Quantifying Systemic Risk: Introduction

Joseph G. Haubrich* Andrew W. Lo

This Draft: 19 April 2011

1 Introduction

In the wake of the Financial Crisis of 2007–2009, many proposals have been put forward for its causes and the appropriate remedies. In response to an impatient and frustrated public, and several months before the Financial Crisis Inquiry Commission completed its analysis, Congress passed the 2,319-page landmark Dodd-Frank Wall Street Reform and Consumer Protection Act of 2010, setting the stage for seismic shifts in the regulatory landscape of the financial industry. Clearly, change is afoot, but are we ready?

In the context of such sweeping regulatory reform, one of the most urgent priorities is establishing the means to measure and monitor systemic risk on an ongoing basis. Even the most cautious policymaker would agree that attempting to eliminate all systemic risk is neither feasible nor desirable—risk is a necessary consequence of real economic growth. Moreover, individual financial institutions do not have the means or the motivation to address systemic risk themselves. Because risk is closely tied to expected returns in this industry, as both theory and practice suggest, in competing for market share and revenues financial entities will typically take on as much risk as its shareholders will allow, without considering the consequences for the financial system as a whole. In much the same way that manufacturing companies did not consider their impact on the environment prior to pollution regulation, we cannot fault financial institutions for ignoring the systemic implications of their risk-taking in the absence of comprehensive risk regulation. Unless we are able to measure systemic risk objectively, quantitatively, and regularly, it is impossible to determine the appropriate trade-off between such risk and its rewards and, from a policy perspective and social welfare objective, how best to contain it.

This is the current challenge that faces policymakers and regulators—even after the passage of the Dodd-Frank bill—and the focus of this NBER conference volume on Quantifying

*The views expressed are those of the authors and not necessariy those of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System

Systemic Risk. The chapters are based on papers presented at an NBER conference held in Cambridge, Massachusetts on November 6, 2009, and jointly sponsored by the Federal Reserve Bank of Cleveland and the NBER. We were fortunate to have a remarkable and diverse array of participants drawn from academia, industry, and government agencies, and the breadth and depth of ideas contained in this volume is a clear testament to their unique expertise. Each paper presented at the conference was assigned two discussants, one from academia and the other from either industry or government, and we have included summaries of the discussants' remarks as well.

In "Liquidity Risk, Cash Flow Constraints, and Systemic Feedbacks," Sujit Kapadia, Matthias Drehmann, John Elliott, and Gabriel Sterne introduce a theme that reappeared in several other conference papers: while outside shocks may touch off a financial crisis, the reaction of market participants determines the course of the disaster. In the model they develop, solvency concerns at one bank lead to liquidity problems, as funding becomes more difficult. This forces the bank to take defensive actions, hoarding liquidity and reducing lending to other banks. In certain cases, the problem snowballs (or becomes contagious) and a crisis looms. As other banks finds it harder to obtain liquidity, the problem can become systemic. The process illustrates, as do several other papers in this volume, how the fallacy of composition can hold in the financial markets: individual defenses against risk lead to greater risk overall.

The paper emphasized the cash flow constraint: banks must have cash inflows that cover their cash outflows. Kapadia *et al.* go further, however, and quantitatively evaluate the systemic effects of this funding liquidity risk. To do so, the work builds on a broader project (RAMSI) underway at the Bank of England, using detailed balance sheet information from UK banks encompassing macrocredit risk, interest and non-interest income risk, network interactions, and feedback effects. Funding liquidity risk is introduced by allowing for rating downgrades and incorporating a simple framework in which concerns over solvency, funding profile and confidence trigger the outright closure of funding markets to particular institutions. The detailed look at the network of counterparty transactions demonstrates how defensive actions on the part of some banks can adversely affect others. The model can accomodate both aggregate distributions and scenario analysis: large losses at some banks can be exacerbated by liability-side feedbacks, leading to system-wide instability.

