non-negative end-of-day balance ask other profit centers of their institutions to avoid large, unscheduled payments, such as settlement of currency trades, near the end of the day. Once a fed funds trader has a reasonable estimate of the extent of current and yet-to-be-executed send and receive transactions, he can adjust pricing and trading negotiations with other banks so as to push his bank's balances in the desired direction. Ashcraft and Duffie uncovered empirical evidence of this behavior; furthermore, they found that such behavior is more pronounced following increases in intraday rate volatility.

In addition, the authors raised the issue of, but did not resolve, whether precautionary balance targeting by banks in the fed funds market, coupled with a regime in which banks forecast the targeting policies of other banks, could have systemically destabilizing consequences. A potential systemic problem could arise, for example, if several large institutions during a day of extreme misallocation of reserves individually "hoarded" reserves, given the heightened risk of other banks doing the same or the institutions' forecasts that other banks are incapable of releasing excess reserves quickly to the market. For instance, Ashcraft and Duffie reported traders' accounts of rumors that this type of behavior was initially feared on September 11, 2001, with the communications disruptions that day affecting the Bank of New York, a large clearing bank.⁵ Any such gridlock was in the end averted by energetic liquidity provision by the Federal Reserve.

Without significant liquidity provision by a central bank during such an event, "a run on reserves" could stress the ability of the fixed intraday supply of reserves to be sufficiently reallocated quickly to meet requirements. (The total amount of reserves in the system is relatively small compared with the total daily volume of transactions.) Even in an extreme

⁴One may think in terms of the usual "eigenvalue" or "mean-reversion" conditions for dynamic stability of a multivariate dynamic system. In this case, the coordinate processes of the system are the current balances of each bank in the system.

⁵McAndrews and Potter (2002) describe how the disruption of the regular timing of incoming payments made a bank's liquidity management more difficult, and for some banks the "increased uncertainty (regarding which payments they might receive later in the day) led them to have higher precautionary demand for liquid balances."

scenario, however, access to the discount window as well as infusions of liquidity by the Federal Reserve and other central banks would mitigate adverse systemic effects, as they did on September 11.⁶

NETWORK-BASED MODELING OF SYSTEMIC RISK IN THE INTERBANK PAYMENTS SYSTEM

The cross-disciplinary team of Bech, Beyeler, and Glass began by discussing two new approaches for characterizing the nonlocal, systemwide interconnections that may lead to systemic risk, and then combined the approaches to analyze the Fedwire interbank payments system. Their research first examined the global structure of interconnections in Fedwire by representing bilateral interbank relationships as a network. This approach allows one to quantify the overall pattern of interaction and interdependencies using well-defined measures applied to other complex networks, as well as to observe how those measures change during disruptions. Bech et al. next presented a model with simple agents interacting within a payments system network; the model exhibited a transition from independent to highly interdependent behavior as liquidity was reduced. When the authors applied the model to a liquidity market, they demonstrated that a reduction in market frictions can reduce this interdependence and thus the likelihood and size of congestive liquidity cascades.

Network Topology of Interbank Payment Flows

A payments system can be viewed as a specific type of complex network.⁷ In recent years, many fields of science have sought to characterize the structure of networked systems and the relationship between network topologies and stability, resiliency, and efficiency. From a communications or information-technology perspective, Fedwire is a “star” network, in which all participants are linked to a central hub: the Federal Reserve. Because of its ability to wire funds, Fedwire is a complete network, as all nodes (participants) can communicate (send and receive payments) with all others. However, the actual behavior of participants, the flow of liquidity in the system, and thereby the contagion channel for financial disturbances are not captured by these network representations.

The graph of actual payment flows over the Fedwire system includes more than 6,600 nodes and more than 70,000 links, but a subgraph, or core, of just 66 nodes and 181 links accounts for 75 percent of the value trans-

⁶For example, see McAndrews and Potter (2002) and Lacker (2004).

⁷See, for example, Newman (2003).

ferred. A prominent feature of Fedwire is that 25 nodes form a densely connected subgraph to which all the remaining nodes connect. By itself, this core is almost a complete graph, and this small number of banks and the links between them process the large majority of all payments sent over the network. Soramäki et al. (2006) show that the network shares many characteristics with other empirical complex networks. These characteristics include a scale-free degree distribution, a high clustering coefficient, and the “small-world” phenomenon. Apart from the core, the network, like many other technological networks, is disassortative. That is, large banks tend to connect to small banks and vice versa.

Bech et al. showed that the topology of the network was altered significantly by the attack on the World Trade Center on September 11, 2001. First, the massive damage to property and communications systems in Lower Manhattan made it more difficult—and in some cases impossible—for certain banks to execute payments to one another, as some nodes were removed from the system or had their capacity reduced. Second, the failure of some banks to make payments disrupted the coordination upon which banks rely when they use incoming payments to fund their own transfers to other banks.

A Payments System Model

In a paper by Bech et al. (2006), the authors use an elementary, agent-based model in the spirit of models applied to understanding self-organized criticality.⁸ Physicists have used these models to study cascading phenomena in a variety of systems (for instance, Jensen [1998]), where models made of very simple agents, interacting with neighboring agents, can yield surprising insights about system-level behavior. In the Bech et al. model, interbank payments occur only along the links of a scale-free network based on the authors’ analysis of Fedwire data; the model thus shows that only a very small fraction of the possible interbank exchanges tend to be active. Bank customers randomly instruct the institutions to make payments, and banks are reflexively cooperative: They submit payments if the balance in their payments system account allows them to; otherwise, they queue the instruction for settlement later.

If a bank receiving a payment has instructions in its queue, the payment enables it to submit a payment in turn. If the bank that receives the payment is also queuing instructions, it can make a payment, and so on. In this way, a single initial payment can cause many payments to be released from the queues of the downstream receiving banks. This is

⁸For example, see Bak, Tang, and Wiesenfeld (1987).

an example of the cascade processes typically studied in other models of self-organized criticality.

In the Bech et al. model, a single parameter—overall liquidity—controls the degree of interdependence of interbank payments. Abundant liquidity allows banks to operate independently; reduced liquidity increases the likelihood that a bank will exhaust its balance and begin queuing payments. When liquidity is low, a bank's ability to process payments becomes coupled with the ability of other banks to process and send payments. In this instance, the output of the payments system as a whole is no longer determined by the overall input of payment instructions but instead is dominated by the internal dynamics of the system, and the correlation between arriving instructions and submitted payments degrades as liquidity is reduced. The model does not exhibit a phase transition between completely noncongested and completely congested states. The distribution of congestive events, or liquidity cascades, is close to a geometric distribution, and it has an inherent length scale independent of network size. This result differs from the finding obtained in systems that exhibit self-organized criticality (such as in Jensen [1998]); the result is directly attributable to differences in the underlying relaxation process that set the simple payments system apart from systems that display self-organized criticality.

The Market for Interbank Payment Funds

Going beyond the study of the payments system in isolation, Bech et al. combined it with a simple model of a liquidity market. Liquidity market transactions were represented as a diffusive process, where liquidity flows are not confined to the links of the payments network. In the model, each bank is directly connected to a central node representing the market, and this connection is characterized by a conductance parameter that reflects the cost or friction associated with market transactions. The market creates a separate network of liquidity flows that, having a star topology, operates parallel to the network of payment flows.

The inclusion of a liquidity market weakened the coupling between banks and reduced the size of settlement cascades. Bech et al. identified trade-off functions that relate different levels of system performance, in terms of cascade size and queue size, to the value of initial system funding and market conductance needed to achieve that performance. The liquidity market is very effective in reducing cascade size and queue size. For a given level of performance, the rate of liquidity flow through the market relative to the rate of flow through the payments system was surprisingly small. The performance of the system can be greatly improved by the market even though less than 2 percent of system throughput flows

in the market. A liquidity market allows a system to achieve a specified level of performance with much less total liquidity than a system without a market would require. Conversely, the performance of such a system is highly dependent on the operation of the market. Disruptions to the market would greatly increase congestion and cascades unless they were mitigated, for example, through the addition of liquidity.

According to the authors, if the combined payments and liquidity market system is modeled with simple processes, the boundary between noncongested and congested states can be explained in terms of the relative values of three time constants: a liquidity depletion time, which is governed by the initial liquidity in the system; a net position return time, which increases with total deposits in the system; and a liquidity redistribution time through the market, which is associated with the market conductance parameter. An understanding of this boundary has significant practical applications because this understanding allows for direct consideration of the trade-offs between liquidity within the payments system and friction within the liquidity market, both of which are modified by central bank policies. While the Bech et al. model does not yet include the behavioral feedbacks that likely factor into the decision making of banks, those feedbacks can be included in order to consider how the congestion boundary may move. A goal of future studies will be to introduce into models this complexity as well as variability in the size of bank payments.

Wide-scale disruptions may not only present operational challenges for participants in a payments system, but they may also induce participants to change the way they conduct business, with the potential to either mitigate or exacerbate adverse effects (Bech and Garratt [2006]). An understanding of how participants interact and react when faced with operational adversity will assist payments system operators and regulators in designing countermeasures, devising policies, and providing emergency assistance, if necessary. Accordingly, as Bech et al. argue, it is not sufficient just to understand the topology of the network; knowledge of the processes operating in that topology is as important.

DISCUSSION

Robustness Issues

Much of the discussion following the presentations centered on questions concerning the robustness of the fed funds market and the interbank payments system. In particular, how much shock can the system tolerate—and would a shock move the system to a new, less desirable equilibrium? Moreover, if there are multiple equilibria, which are stable and which are

unstable? Although the models presented in this session suggested the resiliency of the payments system when combined with a liquidity market for interbank payment funds, the models were not complete enough to assess these issues definitively.

The models do not have multiple equilibria because they do not yet include the anticipatory behavior that could produce feedback effects leading to system gridlock. In the combined payments system and liquidity market model, congestive liquidity cascades within the payments system are mitigated through transfers within the fed funds market. Each bank is trying to mean-revert toward a target level of reserves, and the more quickly they adjust toward the target, or the lower the friction, the faster the system moves to a stable and stationary equilibrium. If expectations and anticipatory behavior are introduced, such behavior could produce a model with explosive or unstable equilibria—the “gridlock scenario” that so alarms payments system operators. While the model has not been extended in that direction, it is conceptually feasible to model the self-sustaining volatility that would prevent banks from reaching a stable equilibrium. Doing so empirically, however, would pose a challenge because the rarity of systemic crises severely limits the historical data on the behavior being modeled.

Another challenge in modeling empirically the fed funds market and the payments system is to account for the central bank’s deep involvement in the system as well as its ability to take corrective action to stabilize the system.⁹ Because of this stabilizing role, no data exist on the extreme states of the system where instability could be found. Those states never occur, or are extremely rare, because of stabilizing intervention. This consideration also complicates the formulation of the anticipatory behavior discussed above: Participants understand the stabilizing role of the central bank and may come to rely on it when forming their expectations—a point we consider later in this summary.

Large-Scale Simulation Models and Policy Decisions

Some of the discussion focused on what lessons might be learned from simulation models and how the integration of models of different scales might inform policy decisions. The models presented analyzed liquidity in terms of the liquidity properties of the network, such as the transition between congested and noncongested states, and in terms of the functioning of the marketplace for liquidity redistribution. These models allow

⁹As an analogy, consider the electrical power distribution system: What if the government instantaneously opened a new power line that was not there before whenever a failure occurred in the grid?

for simulation of the network's responsiveness to liquidity injections by the central bank. What is striking about them is their sensitivity to these injections. A useful endeavor, suggested during the discussion, would be to construct simulations that shed light on the amount of liquidity needed by the system in different types of stress scenarios. That is to say, how much liquidity should be provided in response to different physical disruptions to the network, and to which institutions should it be provided? For instance, while we tend to focus on the largest institutions—perhaps on the assumption that they are the most connected—a medium-size network participant might also be important because of the high degree of its connectedness in the network.

Another issue raised was whether the simulation models' degree of resolution would be comparable to the type of information policymakers may actually have in a crisis. For example, is it practical to build simulation models with a high level of resolution? And at what level of policy can simulation exercises be feasibly aimed—at a high level to inform policymakers of trade-offs and general principles, or at a more granular level to emphasize specific steps to take in a crisis?

With regard to these questions, conference participants would like to see closer integration between simulation models at different scales—between the small-scale behavior of individual decision makers and the large-scale aggregate behavior of the system. Accordingly, a hierarchy of models with different scales is needed. The general challenge here, common to multiscale modeling in most domains, is building realistic links between the small-scale behavior at the level of individual agents and the large-scale aggregate behavior of the system. A particular challenge in this application is likely to be the range of behaviors and expectations at the micro level of individual decision makers that can be feasibly modeled in simulations. Thus, to what degree can anticipatory behavior and expectations of economic agents be realistically represented in simulation models, especially in simulations of systemic crises that are not modeling "business as usual"? While the issue has both conceptual and empirical dimensions, the empirical basis for this behavioral component is apt to be limited because systemic crises occur so infrequently.

Behavioral and Mutable Aspects of Network Connections

John O'Brien of the Haas School of Business at Berkeley commented on the application of the web-and-node model to the events of 1987. As O'Brien explained it, the nodes represent the individual investment banks and brokers, and they were all under pressure to be more conservative. Therefore, each node was trying to eliminate its assets and reduce its leverage. The system overall would be harmed by that activity, however,

because as each node protected itself it would put pressure on the others. At the system level, the Federal Reserve was trying to make credit more readily available, but that change could not be pushed down to the node level because the top priority of the brokers in the investment banks was to meet margin calls and reduce their leverage. Since the situation in 1987 seems to fit the web-and-node model, O'Brien observed, what does the model suggest about policy if a similar situation were to develop? Robert Litzenberger of Azimuth Trust explained that the answer depends on knowing whether the shock will dissipate or be self-reinforcing. Because many events will cause only a mild market decline, we need to learn what types of conditions will exacerbate problems. One relevant insight is that economists sometimes partition the market into informed traders and liquidity traders. If the shock is affecting mostly liquidity traders, policymakers might decide to act, Litzenberger said. However, if some of the selling is attributable to informed traders, policymakers might instead opt to refrain from taking action.

Douglas Gale of New York University added that in 1987, some investors could not get in on the other side of the market because dealers essentially would not answer their phones. This problem illustrates one of the important differences between contagion in a complex system that includes rigid links between nodes (such as the electric grid), where the collapse of a node has to be triggered by an event that hits it in a physical sense, and the banking system, where links are somewhat fluid depending on the perceptions and expectations of the nodes. When the links are choices made by people anticipating contagion, the links might start breaking before the contagion reaches them. Gale's comment points to the central importance of human expectations and decisions in the operation of financial markets. Risk analysis of financial systems is more dependent on human behavior than is risk analysis of engineered systems.

Evolution of the Payments System

The session concluded with a discussion of whether the payments system could evolve into a less stable configuration. Participants considered constraints that could be placed on the evolution of the payments system to maintain its stability in the face of ongoing changes in the financial system.

An important issue raised in this context was moral hazard. Whenever the central bank intervenes to mitigate or reduce the effects of a systemic shock, it can influence the future risk-taking incentives of private sector decision makers by weakening the perceived need to plan for the occurrence of extreme shocks. This moral hazard issue can potentially influence the state of the system and its vulnerability to shocks through

the risks taken by investors and financial institutions. For researchers, this issue could be a factor when choosing the behavioral features of their models. Consideration of such behavioral issues, more generally, can add to the richness of the results obtained from simulation models.

One notable issue that might be illuminated by the models is the current tendency of payments in the system to migrate toward the end of the day. This shift in payments timing would appear to raise the likelihood of a congested state, as system participants seem to hold back payments until late in the day. Conference discussants noted that the Federal Reserve is looking at how payments system policies affect use of the liquidity pool. For example, which policies might be inducing participants to shift their payments to later in the day, and which might reverse that practice so participants can use the available liquidity in the system more efficiently? Two changes that might have played a role in the shift in payments timing in recent years are the reduction in reserve balances, as banks make more efficient use of reserves, and a higher demand on the pool of liquidity attributable to increases in the volume of obligations to be settled.

Another issue raised was whether the financial system is moving toward increased homogeneity and greater connectivity. Greater connectivity seems to be self-evident, but in some ways we are actually seeing an increase in heterogeneity. A shift has occurred from a bank-oriented financial system to a securities-market-oriented system in which a more diverse population of financial institutions interact. The diversity of investors in today's financial system—such as pension funds, insurance companies, mutual funds, private equity and hedge funds, investment banks, and commercial banks—makes for a more heterogeneous set of financial system participants than in the traditional bank-oriented models of the financial system. In this sense, the system has become more robust. That said, conference participants observed that a handful of very large financial firms play central roles in the financial system, and we need to look carefully at ways to increase the robustness of the systems and utilities that tie these firms together.

REFERENCES

- Bak, P., C. Tang, and K. Wiesenfeld. 1987. "Self-Organized Criticality: An Explanation of $1/f$ Noise." *Physical Review Letters* 59, no. 4 (July): 381-4.
- Bech, M., and R. Garratt. 2006. "Illiquidity in the Interbank Payment System following Wide-Scale Disruptions." Federal Reserve Bank of New York *Staff Reports*, no. 239, March.
- Beyeler, W., R. Glass, M. Bech, and K. Soramäki. 2006. "Congestion and Cascades in Payment Systems." Federal Reserve Bank of New York *Staff Reports*, no. 259, September.
- Jensen, H. J. 1998. *Self-Organized Criticality: Emergent Complex Behavior in Physical and Biological Systems*. Cambridge: Cambridge University Press.

- Lacker, J. M. 2004. "Payment System Disruptions and the Federal Reserve Following September 11, 2001." *Journal of Monetary Economics* 51, no. 5 (July): 935-65.
- McAndrews, J. J., and S. M. Potter. 2002. "Liquidity Effects of the Events of September 11, 2001." Federal Reserve Bank of New York *Economic Policy Review* 8, no. 2 (November): 59-79.
- Newman, M. E. J. 2003. "The Structure and Function of Complex Networks." *SIAM Review* 45, no. 2: 167-256.
- Soramäki, K., M. L. Bech, J. Arnold, R. J. Glass, and W. E. Beyeler. 2006. "The Topology of Interbank Payment Flows." Federal Reserve Bank of New York *Staff Reports*, no. 243, March.

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 and potential applications to policy.

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 Haimès of the University of Virginia and Massoud Amin of the University of Minnesota is in identifying extreme events. The conference background paper¹ 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 Chapter 1 of this volume vividly described as the feeling of regime shift during the events of 1997-98: the Asian currency crisis, the Russian default, and the Long-Term Capital Management collapse.