In "A Tax on Systemic Risk," Viral V. Acharya, Lasse H. Pedersen, Thomas Philippon and Matthew Richardson take the important step of tying a specific regulation to a quantitative measure of systemic risk. They propose taxing each firm based on its contribution to systemic risk. Specificaly, the tax would depend on a firm's expected loss conditional on the occurrence of a systemic crisis. Note the dual trigger: both the individual firm and the financial sector must become undercapitalized. The tax is then just the fair-value premium of insurance against this event. Although they derive the pricing for such insurance, their preferred solution involves letting the market set the price. Individual firms would be required to purchase contingent capital insurance, that is, insurance against the losses they incur during systemic crises. The cost of this insurance determines the firms systemic risk tax. In a true systemic crisis, however, it is not clear that private firms would be in a position to provide the insurance. Rather, joint private-public provision of such insurance (say 5%-95%) lets the government piggyback on the market's superior price-setting ability. The total insurance premium, or tax, should induce the financial sector to internalize the systemic risk. A further element of the design addresses the moral hazard problem: If the firm has insurance, why should it avoid the risk? In this paper, the payoff goes not to the firm, but to the regulator. This adds a measure of pre-commitment to the government rescue policy.

Applying this measure of systemic risk to the recent crisis provides some encouraging results. The paper calculates both the tax and the insurance premium for major financial firms prior to the crisis, and Bear Stearns, Lehman Brothers, Fannie Mae and Freddie Mac show up high on the list, although AIG is prominently missing. This suggests the intriguing possibility of an early warning system, but it is an entirely different question whether the tax would have been enough to reduce systemic risk in these firms—or the market—to a manageable level. A further consideration is how this contingent capital proposal compares with other related proposals such as forced debt-for-equity conversions.

In "Systemic Risks and the Macroeconomy" Gianni De Nicolò and Marcella Lucchetta make a distinction between real and financial risk, and present a modeling framework that jointly forecasts both sorts of systemic risk. They emphasize that lost output and unemployment constitute the true costs of financial crises. Thus, their systemic version of Value at Risk (VaR) has two components: the 5 percent tail of a systemic financial indicator(market adjusted return for the financial sector), and GDP at Risk, the 5 percent tail on real GDP growth. This framework is implemented using large sets of quarterly time series of indicators of financial and real activity for the G-7 economies for the 1980Q1-2009Q3 period. They first use a dynamic factor model to check forecasting power, and then impose sign restrictions from a simple macromodel to identify the shocks. For example, an aggregate supply shock should increase output but decrease inflation.

They obtain two main results. First, the model can, with some accuracy, forecast large declines in real activity, showing promise as an early warning system or a risk monitoring tool. Second, in all countries aggregate demand shocks drive the real cycle, and bank credit demand shocks drive the bank lending cycle. These results challenge the common wisdom that constraints in the aggregate supply of credit have been a key driver of the sharp downturn in real activity experienced by the G-7 economies in 2008Q4-2009Q1.

In "Endogenous and Systemic Risk," Jon Danielsson, Hyun Song Shin, and Jean-Pierre Zigrand explore the feedback between market volatility and traders' perception of risk. Trading activity sets and moves prices, but traders gauge risk from the resulting price volatility. Equilibrium requires a consistency between the perceived and the actual risk. In a setting where traders operate under Value-at-Risk constraints (although the logic carries over to risk-based capital requirements and more), volatility can become stochastic, even as fundamental risk remains constant. Trader reactions amplify fluctuations, creating a spiral of even greater response. If the purpose of financial regulation is to shield the financial system from collapse, then basing regulation on individually optimal risk management may not be enough: in this case, the prudent behavior of individuals increases the aggregate risk.

Roughly speaking, a market shock, say a decrease in prices or an increase in volatility, now makes the asset look riskier according to risk management rules, be they Value at Risk or some other method. This forces the firm to reduce risk by selling the asset. But of course other firms, also noting the increase in risk, do the same, leading to an even larger drop in price, starting a spiral to even more risk. A crisis can arise quickly, because the process is highly non-linear, with larger movements appearing suddenly. The critical threshold depends on the specifics of each market: risk management strategies, leverage, and capital plans. The paper applies this insight ot a variety of markets, from explaining the implied volatility skew for options, the procyclical impact of Basel II bank capital requirements, to the design of lenders of last resort and of centralized clearing for derivatives. Spelling out the precise mechanism, though a challenge, also is a vital first step in the design of more robust institutions and policies.

In "Hedge Fund Tail Risk" Tobias Adrian, Markus K. Brunnermeier, and Hoai-Luu Nguyen estimate the tail dependence between the major hedge fund styles, such as long/short equity and event driven funds. They use quantile regressions to document how the return of one strategy moves with the return on another. Quantiles can explicitly compare the dependencies between normal times (50 percentile) and stress periods (5 percentile). The tail sensitivities between hedge funds increase in times of crisis, some more than doubling.