Chapter 3 of this volume explains the approaches followed by Haimès and Amin for identifying possible extreme events—for instance, a shut-down of the electric grid—and considering what set of circumstances could produce the failure. Haimès 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.”

¹The background paper can be found in Appendix B.

While Haimés’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.

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 Chapter 2 of this volume, and the research carried out in ecology and engineering, as glimpsed in Chapter 3. Economists were impressed by the quantity and quality of data available to researchers in the examples cited by Levin and by Haimés 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 Haimés 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.

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 Chapter 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 Chapter 4—economists and financial

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

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, 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 Research Council'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.

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 be either strong enough to resist disturbances or 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 finan-

cial 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.

REFERENCE

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

Appendixes

Appendix A

Conference Program

NEW DIRECTIONS FOR UNDERSTANDING SYSTEMIC RISK

Federal Reserve Bank of New York
33 Liberty Street, New York, New York
May 18-19, 2006

Thursday, May 18

8:30 a.m.-8:45 a.m. Welcome and Overview of Conference Goals

Speakers: *Christine M. Cumming*, First Vice President and Chief Operating Officer, Federal Reserve Bank of New York
Charles Lucas, Member, Board on Mathematical Sciences and Their Applications, National Research Council

Industry consolidation, global networking, terrorist threats, and heavy dependence on computing: these and other trends introduce the possibility of new systemic risks, perhaps with increased ramifications. Ms. Cumming and Mr. Lucas will discuss the importance of understanding and managing systemic risk and set the stage for the conference as an opportunity for the central banking community to take a fresh look at systemic risk, aided by the insights of researchers who study similar risks in other complex systems.

8:45 a.m.-9:45 a.m. Background on Systemic Risk in the Financial Sector

Speakers: *Darryll Hendricks*, UBS

Thomas Daula, Morgan Stanley

D. Wilson Ervin, Credit Suisse 1

Three presentations will help to identify key issues relating to systemic risk in financial markets and institutions, describe the structure of financial markets, and give historical examples of risks that concern central bankers and market practitioners.

10:00 a.m.-12:30 p.m. Presentations on Current Research Directions

Speakers: *Roberto Rigobon*, Massachusetts Institute of Technology

Hyun Song Shin, Princeton University

Markus K. Brunnermeier, Princeton University

Chair: *Franklin Allen*, Wharton School, University of Pennsylvania

A panel of experts in systemic risk in the financial sector will present a cross section of current work in this area. The chair will raise discussion topics for the panel and serve as moderator during a question-and-answer session with the audience.

12:30 p.m.-1:45 p.m. Welcome and Keynote Speaker

Welcome and Introduction: *Timothy Geithner*, President, Federal Reserve Bank of New York

Keynote Speaker: *Donald L. Kohn*, Governor, Board of Governors of the Federal Reserve System

1:45 p.m.-4:00 p.m. Panel Discussion: Models of Systemic Phenomena in Other Complex Interactive Situations

Panelists: *Yacov Haim*, University of Virginia

Massoud Amin, University of Minnesota

Chair: *Charles R. Taylor*, Risk Management Association

Two researchers who model systemic phenomena in nonfinancial systems will each make a presentation. The first explores concepts and methods for analyzing risk in complex engineered systems. The second outlines concepts and methods for modeling systems, infrastructure reliability, and catastrophic failures in complex networks such as power grids. The speakers will address the treatment of heavy-tailed events, the modeling of networks, and the identification of vulnerabilities. The session chair will compare and contrast the approaches with those typically applied in studies of systemic risk in the financial sector and then open the floor for discussion.

4:15 p.m.-5:30 p.m. Presentation on Systemic Dynamics in the Federal Funds Market

Speakers: *Darrell Duffie*, Stanford University

Adam Ashcraft, Federal Reserve Bank of New York

The presenters will discuss preliminary results of a simulation of the systemic risk arising from settlement flows in the federal funds market.

Friday, May 19

8:00 a.m.-10:15 a.m. Panel Discussion: Models of Risks Facing Complex Systems

Panelists: *Simon Levin*, Princeton University

Morten Bech, Federal Reserve Bank of New York

Walter E. Beyeler, Sandia National Laboratories

Robert J. Glass, Sandia National Laboratories

Chair: *George Sugihara*, University of California at San Diego

This session will present two talks about the risks facing complex systems. The first speaker will explore concepts and methods for analyzing behaviors of ecosystems, especially as they adapt to or approach precipitous changes. The second talk, by a cross-disciplinary team of researchers, will present a pilot attempt to analyze critical nodes in the financial transaction system using tools and concepts that are not in common use in the central banking community. A discussion with conference participants will follow.

10:30 a.m.-12:00 p.m. Wrap-up Panel Discussion: What Has Been Learned?

Panelists: *Douglas Gale*, New York University

Robert Litzenberger, Azimuth Trust

George Sugihara, University of California at San Diego

Vincent Reinhart, Board of Governors of the Federal Reserve System

Chair: *Timothy Geithner*, President, Federal Reserve Bank of New York

Panelists from the fields of finance, economics, and science will share observations on the conference findings and offer thoughts for the road ahead. Conference participants will be invited to respond.

Appendix B

Background Paper

SYSTEMIC RISK AND THE FINANCIAL SYSTEM

*Darryll Hendricks, John Kambhu, and Patricia Mosser*¹

Introduction

This paper is intended as background material for a cross-disciplinary conference, sponsored by the Board on Mathematical Sciences and Their Applications of the National Academy of Sciences and the Federal Reserve Bank of New York, on new approaches to evaluating systemic risks and managing systemic events in the global financial system. A key objective of the conference is to bring together a diverse group of leading researchers who have developed analytical tools for the study of complex systems in a range of fields of inquiry.

The stability of the financial system and the potential for systemic risks to alter the functioning of that system have long been important topics for central banks and for the related research community. However, recent experiences, including the market disruption following the attacks

¹Darryll Hendricks was a senior vice president at the Federal Reserve Bank of New York when this paper was prepared in May 2006; he is now a managing director and the Global Head of Quantitative Risk Control at UBS Investment Bank; John Kambhu is a vice president and Patricia Mosser a senior vice president at the Federal Reserve Bank of New York. The views expressed are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

of September 11, 2001, suggest that existing models of systemic shocks in the financial system may not adequately capture the propagation of major disturbances. For example, current models do not fully reflect the increasing complexity of the financial system's structure, the complete range of financial and information flows, and the diverse nature of the endogenous behavior of different agents in the system. Fresh thinking about systemic risk is therefore desirable.

This paper describes the broad features of the global financial system and the models with which researchers and central bankers have typically approached the issues of financial stability and systemic risk—information that can serve as a shared reference for conference participants. The conference itself will provide an opportunity for participants to discuss related research in other fields and to draw out potential connections to financial system mechanisms and models, with the ultimate goal of stimulating new ways of thinking about systemic risk in the financial system.

Systemic risk is a difficult concept to define precisely. A recent report by the Group of Ten (2001) on financial sector consolidation defined systemic risk as “the risk that an event will trigger a loss of economic value or confidence in, and attendant increases in uncertainty about, a substantial portion of the financial system that is serious enough to quite probably have significant adverse effects on the real economy.” This definition is broad enough to permit different views on whether certain recent episodes within the financial system constituted true systemic risk or only threatened to become systemic if they had a significant adverse impact on the real economy.

Some argue that even damage to the real economy is not sufficient grounds to classify an episode as systemic; rather, the key characteristic of systemic risk is the movement from one stable (positive) equilibrium to another stable (negative) equilibrium for the economy and financial system. According to this view, research on systemic risk should focus on the potential causes and propagation mechanisms for the “phase transition” to a new but much less desirable equilibrium as well as the “reinforcing feedbacks” that tend to keep the economy and financial system trapped in that equilibrium.

While differences in the definition of systemic risk are clearly important from a policymaking perspective, this paper includes discussions of episodes that not everyone would agree were systemic in nature. This is because our primary interest is in stimulating further research on the types of propagation or feedback mechanisms that might cause a small financial shock to become a major disturbance, allow a financial shock to have a material impact on the real economy, or mire the financial system in a suboptimal equilibrium. In this regard, the dynamics of nonsystemic episodes may still be very relevant to the modeling of financial market

behavior. Moreover, as noted by a recent private sector report on risk management practices (Counterparty Risk Management Policy Group II [2005]), “Unfortunately, in real time it is virtually impossible to draw such distinctions.”

The remainder of this paper is divided into three sections. The first describes the classical case of systemic risk in a banking-dominated financial system and provides some background information on the current workings of the international banking system. The second section focuses on issues that arise in market-oriented “panics” (for example, the October 1987 stock market crash) and again seeks to provide some relevant background information. The third section discusses the challenges in understanding the full nature of systemic risk posed by events of the last decade (for example, the Asian currency crises and 9/11 payments system disruptions) as well as key ongoing trends in the financial markets generally. Significantly, this paper is meant primarily to stimulate discussion of relevant issues at the conference; it is not intended to provide a comprehensive overview of the substantial economic literature on systemic risk and financial instability.

Systemic Risk in Banking Systems

Banks have long been at the center of financial activity. They remain so today, even though their share of financial intermediation has been reduced by the growth of capital markets and mutual funds and other developments of the last few decades. The largest commercial banks have balance sheets in the \$1 trillion range, engage in extensive international operations, and maintain a presence in a wide variety of retail and wholesale financial business activities. These activities include making loans to corporations and individuals; underwriting debt and equity securities offerings; acting as dealers in foreign exchange, securities, and derivatives markets; providing asset management services; providing payments, settlement, and custodial services; and taking deposits.

The classical model of a commercial bank is a firm that makes loans on the asset side of its balance sheet and takes demand deposits (checking and savings accounts) on the liability side.² The loans are typically perceived as being long-term “illiquid” assets in the sense that efforts to liquidate them prior to maturity will yield a reduced value relative to their intrinsic worth if held to maturity. However, the bank is obligated to pay back demand deposits at any time the depositor requests. Thus, banks are seen as providing a fundamental maturity and liquidity transformation that is both beneficial and inherently unstable.

²The following discussion draws heavily on Diamond and Dybvig (1983).

This instability can be seen by considering a case in which each depositor at a particular bank would be willing to leave his or her funds on deposit, but believes that other depositors are likely to withdraw their funds, thus making it necessary for the bank to call in its loans and suffer the associated losses. In this case, all rational depositors will seek to withdraw their funds as quickly as possible, producing a “run” on the bank. In this simplified model, bank runs can be caused by concerns over liquidity even if the bank’s assets are fundamentally sound on a going-concern basis (that is, the bank is solvent). The distinction between illiquidity and insolvency is one that occurs repeatedly in discussions of systemic risk.

Moreover, in this model, bank runs can be contagious. The contagion can arise simply as a result of a self-fulfilling prophecy if depositors believe that other depositors will regard a run at one bank as an indication that runs are now more likely at other banks. Somewhat more concretely, contagion may be more likely to occur if the issue that sparked the original run—excessive loan exposure to real estate or the oil industry, for example—is perceived potentially to affect other banks, or is the result of concerns about significant interbank exposures (that is, runs at banks seen as having large exposures to the bank subject to the original run). Naturally, in this model, runs are more likely at banks perceived to have a smaller equity capital cushion to absorb declines in asset values and at banks whose financial condition is difficult to assess in the first place.

Although the model just described is highly simplified, it nevertheless captures the essence of past bank runs, which occurred with some frequency before the 1930s. The primary approaches to dealing with the risks inherent in banking activity have included (1) controlling the relative risk of the loans extended—for example, through regulation, (2) requiring that bank balance sheets contain a larger share of equity capital and a smaller share of demand deposits, and (3) ensuring that government provides a “lender of last resort” function and/or deposit insurance.

The lender of last resort role is one of the most distinctive functions of central banks. In this role, central banks such as the Federal Reserve typically have the authority to provide short-term loans to banks against collateral. For example, a bank could pledge some of its loans to the central bank and obtain cash on a short-term basis. In determining whether to provide funds, the central bank must make a judgment about the bank seeking funds. The conventional wisdom that emerged in the nineteenth century was that central banks should “lend freely at a penalty rate” when they believe that the bank needing funds is illiquid but not insolvent, but should not lend at all to a bank that is truly insolvent. Of course, there is often substantial practical difficulty in distinguishing illiquidity from insolvency.

The provision of deposit insurance in the United States followed the bank runs of the early 1930s. Deposit insurance aims to eliminate the threat of a bank run directly, by assuring depositors that they will be paid regardless of whether the bank ultimately fails. While clearly effective in discouraging bank runs, deposit insurance further reinforced the need for bank regulation to limit the extent of banks' risk taking. Economists refer to the incentive problems created by the presence of deposit insurance as an instance of "moral hazard." That is, bank managers will want to take on risk to increase their upside potential, but insured depositors will have no incentive to monitor or constrain their behavior. Thus, the bank runs of the Great Depression served to shape the institutional framework in which banks operate today—a framework that emphasizes official regulation and supervision of banks.

In considering the systemic risk associated with banking crises, one should also bear in mind the social costs of such episodes. On balance, the economic literature on the Great Depression in the United States concludes that much of the social cost of this episode stemmed from the interruption to credit allocation that occurred as a result of the bank runs and contraction of the money supply. That is, the broader, nonfinancial portion of the economy was seriously hurt by the interruption in the financing of its activities and by the reluctance of banks to extend new financing amid a series of bank runs. Concern that financial sector crises may adversely affect the functioning of many other parts of the economy is a recurrent theme in discussions of systemic risk.

Although the overall importance of banks within the financial system has declined in the last few decades, the largest banks in the key financial centers remain sufficiently important that their failure to function normally would raise questions of a systemic nature. Significantly, these institutions exhibit several of the characteristics discussed above.

- They are highly leveraged, with equity-to-total-asset ratios ranging between 5 percent and 10 percent.
- While banks are less reliant on short-term-deposit funding than the stylized model just outlined would imply, such funding remains a material part of the liability structure for the largest institutions.
- The scope and complexity of their activities and legal/organizational structures make assessments of their true financial condition by outsiders difficult, while also posing significant management challenges for the banks themselves.
- The largest banks typically have significant exposures to one another, for example, through interbank deposit markets, interdealer transactions in over-the-counter derivatives, and wholesale payment and settlement arrangements.

- According to some commentators, banks remain particularly prone to cyclicality and myopia in their credit processes, tending to forget the last cycle of bad lending too rapidly when economic conditions brighten. The old saying that “bad loans are made in good times” captures the essence of this concern.

- Finally, the largest banks appear to be increasingly subject to legal and regulatory risks stemming from actions of their employees, risks that in some cases could result in sudden adverse impacts.

Nevertheless, despite the vulnerabilities just outlined, the financial system today does not seem highly prone to contagious runs on very large banks. This reflects the relative profitability and health of banks in many countries, their risk management discipline, and the perception that the largest banks would benefit from liquidity provision or other forms of official assistance should runs appear imminent. In Japan, for example, official intervention following the emergence of significant banking sector problems in the 1990s largely forestalled major bank runs. Interestingly, however, Japan’s policies to prevent runs did not prevent economic weakness associated with a banking sector too fragile to play a full and vibrant role in financing broader economic activity. In the United States, policy-makers have indicated that large banks are not “too big to fail” and have worked to ensure that such banks maintain a strong financial condition and adopt rigorous risk management policies and procedures.

This last point reflects a concern that inappropriate public policy choices can serve to generate systemic risk. For example, many observers note that the U.S. savings and loan crisis of the late 1980s resulted in part from policies that paid insufficient attention to “moral hazard” concerns. In addition, supervisors failed to deal with insolvent firms promptly, creating strong incentives for the management of such firms to invest in high-risk projects in an effort to restore solvency. The downside risks of these (frequently suboptimal) investments ended up being borne by society at large, both in the cost of government bailouts of depositors and in the opportunity loss of numerous investments yielding little or no return.

While the liquidity-based contagious run model of systemic risk applies very directly to banks, it also has relevance to other kinds of institutions. The largest securities firms rely on debt rather than bank deposits as a significant funding source and hold a greater share of their assets in the form of marketable securities than do banks. Nevertheless, some analysts have argued that securities firms may be vulnerable to contagious runs because of their reliance on short-term funding sources such as commercial paper, the complexity of their transactions in less liquid securities markets, and their derivatives businesses. As leveraged institutions,

hedge funds that do not effectively manage their liquidity risks could also be subject to runs by their investors and creditors. Indeed, liquidity risk management failures contributed to the problems experienced by the Long-Term Capital Management (LTCM) hedge fund in 1998. The case of LTCM is discussed further in this paper's final section, as it raises other issues about the sources and propagation of systemic risk.

Significantly, a run on an individual firm alone might not be enough to create systemic risk according to the definition outlined above unless the liquidation of assets by the firm or an associated reduction in the firm's underwriting activities were to have a material impact on economic growth. For example, in 2001, Enron suffered what amounted to a run on its short-term liabilities in the period immediately preceding its bankruptcy filing, but there appeared to be very limited systemic contagion to other energy-trading firms and very little impact on the broader economy.

Systemic Risk in Financial Asset Markets

While the bank run model of systemic risk has been studied fairly widely in the financial economics literature, more recent examples of events in which concerns about systemic risk arose have often been associated with disruptions to financial markets, rather than runs on particular financial institutions. For example, the 1987 stock market crash was not precipitated by concerns at an individual institution, nor was it the proximate cause of the failure of any large bank. Nevertheless, it was clearly viewed—at the time and since—as an event with potentially systemic consequences that warranted official sector intervention.

A market-oriented systemic crisis typically manifests itself as a breakdown in the functioning of financial markets for traded assets such as stocks and bonds, and it may develop in response to a sharp decline in the value of one particular type of asset. In addition to the 1987 stock market crash, examples of such crises might include the widening of interest rate spreads and decline in liquidity following the collapse of LTCM in 1998 and the collapse of the junk bond market in 1989-90. In the more distant past, the Dutch tulip mania of the 1630s and similar episodes could, in their end-stage, be viewed as additional examples of this type of crisis.

Consider first the characteristics of the 1987 stock market crash. Two aspects of this systemic market episode are particularly important to highlight. First, the episode suggests that asset price declines can in some circumstances become self-reinforcing and even feed into a reduced willingness on the part of major financial institutions to bear risk across the full range of their activities. Second, the episode underscores the potential importance of not only the specific institutional arrangements that are in

place for clearance and settlement of transactions but also the credit and liquidity exposures arising from those arrangements.