The paper identifies seven factors that explain this tail dependence; these risk factors include the overall market excess return, a measure of volatility, and the slope of the yield curve. Because the seven factors are effectively tradeable in liquid markets, it is possible to hedge, or offload that risk, which significantly reduces tail dependence. The paper thus provides a built-in solution to the problem it uncovers. Implementing this solution may not be easy, however. In fact, the paper demonstrates that individual hedge fund managers have no incentive to offload the tail risk, as funds that increase their exposure to the factors also increase their returns and their assets under management. Offloading the risk then lowers both sides of managers' expected compensation (the famous "2 and 20" rule).

In "The Quantification of Systemic Risk and Stability: New Methods and Measures," Romney Duffey approaches the problem of predicting financial systemic risk from the standpoint of a general theory of technical systems with human involvement. The discussion about the financial crisis often borrows terminology from meteorology or other physical sciences: we hear about "hundred-year floods " or "perfect storms." The analogy can be misleading, not only because it neglects the rich analysis of risk quantification, minimization and management within the engineering profession, but also because it ignores the human element. Among other problems, the meteorological terminology puts an undue emphasis on calendar time. In human systems, failure instead depends on experience time. Airline crashes and automobile deaths, for example, depend on miles traveled. Just what best captures the experience time for financial markets is unclear, but quite likely involves something like volume or the dollar value of transactions, and those have increased. Between 1980 and 2009, monthly trading volume on the New York Stock Exchange increased by a factor of 100, from 1 billion shares to 100 billion.

Accumulating experience has contrasting effects on the probability of major failures, sometimes known as the learning paradox. Learning reduces risk, but learning requires taking the risk and experiencing the very events you seek to avoid. As learning brings risk down to acceptable levels, there is more time for the unknown and rare events to manifest themselves. Indeed, risk often looks low before a major crisis, as the obvious problems have gotten resolved, but not enough (experience) time has passed for the new, rare problems to occur. This interaction often makes it difficult for simple statistical models to capture the distribution of losses.

A related theme was emphasized in the Keynote address by Henry Hu on "Systemic Risk and Financial Innovation: The Importance of the Underlying Process." Hu argues that a proper understanding of systemic risk requires understanding financial innovation as a process, focussing less on particular products and more on how products are invented, introduced, and diffused through the marketplace. Any fixed classification or regulatory scheme quickly becomes obsolete, both because firms find ways around regulation and because the marketplace continually evolves. Such rapid evolution makes mistakes inevitable, because learning takes time, and while that occurs the heuristic approaches and cognitive biases of market participants have room to operate. This human element again emphasizes the dangers of taking physical models of the market too literally: a market crash, the net result of many voluntary trades, is not a meteor strike, and indeed financial markets have an element of a self-fulfilling prophesy: if everyone trades according to a price rule, that rule really works, even if it is flawed.

As an example of this evolution, Hu emphasized financial decoupling: the ability of firms to separate the economic and legal benefits, rights and obligations that standard debt and equity bundle together. For example, a fund may buy stock and obtain voting rights in a corporation, but hedge the financial exposure with offsetting credit default swaps. Conversely, selling a CDS can allow economic exposure without voting rights, and more complicated examples abound. Reckoning with such possibilities clearly requires more than even the most sophisticated economic analysis, needing a more unified, interdisciplinary approach drawing on both law and economics, each situated in the proper dynamic context.

Some of the most important themes of the day arose not from the paper presentations but from the discussions, both from the asiigned discussants and comments from the floor. There were philosophical discussions about what it meant to understand: in biology, the question as to why polar bears are white has an answer from an adaptive/evolutionary standpoint (they blend in with the snow) or from a developmental standpoint (which genes create white fur). Others considered the differing roles of models used for description or for prediction. Regulators from different jurisdictions considered the merits of systems that discouraged risk as opposed to early warning systems, and of deeply understanding one market versus testing across many markets. Others argued over the relative merits of different risk measures: value at risk, simple leverage, even instinctive feelings of discomfort among traders.

However, there was widespread agreement that any serious effort at managing systemic risk must begin with measurement—one cannot manage what one does not measure. In the very best tradition of the NBER, these discussions, and the analytical foundations that the following chapters have begun developing, represent an important first step in our attempt to better understand the nature of financial crisis and systemic risk.