Market-Based Financial Crises, Liquidity, and Self-Reinforcing Price Movements

The shift of emphasis from bank runs to “market gridlock” as a source of systemic risk has arisen from a number of factors, not least the success of policies aimed at preventing bank runs mentioned earlier. In addition, financial crises now manifest themselves in markets rather than in institutions because financial intermediation has moved into markets and away from institutions. This “disintermediation” in financial activity has been particularly pronounced in the United States in the last thirty years. For example, in 1975, commercial banks and thrifts held 56 percent of total credit to households and businesses; by 2005, this figure had dropped to 33 percent. A large fraction of financial assets—both equity and debt—is sold directly by issuers/borrowers to investors, especially institutional investors, via stock and bond markets, with traditional banking effectively bypassed.

The shift from a bank-based to a market-based financial system has expanded the types of activities that banks and other financial intermediaries engage in and the assets that they invest in. The large financial institutions at the core of the system have expanded their activities to intermediate the movement of capital among the various other participants in multiple ways. They assist businesses in the issuance of new stocks and bonds directly into the market (investment banking), intermediate to buy and sell stocks and bonds (market making) after issuance on behalf of clients (broker-dealers/trading desks), manage asset portfolios on behalf of individuals and institutions (asset management), and lend directly to households and businesses (traditional commercial banking). A general trend toward consolidation of financial activity has led to the formation of large, complex institutions at the core of the financial system. At the same time, however, disintermediation has increased the importance of “end-user” financial institutions that invest in securities on behalf of households and firms. These include mostly unleveraged institutional investors (mutual funds, pension funds) as vehicles for household savings as well as more lightly regulated and more leveraged risk-bearing entities (hedge funds).³

³Note too that any large market participant itself consists of a very large number of separate legal entities, with many different charters, incentive structures, constraints, and regulations.

Market-based financial intermediation has a number of advantages over a banking-oriented financial system. One important advantage is that the investment risk in holding securities is dispersed broadly among investors instead of being concentrated in financial intermediaries. For example, debt instruments issued by ultimate borrowers are held directly by savers/investors to a greater degree than in a banking-oriented financial system. Another feature of today's financial system that works to reduce systemic risk is the replacement of bank deposits by mutual fund shares as an investment vehicle for households. While the fixed face value of a bank deposit is inherently fragile, the value of a mutual fund share fluctuates with market prices daily. As a result, the mutual fund model is better able to absorb and disperse shocks across a wide set of investors.⁴

Although superior to a banking-oriented financial system in some respects, market-based financial intermediation carries its own vulnerabilities. The capital markets work best when key markets are liquid. In this context, the term *liquidity* refers to tradability. When a market is liquid, any single trade to buy or sell a particular asset is unlikely to have a major effect on the price of the asset because of the large number of willing transactors on both sides of the market. Market liquidity also ensures that investors can buy and sell securities without undue delay or loss in value from the price impact of the transaction. Almost by definition, then, a market-gridlock systemic crisis is a period when market liquidity is absent.

In normal circumstances, market liquidity rests on a number of foundations. Foremost among them are market making, trading, and arbitrage. Market makers buy and sell out of inventory they maintain to meet customer demand. They smooth out short-run imbalances in market supply and demand, and profit from the bid/ask spread. Traders also contribute to market liquidity by trading on bets that prices will converge to long-run fundamental levels. These traders typically take positions that they hold for potentially long periods of time until prices converge to their long-run norms. Traders play an important role in maintaining the stability of markets and speeding up the convergence of prices to their fundamental values.

Market-oriented crises tend to begin with a large change—usually a decline—in the price of a particular asset; the change then becomes self-sustaining over time. When asset prices drop sharply, there are generally some participants willing to “swoop in” and buy assets that have

⁴Money market mutual funds raise some of the same issues as bank deposits because of their limited ability to bear credit losses; historically, parents of such funds have absorbed impaired money market instruments rather than allowed a credit loss to reduce the fund's share value below \$1.

declined sufficiently in price—an action that largely prevents the stress from becoming worse. In systemic crises, however, investors and traders are either unable or unwilling to step in, perhaps because their own losses have limited their trading capacity or because an infrastructure failure in, say, payments or settlement systems has made such a step difficult. As prices decline, more and more market participants sell, pushing prices lower. Eventually the price declines become so large and persistent that no buyers emerge, market liquidity dries up, market participants reduce their intermediation activities and their risk taking, and market gridlock takes hold. This sequence of events is in some measure self-reinforcing: if price declines are sufficiently large to create losses for traders and market makers, these participants may cease providing liquidity to the market, thereby exacerbating the price declines.

Market-based crises are often characterized by a coordination failure in which a wide cross section of participants in financial markets, including market makers, simultaneously decide to reduce risk taking and effectively pull back from financing activities (trading stocks, issuing new stocks and bonds, lending, and so forth). While no one institution is necessarily insolvent or illiquid, each firm reduces its activity and risk to protect capital and profits. In aggregate, however, the firms' actions combine to slow down or stop financial market activity. In severe cases, the financial system could become almost paralyzed and unable to perform its core functions of channeling capital to investment opportunities. The period immediately following the 1987 stock market crash is an example of this type of coordination failure, although its consequences were contained.

The potential for self-sustaining dynamics in financial price movements has been studied extensively in the finance and economics literature. Minsky (1977) and Kindleberger (1978) have advanced theoretical explanations for many varieties of financial crisis in which an exogenous change in the economic environment leads to the creation of new profit opportunities that attract capital fed by an expansion in credit.⁵ For a time, these investments give rise to even more profit opportunities, leading to a speculative euphoria that, by involving segments of the population typically not involved in such ventures, becomes a "mania" or a "bubble." However, at a certain point, knowledgeable insiders begin to take profits and sell out. Prices begin to level off and some financial distress may ensue. A crisis occurs when a specific event precipitates the equivalent of a run on the asset class that was the subject of the speculative frenzy. Aversion develops toward that asset class by banks and others that had previously lent against it, and with this aversion arises a desire to obtain

⁵The discussion in this paragraph draws heavily on Kindleberger (1978).

liquidity at nearly any cost. The resulting panic culminates when (1) asset prices fall so low that investors are tempted back, (2) trading is cut off, perhaps by the closing of the relevant exchange, or (3) a lender of last resort succeeds in convincing the market that sufficient liquidity will be available if necessary.

Although not all economists would subscribe to this broad theory of speculative financial crises, the theory is useful to keep in mind, especially in relation to those features of the financial system that could make the system particularly vulnerable to large, self-sustaining changes in asset prices, and thus to market gridlock.

In modern financial systems, debt and leverage are necessary and pervasive. Many market participants, including the largest intermediaries, borrow funds in order to expand their balance sheets and thus increase their ability to invest and trade in financial assets. They adopt this strategy to increase their return on equity capital invested (that is, by holding assets expected to yield returns exceeding the cost of the funds borrowed). As noted earlier, the largest banks are nearly all leveraged more than ten to one, implying that such institutions cannot afford to realize losses greater than 10 percent of the value of their assets if they are to remain solvent.

The obvious implication of leverage is the need for financial institutions to control their losses carefully and to take steps to reduce their risk taking in the face of declining asset values. In other words, leverage creates an incentive to sell assets whose prices are declining, particularly if further price declines are expected in the future. For example, if a firm is leveraged ten to one, then even a 1 percent realized loss in asset value translates to a 10 percent loss in the firm's capital value. Collectively, of course, widespread selling after an asset price decline will likely push prices even lower and losses higher. This scenario raises the obvious concern that such liquidations would further amplify the underlying price movements.

Moreover, in some markets, liquidations after losses can be automatic. For example, when an investor trades stocks on margin accounts (by borrowing a percentage of the stock value), a subsequent decline in the value of the stock requires that the investor post (add) collateral—usually cash—in order to bring the margin account back into compliance with the margin rule of the stock exchange.

In the 1987 stock market crash, large margin calls required investors to sell stock, thus putting further downward pressure on stock prices. The sudden and large fall in stock prices created large debits in the accounts of investors that had purchased stock on margin at brokers or held long positions in equity-linked derivatives contracts on futures exchanges. These margin account debits created a need to transfer large sums of

cash that many investors were not able to provide within the time frames required by brokers and the futures exchanges.

An additional feature of contemporary financial markets that can create self-reinforcing asset price dynamics relates to financial products that exhibit convexity in their price behavior. Assets (or derivatives) with convexity are those that become more or less sensitive to changes in an underlying asset price (or other variable) as that price or variable changes.

The classic example of convexity is an option. The buyer of a call (put) option has a right, but not an obligation, to buy (or sell) a particular asset (for example, 100 shares of IBM) at a particular price at some point in the future. Conversely, while the option buyer has a right to exercise, the writer or seller of the option has an obligation to perform. Those who have sold options to others are exposed to what market participants call negative convexity: as the underlying asset price moves against the seller of the option, the value of the option position becomes increasingly sensitive to further changes in the price of the underlying asset. In the case of a put option, if the seller of the option should try to compensate for this increased sensitivity by selling the underlying asset as a hedge against further price declines, it will put additional downward pressure on the underlying asset price. But a further decline in the underlying asset price simply increases the option sensitivity again, prompting even more selling. Thus, what appears to be a risk-mitigation strategy by the option seller locally is, in fact, a strategy that can reinforce adverse asset price dynamics when undertaken by a large number of sellers.

This phenomenon was evident in the 1987 stock market crash. At that time, many institutional investors had purchased portfolio insurance from intermediaries or were attempting to replicate such insurance through dynamic hedging strategies. Portfolio insurance is nothing more than a put option on the underlying asset; it exhibits exactly the characteristics outlined above. The seller of the insurance (or the firm trying to replicate it) must hedge its position by selling in greater amounts as prices decline, creating even further downward price pressure. Although the extent to which such activity was responsible for stock price declines in October 1987 is heavily debated, there is little doubt that such strategies—if widespread—could create self-reinforcing market movements.

Importance of Clearance and Settlement Arrangements

Clearance and settlement mechanisms contributed greatly to the liquidity strains created by the large price declines across cash, futures, and options markets, and the resulting margin calls in the 1987 stock market crash. The different settlement arrangements and time frames for different products (that is, T+5 for stocks traded on the exchanges at that time in

contrast to same-day settlement for stock index futures) meant that even investors that were hedged across the different markets could face large cash demands during the interim period.

This sudden need for large cash transfers threatened to create gridlock in the payments system and in the stock and futures markets. Securities firms did not have the funds to make margin payments at futures exchanges because their customers had not made margin payments to them. The futures exchanges' credit risk management practices required that positions be closed out when margin payments were not made. This unwinding of futures positions would likely have triggered further massive selling pressures in the stock market, exacerbating what had already occurred. However, the concentration of risk in the clearinghouses used to guarantee settlement of both securities and futures transactions meant that if positions were not closed out and markets fell further, the integrity of the clearinghouses themselves could be threatened. Since these clearinghouses form a core part of the infrastructure supporting the relevant trading activities, such an outcome could have significantly impaired market functioning for a sustained period of time.

In the end, large commercial banks were persuaded of the need to supply liquidity to those firms most heavily exposed to equities (that is, by lending against the value of those portfolios), and the most severe consequences were averted. However, the banks' action was due in part to official sector appeals to their collective desire to avert a further deepening of the crisis, as well as a stated willingness by the Federal Reserve to make more liquidity generally available to the banking sector.

This example illustrates the presence of systemic risk in wholesale market payments, clearance, and settlement, owing to the very sizable credit and liquidity exposures that typically characterize such arrangements, particularly on an intraday basis. In normal circumstances, the extension of such large amounts of intraday credit and liquidity between the major participants in these mechanisms facilitates more rapid settlement of the transactions. During a crisis, however, the reluctance of participants to continue doing business in this fashion can potentially lead to gridlock.

A case in point is the 1974 failure of Herstatt Bank, a midsize German bank that was closed down after it received the deutsche mark leg of its deutsche mark-U.S. dollar currency trades but before its pending U.S. dollar payments were completed in the United States. This created a short-term gridlock in the foreign exchange market that remained a source of systemic concern until the mid-1990s, when central banks made clear that the amounts of such "payment versus payment mismatch" were too large to be tolerated indefinitely and the large commercial banks invested in

the CLS Bank, a system for simultaneously settling both sides of foreign exchange transactions.

Broadly speaking, central banks and others in the official sector have been pushing for continuing improvements in the robustness of payments, clearance, and settlement mechanisms. These improvements provide greater assurance to investors that their transactions will settle, and that the mechanisms for payment and settlement will not themselves become a channel for propagating systemic disturbances. Nevertheless, these arrangements remain highly complex and are increasingly concentrated. For example, most market participants effectively outsource payments, clearance, and custodial functions associated with their transactions to an increasingly small number of global banks that specialize in those activities.

In turn, these banks at the core of the financial system interact with a relatively small number of specialized organizations that actually provide the central settlement functions for specific assets. For example, the Federal Reserve provides settlement services for U.S. dollar wholesale payments through its Fedwire service while the European Central Bank does the same for euro-denominated payments. The Federal Reserve is involved in settlement services for U.S. government bonds through its book-entry and transfer services for those securities, while the Depository Trust and Clearing Corporation provides clearance and settlement services for a wide range of securities, including all equities traded on U.S. stock exchanges and corporate bonds. Significantly, all of these systems continue in one form or another to provide large amounts of intraday credit to their major participants.

Many financial markets (especially securities and futures markets) have, in addition to the settlement mechanism, a clearinghouse that provides further assurance that transactions will settle by interposing itself as the legal counterparty to both sides of the original transaction. The clearinghouse typically imposes margin requirements or other controls on member transactions while also maintaining its own financial resources and/or the ability to call on its members' resources. Although clearinghouses have the ability to contain financial distress, the concentration of settlement risk in an exchange has the potential to focus it—as the 1987 stock market crash vividly illustrates—if problems are severe enough to call the integrity of the clearinghouse into question.

The Role of Central Banks

Central banks have historically played a key role in ensuring that financial markets have sufficient liquidity to function effectively. They have several tools that can be used in this regard. First, they control the aggregate

supply of bank reserves—the ultimate unit of account. By increasing the supply of reserves, central banks can increase the aggregate amount of liquidity in the financial system. Second, central banks function as the lender of last resort, a role that gives them the ability to lend directly to individual commercial banks. In extraordinary circumstances, the Federal Reserve System also has the power to lend directly to any individual or corporation, although this power has not been exercised since the 1930s. Third, the central bank typically possesses sufficient influence to persuade market participants that a collective decision to make liquidity available in particular circumstances will produce a better outcome than if individual market participants all seek to “free ride” on the actions of others. Largely because these tools are so effective, the central bank can often forestall liquidity pressures simply by announcing its willingness to make liquidity available should the need arise. Such announcements were made by the Federal Reserve in the wake of the 1987 crash as well as after the events of September 11, 2001. Elaborate assurances of this kind were also given in advance of the Y2K rollover.

Of course, central bank actions to forestall financial crises may themselves have a cost. In line with the moral hazard argument discussed in the section on banking-oriented crises, it is important that market participants not become so complacent that they count on the central bank to defuse any potential market-oriented financial distress and thus underinvest in their own management of market and liquidity risks.

New Sources of Systemic Risk

In the last ten to twenty years, financial markets have evolved significantly. They are more global and involve a wider range of more complex products than ever before. In some areas, market activities have become increasingly concentrated in a handful of very large firms. In other areas, the role of smaller, more specialized entities has grown significantly. From a policy perspective, there does not seem to be a clear consensus on whether the financial system today is more or less vulnerable to systemic disturbances than it was in, say, 1990.

Moreover, several of the most significant financial market disturbances of the last decade manifested features that, though present in earlier financial crises, have become more prominent. As the “supply chain” has evolved from the simplicity of a bank’s making and servicing a loan over its life to the complexity of securitization (involving originators, holders, servicers, trustees, and hedging markets), the focus on core banks and securities firms and major markets must expand to include other potential single points of failure. In addition, the economic forces leading to consolidation have included economies of scale in risk and

liquidity management. The liquidity needed in key market-risk-management markets and in the processing of high-value dollar payments derives in substantial part from the natural offsetting of risks or payments when volumes are high. Finally, the global scale of large banks and securities firms and some major investors has expanded the channels that can transmit systemic risk.

These new features raise interesting questions about whether the kinds of conceptual models outlined in the preceding two sections fully capture the range of possible causes and propagation channels for systemic risk. The discussion below addresses two cases: (1) the events of 1997 and 1998 that involved currency crises in several Asian countries, the Russian debt default, and the collapse of the Long-Term Capital Management hedge fund, and (2) the disturbances in payment and settlement arrangements following operational disruptions resulting from the terrorist attacks of September 11, 2001.

Asia, Russia, and LTCM

This sequence of events began in the summer of 1997 as certain Asian countries faced a substantial change in market sentiment that exposed relatively fragile macroeconomic conditions. In particular, several countries had short-term foreign currency debts that far exceeded their international reserves. The countries were thus susceptible to a run on their currencies, with generally negative consequences from a macroeconomic point of view. While currency crises are an extremely well-studied subset of economic crises, the Asian episode was notable in several respects.

First, the Asian crisis was characterized by a significant interplay between macroeconomic and financial sector factors. This interplay reflected weakness in the banking sectors of some countries that, while not the root cause of the crisis in all cases, clearly affected how the crisis played out and how well each country absorbed the macroeconomic impact of the crisis.

Second, consistent with the model of bank runs outlined earlier, contagion figured very prominently in the Asian crisis. Indeed, the events demonstrated a new mode of contagion. Various trading and risk management strategies now commonly used by market participants created linkages between different assets and activities that may not have previously existed, in some cases requiring positions in one currency to be adjusted largely as a result of movements in another. In some instances, a problem triggered by a currency or maturity mismatch in one country or market would lead global investors seeking to reduce risk to identify similar vulnerabilities in other markets.

A year later, contagion figured in the relationships between Russian

debt and the debt of Brazil and other emerging economies. Although the economies of Russia and Brazil are not themselves closely integrated, the prices of their debt fluctuated largely in tandem. In part, these parallel fluctuations reflected the fact that many of the holders of this debt specialized in holding the debt of emerging market countries, regarded these countries as proxies for each other, and needed to maintain some stability in their overall risk profile. Thus, when Russian debt began to be perceived as increasingly risky and to lose liquidity, some of these participants began to sell their Brazilian debt to reduce their risk profile and to take advantage of the Brazilian debt's greater liquidity. Ultimately, of course, the correlation between these two assets broke down as Russia defaulted while Brazil did not.

The Russian government default of August 1998 occurred against the backdrop of the Asian crisis that had been playing out over the preceding year, but otherwise took place in a period that was characterized both by the strong macroeconomic performance of the United States and by the strong financial condition of the major financial intermediaries. Nevertheless, the Russian default set in motion a chain of events that created significant fear among the leadership of those same intermediaries and served to reduce liquidity across most of the world's capital markets for some months.

LTCM was a hedge fund that conducted leveraged trades involving both securities and derivatives on a large scale and used highly sophisticated mathematical approaches to manage its risk. The firm suffered a severe loss of capital when prices moved against its positions following the Russian default. While LTCM's uniquely high leverage made it a fragile enterprise, it may not have been the only leveraged investor to be vulnerable, and this broader vulnerability may have played a role in amplifying the price shocks that occurred in a number of markets following the Russian default. For a year or two before the crisis, the liabilities of financial intermediaries had increased substantially relative to the liabilities of the nonfinancial sector, suggesting that others besides LTCM had also taken on more debt and were similarly vulnerable to price volatility and liquidity shocks. At the onset of the crisis, however, signs of an abrupt scaling back of leverage in trading activity emerged. For example, the repurchase contracts that securities dealers use to finance their own and customers' trading positions showed a sharp and unusually sustained decline in volume. An implication of the deleveraging was that other traders that might have speculated against the fall in asset prices and thereby stabilized the markets were no longer a support in the markets.

As the ensuing market liquidity crisis unfolded during August and September 1998, growing risk aversion made ever larger numbers of investors seek out low-risk assets and retire to the sidelines, and credit

spreads widened sharply beyond what had already occurred following the Russian default. To avoid a disorderly default, and the potentially adverse consequences of the further selling pressures it might have incited, a consortium of LTCM's trading counterparties undertook a recapitalization of the hedge fund in what was essentially an informal bankruptcy procedure conducted by the creditors with the cooperation of the fund's management.

Even after the LTCM recapitalization, however, spreads in many markets continued to widen as participants showed an ongoing aversion to risk. Other hedge funds in particular saw dramatic changes in the willingness of major intermediaries to finance their activities—a development that prompted further selling and spread widening. By mid-October, reports had grown that the situation was hindering the ability of nonfinancial businesses to raise capital and that risk aversion was beginning to manifest itself in payment and settlement procedures. Only after the Federal Reserve surprised markets with an intermeeting rate cut did the markets gradually return to normal.

While analysts differ in their views on whether the disorderly collapse of LTCM would have been a systemically significant event, the episode nevertheless signals the need to think broadly about the potential sources of systemic risk. In particular, how has the growing emphasis on trading activities—which are increasingly conducted through hedge funds—affected the potential for systemic risk? Does this emphasis create mechanisms for propagation that did not exist previously? Can these mechanisms be fully captured by the classical models associated with bank runs or market gridlock, or do they introduce fundamentally new elements?

Several recent trends in the financial markets bear on these questions. One trend relates to the blurring of distinctions between types of financial firms. Commercial banks that have traditionally focused on making loans have increasingly removed loans from their balance sheets through securitization (pooling loans such as mortgages into securities sold to investors) or outright trading of loans and securities; at the same time, they have increased their investment banking securities underwriting and trading activities. Conversely, some of the largest investment banks and trading houses now lend directly to businesses and households.

One result of this broadening of activities has been an increased volume of trading in asset types that have in the past been regarded as illiquid. Traditionally, financial assets have been separated into liquid and illiquid assets: liquid assets (such as stocks and government bonds) are priced and traded regularly after issue on exchanges or in large interdealer markets, while illiquid assets (such as bank loans) are held by financial institutions, particularly commercial banks, over long periods

of time and are rarely traded or priced after origination. In recent years, however, the sharp distinction between liquid and illiquid assets has eroded, and liquidity, or tradability, has become a continuum. While some types of assets still trade very little after issuance, there is a trend toward trading asset types that have traditionally been regarded as illiquid—for example, bank loans, debt and equity of small firms, and debt of bankrupt or distressed firms.

Moreover, financial institutions now securitize many previously illiquid assets. Securitization involves pooling together collections of illiquid assets such as mortgages, auto loans, or credit card loans and creating a relatively standardized security that pays investors the cash flows from these assets. As a result of these changes, market participants today trade and price a much wider array of risky assets—at least when markets are functioning normally. During times of financial market distress, however, the liquidity of many assets can drop sharply, and differences in liquidity across asset types can widen dramatically.

Similarly, the tremendous growth in the use of financial derivatives reflects the increased tradability of financial risk. A substantial amount of current financial market activity involves the repackaging of claims on underlying assets and the redistribution of the underlying risks. This last activity has spawned enormous growth in the trading of derivatives, which are contingent claims in which payoffs are conditioned on the behavior of underlying variables such as interest rates or equity prices. The institutions at the core of the financial markets not only participate in these various activities, but also frequently serve as market-making intermediaries.

Derivatives offer a number of advantages in the trading and hedging of the price risks in underlying assets. First, because they are equivalent to a leveraged trading position, derivatives contracts can often be entered into with very little capital up front. Thus, they are an ideal hedging instrument because the underlying risk can be hedged without the cost of committing a substantial amount of capital. At the same time, however, the leveraged nature of derivatives contracts makes them risky trading instruments, and traders that use these instruments to speculate can lose large sums very quickly. Second, the ability to structure and specify the particular underlying risk that a derivatives contract is exposed to enables users to unbundle a collection of risks embodied in an asset and trade the components separately. This precision also makes derivatives an ideal hedging and trading tool, since a hedger can choose which risk to hedge and which to leave uncovered.

An important consequence of the widespread use of derivatives contracts is the parsing and dispersal of the risks embodied in underlying assets. Overall, this has provided a net benefit to the economy, because

risks that would have remained locked up and concentrated in underlying assets are now spread out and allocated to those more willing to bear them. This ability to transfer unbundled risks through derivatives contracts separately from the aggregates in underlying assets enables investors to better select which risks they are exposed to, providing two important benefits: lower risk premia in asset prices because investors are no longer locked into bearing unwanted risks, and the potential for a better allocation of risks to those more able to bear them.

Accompanying the growth of trading in less liquid assets and derivatives has been the general trend toward fair value accounting for more types of instruments and positions. Fair value, or mark-to-market, accounting imposes a discipline and transparency that can force institutions to take action to address emerging problems that might not occur under historical cost accounting. By contrast, historical cost accounting is more likely to allow serious problems to go undetected and unaddressed for longer periods of time.

A second significant trend, alluded to earlier, is the increasing role played by a broader range of market participants—not only hedge funds but also other forms of specialized vehicles such as private equity firms and collateralized debt obligation managers. These new agents for risk bearing have the potential to alter the dynamics of how the financial system as a whole manages risk. By allowing risk to be spread more widely, they have the potential to help insulate the financial system against external shocks. In the view of some analysts, however, a greater capacity for risk bearing may lead the system to become even more inclined to cyclical behavior. The extent to which these new entrants are stabilizing or destabilizing depends in part on whether the extent of aggregate leverage in the financial system is greater today than in the past, since more highly leveraged institutions are more susceptible to large shocks that erode capital. Another critical question relates to the linkages between these new entrants and the traditional financial intermediaries. For example, in a financial crisis, it may not be sufficient for banks to have transferred risks to hedge funds if the ultimate source of financing and liquidity for those hedge funds remains the banks themselves. Again, the overall impact will likely depend on whether the new arrangements increase or decrease the amount of total equity capital at stake (including both bank equity and hedge fund investors' equity) relative to the size of the risks being taken.

A third trend is the strong emphasis that leveraged institutions—not only the large banking and securities intermediaries but also the majority of hedge funds—put on quantitative models for the pricing and risk management of their activities. Risk management practices at such organizations owe a significant debt to the efforts over the past fifty years of

many academics and practitioners to apply statistical and mathematical techniques to the problem of analyzing movements and comovements in market prices and other relevant variables. Such analysis, leavened in most cases by market experience, is used to help assess a firm's ability to operate safely with different combinations of assets and leverage. Risk management strategies are also obviously critical in influencing how financial market participants will react to changes in market conditions. To the extent that there is commonality in risk management models and strategies, there is potential for a broad cross section of market participants to react similarly to changes in asset prices.

In valuing complex derivatives transactions, it is often necessary to interpolate or extrapolate the fair value of such instruments using mathematical models calibrated to the observed market values of other, simpler instruments. In some cases, these models are very difficult to test against an objective reality beyond the fact that other participants are using similar models. It is no accident that models are most commonly used to price relatively illiquid assets; thus, during periods of financial distress, actual prices are most likely to differ substantially from modeled prices. A related issue is the degree to which the positions and strategies of the diverse participants in various markets are correlated. To the extent that many participants are pursuing very similar strategies and will behave very similarly in response to market shocks, the diversification of the system as a whole may be less than it appears during more benign periods.

All of these trends—a substantial emphasis on trading, risk transfer, and derivatives; greater market involvement by hedge funds; and a heavy reliance on quantitative risk management models—were at work to some extent in the LTCM episode. While the classical models of bank runs and market gridlock were undoubtedly also relevant to LTCM, the episode highlights the need to expand these models to incorporate more fully the potential endogeneities and feedback effects generated by the trends discussed here.

September 11, 2001, and the Reliance on Critical Infrastructure

While the growth of hedge funds underscores how financial market activities have expanded beyond the major commercial and investment banks, the financial sector events following 9/11 emphasize the reliance of the financial sector on certain core elements of infrastructure and on a relatively small number of organizations. Two related aspects of the post-9/11 period merit discussion in this regard.

First, the terrorist attacks of that day did widespread damage to both

property and communications systems in Lower Manhattan.⁶ Because many of the largest commercial banks had operating facilities in this area (or had electronic communications routed through hubs in the area), they were unable to make payments as they normally would. Since most large banks normally both send and receive a large volume of Fedwire payments every day, relying heavily on incoming payments for the liquidity to make their own payments, the normal coordination of payments broke down and liquidity shortages developed at many banks.

From a systemic perspective, the Federal Reserve attaches extreme importance to keeping the Fedwire system open; otherwise, this central aspect of the financial system nationwide would not be able to function at all. Indeed, in the wake of 9/11, the Federal Reserve extended the operating hours of the system to help provide more time for banks to execute their transactions. In addition, the Federal Reserve made more liquidity available, both to individual banks through its discount window operations and to the system generally through open market operations. These measures, along with the willingness of the Federal Reserve to permit sizable intraday overdrafts, helped restore normal functioning to the payments system.

A second set of issues arose in the market for U.S. government securities. The clearance and settlement of these securities (as well as a number of other fixed-income securities) are concentrated in two commercial banks. These same two banks provide the primary mechanism through which the securities portfolios of the major securities firms are financed on a daily basis (the “tri-party repo market”). The financing itself is provided by money market mutual funds and pension funds primarily, but the two banks provide the systems, services, and intraday credit on which this nearly \$1.5 trillion market critically depends.

Following the 9/11 attacks in New York City, one of these two clearing banks suffered very significant operational disruptions, reflecting the proximity of its primary as well as back-up operating sites to downtown Manhattan. Although these disruptions did not completely obstruct the processing of securities transactions, the processing slowed considerably. Further, the destruction of brokers’ offices obstructed the clearing and reconciliation of trades, and trade records were not fully reconciled for several weeks. In the meantime, the uncertainty arising from the disruptions contributed to a significant increase in the number of trades that failed to settle. This “fails” problem became so serious that the U.S. Treasury conducted an unprecedented reopening of the auction for the ten-year note in order to increase the supply of that security in the marketplace.⁷

⁶This discussion draws heavily on McAndrews and Potter (2002).

⁷See Fleming and Garbade (2002).

Although the systemic financial consequences of the events of 9/11 are probably best described as a “near miss,” they do demonstrate the global financial system’s vulnerability. Investigation of the possible outcomes of the attacks indicates that if one of the two clearing banks had not, in fact, been capable of operating for a sustained period of time, the task of replicating such functionality elsewhere would have taken considerable time, possibly as much as a year or more. In the meantime, the underlying securities markets that are supported by the financing activities that clear through these banks would be disrupted. In particular, the U.S. government securities market that forms the basis for the implementation of U.S. monetary policy and the financing of U.S. government activities and that is used as “riskless” collateral in countless financial transactions worldwide could be impaired.

While this particular vulnerability was highlighted by the events following the 9/11 attacks, it is almost certainly not the only critical “choke point” in the global financial system today. That is, the operational disruption of other relatively modest organizations or physical facilities could significantly damage the functioning of the overall financial system. Indeed, the last decade has seen increased concentration in the provision of critical infrastructure services such as payments, settlement, and custody activities. Not surprisingly, the potential systemic risk associated with threats to such critical infrastructure has since 9/11 spurred a significant amount of effort by both the public and the private sectors to increase the resiliency of that infrastructure.

Clearly, traditional financial models of systemic risk cannot readily capture the type of systemic risk that arises from the potential for critical points of failure to lead to broader disruptions in the system. For one, the proximate cause is more operational than financial in nature. Nevertheless, the financial aspects cannot be ignored. As the example of the breakdown in payments flows illustrates, even if the initial disruption stems from physical damage or computer malfunction, the methods of propagation may still be financial. Thus, there is a strong need for models that are more capable of capturing the complex interactions between operational infrastructure and the financial flows that the infrastructure supports. Similar models would be helpful in understanding the consequences of a pandemic event that made it impossible for large numbers of urban employees to work from their offices. Is the existing financial system capable of a smooth transition to a temporarily reduced level of activity? Current models cannot readily even frame such a question.

Implications for Systemic Risk

Three interesting themes emerge from the events discussed in this section. First, the number of relevant points of failure has increased with the

growing complexity of the financial system. Large financial institutions such as banks and major financial markets such as the U.S. equity market continue to be focal points in any assessment of systemic risk. But new sources of risk have arisen with the growth of risk transfer through securitization and derivatives as well as the increasing use of central counterparties and other specialist financial institutions that fill specific roles in the financial market infrastructure. One further implication is that when individual institutions have problems, the number of business relationships and elements of risk has expanded dramatically.

Second, as the volume of transactions—payments, derivatives, and secondary-market trading—has increased, the apparently strong economies of scale in risk and liquidity pooling have led to consolidation, typically into a subset of the larger financial institutions. The high velocity of transactions creates substantial efficiencies that are reflected in timing and pricing. However, sharp slowdowns in transaction volume, such as those occurring in the payments system after 9/11, can reverse these efficiencies and potentially impair the performance of the financial system when key parts of the system are under stress. Similarly, a key institution's loss of credit standing can diminish the flow of business substantially and increase the cost of managing its derivatives or payments books.

Third, in the information-rich global environment that has emerged over the last few decades, the potential for contagion has changed. That potential continues to include direct linkages among large institutions through common credit and market exposures or exposures to one another, although many policy changes and enhancements to private risk management have sought to reduce that potential. Now, however, the potential for contagion has expanded to include associations between risk dimensions created through common investors, similarities in risk profiles and risk appetites, and common exposures to macro-level risk factors such as geopolitical risk. In periods of distress, such as the Asian currency crises, the Russian debt default, or the LTCM collapse, such associations may lead to the propagation of market disturbances in hard-to-predict and probabilistic ways, and therefore make crises more difficult to anticipate and manage.

Questions for Discussion

This background paper covers many different subjects at a relatively high level. Some key questions that conference participants might pursue are as follows:

- What types of models of systemic failure or collapse have proved useful in other disciplines? How applicable are these models to the kinds of issues discussed above?

- Which aspects of the financial system seem most important and/or challenging to capture in considering the potential for systemic risk in the financial sector?
- What potential avenues for future cross-disciplinary collaboration on systemic risk issues seem most promising?

References

- Counterparty Risk Management Policy Group II. 2005. *Toward Greater Financial Stability: A Private Sector Perspective*. Available at <<http://www.crmpolicygroup.org>>.
- Diamond, D. W., and P. H. Dybvig. 1983. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy* 91, no. 2 (June): 401-19.
- Fleming, M. J., and K. D. Garbade. 2002. "When the Back Office Moved to the Front Burner: Settlement Fails in the Treasury Market After 9/11." Federal Reserve Bank of New York *Economic Policy Review* 8, no. 2 (November): 35-57.
- Group of Ten. 2001. *Consolidation in the Financial Sector*. Available at <<http://www.bis.org/publ/gten05.html>>.
- Kindleberger, C. P. 1978. *Manias, Panics, and Crashes: A History of Financial Crises*. New York: John Wiley & Sons.
- McAndrews, J. J., and S. M. Potter. 2002. "Liquidity Effects of the Events of September 11, 2001." Federal Reserve Bank of New York *Economic Policy Review* 8, no. 2 (November): 59-79.
- Minsky, H. P. 1977. "A Theory of Systemic Fragility." In E. I. Altman and A. W. Sametz, eds., *Financial Crises: Institutions and Markets in a Fragile Environment*. New York: John Wiley & Sons.
- _____. 1982. "The Financial-Instability Hypothesis: Capitalist Processes and the Behavior of the Economy." In C. P. Kindleberger and J.-P. Laffargue, eds., *Financial Crises: Theory, History, and Policy*, 13-39. Cambridge: Cambridge University Press.

Appendix C

About the Rapporteurs

John Kambhu is a vice president in the Financial Intermediation Function of the Research and Statistics Group of the Federal Reserve Bank of New York, where he has worked since 1988. He focuses on issues relating to market liquidity, risk management, financial derivatives, and public disclosure of market and credit risks. Prior to joining the Bank, John was an assistant professor of economics at Columbia University. He received a Ph.D. in economics from New York University in 1981.

Scott Weidman is the director of the National Research Council's Board on Mathematical Sciences and Their Applications. He joined the NRC in 1989 with the Board on Mathematical Sciences and moved to the Board on Chemical Sciences and Technology in 1992. In 1996, Scott established a new board to conduct annual peer reviews of the Army Research Laboratory, which conducts a broad array of science, engineering, and human factors research and analysis, and he later directed a similar board that reviews the National Institute of Standards and Technology. Scott has been with the BMSA full-time since June 2004. During his NRC career, he has staffed studies on a wide variety of topics related to mathematical, chemical, and materials sciences, laboratory assessment, risk analysis, computational science, and science and technology policy. His current focus is on building up the NRC's capabilities and portfolio related to all areas of mathematical analysis and computational science. Scott holds bachelor's degrees in mathematics and materials science from Northwest-

ern University and an M.S. and a Ph.D. in applied mathematics from the University of Virginia.

Neel Krishnan received a B.A. in economics and anthropology from Columbia University in 2005. He is currently a research associate in the Capital Markets Function of the Research and Statistics Group of the Federal Reserve Bank of New York.